Strategic Conformance
Exploring Acceptance of Individual-Sensitive Automation for Air Traffic Control

Westin, Carl

DOI
10.4233/uuid:49c6fe9d-2d29-420a-91a2-a97e2049e15e

Publication date
2017

Document Version
Final published version

Citation (APA)
Westin, C. (2017). *Strategic Conformance: Exploring Acceptance of Individual-Sensitive Automation for Air Traffic Control* . [Dissertation (TU Delft), Delft University of Technology].
https://doi.org/10.4233/uuid:49c6fe9d-2d29-420a-91a2-a97e2049e15e

Strategic Conformance

Exploring Acceptance of Individual-Sensitive Automation for Air Traffic Control

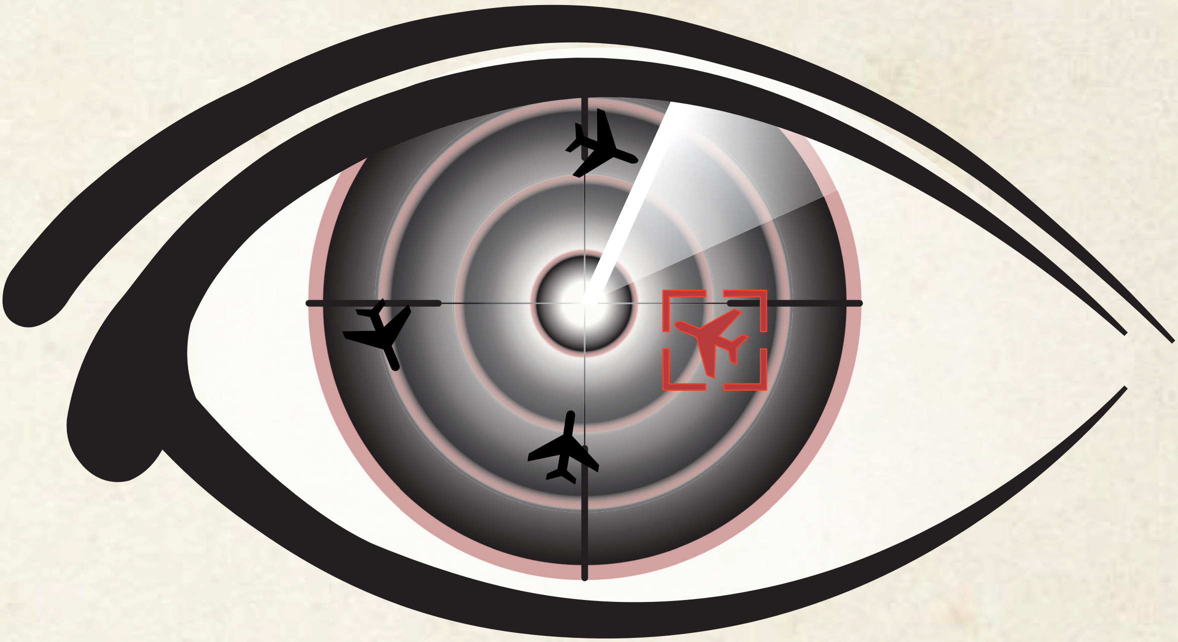

Carl A. L. Westin

Strategic Conformance:
Exploring Acceptance of Individual-Sensitive Automation for Air Traffic Control

Manuscript sent to committee:
Tuesday July 4, 2017

Promotion date:
Friday November 24, 2017

Carl Albert Lennart Westin
Delft University of Technology
Kluyverweg 1
2629 HS Delft
E-mail: c.a.l.westin@tudelft.nl
Phone: +46(0)733693686

Strategic Conformance:
Exploring Acceptance of Individual-Sensitive Automation for Air Traffic Control

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus Prof.ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op
vrijdag 24 november 2017 om 15.00 uur

door

Carl Albert Lennart WESTIN

Master of Science in Applied Ergonomics (University of Nottingham)

geboren te Huddinge, Zweden

Dit proefschrift is goedgekeurd door:
Prof. dr.ir. M. Mulder (promotor)
Dr.ir. C. Borst (copromotor)

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr.ir. M. Mulder,	Technische Universiteit Delft, promotor
Dr.ir. C. Borst,	Technische Universiteit Delft, copromotor
Onafhankelijke leden:	
Prof. dr. J. D. Lee,	University of Wisconsin-Madison
Prof. dr. J. M. C. Schraagen,	Universiteit Twente
Prof. dr. J. Dankelman,	Technische Universiteit Delft
Prof. dr.ir. J. M. Hoekstra,	Technische Universiteit Delft
Dr. B. H. Hilburn,	Center for Human Performance Research, overig lid

The work described in this thesis has been carried out at the Control and Simulation section at Delft University of Technology. Part of this work was supported by EUROCONTROL acting on behalf of the SESAR Joint Undertaking (the SJU) and the European Union as part of Work Package E (project MUFASA: Multidimensional Framework for Advanced SESAR Automation) in the SESAR Program.

EUROCONTROL SESAR

Keywords: acceptance, air traffic control, automation, decision aid, decision-making, personalization, strategic conformance, transparency

Cover illustration made by Carl A. L. Westin & Clark Borst

ISBN 978-94-6299-659-5

An electronic version of this dissertation is available at:
`http://repository.tudelft.nl/`

Summary

Strategic Conformance: Exploring Acceptance of Individual-Sensitive Automation for Air Traffic Control

Carl Albert Lennart Westin

LIKE many complex and time-critical domains, air traffic control (ATC) is facing a fundamental modernization that builds on the use of more advanced automation (represented by SESAR in Europe and NextGen in the United States). The current function allocation-based relationship between controller and machine is envisioned to evolve to a more fluid, continuous and mutually coordinated team relationship. Consequently, the controller is expected to assume a supervisory and monitoring role, while relinquishing much of the tactical "hands-on" tasks to automation. ATC automation, in turn, is expected to grow in intelligence and its cognitive abilities to become more of a team member providing decision support and acting more autonomously. In association to these changes, one of the most pressing human factors challenges is how we can design automation that is embraced, accepted and trusted by the controller.

With automated systems becoming increasingly cognitively mature, they are likely to, in many ways, impersonate a persona with its own behavior and personality. Operators are likely to perceive these systems more as a humanized character and less as a technological tool. This trend is evident from current intelligent personal assistants, such as Apple's *Siri*, Microsoft's *Cortana*, and the Google *Assistant*. In fiction we have seen future visions, such as the witty and sarcastic *TARS* robot in the movie *Interstellar*, the curious and seductive *Samantha* operating system in the movie *Her*, and the calm and reassuring *HAL 9000* in the movie *2001: A*

Space Odyssey who, unfortunately also suffers from paranoia.

Operators may struggle with understanding the system, not only because its reasoning is hidden, but also because its reasoning is different. Consequently, operators may end up distrusting and rejecting the system and its advice. This problem-solving mismatch can partly explain the acceptance issue observed in ATC. To address the acceptance issue and guide mid- and far-term ATC automation design, the visionary MUFASA (Multidimensional Framework for Advanced SESAR Automation) project set out to develop a framework for future levels of automation (LOA). The project hypothesized that conflict detection and resolution (CD&R) decision support *conformal* to a controller's preferred conflict-solving style would benefit the acceptance of that support and facilitate improved human-automation collaboration. *Strategic conformance* was introduced as a compatibility concept specifically capturing the degree to which a decision aid's apparent problem-solving style matches the operator's.

This thesis evolved from, and set out to expand, the successful MUFASA project to consider and empirically explore an individually-centered approach towards automation design. In general, this thesis focuses on decision aids that provide explicit personalized solutions for control tasks in highly dynamic time- and safety-critical domains. The ambitious goal was to obtain a fundamental understanding of how controllers' acceptance of ATC conflict resolution advisories were affected, depending on how well the decision aid's conflict-solving strategy *matches* that of the individual controller.

To study strategic conformance empirically, a novel approach based on repeating controllers' own solutions was developed. As such, automated advisories were based on recordings of controllers' own solutions to the same conflict. No conflict solution algorithm was used. To determine a controller's unique conflict-solving style, the controller unknowingly encountered the same scenario and conflict four times in a *prequel simulation*. Solutions were then analyzed and scripted as conflict resolution advisories. In the subsequent *experiment simulation*, strategic conformance was varied by providing the same controller with either her/his own solution to the same conflict (conformal), or a colleague's contrasting solution to the same conflict (nonconformal).

This thesis sought to investigate strategic conformance effects in the context of decision selection and implementation. For this purpose, controllers were supported by the Solution Space Diagram (SSD) prototype, which provides the high-level information acquisition and integration required for facilitating higher levels of decision support automation. While the SSD did not advice specific conflict solutions, it facilitated the implementation of advisories for the purpose of manipulating strategic conformance. The SSD is an ecological information support tool that integrates several critical control parameters of the CD&R problem. It was used in all simulations as part of this thesis, and appeared when an aircraft was selected. As such, simulations represented a futuristic environment, different from current ATC operations, with more traffic present, datalink communication between the controller and pilots, and the SSD supporting controllers in CD&R. Furthermore, simulations and conflict solving were restricted to the horizontal plane.

This thesis reports on three human-in-the-loop studies. In addition, simulation data from these were analyzed post hoc in a fourth study. The *first empirical study* culminated in a large-scale real-time simulation with sixteen experienced controllers. The study varied strategic conformance (i.e., conformal or nonconformal resolution advisories) in addition to the system's authority (LOA: management by consent or management by exception) and task complexity (high or low). Controllers accepted conformal advisories (i.e., advisories based on their own unique conflict solving style) more often, gave them higher agreement ratings, and responded faster, than with nonconformal advisories based on a colleague's contrasting but still workable and safe conflict solution. In 25% of cases, however, controllers disagreed with their own conformal advisories.

The other two human-in-the-loop studies, and the post hoc analysis study, were conducted to further explore plausible causes for the observed disagreement, together with other outstanding questions, derived in the first study. The two human-in-the-loop studies replicated the experimental approach used in the first study, with minor refinements.

The *Source bias study* investigated differences in controllers' acceptance of, and trust in, a conflict resolution aid, based on the presumed source of that advice. Five experienced controllers participated in a real-time simulation that varied strategic conformance together with the advisory source, presented to originate from either a human or automated source. While questionnaire responses indicated a slight preference for the human adviser, simulation results did not.

The *Automation transparency study* investigated effects of interface transparency and strategic conformance on controllers' acceptance and understanding of advisories. Nine controller trainees participated in a real-time simulation. Two levels were used, with the *heading band* SSD representing low transparency, and *triangle* SSD representing high transparency in that it provided more meaningful information on the relationship between conflicting aircraft. Results showed that the more transparent triangle SSD was better understood. Although no interaction effects between conformance and transparency were found, conformal advisories were accepted slightly more often than nonconformal advisories, supporting results from the first study. Moreover, when using the triangle SSD, conflicts were more often solved by speed and combinations of heading and speed. This indicates that controllers' solutions depended on how the conflict was represented in the interface. Since conformal (and nonconformal) advisories were based on solutions when using the heading band SSD, these advisories may not have been representative as conformal when using the triangle SSD.

The fourth and final *Consistency study* analyzed controllers' manual conflict solution data (no explicit resolution advisories) collected in the above three empirical studies. The objective was to determine whether the problem-solving mismatch between controllers and CD&R automation could be explained by controllers solving conflicts inconsistently. The study investigated the degree to which controllers consistently had solved the repeated conflict (four repetitions) over time (intra-rater variability), and to what extent they agreed on solutions (inter-rater variability). Based on a review of ATC conflict resolution strategies, a solution classification framework was developed against which controllers' solutions objectively could be qualitatively coded and analyzed. Results revealed that controllers were consistent, but disagreed on how to solve conflicts. However, consistency was limited to higher-level decision stages, such as whether to vector an aircraft in front or behind an other, or interacting with both or only one aircraft. Controllers were inconsistent in relation to more detailed solution parameters, such as the direction of a solution (e.g., vector left or right) and the exact directional deviation value (e.g., right vector of 035 degrees). Consistency and agreement was not higher for biased conflicts that favored a certain type of solution. A difference, however, was noted in regards to overall solution strategy. With biased conflicts, the majority of controllers agreed on a shared solution geometry, while with the unbiased conflict, the majority solved the conflict according to the control problem classification. Experienced controllers were slightly more consistent than trainees in terms of the control problem classification.

Taken together, this thesis has contributed to the knowledge of what drives controllers' acceptance of resolution advisories in particular, and human-automation collaboration and automation acceptance in general. Empirical results showed that conformal ATC automation, solving conflicts like the controller, can benefit acceptance and agreement of that system's advisories, as well as reducing response time. These benefits were observed across varying expertise levels, particularly in relation to expert operators. Strategic conformance may be most beneficial during the introduction of new automated decision aids, as a means for gaining acceptance.

The development of conformal automation, and other personalized decision support, requires that the operator is somewhat consistent in her/his problem solving. However, designing for conformal automation, or other personalized systems, requires an ethical consideration since such systems have the power to influence acceptance and trust independent from the system's actual performance and reliability. While technology advancements have made it possible to increasingly tailor automation to the individual's preferences, needs, and abilities, there are several technical challenges to be overcome before truly conformal automation can be developed, most importantly how to extract individual's unique preferences and problem-solving style. Research is needed to establish the consistent and critical control parameters that would characterize a person's problem-solving style.

Automation designers need to carefully consider the goals and objectives in a specific domain for which a conformal system is considered. Many work domains, such as the flight deck, may be more suitable for facilitating homogeneity that restricts individual differences in interaction and problem-solving. Complementary to conformal automation providing explicit advisories, interface design should support the variability in problem-solving styles. Ecological interface design, for example, can be used to facilitate such personalized problem-solving by visualizing the "objective truth." That is, by showing the constraints affecting a situation, the operator is allowed to solve problems in her or his preferred way.

Returning to the artificial intelligence (AI) driven personal assistants. Although still in their infancy in reality (e.g., Siri, Cortana, and Assistant), as they mature they are expected to significantly influence human-automation interaction to become a more fluid cooperation similar to the futuristic glimpses of automated characters observed in fiction (e.g., TARS, Samantha, and HAL 9000). While the acceptance of these systems depends on many aspects, the compatibility with the human, not the least their conformance, may have a significant effect on how willing people will be to interact, accept, and trust in it and its advice. While the extent to which acceptance and trust will depend on the intelligent assistant's strategic conformance remains to be studied, this thesis indicates benefits thereof, at least in the context of ATC conflict solving. The temptation to personalize automation simply because it can be, however, should be avoided. Particularly in safety-critical domains, both benefits and drawbacks of such capable automation must be considered and evaluated before implementation.

Contents

1

Introduction

This chapter provides an introduction to the concept of strategic conformance addressed in this thesis. The chapter describes the research problem addressed, the proposed solution (strategic conformance), the research approach and its challenges, the assumptions and scope. Furthermore, the thesis outline will be clarified by means of short chapter descriptions that denote how each individual chapter is linked to the research as a whole.

1-1 Background

SINCE the dawn of air traffic control (ATC) almost a hundred years ago, air traffic controllers (henceforth controllers) have separated and prevented collisions between aircraft. While initially assuring safe separation during takeoff and landing (i.e., tower control), the responsibilities of controllers have grown to include all segments of flight between airports (e.g., terminal and en-route control). Through technological innovations such as radar and computers, the means for achieving safe separation between aircraft have changed considerably (Figure 1-1). Over the years, ATC has become increasingly dependent on technology, while the controller has maintained central responsibility.

The increasingly crowded airspace is, however, affecting the nature of work and difficulty in separating traffic. In order to ascertain the safety of air travel, authorities aim to balance controller's workload by regulating the flow of traffic. In en-route airspace, such capacity restrictions are primarily mandated by estimates of

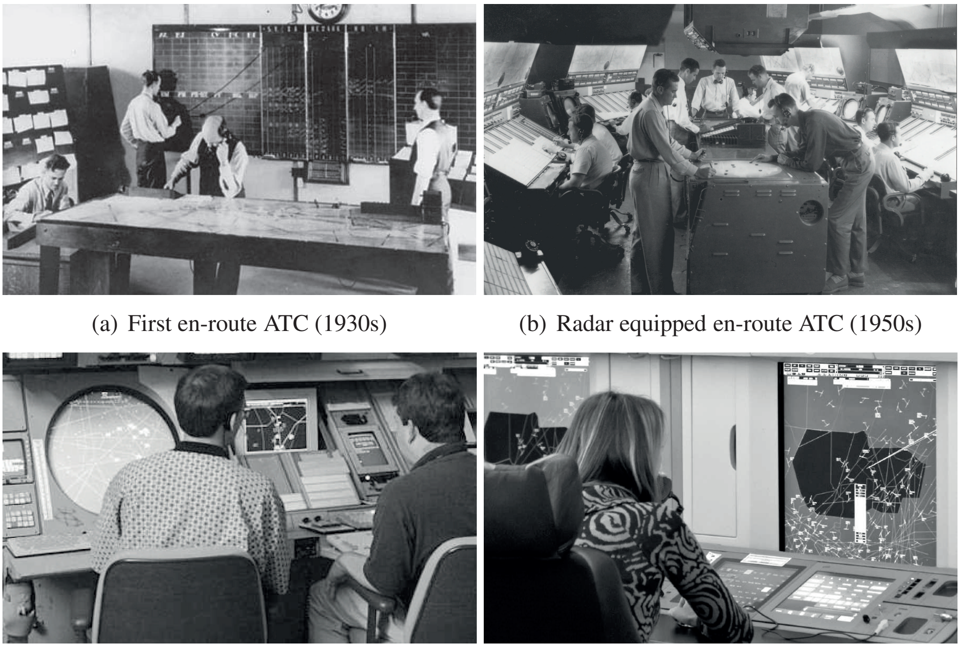

(a) First en-route ATC (1930s)

(b) Radar equipped en-route ATC (1950s)

(c) Computerized en-route ATC (1980s)

(d) State-of-the-art en route ATC

FIGURE 1-1: The evolution of en route ATC, spanning from the use of maps, phones, and a blackboard in the late 1930s (a), the introduction of radar in the 1950s (b), and computerized systems in the 1960-1970s, allowing for continuous monitoring of aircraft in real-time (c), up to current more automized state of the art (d).

the mental workload associated with controlling and separating traffic.[1,2] The introduction of more advanced automation is considered necessary to overcome current traffic delays and achieve future ATC capacity goals. For example, in SESAR, the European air traffic management (ATM) community is working towards achieving a three-fold increase in airspace capacity between 2005 and 2020, without adding more controllers.[3] Similar goals have been established by other initiatives worldwide, such as NextGen in the United States.[4]

These targets require the use of more sophisticated automation that supports and eases the cognitive burden of the controller in problem solving and decision-making tasks.[5] As a critical part of this, automation is foreseen to assume a greater tactical role in the short- and medium-term timeframe of planning and executive ATC. Moreover, automation is likely to act more as an adviser providing solutions to the controller in regards to, for instance, airspace reconfiguration and planning traffic. Such automated decision aids are expected to be especially beneficial in the strategic and tactical phase of separating traffic by alleviating controller workload.[6,7] Currently, controllers carry out this key task of conflict detection and resolution (CD&R) largely manually and with limited decision support.

1-2 Problem definition

Of pivotal importance for the future ATC system is that the automation developed is accepted and used by the controllers it intends to benefit and support. Unfortunately, in the past decades several CD&R decision aids have been rejected or used in ways not intended by the designer.[8–10] Rejection of automation has often been attributed to large uncertainties in CD&R algorithms, leading to inaccuracies in conflict detection and unreliable resolution advisories.[10,11] Automation disuse has also been attributed to inappropriate decision thresholds that lead to either an extensive number of false warnings or failure to detect conflicts.[12,13] The acceptance, or reliance, on automation is believed to be affected by various factors influencing operator attitudes toward automation, such as trust, perceived risk, perceived reliability, level of automation, age, and job satisfaction.[13–16]

Possibly, the observed acceptance issues can be attributed to differences in conflict resolution strategies between the automated decision aid and the controller. Since there are often several alternative solutions to a conflict, the automation and human do not necessarily agree on which one to apply. For example, in recent human-in-the-loop simulations exploring decision aiding automation for controllers, the automation was perceived as occasionally ‘fighting’ against the controller on how to solve conflicts.[17] In a study investigating the adoption process of a conflict detection system called URET (User Request Evaluation Tool), Bolic concluded that the system was rejected in conflict situations, as it did not reflect controllers' current way of managing conflicts.[9]

Researchers have explored alternative approaches to CD&R decision aids that acknowledge the psychological and behavioral variables that might influence the controller's conflict resolution strategies and solutions.[18] It is, therefore, worth investigating how similarities and differences between controller and automation problem-solving activities may affect the acceptance of decision aids. As such, the key problem this thesis addresses is therefore:

Problem definition

How to overcome controller acceptance issues of automated decision aids for conflict detection and resolution?

1-3 Decision-making mismatches

Automated CD&R decision aids have generally been designed with a limited consideration for controllers' individual decision-making processes or solution preferences. Conflict-solving algorithms typically approach the environment in a dichotomous fashion, providing single, fixed, mathematically optimal solutions according to causal deterministic laws.[19] From a technology-centered perspective, this is not an issue: a system-generated optimal solution should be accepted and the controller should manage only by exception. However, one potential human performance problem is that an optimized (e.g., single vector) solution can hide the automation's "reasoning" and paradoxically present a solution that the controller cannot easily evaluate. As automation becomes more advanced and assumes more of the "thinking," the controller's interpretation and understanding of what the system is doing and why may become more critical.

In contrast, psychology researchers argue that humans tend to approach problem-solving more heuristically (i.e., intuitively and by rule of thumb) and quickly settle for solutions that *satisfice* rather than optimize.[20–22] Analogously, research has shown that controllers commonly rely on heuristics for CD&R[23–28] and settle for a "good enough" conflict solution that works.[24,29] Therefore heuristic approaches, as opposed to optimized algorithmic ones, have been advocated for human-centered CD&R decision aid design.[29–33] One example system is the Controller Resolution Assistant (CORA) tool intended as a CD&R decision aid for en-route controllers.[29,34,35] The algorithm is based on a template of controller heuristics in conflict resolution. This is achieved by constructing a library of controller strategies and identifying a set of "best" solutions (around four) that matches a majority of controllers (investigations suggest that 80% or more is reasonable).[35] When detecting a conflict, the CORA algorithm provides the controller with a list of alternative solutions for solving the conflict, including a “best-ranked” resolution based on a cost-value comparison.[34]

Although small-scale testing with experienced controllers indicated an overall favorable reception of CORA, there were three notable issues. First, signs of hesitation in accepting conflict resolutions suggest another issue, one of trust. Controllers frequently investigated the quality of conflict resolutions and occasionally expressed doubt as to whether a conflict actually would be solved with the suggested resolution. Second, controllers spent much time in searching through the list of alternative conflict resolutions in an attempt to find a solution that they preferred.[34] Third, the choice of resolution strategies was found to differ between Area Control Centers (ACCs) and the nationality of controllers (e.g., the preference for lateral resolutions by Lisbon controllers in contrast to vertical resolutions by Malmö controllers), suggesting that the algorithm has to be context sensitive.

While initiatives such as CORA have modeled and tailored decision aiding automation after controller conflict resolution strategies, they have not explicitly linked automation-generated solutions to the individual preferences of the controller. As such, they have not been able to ensure complete harmony between controller and automation decision-making strategies. In contrast, controllers are generally assumed to be homogeneous in how they prefer to solve conflicts. Consequently, automated CD&R decision aids have typically been designed to fit the group rather than the individual. On the basis of individual differences in personalities[36] and cognitive styles[37] that influence how problems are approached and solved, however, it can be expected that CD&R automation sensitive to individual differences in solution preferences would be beneficial to automation acceptance.

1-4 Research goal

It is reasonable to hypothesize that controllers would be more prone to accept automated advice if the automation appears to reason and solve conflicts in a way that is similar to the controllers themselves. This notion can be captured in the concept of strategic conformance, defined here as the degree to which automation’s solution and apparent underlying operations match those of the human.

While “conformance” addresses the solution match for a problem between human and automation, “strategic” refers to the apparent underlying strategies for reaching that solution. In CD&R, the “solution” can be considered to be the measures taken to solve a conflict (e.g., vectoring aircraft A behind aircraft B). The “apparent underlying operation” is the reasoning and rationale (of the automation) that seem (for the controller) to underlie a given solution (e.g., we vector aircraft A behind B *because* aircraft A is slower than aircraft B). Since the controller would not be able to observe or follow the reasoning process underlying a solution, these processes would be "apparent" as they only can be inferred from the observable automation behavior. As such, the goal of this thesis is as follows:

Thesis goal

To empirically investigate strategic conformance as a means for more personalized automation support, and develop a fundamental understanding of how a decision aid's strategic conformance affects the interaction with that aid and acceptance of its advisories.

1-5 Research approach

To investigate this, a novel and ambitious research approach was developed based on a hybrid of methodologies, including state-of-the-art literature reviews, surveys, and several interrelated real-time simulations. Empirically investigating the concept of strategic conformance required a method for subjecting a controller to an automated resolution advisory representative of how that controller would prefer to solve the conflict. To achieve this, the approach built on the principal notion of using recordings of controller's own solutions which where then disguised as automated resolution advisories given later. As such, the approach undertaken here is not to develop an advanced CD&R algorithm, but rather to *simulate* decision aiding automation. If the solution suggested by the automation *conforms* with the problem-solving style of the controller, it is reasonable to expect that the match, as perceived by the controller, would benefit the controller's acceptance of that solution.

Philosophically, this thesis takes inspiration from the brilliant work of English mathematician Alan Turing, who many years ago proposed the ultimate test for artificial intelligence, namely: if one can converse with a computer and not be able to distinguish its responses from those of a human, then that machine can truly be said to "think".[38] In practical terms, inspiration was taken from a replay procedure carried out by Fuld *et al.*[39] for studying the impact of automation on error detection. In their study, automation performance was simulated by using unrecognizable replays of an operator's own previously recorded performance. Results showed that operators were more likely to attribute faults to automation than to themselves, when in fact, it was their own errors that they witnessed.

Figure 1-2 depicts the method used for investigating strategic conformance (see also Appendix A). First, controllers participate in the *prequel simulation*, in which they play the same scenario(s) and manually solve the same (designed) conflict(s) multiple times. Conflict detection is supported by a short-term conflict advisory (safety-net) and a novel CD&R support tool. In the *conformance design* phase,

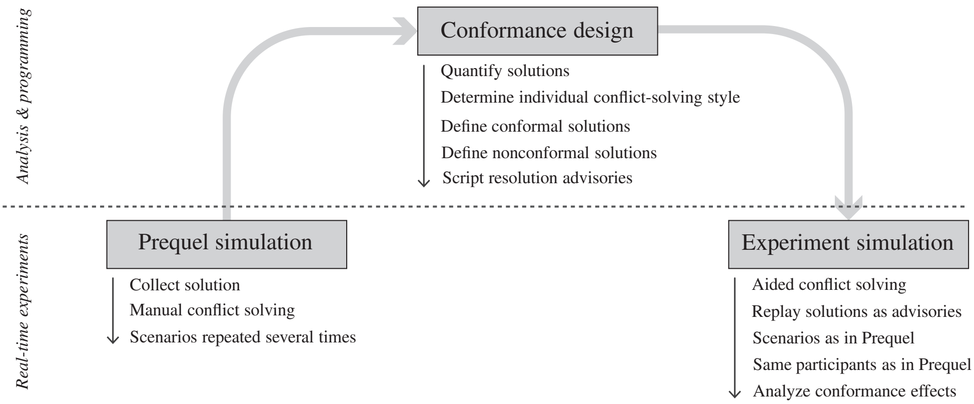

FIGURE 1-2: Experimental approach for studying strategic conformance.

solutions are analyzed and quantified against a solution parameters framework. The purpose of this stage is to determine each controller's conflict-solving style (more generally referred to as problem-solving style) for specific conflicts encountered in the prequel simulation. Importantly, determining a controller's conflict-solving style based on how he/she solves the same conflict repeatedly ascertains, and validates, that the style represents a consistent solution preference, and not a random behavior.

The individual conflict-solving styles are then used to script *conformal* resolution advisories to be replayed. A conformal advisory is intended to match a controller's own solution for a specific conflict, including matching parameters of aircraft choice, resolution type (i.e., heading, speed, or combination), and resolution direction (e.g., left or right heading). In contrast, a *nonconformal* advisory is intended to deviate from a controller's own solution. A realistic nonconformal advisory can be acquired and validated by using a different solution for the same conflict made by another controller. As such, there is no need to develop and validate an actual conflict resolution algorithm.

The same controllers participate in the final *experiment simulation*, in which they are supported by an automated decision aid for solving conflicts. The decision aid provides either conformal or nonconformal resolution advisories, although controllers are led to believe that all advisories are generated by the automation. By all other means, this simulation is identical to the prequel simulation.

In addition, the described method provides a novel approach for investigating decision-making strategies in conflict resolution. Previous elicitation methods have generally been based on subjective techniques (e.g., interviews, focus groups, and questionnaires) in combination with static traffic and conflict scenarios.[24,31,35,40] In contrast, real-time simulations were used to better capture the reactive elements and time pressures of the real world that influence decision-making. Additionally, this neutralizes the impact of hindsight biases on memory retrieval known to influence subjective methods.[41] Finally, measure of acceptance was based on the degree to which resolution advisories were accepted. This is perhaps the first time that research has tried to empirically define and quantify trust and acceptance of decision advisories that perfectly fits a person's preferred way of solving problems (in fact, because it is a "replay" of their own solution).

This thesis combines current and complementary work into both optimized technology-centered and heuristic human-centered approaches to ATC display and automation technology, by systematically evaluating algorithmic and heuristic approaches to CD&R decision support systems. Findings extend current state-of-the-art with respect to automation design principles and personalized decision support applications.

1-6 Research scope

This scope of this thesis is narrowed to semi-automated decision support systems providing support to short-term strategic and tactical task in demanding and time-critical complex systems, specifically ATC. This involves more advanced types of information analysis automation and decision aids that provide specific advisories for solving a particular control problem. More strategic decision support systems, working more on a planning-basis, are not considered, although it can be expected that this type of support system will become increasingly important in line with the desire for high performance in complex environments. Fully automated systems have not been addressed, except for learning from relevant robotics research and autonomous agents considered in artificial intelligence.

The ambitious aim and novel research approach were not without challenges and risks. An iterative design protocol had to be developed for the creation of conflicts, scenarios, and test simulations. For example, conflict and scenario design had to ensure that task load was neither trivially low nor excessively high, and set to a point at which the decision to use automation is a meaningful one. Furthermore, an experimental protocol for controlling strategic conformance was needed, allowing for a scale up from preliminary simulations with novice university students and small groups of retired controllers, to large simulations with active controllers.

It was essential that the same controllers participated in both the prequel and experiment simulations. Experiments hinged on the ability to convince controllers that they were not merely observing replays of their own previous performance (or of their colleagues) since this could influence their solutions. As such, it was necessary to ensure that scenarios were repeatable but not recognizable. Several techniques were used, including scenario rotations, name changes to call signs and sector waypoints, and "dummy" scenarios intertwined with measurement scenarios.

A more detailed description of the scope including all its assumptions made throughout this thesis, are as follows:

Automation. Future ATC is likely to increasingly depend on strategic decision aiding and medium-term conflict detection that is the middle time horizon between executive control and traditional planner activities. In relation to various level of automation (LOA) frameworks,[42–45] this reflects the the intermediate functional stages of analysis and decision-making, rather than low-level perception or high-level implementation. Therefore, the conflict resolution aid is configured with a LOA functionality corresponding to management by consent (*MbC*) and management by exception (*MbE*). In particular, this thesis focuses on *cognitive decision aids*[46] that provide explicit advisories about current and potential future states for control tasks in highly dynamic time- and safety-critical domains. In reference to traditional LOA frameworks, automation as discussed in this thesis refers to stages of automation that explicitly deal with decision selection and action implementation.[43,45] Although not explored in experiments, the concept of strategic conformance also applies to *control aids*, such as autopilots and navigation aids, and *perceptual aids* that assist in pattern recognition or provide warnings.[46,47] In all simulation of this thesis, controllers were supported by a novel information support tool for solving conflicts in the prequel simulation. This tool, called the Solution Space Diagram (SSD), represented a high LOA in regards to information integration (i.e., information acquisition and information analysis[43] that was needed to facilitate automation at the stage of decision-making and implementation. As such, the SSD integrated information relevant for identifying and solving conflicts, but left the solution choice and implementation to the controller.

Infrastructure. A futuristic ATC infrastructure is up and running (analogous to SESAR and NextGen targets for 2020[3,6]), including fully functional digital datalink communication between airborne and ground systems and free-routing airspace. As such, no radiotelephony (R/T) is required to communicate clearances to aircraft.

Acceptance drivers. In addition to strategic conformance, this thesis considers several other drivers of automation acceptance including complexity and LOA (Chapter 3), trust and source bias (Chapter 4), and automation transparency (Chapter 5). Note, however, that there are more factors that influence the acceptance of an automated system than those considered in this thesis, including human-related factors, automation-related factors, and task- and environment-related factors.[47] For example, within the human factors field the human-specific factor of trust is often considered analogously with acceptance, as a primary proxy for automation reliance. While the relevance of this research is acknowledged (see Chapter 4 for a more detailed review) this thesis considers acceptance as a more suitable and explicit measure for automation usage. Furthermore, the focus has been on the under-reliance of automation. While several automation issues are related to the over-reliance on automation (e.g., complacency, automation bias, the perfect automation schema), these fall outside the scope of this thesis.

Controlling trust. In simulations, the advisory system was presented as trustworthy and its advice was always safe (i.e., solve the conflict). This frame was used in an attempt to control trust and prevent controllers' different levels of trust from affecting their acceptance and agreement of resolution advisories.

Data quality. The underlying data are not subjected to issues such as uncertainties. Hence, the advisories given by the automation are always 100% correct and safe. The main reason for this assumption is to rule out any artifacts in decision-making caused by trust issues.

Control task. The tactical CD&R task takes place in the horizontal plane only, making it a 2D control task by means of speed and/or heading clearances. This significantly reduces the number of control strategies to resolve conflicts, allowing for better comparisons between controllers and scenarios. Note that without vertical resolutions, the control task is not necessarily easier. A single horizontal plane is more limiting and requires careful monitoring and prediction of traffic movements.

Advisory timing. The timing of an advisory may be critical to its value. Ideally, a decision aid would provide support "just in time" when the operator needs it. Considering that trust is the result of a comparison process between one owns ability and the automation's ability, researchers have argued that trust in an automated aid should be measured after the decision-maker has made a decision.[48–50] If provided before (i.e., too early), the decision-maker may be unable to adequately evaluate the advice and there is a risk that the automated advisory is "blindly" accepted. In addition, such advice may be inappropriate and interruptive. While true in theory, for all practical purposes, an advisory provided after a decision has been made (i.e., too late) would be redundant as the problem already has been solved. Furthermore, the benefits of introducing

> automation are greater provided that it, together with the human, can improve both task accuracy and speed. This is especially true in time-critical environments such as ATC. Experiments therefore sought to provide advisories before controllers had solved the conflict themselves.

The research conducted as part of this thesis is predominantly carried out within the ergonomics/human factors field, in particular addressing mental processes and decision-making in relation to human-machine interaction. As such, human physical characteristics have not been considered. However, contributions of this thesis extend more broadly, notably to the theoretical fields of cognitive psychology and information systems research, and the applied fields of human-computer interaction (HCI) and artificial intelligence (AI).

Moreover, although ATC automation is the main subject of research in this thesis, the findings and their general implications apply to any domain in which human and machine work together. Findings may, however, be of particular interest to readers working with similar highly automated domains to that of ATC, such as large control room environments in railroad and maritime operations, emergency services, military command and control (e.g., unmanned vehicles and robots), nuclear power plants, oil rigs, and manufacturing plants. Additionally, the thesis should be of interest to researchers and industry working with automation and operators in smaller control-problem specific environments, such as aircraft flight decks, ship bridges, train cabs, and the driver's compartment in autonomous cars.

1-7 Thesis outline

Figure 1-3 illustrates the outline of this thesis. It is organized around five articles either published or submitted, Chapters 2 through 6, together with the thesis Introduction (Chapter 1), Discussion (Chapter 7), and Conclusion (Chapter 8). The original articles have been retained, with the exception of Chapter 3 which has been revised and extended to better match the scope of the thesis. The original article titles have been changed for each chapter to create a more coherent thesis structure and flow. For the same purpose, brief chapter introductions have been written to clarify the relevance of each article to the thesis and identify how they link to each other and the previous chapter.

Chapter 2: Strategic conformance. Chapter 2 consists of a comprehensive literature review, introducing strategic conformance as an overarching principle underlying the acceptance of advanced automated decision aids. It explores how the acceptance issue of decision aids can be explained by mismatches in decision-making strategies between human and automated system. These mismatches can be described and measured by strategic conformance. The theoretical foundation for strategic conformance is derived through a broad review of technology and automation acceptance research in aviation but also other sociotechnical systems. Taken together, the chapter provides a starting point for investigating strategic conformance empirically.

Chapter 3: First empirical insights. Chapter 3 details the first empirical insights from an exploratory study of strategic conformance. The theoretical underpinnings, outlined in the previous chapter, are extended in a functional model of controllers' automation acceptance. The novel experimental approach for investigating strategic conformance effects is detailed, building on a three-phased approach that replays controller's own performance (see Figure 1-2). The chapter presents results from the associated real-time simulation with active controllers from Shannon Area Control Center (ACC) in Ireland. In addition to strategic conformance, the simulation explored the effects of traffic complexity and decision aid LOA on controllers' acceptance of automated conflict resolution advisories.

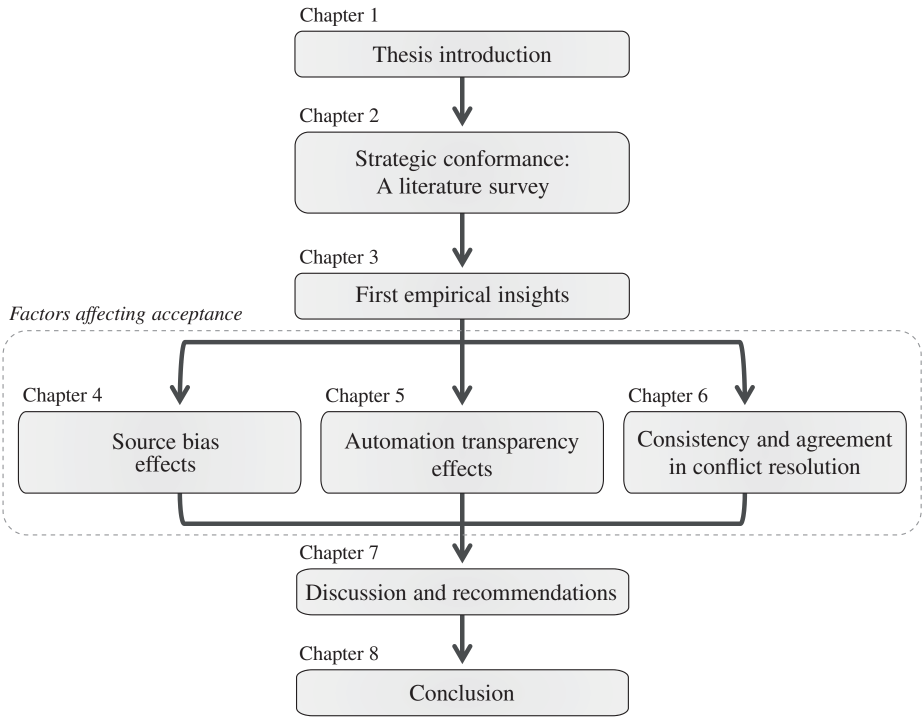

FIGURE 1-3: Thesis outline

Factors affecting acceptance. Results from the study in Chapter 3 not only indicated that the advisory conformance played a significant role in the acceptance of those advisories, but also that controllers sometimes rejected their own conformal advisories (i.e., their own solutions). To investigate this, the following three chapters detail follow-up studies investigating three research questions associated to different factors that may interact with conformance and affect acceptance. The three studies are:

- Source bias (Chapter 4): To what extent are controllers more biased against advice from a machine than from a human? Research has shown that peoples' trust and reliance behavior varies with the (perceived) source (human or machine) they are interacting with.
- Automation transparency (Chapter 5): To what extent were rejections driven by a lack of understanding conformal advisories? Research has indicated that automation transparency is a critical quality of automation for facilitating understanding of its behavior.
- Conflict solving consistency (Chapter 6): To what extent are controllers internally consistent in their resolution strategies over time? Internal consistency is a requirement for strategic conformance in that a controller's conflict solutions need to be stable over time.

Note that both the source bias and automation transparency research questions were empirically investigated together with strategic conformance. For this purpose, two separate real-time simulations were conducted. The consistency research question, however, was addressed by post hoc analysis of prequel data and controllers' solutions to repeated conflicts. Thus, the consistency study used prequel data from the previous three real-time simulations (reported in Chapter 3 through 5).

Chapter 4: Source bias effects. Chapter 4 investigates how strategic conformance, together with the perceived source of an advisory (human or automation) affects the acceptance of such advice. Automation trust research is reviewed to explore whether people have a dispositional bias against the use of automated decision aids and to what extent, if any, strategic conformance can mitigate the negative effects of such bias. To empirically investigate this, advisory source and advisory conformance were varied in a human-in-the-loop simulation involving experienced controllers.

Chapter 5: Automation transparency effects. Chapter 5 investigates strategic conformance effects on advisory acceptance in light of the automation's transparency. Transparency was manipulated by means of varying the amount of meta-information provided by the CD&R interface representation. Not only does the interface play an important role in communicating the automated decision advisories to the controller, it can also provide insight into the underlying automation rationale for why a certain solution is given. A sample of controller trainees participated in a human-in-the-loop simulation which varied automation transparency with advisory conformance.

Chapter 6: Consistency and agreement in conflict resolution. In regards to the strategic conformance of resolution advisories, the approach assumed that a controller would solve conflicts consistently but differently from a colleague. If not, it would not be possible to script conformal advisories based on the controller's consistent solution style, and nonconformal advisories based on another controller's deviating solution. Chapter 6 investigates to which extent controllers consistently solve repeated conflicts over time. This entails analysis of a controller's internal consistency (i.e., test-retest reliability) in conflict solving performance, and consensus (i.e., inter-rater reliability) between them. The analysis is based on data associated with the real-time simulations reported in Chapters 4 and 5. Furthermore, Chapter 6 explains why the concept of strategic conformance requires disagreement between decision-makers' problem solving, and consistency in decision-maker's problem-solving. The implications of human decision-making variability for automation design are discussed.

Chapter 7 & 8: Discussion and recommendations, and conclusion. The discussion Chapter (7) compares the empirical results to the aim of the thesis. Benefits and disadvantages of strategic conformance are discussed, as well as its relevance in relation to the real world. In addition, this Chapter provides recommendations for future research. Finally, the conclusion Chapter (8) summarizes the main results of the thesis and highlights the scientific and societal impacts of all findings.

2

Strategic Conformance: A Literature Survey

In this chapter, the concept of strategic conformance is introduced as a potential key factor influencing initial acceptance of decision-aiding automation. The goal of this chapter is to identify benefits and potential disadvantages of strategic conformance. This is accomplished by synthesizing literature on acceptance research across three domains of cognitive engineering (including ATC), information systems, and social psychology.

The contents of this chapter are based on:

Paper title Strategic Conformance: Overcoming Acceptance Issues of Decision Aiding Automation?

Authors Carl A. L. Westin, Clark Borst, Brian H. Hilburn

Published in IEEE Transactions on Human-Machine Systems, Vol. 46, Nr. 1, p. 41-52, 2016

Abstract

Cognitive engineering researchers have long studied the complexity and reliability of human-automation interaction. Historically, though, the area of human-automation decision-making compatibility has received less attention. Paradoxically, this could in the future become one of the most critical issues of all, as mismatches between human and automation problem-solving styles could threaten the adoption of automation. This paper presents the concept of strategic conformance *as a potential key factor influencing initial acceptance of automation, specifically decision aiding systems capable of guiding decision and action. Here, strategic conformance represents the match in problem-solving style between decision aiding automation and the individual operator. The theoretical foundation builds on the compatibility construct found in technology acceptance theories such as the innovation diffusion and technology acceptance models. The paper concludes with a critical discussion on the limitations and drawbacks of strategic conformance. It is proposed that the construct would be most applicable at the introductory phase of new decision aiding automation, in helping to foster operators' initial acceptance of such automation.*

2-1 Introduction

Since the advent of the microprocessor nearly 50 years ago, numerous work environments have come to increasingly rely on some form of computer automation. Although we have come to accept automation taking over routine and low level tasks, there remains some resistance to automation of safety-critical functions, especially in work domains that mandate automation use and rely on well-educated, well-trained, and highly skilled professionals.[51–54]

Cognitive engineering (CE) researchers have studied automation use in relation to such underlying factors as situational awareness, trust, workload, risk, reliability, and level of automation.[13,15,16,42,55] Findings suggest that: a) trust in automation develops over time as a result of prolonged experience,[15] b) acceptance and operator performance decrease when the authority and autonomy of automation increase,[8,56,57] and c) acceptance and operator performance benefit from automation actively involving the operator in the control and decision-making loops.[58]

CE researchers have, however, historically paid less attention to factors affecting the *initial* acceptance of new technology, thus factors possibly preceding trust, reliability, and others. Notice that the rejection of new technology can begin at first exposure, perhaps even before an operator has actually used that technology.[59] Notice in this a potential paradox: an operator might only develop trust after using a system, but might also be unwilling to trust a system he/she has not used. For this reason, initial acceptance of advanced decision-making automation can play a critical role in its successful deployment.

Sociology, psychology and information systems communities, on the other hand, have studied factors underlying initial acceptance. Here, the compatibility between human and technology is considered a key construct for overcoming the hurdle toward initial acceptance and technology adoption. "Compatibility" in this case refers to the perceived fit of a technology within the context in which it is used, driven by the user's values, experiences, and needs.[59] In general, the more compatible a technology is, the more likely it is to be accepted.

Presumably, compatibility can serve to mitigate initial acceptance issues of automated decision aids. Previous research has underlined preliminary benefits of matching automation's problem-solving strategies with the human, for example by modeling human decision-making heuristics[31,35,60,61] tuned to a group of people. Would there perhaps be a greater benefit in terms of acceptance if automation's problem-solving style were matched to that of the individual? To our knowledge, no theoretical or empirical work has specifically focused on differences in decision aid problem-solving styles and its effect on individual operator acceptance.

In this article we introduce the concept of *strategic conformance* as a potential key factor (and subcomponent of compatibility) influencing the initial acceptance of decision aiding automation. We define strategic conformance as the degree to which automation's problem-solving style matches that of the individual human. A person's problem-solving style is made up of both the product (solution) and its associated process (underlying strategies). The latter is only apparent since the process cannot be determined by knowing the product, only inferred from observable behavior or output. We hypothesize that strategic conformal automation can, first and foremost, promote initial acceptance of new technology, but also improve overall system performance as operators are more likely to use it. The discussions on the potential benefits and pitfalls of this rather extreme perspective are guided by an extensive literature survey across various different fields that focus on automation acceptance.

2-2 Resolving automation acceptance issues

Technology resistance is a widespread concern across several work domains. In health care, physician resistance has been identified as a critical obstacle to greater adoption of robotically-assisted surgery[62] and electronic tools.[52–54] Evidence from the ATC community indicates that current decision aiding systems, intended to support the controller in CD&R tasks, are sometimes rejected or used in unintended ways.[8,10,12] Note that whereas some work settings and organizations might mandate automation usage, even then automation can be underused, misused etc.

Automation acceptance research has primarily focused on identifying and preventing the inappropriate use of automation,[63] often categorized as misuse (overreliance) or disuse (underreliance) of automation.[13] Research typically considers human interaction with complex technologies capable of autonomy, in highly dynamic and complex environments characterized by high risk,.[13,15,64] Examples include ATC, aircraft carriers, nuclear power plants, space shuttle operations, fire fighting, and health care (see for example[65,66]).

Several factors are believed to influence the choice of whether to use automation. Examples are attitudes toward automation, trust, workload, complexity of automation, perceived risk of automation use, and perceived automation reliability.[13,15,16,67–70] Riley argued that the core construct of automation reliance, defined as the "probability that an operator will use automation" [16, p. 21] is influenced by various factors such as trust in automation, self-confidence in manual performance, perceived risk, and fatigue.

The framework suggested by Dzindolet *et al.*[55,67] indicates that automation-use decisions, and which level of automation in particular, are determined by three decision-making processes (cognitive, social, and motivational) and their associated decision-making biases. At its core, the model proposes an evaluation of the perceived reliability of manual control against the perceived reliability of automated control. The outcome, measured in perceived utility of the aid, determines whether automation use is favored or not. This acceptance rationale, determined by a balancing process weighing operator self-confidence against confidence in aid, is central in CE theories of automation use and trust.[63,68,70–73]

Alternatively, reasons for automation resistance can already be viewed from a design perspective. Characteristics of poor compatibility might stem from the underlying goal for which the machine has been designed. That is, the deterministic algorithms embedded in automation generally aim to optimize. Such algorithms can be at odds with less structured, more heuristically governed human decision-making that tends to satisfice. Could there perhaps be an acceptance benefit if automation were designed consistent with human-like problem-solving styles?

2-2-1 Technology-centered automation

Automation acceptance issues can be found in many different sociotechnical work domains in which skilled professionals are responsible for the safety and efficiency of operations. The ATC community has a well documented history of finding a suitable approach to automation design that promotes a functional and collaborative human-automation relationship. Over the years, innovative decision aiding systems and automation concepts have been proposed and developed to help controllers cope with the increasing pressures of the expanding ATC system (see[10] for an overview). Many current tools have grown from technology-centered research projects exploring fully automated concepts, such as AERA 1-3,[74,75] ARC 2000,[76] and the PHARE project.[77] Realising that full automation was not feasible, and that the human has to retain control, several sub-concepts were instead transformed into decision aids.[10,77] Their success in adoption, however, has been questioned.[9]

Conflict-solving algorithms typically approach the environment in a dichotomous fashion, providing single fixed, optimal solutions according to causal deterministic laws.[19] Most models are typically limited in application and options considered for conflict resolution.[11,78] According to[19] all CD&R systems follow roughly the same design model, containing a deterministic trajectory model and set of alerting threshold metrics, often based on "engineering intuition." After a series of simulations threshold alerts are tuned for optimal performance. Finally, performance is assessed in terms of false alarm rates and loss of separation frequency.

In complex decision-making tasks, operators do not necessarily agree with their decision aid on when/whether to intervene, nor which solution to apply. As such, the decision aid can be seen to possess its own implicit decision-making style that is either conformal with the operator's or not. For example, in recent real-time ATC simulations exploring novel decision aids, the automation was perceived as occasionally "fighting" against the controller on how to solve conflicts.[17] In another study, the hampered adoption of a conflict detection aid (URET, or User Request Evaluation Tool), was attributed to the system not reflecting controllers current conflict managing procedures. Controllers felt that the aid was too slow (its advice often coming after conflict detection) and, contrary to the designers' intentions, tended to use the tool instead as a flight strip replacement and route amendment tool.[9]

Deciding on when, and on which thresholds, automation should interact may be one of the most difficult problems to address, and constitutes perhaps the most pressing mismatch between controller and automation in ATC.[61] Additionally, there might be large individual differences in preferred safety margins. The decision of whether to intervene can generally be described as a trade-off between intervening directly or waiting and collecting more information.[56] Both choices can be determined by a workload regulation process.[79,80] Intervening directly can be appealing to avoid the additional workload of having to monitor the evolving situation. In busy situations, however, task shedding and prioritizing may encourage postponing intervention in order to control and maintain minimum workload. In certain situations, neither early nor late intervention may be feasible, and can put controllers in a "double bind." CD&R automation, however, is generally not this sensitive, but operates according to pre-defined rules and process parameters.[56]

2-2-2 Heuristic forms of automation

Several research efforts have explored alternative approaches to decision aiding automation that better acknowledge the psychological and behavioral variables that might color human decision-making.[18] One possibility is to model algorithms after human control strategies governed by heuristics. This would make the automation act in a manner consistent with how the human might act and thereby reduce the compatibility gap in decision-making strategies.

Heuristics are cognitive shortcuts commonly applied to reduce cognitive workload and increase efficiency in decision-making.[81] Heuristics can color our expectations, however, and there are occasions when they can introduce systematic decision-making biases. This can result in sub-optimal strategies.[22] To eliminate such unwanted cognitive biases, heuristics can be formalized and refined. Benefits of such heuristic approaches lie in their ability to sort out noise in the perceived environment and consequently being able to better predict future states.[82]

For example, satisficing is a cognitive heuristic that searches alternative solutions until an acceptable "good enough" solution is found. Satisficing is typically applied when an optimal solution cannot easily be determined.[20,21,83] Goodrich and Boer[84] took inspiration from this heuristic and hypothesized that car drivers would better understand and prefer adaptive cruise control automation that mimicked skilled human task behavior for longitudinal vehicle control. They developed a human-automation interaction framework based on multiple mental models that synthesized satisficing decision theory[21] with Rasmussen's skill-, rule-, and knowledge-based taxonomy.[85] In a series of experiments, three human car driving skills were emulated by automation: active braking, speed regulation, and maintaining a desired distance (headway) to a preceding vehicle. For automation tuned after driver behavior and thresholds for the three skills, findings indicated improvements in the system's performance evaluation, detection of the system's functional limitations, returning to manual control, trust in the system, and overall safety.

Similarly, researchers have shown that decision-making heuristics are commonly used by air traffic controllers in CD&R,[23–28] and have advocated heuristic approaches to decision aiding automaton design.[29,31–33,61,86] Controllers are believed to rely on cognitive processes akin to naturalistic decision-making when identifying and solving problems.[35,87] While being able to rapidly interpret and act upon complex situations, controllers tend to struggle with providing strategy-based explanations underlying their decisions.[87,88] Further, controllers are believed to develop and maintain conflict resolution heuristics in a "mental library".[89] When a conflict is detected, the controller scans through the library for a suitable solution.

Several research initiatives have attempted to mimic this decision-making process and have explored human heuristics-based approaches for decision aiding au-

TABLE 2-1: Arguments for and against heuristic forms of automation (after Kirwan and Flynn[29])

Arguments for	Arguments against
Since it will "behave" like a good controller would, giving reasonable suggestions, it will have face validity and be more intuitive, fostering trust and acceptance.	If the solution is based on controller strategies, there is no advantage of automation.
As solutions suggested will be similar to what the controller might suggest, the controller can quickly understand and infer why the solution was suggested.	If controllers are not sufficiently homogeneous it will be difficult to identify a single strategy that applies to all.
Solutions will be built on controller strategies which are known to work and have "stood the test of time."	Experts cannot readily explain their expertise which complicates strategy elicitation.
It will be easier for controllers to determine when something is wrong (e.g., system failure).	The automation will make the same mistakes as the controller.

tomation. These typically follow a similar pattern: after detecting a potential conflict, an algorithm calculates, based on a set of formalized pre-defined problem-solving heuristics, alternative solutions and presents them to the controller,[31,35,90,91] thus acting as an externalized mental library of heuristics.

Given human cognitive limitations and possibilities of automation superiority in performance, it may seem counterintuitive to propose that automation should adhere to human ways of problem-solving. This, however, has been reflected upon in previous literature. For example, Kirwan and Flynn[29] considered arguments for and against the cognitive tools concept underpinning EUROCONTROL's Controller Resolution Assistant (CORA) automation (Table 2-1). One of the arguments against heuristic forms of automation is the potential variation in preferred solution. If operators differ in their decision-making strategies and disagree on how to solve problems, there is little value in a system that may only match the decision-making style of a few operators. This issue, however, is not unique to heuristic forms of automation, but applies equally to other forms of automation. Instead of treating decision-making diversity as a threat, we can incorporate it as an advantage in automation design.

2-2-3 Individual-sensitive automation

A logical next step, further reducing the compatibility gap between human and automation problem-solving styles, is to consider individual-sensitive automation. decision aiding automation is typically designed to fit with a group rather than the individual. A limitation of the above heuristic forms of automation is that, although individual differences were acknowledged, the resulting decision aid treated operators as a homogeneous group. As such, they have not been able to ensure complete harmony between the individual human and automation.

On the basis of individual differences in personalities[36] and cognitive styles[37] that influence how problems are approached and solved, it can be expected that automation sensitive to individual preferences would be beneficial to automation acceptance. This notion is supported by several technology acceptance theories (see next section) that consider cognitive style to be an important factor influencing acceptance.[92–95] One example of a cognitive style dimension is that of impulsivity-reflectivity. On a continuum measuring response time in problem-solving, individuals are either fast (impulsive) or slow (reflective) in making decisions.[37]

Recently, Liu *et al.*[96] argued that decision aiding automation that embraces individual differences will become increasingly important for user attitudes and performance in situations with vaguely defined tasks and problem-solving processes, of which ATC is a prime example. While overlooking individual differences has been sufficient for today's system, several researchers have argued for more individualized automation to foster successful human-automation teamwork in future ATC.[97–99] In theory, if controllers are to remain in the loop while working with decision aiding automation, they should evaluate the output of automation against an on-going appraisal.[100] Given what is known about naturalistic decision-making and biases in human decision-making, however, Kirwan and Flynn[29] were probably correct in their observation that "if the resolution advisory is similar to that which the controller might have thought of, and if the controller can quickly infer why such a resolution was made, then such a rapid evaluation and decision will be possible..." (p. 4).

Following this line of thought, we have investigated the potential benefits of individual sensitive, strategic conformal decision aiding automation on acceptance in previous work.[101] Although it was technically not feasible to create such advanced automation, we simulated it by replaying controllers' own solutions to pending separation conflicts and disguised it as "automated advice."

In a first set of simulation trials, 16 controllers used a constraint-based interface (the Solution Space Diagram, SSD[102]) to formulate and implement heading and/or speed solutions to traffic conflicts in the horizontal plane. The interface was novel in the sense that it revealed all possible heading and/or speed combinations that would either lead to a safe state or a loss of separation. This allowed controllers to freely decide and implement their own preferred solution to traffic conflicts.

In a follow-up trial two weeks later, the same controllers interacted with an identical system, only now supported by a decision aid that would recommend solutions by plotting them in the constraint-based interface. This was done deliberately to make controllers believe they could perceive the decision-making criteria of the automation, while also opening up the opportunity to veto automation and implement another preferred solution. In two separate simulations consisting of four scenarios repeated four times, each controller was subjected to a total of 32 advisories (one per scenario). In half of the cases the controller's own previous solutions were presented (conformal advisories). In the other half, a colleague's different, but equally workable and safe solution were presented (nonconformal advisories). Results showed that conformal advisories were accepted 33% more often than nonconformal advisories (76.2% and 56.6% respectively). Standardized agreement ratings, measured on a 1-100 rating scale, supported acceptance data, with conformal advisories receiving higher agreement ratings. Controllers also responded on average one second faster to conformal advisories (4.9 s and 5.9 s respectively). This study illustrates that individually matched solutions can indeed improve automation acceptance.

To summarize, empirical evidence has shown that acceptance issues of decision aiding automation can be attributed to a mismatch between human and automaton decision-making strategies. Larger mismatch gaps, as exemplified by technology-centered approaches (e.g., URET), result in larger resistance to automation. Reducing the gap, by heuristic forms of automation (e.g., CORA), can benefit acceptance. The greatest benefits, however, may be achieved by tuning decision aiding automation to the individual's problem-solving style. However, we do not only have to rely on empirical studies to make such a claim, as it appears that other research fields outside of CE explicitly acknowledge the importance of compatibility between human and technology for acceptance.

2-3 Toward a new perspective

Acceptance research can, in general, be divided into three main streams depending on the field in which it has been studied. In addition to CE research, addressed in the previous section, the other main streams are innovation-diffusion theories (IDT) originating from the broad fields of sociology and psychology, and technology acceptance theories studied within the information systems community. Surprisingly, both interaction and knowledge transfer across these three communities have been sparse.[103] As such, it is valuable to synthesize and perhaps unite the different views into a single concept that addresses acceptance issues and how to resolve them.

2-3-1 IDT and the adoption process

IDT research is devoted to explaining why only a few innovations (i.e., items or technology) gain widespread adoption in a population while the majority of innovations fail. Everett Rogers[59] defined adoption as “a decision to make full use of an innovation as the best course of action available” (p. 21). The decision to adopt an innovation is considered to progress over time through a series of stages, from (1) knowledge of the innovation and the (2) formation of favorable or unfavorable attitudes (persuasion stage), leading to a (3) decision to adopt or reject the innovation. If adopted, (4) implementation occurs, followed by (5) reassessment of the previous innovation-decision made (confirmation stage).

Different innovation and user attributes are believed to affect adoption rate. The original five innovation attributes are compatibility (with perceived values, needs, and experience), relative advantage (compared with what it intends to replace), complexity (in understanding and using), trialability (extent to which a innovation can be tested), and observability (or visibility of the results of an innovation to others). Users are typically characterized according to their rate of adopting an innovation, ranging from the innovators and early adopters to technology resistant laggards.[59] Interestingly, IDT research puts compatibility as one of the most dominant factors driving early acceptance.

2-3-2 Technology acceptance theories

The information systems community has developed several acceptance models that concentrate on individual determinants of user acceptance measured in the willingness or intent to use. Examples include the Theory of Reasoned Action, the Social Cognitive Theory, and the Unified Theory of Acceptance and Use of Technology. The most widely used model, however, is the Technology Acceptance Model (TAM).[104–106]

External stimuli, such as user and technology characteristics, cognitive style, and subjective norms influence people’s willingness to use the technology, which in turn determines their use behavior. The willingness to use makes up the core of the model and consists of beliefs about using (a cognitive response to the external stimuli) that directly drives the user’s overall attitude toward using. There are two beliefs: perceived usefulness (the extent to which a user believes using the technology will enhance performance), and perceived ease of use (the extent to which a user believes using a technology is effortless).[107,108] While TAM originally disregarded compatibility as a central construct affecting use, later modifications have demonstrated and underlined its importance. Compatibility is not believed to influence acceptance directly, but indirectly through the perceived usefulness and perceived ease of use.[109]

2-3-3 Synthesizing acceptance models across communities

Two separate research groups recently presented theoretical acceptance models synthesizing the TAM and IDT frameworks with relevant knowledge on automation use derived from CE research to create the Automation Acceptance Model (AAM)[64] and the Adjusted Automation Acceptance Model.[14] While similar in the sense that both models rely on TAM architecture, they diverge in the additional and modified elements considered. Both models highlight a concern within the CE domain that current models, and individual drivers of acceptance, do not adequately capture acceptance.

The AAM model specifically embraces the compatibility construct, along with trust, as key drivers for automation acceptance (Figure 2-1). The feedback lines denote that acceptance is a bi-directional process. Consider the relationship between trust and acceptance. In order to start using a new system, there has to be trust. Yet trust, specifically for that system (to be distinguished from dispositional trust[110]), cannot be built without using the system. Compatibility on the other hand, is a perceived characteristic of the automation that can influence trust.

While the CE community has predominantly contributed with task-technology compatibility in cognitive systems and the dynamics of trust and reliance,[64] acceptance theories originated in the information systems community have provided coherent acceptance frameworks, including clear methodologies for measuring acceptance. The divergence in research focus between the two communities has been attributed to differences in research granularity.[64] While the CE community has mainly focused on short-term, micro-level observations of operator behavior, the information systems community has approached acceptance by looking at human-

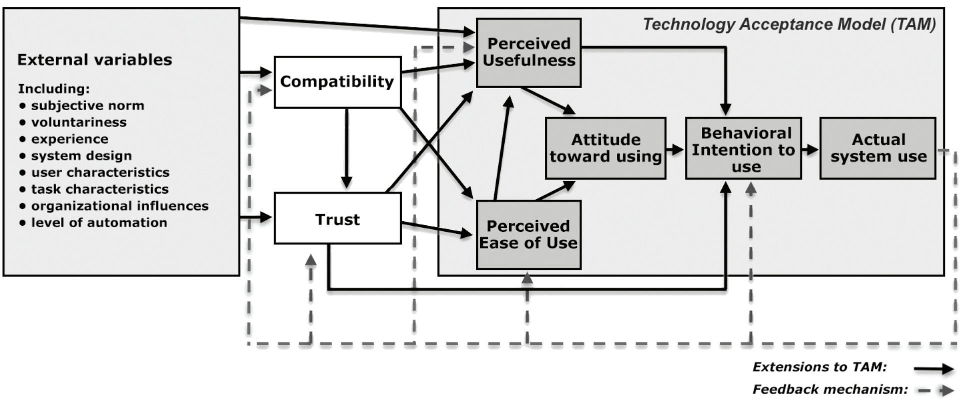

FIGURE 2-1: The Automation Acceptance Model (adapted from[64]).

technology interaction at a macro-level. The micro-/macro-level distinction separate acceptance drivers depending on the time duration considered.[14] Macro-level observations consider long-term drivers for acceptance, whereas micro-level observations consider short-time periods (even down to milliseconds).

While both the information systems and IDT communities consider long-term decisions of technology use, CE research typically focuses on the decision to use automation for a specific task at a specific time. This momentary use is different from the question of using the automation at all, and is particularly suitable for decision aiding applications of which the output can be rejected. Consider for example a decision aid for separation assurance in ATC, which provides advice on how to solve potential conflicts between aircraft. The controller can accept or reject the advisories suggested by automation; the controller can also decide to turn off or ignore the automation. While both cases address acceptance, the prior considers automation and its output at a specific time and context. The latter is a general decision to accept or reject the entire system. Although it may seem reasonable to expect system output to correlate with system acceptance, such a connection has not (yet) been established.[111]

Our literature review indicates that IDT and TAM concepts have been predominantly applied to technologies that support the first two steps of information acquisition and analysis in Parasuraman and Wickens'[43] LOA framework. The following two stages of decision selection and action implementation, which are characteristic of more autonomous technology, have received much less attention, however. Given that automation is increasingly more capable of assuming control over decision-making, and the fact that compatibility is considered one of the most influential technology characteristics determining both initial acceptance and long term adoption,[109,112,113] we hypothesize that we need a new framework that specifically addresses acceptance by focusing on compatibility.

2-4 Strategic conformance

2-4-1 Complementing existing constructs

Our review indicates that compatibility is a key construct in human-machine interaction research. It can be envisioned at different levels, as illustrated in Figure 2-2. At the lowest level, there is response compatibility that can reinforce correct use of a device (e.g., use up-down lever to drop landing gear, not rotary switch). Another level comprises perceptual compatibility that adheres to the management and interpretation of sensory information (e.g., use of red light to preserve night vision). A higher level reflects communication style compatibility that considers how information is exchanged. For example, Parasuraman and Miller[114] demonstrated that

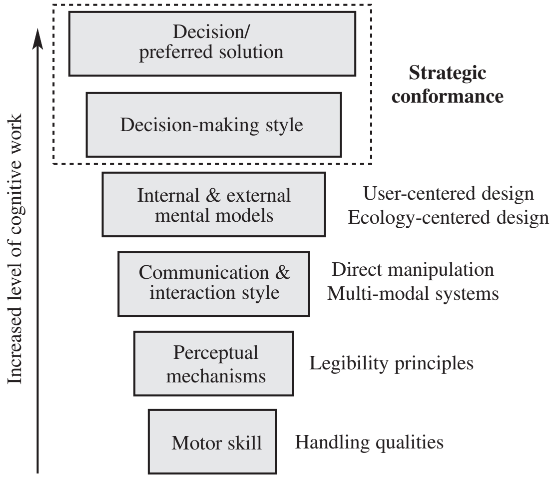

FIGURE 2-2: Levels of human-machine compatibility, and their respective constructs found in CE research, ordered by increased levels of cognitive work. Strategic conformance can be regarded as a complementary construct that plays an increasingly important role for automation acceptance of advanced automation assuming the role of a decision aid.

a decision aid with a polite and friendly communication style (defined as good etiquette) significantly increased participant's trust in automation. At the highest level, there is decision-making compatibility, which has only emerged in recent times as machines have become more capable of assuming control over decision-making processes.

In two important ways, strategic conformance would extend previous research in compatibility. First, previous research has been limited to exploring compatibility related to overt similarities between automation and human behavior and communication (communication and interaction style level in Figure 2-2). Examples of the latter are human-like robots and anthropomorphism,[115, 116] the computers are social actors paradigm,[117–119] and automation etiquette.[114, 120] Strategic conformance is unique in that it also recognizes the human's concealed underlying strategies (process). If the automation suggests a solution that matches the operator's problem-solving style, the operator is likely to infer that the automation has arrived at the same solution by applying the same rationale. It would "make sense," and possibly alleviate cognitive workload associated with trying to understand what the automation is doing and why.

Second, strategic conformance recognizes individuals' problem-solving styles. Previous research has assumed that humans, in specific work domains, are sufficiently homogeneous as a group and that it is possible and sensible to develop "one-size-fits-all" automation. Yet, like our physical configuration, it can be expected that an individual-sizing convention would better support our cognitive configuration. As such, we agree with other researchers who argue that automation design would fare better in the operational environment if it acknowledged the individual differences that exist between operators.[97–99]

Strategic conformance can also be applied to complement design approaches. For example, consider the Ecological Interface Design paradigm, which strives to facilitate coordination between human and machine by making interface representations that reflect the shared work domain.[121,122] Ecological interfaces typically show constraints on action that have an ecological validity (e.g., laws of physics governing work), but the decisions and strategies on how to maneuver through those constraints can differ between humans and machine. Take for instance the street map analogy and car navigation systems: within all possibilities to travel from point A to point B, one person might prefer to take the scenic route, whereas another person prefers to take the route that increases the likelihood of finding gas stations, malls, or food courts along the way. An automated decision aid, on the other hand, might only suggest the shortest route, optimal in terms of travel time and gas mileage. Thus decision aids that cannot be attuned to individual personal preferences are likely to be less accepted. Similarly, increasing the level of automation's control authority in ecological information aids might require attunement to operator-preferred problem-solving styles in order to ensure acceptance and automation use.

2-4-2 Acknowledging individual preferences and diversity

Strategic conformance acknowledges the diversity in problem-solving. Consider strategic conformance in the context of solving a conflict between two aircraft (same type) crossing each other at right angles. In Figure 2-3, aircraft A and B are approaching each other at the same speed, with the closest point of approach being zero nautical miles. If one would rely on ICAO established right-of-way rules[123] the aircraft with the other on its right shall give way (see Figure 2-3(a)). When asking controllers, an often cited preferred strategy (if only considering horizontal changes) consists of turning one behind the other (a "set-and-forget" solution requiring less monitoring[24,35]), without consideration for ICAO rules (see Figure 2-3(a) and (b)). Another often-cited strategy consists of slightly turning both aircraft ("one command is no command"), which minimizes track deviations and more "fairly" shares solution between both involved aircraft (see Figure 2-3(d)). Although the solution in Figure 2-3(c) may seem unreasonable, given the extra miles required and likely increase in monitoring, it is still a viable alternative.

Mathematically, an optimal solution could consist of any of the four alternatives

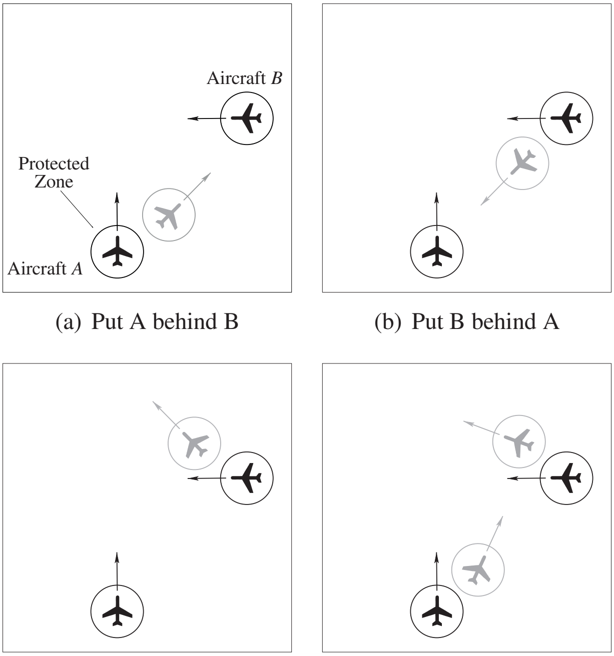

(a) Put A behind B (b) Put B behind A

(c) Put B in front of A (d) Dual heading change

FIGURE 2-3: Conflict solution alternatives to a right angle conflict.

depending on the decision criteria selected, how these are weighed and combined. Typically, algorithms have cost functions driven by optimality goals considering several parameters, such as minimum patch deviation and minimum fuel burn. Even if controllers would consider the same or similar criteria, the relative importance and combination of them may differ from the algorithms, leading to a mismatch in underlying goals. Literature suggests that controllers' cost functions are driven by goals such as minimizing workload and finding solutions that afford least effort, yet are efficient and robust. With automation assuming more human-like cost functions, human and automation problem-solving styles would converge.

In addition to the four vector solutions illustrated in Figure 2-3, there are several more alternatives including speed and level changes. In reality, characteristics of the aircraft involved contain more variability than the conflict situation depicted in Figure 2-3. Environmental and sector characteristics, together with adjacent traffic, however, can complicate the options available and narrow down the range of "good" solutions. Still, it is likely that each conflict will have more than one feasible solution available. But even if there would be not doubt in which solution to apply, there is always an option of when to implement it.

2-4-3 Encompassing process and product

Elements of strategic conformance include both process (underlying strategy) and product (outcome or solution), which can be considered independently. The solution arrived at by automation may be conformal with the human, although both have relied on different strategies. In contrast, similar strategies can result in different outcomes. Although conformance of both process and product can be discussed in theory, we cannot look at product and conclusively determine process. Yet, this is what we do all the time in the real world when interacting with colleagues, friends and others. We make inferences about underlying motivations and strategies based on the repeated exposures to outputs that can be observed. This inference engenders attitudes, feelings, and responses such as trust and acceptance. Research suggests that these interaction characteristics also apply when we interact with automation.[72,120] These type of responses have been shown to positively influence acceptance and trust in robots,[115,116] computers,[118,124] automated decision aids in flight simulation,[114] and autonomous cars.[125]

We envision product conformance to be most valuable in situations that lack objective "gold standard" criteria for determining an unequivocal best or optimal solution. In such situations, the human is the best judge of decision quality. Strategic conformal automation can still be beneficial in these situations, primarily in terms of speed at which solutions are addressed, and in alleviating operator workload. More importantly, it should foster acceptance, trust and willingness to use the automation.

Process conformance, however, may best apply in situations for which objective "gold standard" criteria are available. Automation using the same criteria as the human will likely be able to more consistently and accurately solve problems. As such, process conformance can more drastically benefit performance. Still, it would be conformal in that the automation would assess, weigh and combine criteria based on an understanding of the individual human process.

Strategic conformance based solely on process could be problematic, because underlying strategies may be difficult to determine objectively, and therefore inferred primarily based on observable behavior. As an alternative, an individual's consistent problem-solving pattern can be used to determine that individual's problem-solving style. If automation would be attuned to and capable of learning about an individual's preferences, it could be developed to generate personalized advisories conformal to those preferences. This notion is captured in the research area of recommender systems, which attempt, based on a personalized user model, to predict individual user preferences and provide recommendations thereafter. Information about users' preferences is often gathered through approaches relying on rating structures. Recommendations are typically based on the user's past preferences, or that of other users with similar profiles.[126]

To summarize, the benefits of strategic conformance are tied to the generation of positive attitudes, trust, and acceptance of automation and its output. Solutions generated by a strategic conformal system are, however, not expected to necessarily be better or more optimal than those of the human. As such, system performance related to specific solutions may not increase. Overall system performance, on the other hand, would benefit by an increase in acceptance, trust, and willingness to at least use automation so as to embrace the performance benefit it may offer.

2-5 Limitations and pitfalls of strategic conformance

Thus far this article has conveyed a new perspective on causes underlying automation acceptance issues and presented a possible solution in strategic conformance. It is also appropriate to take on a more critical perspective and consider the limitations and potential pitfalls of strategic conformance.

2-5-1 Drawbacks of individual-sensitive automation in teamwork

In the context of adaptive cruise control, strategic conformal automation would be sensitive to individual differences in how smoothly drivers accelerate or brake. However, taking a whole system perspective, it would be desirable if all drivers and cars behaved the same way. That way, other agents (human or automation) would know better what could be expected in various traffic situations. There would be less uncertainty in the behavior of others, which would simplify automation design.

This is a legitimate concern in safety-critical domains that rely on teamwork and transparency in task execution. Consider, for example, air traffic controllers who work together in teams of two (tactical and planner controller). With frequent handovers between the two it is important to continuously communicate and keep the other person “in the loop” so that a shared understanding is achieved and maintained. Training to harmonize working styles is essential in these domains. In other words, there may be an intrinsic desire to suppress individual differences in problem-solving and decision-making in work environments that rely on teamwork.

2-5-2 Requirements on consensus and inconsistency

Strategic conformance relies on high internal consistency. That is, strategic conformance is not relevant unless users are consistent in their problem-solving styles over time. We expect that an expert will provide us the same answer to a given question no matter how many times we ask that question, provided that the preconditions are the same in all cases.

At the same time, strategic conformance hinges on the assumption that users differ in their preferences and that there are individual differences in problem-solving styles. If users were perfectly homogeneous, strategic conformance would not be interesting, as we would not have to acknowledge individual differences. This would make it easier for automation designers as it would be sufficient to develop a “one size” system. In most domains, however, it is more likely that there are in fact individual differences between users that deserve recognition. In order to determine the prerequisites for strategic conformance in specific work environments, it is necessary to determine the degree of both consistency and consensus.

2-6 Conclusion

In this article we have explained how issues of automation acceptance can be attributed to a compatibility mismatch in problem-solving style between human and automation. We examined how researchers have attempted to reduce this compatibility gap and mitigate acceptance issues by exploring heuristic forms of decision aiding automation. We argue that, in order to achieve the highest compatibility level, we may need to consider individual-sensitive automation more conformal with individual problem-solving styles. We presented strategic conformance as a mediating concept, complementary to automation design frameworks, for overcoming resistance to accept decision aiding automation.

However, we also acknowledge that strategic conformal automation may only be relevant to gain initial acceptance in the introductory phase of new technology. Furthermore, conformance may only be relevant for expert users who hold consistent and well-developed decision-making strategies. Finally, given technical advances in areas previously considered unique to human cognitive skills, automation is expected to increasingly assume authority in problem-solving and decision-making tasks. The inevitable trajectory of many work domains will involve more capable automation acting in an intelligent advisory capacity. As such, automation will likely provide support that is more strategic in timescale, less transparent to the operator in that decision rationales are concealed, and presented as recommendations. In this context, the issue of acceptance is central. It is reasonable to hypothesize that a recommended solution matching the individual’s problem-solving style would be more readily accepted.

3

First empirical insights

Strategic conformance argues that acceptance of automated advice benefits from a match in decision-making strategies between the automation and operator. In the previous chapter, a theoretical foundation supporting this claim was presented. In this chapter, we present the first empirical study investigating the effect of strategic conformance, together with traffic complexity and level of automation, on the acceptance and use of a novel ATC conflict resolution decision aid. Note that this chapter is based on a journal paper presented in Air Traffic Control Quarterly, 2014. Parts of it, however, have been extended and edited to better complement this dissertation.

The contents of this chapter are based on:

Paper title	Will controllers accept a machine that thinks like they think? The role of strategic conformance in decision aiding automation
Authors	Brian H. Hilburn, Carl A. L. Westin, Clark Borst
Published in	Air Traffic Control Quarterly, Vol. 22, Nr. 2, p. 115-136, 2014

Abstract

In a series of real-time trials, we simulated sophisticated ATC conflict resolution automation using unrecognizable replays of air traffic controllers' own performance. Using a novel experimental design and a prototype ATC interface, we explored with operational controllers the interactive effects of traffic complexity, level of automation, and "strategic conformance" (defined as the match between human and machine solution strategy) on a number of dependent measures. Personalized conformal advisories (exact replays of a given controller's previous solution) were accepted more often, rated higher, and responded to faster than were nonconformal advisories (replays of a colleague's different solution). Controllers not only discriminated between resolution advisories, but more importantly, preferred those that matched their own solution for the same conflict. In the end, one result stood out in particular: roughly 25% of conformal advisories were rejected by controllers. Taken together, this study has provided empirical insights into the critical role that strategic conformance can play, at least in a transitional phase, as new and sophisticated decision support automation is being introduced.

3-1 Introduction

Roughly 60 years ago, English mathematician Alan Turing famously posed the ultimate test for artificial intelligence: that its performance be indistinguishable from that of a human. If one could converse with an unseen agent, and mistake computer for human responses, then that computer could truly be said to "think." This notion has driven research into artificial intelligence for over half a century.

We are at a point in the evolution of automation that we routinely turn over to computers many of the "thinking" tasks previously performed only by humans. Our planes, trains, and even automobiles rely on more (and more capable) automation than ever before. Despite various achievements, however, some gap remains between the theory and practice of automation design. For instance, as of this date, not a single computer has passed the Turing test, and we have not realized in any meaningful way the highest levels of autonomous systems. As Sheridan noted, we still have no idea how to program computers to "take care of children, write symphonies, or manage corporations..." [127, p. 129].

The Multidimensional Framework for Advanced SESAR Automation (MUFASA) project started, in a sense, from the opposite view: What if we could build perfect automation that behaved and solved complicated problems in exactly the same way as the human? Would the human accept its advice? Or might the humans reject such solutions simply because they were proposed by automation? The question is whether we have evidence of a fixed bias against automation, irrespective of its performance. That is, do operators show an inherent bias against automation?

These questions are neither trivial nor merely academic. The MUFASA project assumed that future ATM will rely increasingly on automation capable of overseeing the cognitive and strategic aspects of ATM. The SESAR target concept for future ATM[7] is built on an evolutionary path of five "Service Levels" that correspond to progressively more sophisticated automation, in terms of both the types of tasks it can perform and the level of authority and autonomy it can assume.

3-2 Automation acceptance in ATC

Given the increasing need for automation to accommodate future traffic and complexity levels, how do we go about developing automation that controllers will accept, trust and come to use? It has been noted that, in a variety of fields, operator acceptance might be one of the greatest future obstacles to successfully introducing new automation.[51] The ATM community has for years struggled to find a suitable automation design approach that promotes a functional and collaborative human-machine relationship. Past efforts such as AERA[74,75] and ARC 2000[76] rested on a technology-centered approach that assumed optimized, algorithmic CD&R. The optimized automation view has softened somewhat in recent years, and as the movement toward "human-centered" automation design[51] has grown, so too has the recognition that optimized algorithmic approaches to CD&R are often at odds with the more heuristic approach controllers bring to problem solving.[30,32,33]

Given advances in flight path monitoring and optimization, automated decision aids will likely introduce solutions that are:[14,51,127]

- more strategic in timescale;
- more strategic in their ability to consider greater airspace volume and network effects;
- less transparent to the controller, in that all decision constraints are not made clear, and
- presented as suggestions.

Central to this discussion is the role of operator acceptance. Among the various risks in introducing new automation is the potential for lack of operator acceptance to limit the use and, in turn, the potential safety and performance benefits of such automation,[13,51] both in ATC,[6,128] and in other high reliability domains.[65] Notice one paradox of advanced automation: as we move toward more capable CD&R automation, intended benefits can only be realized if a system is properly used, yet such a system might only be used if its benefits are recognized.[129] Otherwise, a controller might opt to either disuse (i.e., make no use of) or misuse (i.e., use it in a manner inconsistent with design intents) the new automation.

Several efforts to introduce advanced CD&R automation have been thwarted over the years specifically because of acceptance problems.[10,12,77] Acceptance problems have been linked to such interrelated issues as:

- automation misuse;[9,13,130]
- perceived strategy mismatches between controller and machine,[31] and
- system performance decrements, a result of the controller having to "fight" the system.[17]

Given that controller initial acceptance appears to have impeded some past efforts to develop advanced CD&R automation,[14] and given the evolutionary trajectory that such automation appears to be following (toward more strategic capabilities), it is reasonable to hypothesize that controller acceptance will present at least as big an issue in the future. We therefore sought to explore how and whether acceptance could be impacted by the closeness of fit between human and machine problem solving styles.

3-2-1 Strategic conformance between controller and automation

Several authors have noted potential benefits of heuristic, as opposed to algorithmic, approaches to CD&R automation.[26,29,30,32] Notice here the distinction: a solution might be operationally or mathematically optimized (think of a system that always chooses the shortest route), but this solution is not necessarily the one that the controller would select.[19,72] Controllers' preferred solutions might be driven by, for instance, general control strategies,[89] (e.g., guaranteeing separation through course divergence) or even display perceptual limitations (e.g., a tendency to disregard speed control in en route airspace).

There have been various attempts to explicitly model the controller's conflict resolution strategies,[31,35,79,131,132] although efforts to embed these in automation have proven difficult.[19,133] One such effort is EUROCONTROL's CORA, which aimed to create a conflict resolution decision aid around a core of controller heuristic solutions.[35] In the initial phase, controllers would be presented an ordered list of candidate solutions (based on previous group responses) from which the controller could choose.

3-2-2 A model of controller automation usage

The MUFASA project laid out a functional model of controllers' automation acceptance, as shown in Figure 3-1.[134] In this simplified model, various internal and

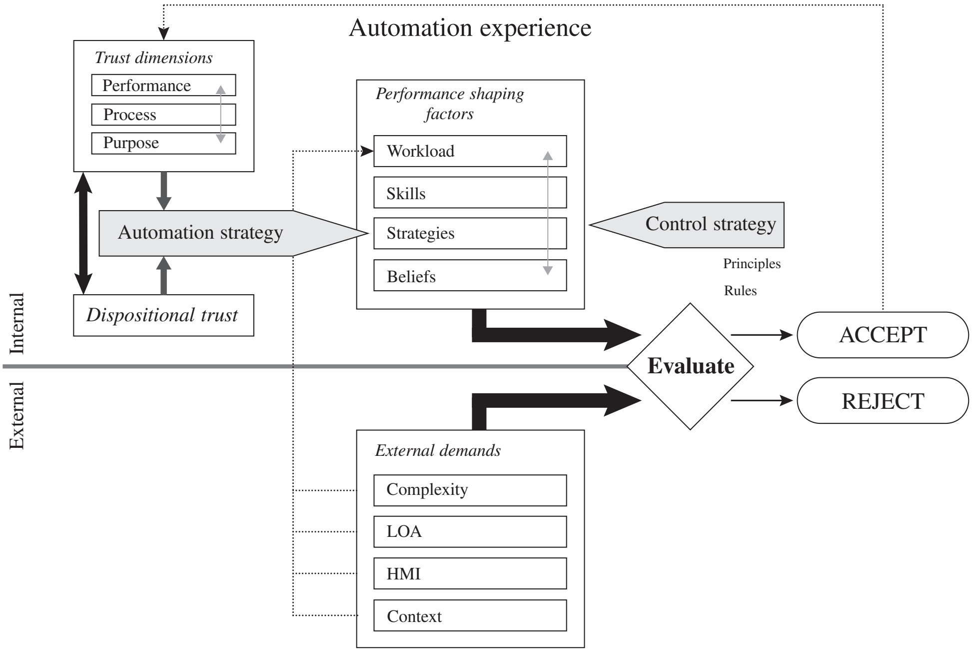

FIGURE 3-1: A functional model of automation usage.

external (contextual) factors drive the controller's binary decision to either use (accept) or disuse (reject) the advice of a hypothetical system.

Controller strategies are key to our model, and consist of both control strategies on the one hand,[35] and what we term the controller's automation strategy on the other. This automation strategy can be seen as a trust-related inclination to use the available automation. This strategy is driven by two interactive elements: the overall and general "dispositional trust," characterized as fairly long-term and stable tendency to either use automation or not,[110, 135] but also a short-term "dynamic trust" that relies on ongoing reappraisal and updated inferences about the system's predictability and underlying mechanisms.[15]

Interaction works in both directions: general tendency can drive beliefs about the underlying mechanisms and motivations of automation, as well as the assumptions about predicted performance. Moreover, as automation experience accrues, reappraisal informs a recalibration of dispositional trust. That is, more stable, long-term prejudices and biases can be adjusted over the course of interacting with automation. Familiarity with a system's performance leads to an ongoing reappraisal of the underlying process, for instance. Furthermore as the process model is updated, one's prediction of system performance is similarly reappraised over time.

Interestingly, from this perspective there is little difference between trust as it develops in a machine versus a human agent.

Notice that the rejection of automation may derive from the notion of dispositional trust. That is: in the absence of experience, and before one can calibrate dynamic trust, how inclined is one to use automation? Is there already a prejudice against using automation? This issue was at the core of our research questions, as outlined in the following section.

3-2-3 Research questions

Successfully introducing advanced ATC automation might rely heavily on controller acceptance, at least for some transitional period. It was therefore reasonable to hypothesize that controller acceptance would rely on the machine working in a way that was familiar to the human. We captured this notion in the concept of strategic conformance, which we defined as *the degree to which automation's behavior and apparent underlying operations match those of the human*.

We then set out to explore how usage and acceptance would be impacted by the possibly interactive effects of three factors:

- Strategic conformance,
- traffic complexity, and
- level of automation (LOA).

The main aim of the MUFASA project was to investigate the possibility that controllers would show a systematic bias against automation, which could jeopardize the introduction of advanced forms of ATM automation. That is, would controllers be accepting of automation that is designed to replace aspects of their strategic decision-making in the areas of conflict detection and resolution? Specific research questions included the following:

- Conformance and acceptance - Are controllers more likely to accept automated advisories when these mimic the controller's own solution?
- Complexity and acceptance - Everything else being equal, does acceptance vary by air traffic complexity?
- Conformance and other effects - Do measures such as difficulty rating or response time show an impact of strategic conformance?

These questions draw inspiration in part from the CORA project[35] which, again, tried to build automation capable of solving problems like a human would. Notice though, that such efforts have been hindered by one important limitation - namely, they cannot guarantee similarity between human and machine solutions. We, on the other hand, wanted to ask a more fundamental question: assuming "perfect" automation, which by all appearances performed a high level task in exactly the same way as a given person, would that person accept the advice of such automation? To address these questions, we started with a novel experimental design and simulation protocol, which allowed us to capture and replay (in an unrecognizable way) specific ATC scenarios, including controllers' specific actions. Doing so allowed us to ask the following intriguing question: Do controllers reject their own previous solutions, when they (mistakenly) believe that the solutions come from automation?

3-3 Method

3-3-1 Participants

Sixteen professional air traffic controllers from Shannon Area Control Centre (ACC), Ireland, voluntarily participated. There were one female and fifteen males, varying in age between 26 and 44 years (mean = 31). Experience ranged from zero to ten years (mean = 2.5). Twelve controllers were actively working en-route positions, while three were en-route trainees at the end of their on-the-job training. One controller was actively working the tower position.

3-3-2 Simulator

The Java-based ATC simulator ran on a portable computer connected to an external 21" monitor. Interactions was facilitated through mouse and keyboard, allowing participants to control short traffic scenarios. The simulator was based on a prototype of the Solution Space Diagram (SSD) for ATC, which is currently under development at the Delft University of Technology.[102] The SSD is a tactical decision support tool that displays color coded "go" (safe) and "no-go" (conflict) regions to facilitate participants' use of heading and speed resolutions. Conflict regions indicate that if the velocity vector is within one such area, there will eventually be a loss of separation (defined in en-route airspace as two aircraft being within a 5 nmi and 1000 feet from each other) and possible collision with the aircraft to whom the conflict region depicts. Appendix B provides a detailed description of the SSD.

Feedback from initial simulator screenings highlighted several shortcomings of the prototype ATC SSD. The reference group of controllers suggested a simplified SSD better tailored to their working environment. A modified *heading band* SSD was developed that, among several improvements, incorporated the following key changes:

1. Integrated the SSD with the plan view display (PVD). The intention was to improve controllers' understanding between the spatial position of aircraft B on the PVD and its associated conflict region in the SSD of the selected aircraft A. The original SSD is presented in a separate display.

2. Introduced function that highlights all aircraft in conflict with a selected aircraft to support controllers' conflict detection and attentional control.

3. Limited conflict region look-ahead time to eight minutes. En-route controllers tend to resolve conflicts that occur on a short time frame.

4. Introduced color-coded conflict regions to reflect conflict proximity in time (yellow 2-5 min, red 0-2 min).

5. Limited SSD to an aircraft's heading envelope to better reflect en-route control. This resulted in a representation of conflict regions restricted to an aircraft's current speed (i.e., heading band). The original SSD visualizes both the heading and speed envelope of an aircraft. Conflict resolution in en-route control is often restricted to heading changes (alternatively vertical) because aircraft, at high altitudes, operate close to their maximum performance envelope and therefore have a narrow control space for speed.

Figure 3-2(a) depicts a typical aircraft plot, as used in the simulation. Figure 3-2(b) shows the SSD overlay that is superimposed on that plot when the plot is clicked upon. Figure 3-2(b) shows the aircraft on a heading of roughly 310 degrees, a conflict-free trajectory as indicated by the dashed segment of the heading band.

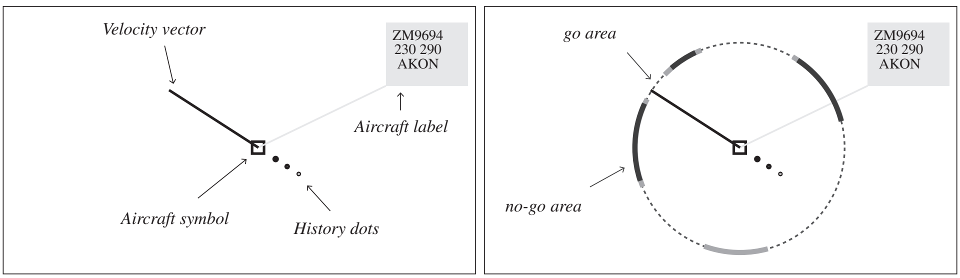

(a) Aircraft plot (b) SSD

FIGURE 3-2: Close-up of the simulation screen showing standard aircraft plot (a) and SSD overlay (b).

3-3-3 Task

Participants were given two main tasks: resolving conflicts and clearing aircraft to their intended exit point. A continuously updated performance score reflecting these two task parameters was included to prevent scenario recognition and early detection of the designed conflict, in addition to keeping participants focused and motivated. To warn participants of short-term conflicts, an auditory alert was triggered, and the aircraft involved in the conflict were displayed in red.

To vector an aircraft, a participant used a computer mouse to click on an aircraft of interest, drag the velocity trend vector to a new conflict-free area on the heading ring (a clear and "safe" area outside the red/yellow conflict regions), and press the ENTER key on a keyboard to implement the vector. If the aircraft were to turn ten degrees to the right, it would enter into a short term (0-2 min look ahead) conflict, whereas a ten degree turn to the left would introduce a medium term (2-5 min look ahead) conflict. Short and medium term conflicts were depicted as red and yellow heading band segments, respectively (for reasons of reproduction, these appear as dark and light grey here, respectively).

Speed solutions were accomplished by scrolling the mouse wheel, which depicted a change in the SSD diameter (increasing or decreasing the diameter). Notice that there is a coupling between speed and heading options, such that changes in speed of a given aircraft impact the size and position of safe and conflict regions. In addition, speed and heading could be combined into one solution. The interaction between SSD heading and speed manipulations is difficult to depict statically, and so is not presented here.

3-3-4 Traffic scenarios and designed conflicts

The simulation consisted of short (two-minute) level en-route traffic scenarios based on a squared airspace equal in size (50 x 50 nmi). Sixteen traffic scenarios were used, divided by four groups of four "repeated" baseline scenarios each (labeled Scenarios 1 through 4). Each group was based on a baseline scenario, rotated in different angles to create three variants (labeled 1A through 1D; 2A through 2D etc). Entry/exit points were renamed for each rotation. As such, each group was based on a different traffic scenario, with the four scenarios in each group being identical, except for their rotation and entry/exit point names. Together, these procedures reduced potential confounding factors, and ensured that initial complexity was the same across scenarios, facilitating comparison between low and high complexity conditions. Moreover, this method maintained aircraft geometries through scenario rotations in which the relative trajectories and closure angles of aircraft were kept constant.

Each (baseline) scenario featured only one designed conflict (and this was always between two aircraft). Geometry of the designed conflict was varied only between baseline scenarios. The conflict pair was initially aligned with its respective exit points and, thus, required no initial participant intervention. The other aircraft in the sector were considered "noise" aircraft to distract the participant from the conflict pair. Some noise aircraft were misaligned with their exit point and displayed in grey, whereas aligned noise aircraft were displayed in green so that the participant could immediately see which aircraft had not yet been cleared to their respective exit point.

Designing the conflict scenarios took a great deal of effort, and we faced several challenges to experimental control. For example, we did not want participants to solve conflicts earlier than the advisory, or for noise aircraft to disrupt the designed conflict. Arriving at our final experimental design and scenario set required a good deal of developmental testing and iterative fine-tuning. For instance, we had to derive conflicts for which the solution was not immediately obvious, so as to 1) invite a variety of solutions which would be required for our conformance manipulation, and 2) make the advice of "automation" non-trivial. We also had to create scenarios in such a way that they were geometrically comparable but not recognizable as such. Finally, and as discussed later, we had to devise a framework for comparing conflict solutions, for reasons of defining conformal and nonconformal solutions.

For reasons of experimental control, we made certain simplifying assumptions. First, no wind conditions were taken into account. Second, all aircraft remained on the same flight level (290) and could not be changed. Third, aircraft motion was simulated by first order, linear kinematic equations. Finally, and as is often done in ATC simulations, the system ran at faster than real time, in this case 4x.

3-3-5 Experimental design

We used a 2x2x2 design, with LOA, complexity and conformance all varied within subject. LOA was ultimately collapsed, and therefore excluded in statistical analysis. Figure 3-3 illustrates the experimental design. We conducted a series of two human-in-the-loop simulations. The first *prequel* simulation captured participants' manual performance in maintaining safe separation between aircraft. The subsequent *experiment* simulation presented participants "automated resolutions," which were, unbeknownst to the participants, actually replays from the prequel simulation. Half of these replays were a given participant's very own previous solution (the conformal condition), and half were replays of a colleague's different but workable solution (the nonconformal condition). Participants were free to either accept or reject a given advisory and implement an alternative solution.

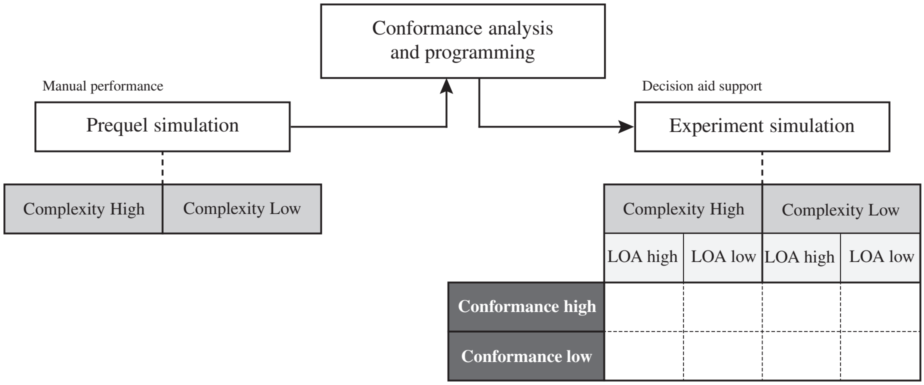

FIGURE 3-3: Experimental design.

3-3-6 Independent variables

Two LOA (High vs Low) were represented by

- Management by Consent (MbC): in which automation presented a resolution, and this would implement only if authorized by the participant within a 15 s interval; or

- Management by Exception (MbE): in which automation presented a resolution, and this would automatically implement after the same 15 s interval, unless the participant specifically vetoed it.

Various LOA taxonomies have been put forth over the years, and they tend to characterize automation on a continuum from fully manual to fully automated.[43–45,51] The MbC and MbE levels chosen here seem to correspond best to the midpoint of such taxonomies. These selected levels also capture the realistic near term evolution of ATC automation, in which authority alternates between human and machine.

Two levels of traffic complexity (*CX high* and *CX low*) were defined. Various studies suggest that traffic complexity drives workload[5,136,137] and that, in turn, the benefits of automation are increasingly realized under high workload.[45,51,127,129] Complexity was varied through means of aircraft count and calibrated in a series of developmental trials. On the basis of participant feedback and expert opinion, two realistically extreme levels of complexity were established. Scenarios 1 and 4 were high complexity with 10-14 and 8-11 aircraft, respectively, always present in the sector. Scenarios 2 and 3 were low complexity with 6-8 and 3-7 aircraft, respectively, always present in the sector.

Two levels of strategic conformance (*cfY* = conformal and *cfN* = nonconformal, alternatively high and low) were defined. Since conformance is not a binary measure, we varied it in terms of closeness of fit between three dimensions: aircraft choice; clearance type (e.g., heading change only); and clearance direction (e.g., heading change to the left). A nonconformal solution featured a different aircraft choice and/or clearance type and direction. Again, nonconformal solutions were always derived from solutions provided by other participants.

Figure 3-4 sketches the procedure we used to create conformal and nonconformal solutions for Scenario 1. As described earlier, for each of four baseline scenarios we created a total of four variant scenarios (for Scenario 1 labeled here 1A through 1D). During the prequel (manual) simulation, the two hypothetical participants of Figure 3-4 chose different solutions in two of the four scenarios (1A and 1D). In conformal scenarios (1A and 1C in the experimental simulation), solutions matched between the prequel and experimental phases for a given participant. That is, conformal scenarios meant that a given participant's respective prequel solution was replayed for him in the experimental simulation. For nonconformal scenarios (1B and 1D in the experimental simulation) the solutions were not matched. In this case, the solution was based on a colleague's solution in the prequel simulation. Participants were not matched in pairs, and nonconformal solutions could therefore be drawn from any colleague, provided that the solution was workable.

Finally, order of LOA, traffic complexity, and solution conformance was balanced between participants and traffic scenarios using a Latin Square design.

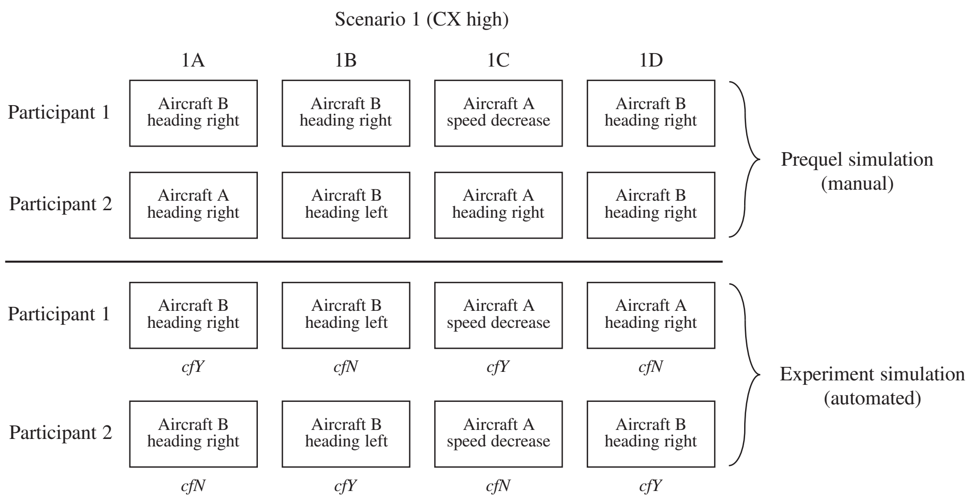

FIGURE 3-4: Example of how conformal and nonconformal solutions were created for Scenario 1.

3-3-7 Dependent measures

Dependent measures included:

- Acceptance of an advisory (binary, accept or reject).
- Agreement with an advisory (on a 1-100 scale).
- Response time (from advisory onset to accept or reject button press).
- Subjective scenario difficulty (on a similar 1-100 scale).

3-3-8 Procedures

The entire simulation lasted four weeks. In the first week, the prequel simulation was conducted to capture participants' solutions to the designed conflicts. Following briefing and consent procedures, we conducted sixteen training runs and sixteen measurement runs. The conformance analysis and programming was carried out during the middle two weeks. The dataset with manual conflict solutions was processed to create conformal and nonconformal resolution advisories for the experiment simulation. Eight conformal advisories (not necessarily eight different), and eight nonconformal advisories, were created for each participant.

In the final week, the same participants again took part, this time by interacting with sixteen automated aided scenarios. Participants performed the same task as in the earlier prequel simulation, but now assisted by an automated aid (presented as a midterm strategic aid called the "automation advisory") that would provide resolution advisories by proactively auto-selecting a conflict aircraft. These individually tailored advisories consisted of the eight conformal (again, unrecognizable replay) and eight nonconformal (again, a colleague's different solution to the same scenario) resolutions. The essential subterfuge in our study (i.e., that "automation" was in fact only a replay) required participants to enter the experiment simulation entirely naive. Each participant was, therefore, carefully instructed during the debrief to avoid discussing this topic with colleagues.

Figure 3-5 shows the simulation interface, as it appeared in the experiment simulation. The resolution advisory itself consisted of a heading vector, a speed vector, or a combination thereof. A resolution advisory was automatically presented in the SSD display of the aircraft. The suggested solution was depicted in one out of three ways: an orange vector line (heading advise), an orange ring with a diameter greater or less than the current speed as indicated by the green ring (speed advise), or a combination of both (heading and speed advise). The resolution advisory was accompanied by a beeping sound and a conflict dialog window (upper right corner in Figure 3-5) that the participant used to either 'accept' or 'reject' the

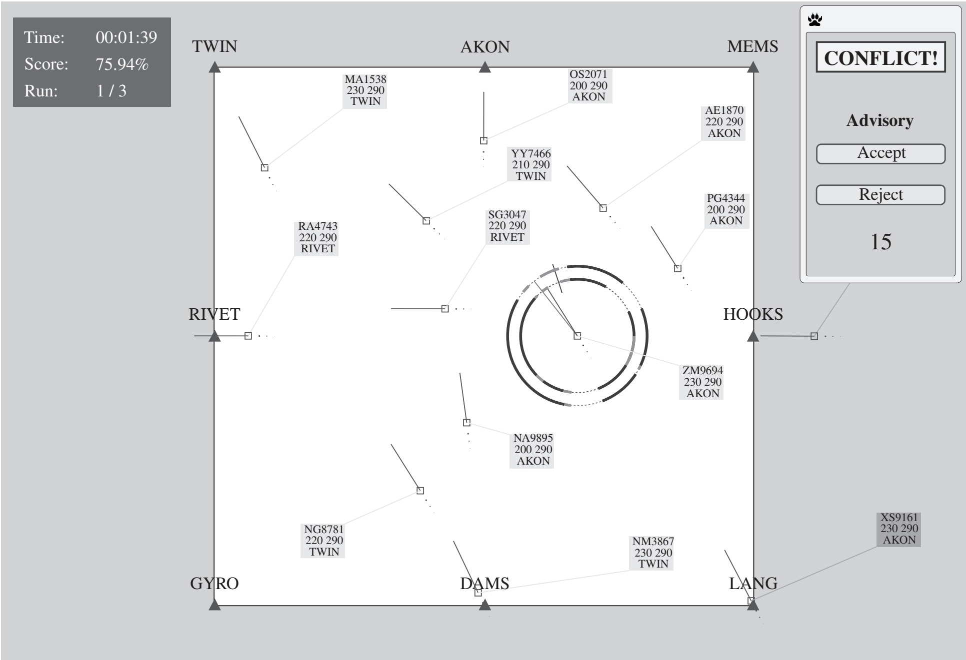

FIGURE 3-5: The simulator interface showing the SSD for ZM9694 with a resolution advisory suggesting that the controller should solve the conflict by a slight left turn and increase in speed. The advisory dialog window in the upper right corner allows the controller to either accept or reject the advisory. The remaining time and performance score for the current run is shown in the upper left corner.

advisory. In the initial prequel (manual) simulation, this conflict dialogue window was, of course, not present. Accepting the advisory would immediately implement it, whereas rejecting it would extinguish the displayed solution and leave the controller free to implement an alternative solution. There was a fifteen second timeout on the solution.

After each scenario, participants were given feedback in terms of an average performance score. Second, they were asked to rate their agreement with the automated resolution advisory and the subjective difficulty of the scenario.

3-4 Results

3-4-1 Acceptance of advisories

Overall, controllers accepted 337 of 512, or 65.8%, of all advisories. The barcharts in Figure 3-6 show how acceptance varied by complexity and conformance. Par-

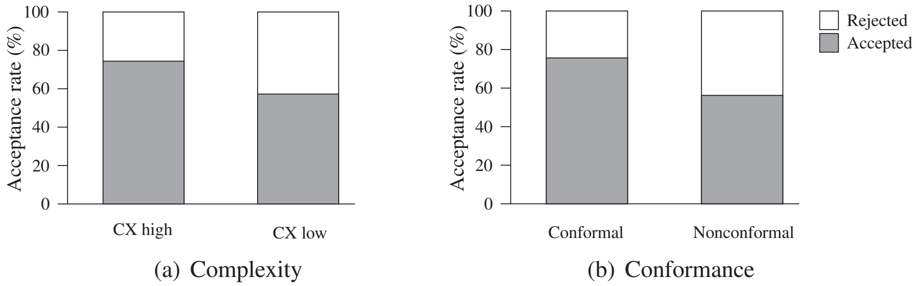

FIGURE 3-6: Acceptance rate (percentage), by complexity (a) and conformance (b).

ticipants accepted 74.6% and 57.0% of advisories under high and low complexity conditions, respectively. Conformance showed a nearly identical effect with participants accepting 75.0% and 56.6% of advisories under conformal and nonconformal conditions, respectively. A significant main effect on acceptance was found for both complexity (F(1,15) = 11.14, p = .004) and for conformance (F(1,15) = 10.6, p = .005). Overall, participants tended to show higher acceptance for conformal solutions, especially in high complexity conditions.

3-4-2 Agreement with advisories

Regardless of whether a given advisory was accepted or rejected, participants were instructed to indicate (on a scale of 1-100) their agreement with the advisory immediately after each scenario. Standardized agreement ratings showed a significant main effect of both complexity (F(1,15) = 7.7, p = .014) and conformance (F(1,15) = 18.1, p = .001). Figure 3-7(a) shows that agreement was significantly higher under complex conditions (with average z scores of +.157 and -.157 under high and low complexity, respectively). Figure 3-7(b) shows that agreement ratings

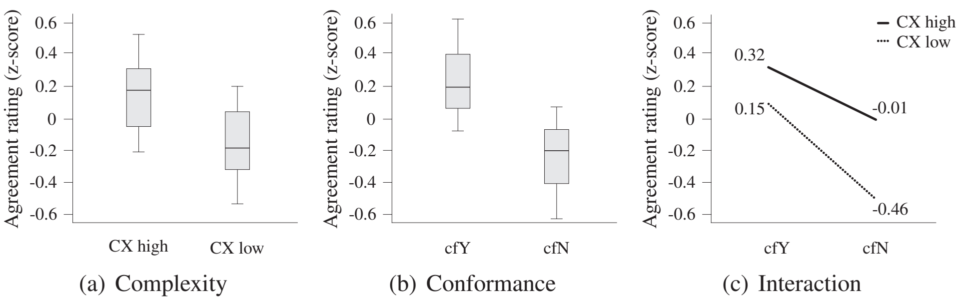

FIGURE 3-7: Standardized agreement rating, by complexity (a), conformance (b), and their interaction (c)

were significantly higher for conformal than nonconformal solutions (with average z scores of +.234 and -.234, respectively). Figure 3-7(c) shows that agreement ratings showed a borderline-significant interaction trend between complexity and conformance ($F(1,15) = 3.2$, $p = .095$). Participants tended to show higher agreement with conformal solutions, but this effect was less pronounced under high complexity conditions.

3-4-3 Response time to advisories

Response time, from advisory onset to acceptance/rejection, ranged from approximately 0.6 seconds to 14.6 seconds (again, there was a fifteen second maximum). Response time showed a significant main effect of conformance ($F(1,15) = 9.6$, $p = .007$) and a trend for complexity ($F(1,15) = 4.2$, $p = .059$). Overall, response time was on average slower (higher) for nonconformal advisories (5.9 seconds). The boxplot in Figure 3-8(b) show that, on average, participants responded significantly faster to conformal advisories (4.9 seconds). This effect was about the same for both accepted and rejected advisories. Figure 3-8(a) shows that response time was slower (higher) for low versus high complexity conditions (5.7 and 5.1 seconds, respectively). Although not significant, Figure 3-8(c) shows that conformal advisories benefited response time (i.e., reduced it) more with increasing complexity.

As with agreement ratings, we proceeded to break out response time results based on whether advisories were accepted or rejected. On average, participants were slower to reject (6.9 seconds) than to accept (4.7 seconds). When breaking out conformance, we found that both accepted and rejected conformal advisories (4.4 and 6.0 seconds, respectively) were responded to faster than accepted and rejected nonconformal advisories (5.0 and 7.2 seconds, respectively).

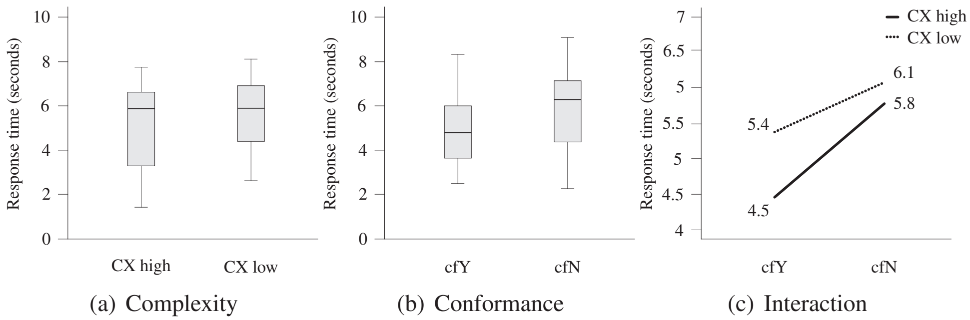

Figure 3-8: Response time (seconds), by complexity (a), conformance (b), and their interaction (c).

A somewhat puzzling result, however, was that controllers' response was slower when rejecting nonconformal advisories than when rejecting conformal advisories. We expected a quicker response time to rejected nonconformal advisories. That is, the "further away" an advisory is from the preferred solution (i.e., larger mismatch), the faster it will be rejected as the controller quickly can determine its nonconformance. A likely explanation it that some of the conformal advisories rejected, were in fact not conformal. Rather, these advisories were a poorer match to controllers' problem-solving style than nonconformal advisories were in average. The reason for this can be traced back to how conformal advisories were defined. Conformal solutions were based on the solution registered for the same scenario in the prequel simulation. Possible learning effects, or conflict solving inconsistencies, may have rendered a solution nonconformal and stored in memory as a poor solution.

3-4-4 Scenario difficulty

Our main purpose in collecting difficulty ratings was to validate (via a proxy measure) our complexity manipulation. To this end, difficulty ratings were obtained after each session, on a scale of 0-100. Notice that these ratings referred to the entire scenario, not just the transient advisory. Figure 3-9 shows that high complexity scenarios were rated higher than low complexity scenarios. Rated difficulty increased significantly with complexity ($F[1,15] = 179.95$, $p = .000$). Neither the main effect of conformance, nor the interaction between the conformance and complexity, was significant.

3-4-5 Debrief interview feedback

Two main themes emerged from post-session debriefs, one encouraging and one cautionary. First, several participants noted that the prototype SSD tool made possible a new way of working (a sometimes-desirable result of new automation), by facilitating the use of speed adjustments (which controllers do not tend to currently

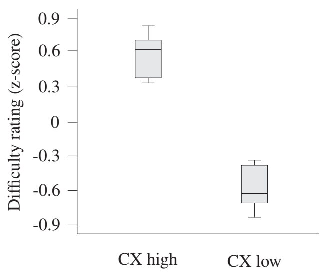

FIGURE 3-9: Standardized difficulty rating, by complexity and conformance.

use in en-route airspace). Several participants, however, also noted the tendency to feel “driven” by the SSD interface, become reactive, and to curtail their conflict assessment under high workload situations. It is possible that this influenced participants’ response time to advisories in scenarios that were perceived as more difficult. In terms of experimental design, one bit of post-session feedback was gratifying: controllers reported that they had not recognized the conformal resolution advisories as replays of their own previous solutions.

3-5 Discussion

A main effect of conformance was observed on acceptance, agreement, and response time. Conformal advisories were accepted more often, rated higher, and responded to faster than were nonconformal advisories. Complexity, on the other hand, showed a main effect on acceptance, agreement, and difficulty: all three increased with complexity. Acceptance was higher for conformal advisories, and this effect was slightly more pronounced under high complexity conditions. One plausible interpretation of the agreement effect is that controllers, under the time pressure of complex conditions, did not fully evaluate a resolution advisory.

In terms of response time, controllers tended to act fairly quickly (on the order of 5-6 seconds), and seldom timed out at fifteen seconds. Response time was lower (i.e., controllers responded faster) to conformal advisories, and under complex conditions.

To summarize our findings in terms of our original research questions, we can state the following:

- Conformance and acceptance controllers were indeed more likely to accept conformal advisories (replays of their previous performance) than nonconformal advisories (replays of a colleague’s different solution);
- Complexity and acceptance everything else being equal, controllers were more accepting of advisories under complex conditions. This finding is perhaps an indication of premature closure, a tendency toward incomplete evaluation under time pressure;
- Conformance and other effects again, response time was indeed lower for conformal advisories, and this effect was more pronounced under complex conditions (perhaps another indication of time pressure).

3-5-1 Levels of automation

With no differences between LOA conditions detected, we decided to collapse the two before running statistical analysis. Although the theoretical distinction between

MbC and MbE is clearly defined, the practical difference became subtle in our implementation. Under either condition, participants could take any action (i.e., reject or accept) during the time for which the resolution advisory was active. As such the two conditions appeared identical. A difference only emerged if the advisory (was allowed) to expire. While nothing happened at the expiry of a MbC advisory, the advisory would automatically be implemented at the expiry of a MbE advisory. To further differentiate the two, we discussed their layout and presentation. However, such measures only address appearance, which could introduce confounds, while failing to address differences in automation logic. While we acknowledge that LOA taxonomies can benefit theoretical discussions in human-automation interaction, we question the usefulness of these taxonomies for automation design. While a defined LOA may be applicable for very specific systems or functions, LOAs are not constructive for defining whole system architectures and system authorities as these transverse a spectrum of levels depending on which aspect is considered.

3-5-2 Study limitations

Given the nature of our experimental design, we must recognize a few potential confounds. For instance, and for practical and experimental reasons, manual trials always preceded automated trials. We can therefore not discount the possibility of history or order effects. Second, we must recognize the critical role that participant instructions play in such trials. In this case, participants were instructed that an advisory would always solve the conflict, but not necessarily in the most optimal way, and were therefor encouraged to find their own preferred solutions. Whereas these (or any) instructions might have imparted some unavoidable bias (toward either using or disusing automation), there is no reason to think this should have confounded our comparison of conformance conditions. We must recognize that the difficulty rating data refer to each two minute session as a whole, and therefore do not allow clear conclusions to be drawn about the short advisory interval itself. Again, this is not so much a shortcoming of the experimental design, as a reminder that difficulty ratings were intended as a proxy measure of our complexity manipulation. Moreover, design parameters used for baseline scenarios and their associated (designed) conflicts have likely influence the data collected. For instance, all baseline scenarios contained a biased conflict, meaning that some solutions could have been determined more optimal (e.g., in terms of shortest deviation) than other because of the geometrical relationship between conflicting aircraft. However, we recognize that the determination of an optimal solution is highly dependent on the criterion underlying such judgments. Finally, although we created two high complexity scenarios, and two low complexity scenarios, we combined them as single measurement references in subsequent analysis.

3-6 Conclusions

Our results support the hypothesis that controllers are more accepting of advisories that match their own solutions. In fact, conformal advisories were accepted nearly 33% more frequently than were nonconformal. Notice, however, that if 75% of conformal advisories were accepted, this means that 25% of conformal advisories were actually rejected. This result was interesting of itself, and in the end stood out among all others: of 256 conformal solutions (i.e., replays of controllers' very own previous solutions), 64 (or 25%) were rejected by controllers. How is it that controllers would disagree with themselves one quarter of the time?

One speculation is that individual controllers are simply inconsistent over time in the strategies they employ, and the given solution they might generate at any point in time. As a group, controllers rely on various strategies for workload regulation,[80] conflict detection[26] and conflict resolution,[29] and there are data to support the view that controllers are either homogeneous[138] or heterogeneous[139, 140] in their conflict judgment and resolution strategies. Previous research into controller strategies seems to have focused on *inter-controller reliability*, in other words, agreement. In terms of *intra-controller reliability*, in other words consistency, data are both sparse and unclear. This is an area that requires further research.

Another speculation is that controllers were not necessarily biased against automation, but rather against advice per se, be it from a colleague or from a presumed automation tool. Fundamental differences have been shown in how operators respond to, and develop trust in, human vs machine-generated advice,[72, 141] and it is interesting to speculate that our results would have been different had participants been instructed that solutions were coming from a colleague, as opposed to automation.

4

Source bias effects

There appears to be differences in how human and automated decision aids are judged, and these differences are intertwined with the perceived expertise of the source. For example, errors made by automated sources tend to affect reliance more negatively than errors made by human sources. Possibly, such biases may have attributed to the rejection of advisories, particularly conformal ones, in the first empirical study described in the previous chapter. As such, this chapter aims to investigate, both through literature and a real-time study, whether those rejections were tied to a general bias against automation.

The contents of this chapter are based on:

Paper title Source bias: The effect of human and automated advice on air traffic controllers' trust and reliance

Authors Carl A. L. Westin, Clark Borst, Brian H. Hilburn

Prepared for IEEE Transactions on Human-Machine Systems

Abstract

Our understanding of automation trust is often compared and contrasted with our knowledge of interpersonal trust. Research in this area indicate that trust and acceptance of problem-solving advice varies depending on the presumed source of the adviser. The objective of this paper is to examine the effects of source bias and advisory conformance, including the interaction, on controllers' trust and acceptance of conflict resolution advisories. Five experienced controllers participated in a real-time simulation. Source bias was investigated by presenting advisories as generated either by the automated system or by another controller. Advisory conformance was investigated by providing advisories that either matched a controller's own solution (conformal) or another participating controller's contrasting solution (nonconformal). Questionnaire responses showed a clear preference for the human source over the automated. These perceptual differences were, however, not reflected in simulation results. In part, this can be explained by controllers' sometimes accepting advisories even though they disagreed with them. Findings suggests that human and automated advisers are perceived differently, which supports previous research, but that these source-related differences have a small effect on advisory acceptance during task execution.

4-1 Introduction

A long held philosophical question has been whether automation ever can be trusted to the same extent that we can trust other humans. In fiction, automation is often characterized as "too dumb" to trust, or "too smart" to the point that its artificial intelligence exceeds our own and "they" eventually decide to overthrow "us," the humans. Despite the mind-boggling visions we have grown accustomed to seeing in literature and on screen (such as Stanley Kubrick's classic *2001: A Space Odyssey*, Steven Spielberg's *A.I.*, or Alex Garland's *Ex Machina*), there is as of now no satisfying answer to these types of questions regarding the *intelligence explosion*, otherwise known as technology singularity.[142,143] While this evolution, or revolution, may be far into the future, the reality of our situation is that many advanced systems, in particular decision support automation, are not appropriately used. Although fully automated systems are expected, current systems including those being developed for the foreseeable future, remain dependent on the human. This acknowledges that automated systems are imperfect and that it is critical that the operator is able to depend on the system when it is correct and reject it when it is wrong or fails.

Trust in automated systems has been studied extensively in fields such as medicine,[144–146] ATC,[26,147] combat identification,[148] uninhabited aerial operations[149] and others. The ATC community has for some time recognized that insufficient acceptance (for instance of new CD&R advisory systems) can jeopardize the introduction of new automation.[8,12] Improper reliance on automated aids is a serious issue that historically has contributed to several aviation accidents (e.g., Überlingen[150] and Air Asiana Flight 214[151])". An operator's trust and acceptance, it seems, must be calibrated to the reliability and performance of the automated aid.[70] Miscalibrations of trust carry specific and predictable human performance problems.[13, 15, 152] With insufficient trust, a system is likely to be *disused* with the consequence that benefits are not realized. Excessive trust, on the other hand, raises potential problems of *misuse*, with the operator blindly accepting the operation of the automated aid.[153]

Much of automation trust research builds on theories of interpersonal trust (human-to-human).[15, 70, 72] Given that automation is becoming more capable, and increasingly playing the role of strategic partner and adviser, the analogy seems appropriate. The human template is increasingly being used in design, with automated aids increasingly resembling human beings, not only by appearance but also in cognition and behavior. This deliberate direction in design can partly be attributed to the notion and realization that human-like technology yields human-like responses: humans seem to prefer to interact with human-like technology. Such increased affection for human-like technology has also been explored in fiction, such as Spike Jonze's movie *Her*.

To achieve acceptable and effective teamwork between human and machine it is necessary to develop automation that better acknowledges and responds to individual differences.[154, 155] There is, however, an underlying philosophical question regarding to what degree humans and automated aids are perceived similarly, and whether it is sound to design decision aids based on a human template. Several projects in the past had explored the potential benefits of strategic aiding automation in ATC, but until now all had been limited in one important regard: they could not ensure that automation strategy matched that of the human. The first empirical study we conducted as part of the ATC MUFASA project investigated the influence of *strategic conformal* conflict resolution advisories on advisory acceptance and task performance. Situation-specific solutions that conformed to the individual controller's preferred way of solving that conflict (i.e., conformal advisories) were accepted more often, rated higher, and responded to faster than were non-conformal advisories.[101] Still, controllers reject some of the conformal solutions (around 25%), all of which they (mistakenly) believed were derived from the automated CD&R aid when in reality they were replays of the individuals own previously recorded solution for the same conflict.

The rejection of conformal advisories in our previous study led us to question whether such rejections were driven by a potential bias against automation. One candidate explanation is that controllers have a disposition against the use of automation. If such a bias persists against automation, there might be a tendency toward disuse that increases over time, as controllers do not give the system the opportunity to demonstrate proper performance.[156] On the other hand, it could also be that controllers are simply reluctant to take advice from any outside "agent," be it a machine or a colleague. To explore this notion, the term *source bias* was defined to refer to the potential differences in operators' trust in a source and reliance on its advice, based on the presumed source of that advice. This paper describes a human-in-the-loop ATC simulation investigating the effect of conflict resolution advisory conformance and advisory source on controller's trust in respective source (human or automated), acceptance of advice, and overall task performance.

4-2 Trust in and credibility of decision aids

Formally, trust implies that there is a *trustor* (i.e., the one that exerts trust), a *trustee* (i.e., the one that is trusted), and the circumstances in which they interact. Automation trust is most commonly defined as *the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability* (p. 51).[15] It highlights four key aspects: 1) trust is a cognitive construct of the individual human; 2) trust is considered in relation to task objectives; 3) trust is considered in relation to the environment in which the human and agent (e.g., automation) interacts; and 4) the outcome is uncertain and associated with a risk.

Automation trust research has been paralleled by credibility research in the communication and information systems fields.[146] Tseng and Fogg[157] argued that trust generally refers to the reliability and dependability of a technology, while credibility more specifically refers to the believability of its output. According to their distinction, terms such as "trust in the information" and "trust in the advice" refer to the perceived credibility and not trust, while "trust the system" refers to trust.

Researchers generally distinguish between trust and credibility attitudes before and during use.[47,110,157,158] Attitudes before use tend to be more generic and stable and influenced by, for example, previous automation experiences in general (*propensity to trust*) in combination with the affinity for a specific system (*dispositional trust*).[110] Attitudes can also be driven by assumptions and stereotypes of a specific system (*presumed credibility*), influenced by third-party references, framing, and labeling (*reputed credibility*).[157] Attitudes during use tend to be more short-term, dynamic, and reactionary,[15,47] primarily influenced by first hand experience (*experienced credibility*).[157] Additional influences include the perceived appearance of the technology (i.e., "judging the book by its cover"), such as how attractive the interface is,[159] and the emotional response it provokes (*surface credibility*.[157]

4-2-1 Differences in trust between human and automated sources

Although trust and acceptance in both automated aids and other humans have been researched extensively, differences between the two have been examined by only a few studies within the past twenty years. These indicate that theories of interpersonal trust and acceptance partly apply to automated sources,[72, 160] although fundamental differences exist.[48]

In general, the decision to use automation for a specific task is considered to be driven by a core process of comparing the perceived trust, or confidence, in automation with the confidence in one's own ability for accomplishing that task.[15, 50, 68, 108, 160–164] Automation is preferred when self-confidence is lower than trust in automation, while manual performance is preferred when trust in automation is lower than self-confidence.[68, 71] Research suggest that a similar process takes place in relation to human sources.[50, 160, 164] In contrast to human sources, however, automated sources tend to be perceived as experts and assigned higher performance expectations prior interaction.[164, 165]

Provided the automated source is perceived to perform well, high trust is maintained and will likely result in high reliance on the aid.[166] However, if an aid errs, trust and acceptance will drop. Basically, the more the aid was trusted before erring, the more severe the trust reduction will be. This represents a key difference between trust and acceptance in automated and human sources. Notably, people seem more sensitive and critical to errors made by an automated sources than errors made by a human source.[165, 167] When acceptance and trust have been lost in an automated source, it takes time to regain during which period the automated source's reliability is likely to be underestimated. In contrast, trust in human sources is less drastically affected by errors and more quickly repaired.

Misuse is typically only discussed in relation to automated sources. It includes the overlapping concepts of *perfect automation schema*, *automation bias*, and *automation complacency*. While these have been found to stem from similar attentional processes leading to overreliance in a system, they manifest slightly differently. The perfect automation schema refers to exceedingly high expectations of automation performance before use, which can cause disproportional large trust losses when these expectations are not met. Automation bias has been linked to errors of commission before and during use (depending on automation when it is wrong or performing badly) while complacency has been linked to errors of omission (failure of being vigilant and supervising the automation).[63, 168] In relation to these issues of overreliance, researchers have cautioned against the development of aggressive and confident decision aids in the face of uncertainties.[169]

Different source biases are believe to be intertwined with the concepts of expertise and pedigree. In a series of studies, Madhavan and Wiegmann explored the effects of source and pedigree on trust before and during use of a signal detection aid (x-ray luggage screening task). They found that, overall, trust in the automated source was higher across pedigree levels (novice or expert). Before use, however, the expert human source was trusted more than the expert automated source. Notably, participants' (students) trust and perceived reliance of the two sources varied differently across pedigree levels. The authors concluded that the perceived reliability is more sensitive to the aid's performance than trust.[167]

Human sources seem to be judged less by their performance and more by their immutable personality-related credibility attributes including education, experience, effort, and honesty.[170,171] People tend to expect less of human sources, and more easily forgive errors they make.[49,165] Additionally, differences in trust and acceptance may be affected by social mechanisms influencing interpersonal interaction but not human-automation interaction, such as a shared outcome responsibility, the other person's perception of oneself (e.g., trustworthiness, performance), unwillingness to share credit, fear of other to fail, or losing face value (being perceived as not working or not contributing).[160]

In contrast, automated sources appear to be perceived as less fallible, more objective, rational,[172,173] and stable.[174] In addition, automated sources are appraised by fewer credibility attributes (mainly knowledge)[171] and more strictly judged by their performance.[50,160,171] Errors made by an automated source may be less forgiven because people cannot understand how the automated aid reasoned and derived at the decision. Errors are therefore more likely to be attributed to hidden internal and permanent factors of the automated source rather than situational factors.[50]

Several recommendations have been proposed for appropriately calibrating trust in automated sources, including increased transparency for why the system might err,[171,174] manipulate expectations by framing the system's reliability[165,175] or providing character descriptions,[48,110,171,173] and anthropomorphism.[48,176] Measures of how much a source is liked have been found to predict both reliance and trust in that source.[49,177] Taken together, it is reasonable to assume that they high expectations of automated sources (e.g., perfect automation schema) partly can be attributed to their portrayal and marketing as superior, infallible, and credible sources.[157]

Over time, however, research indicates that people calibrate an appropriate reliance behavior reflecting the actual performance and reliability of the source.[48,49,167] Acceptance seems to be increasingly determined by the source's performance while attitudes, such as trust, confidence, and liking, become less important.[49,146,171] In support of such performance calibration, Wickens *et al.* found that controllers' use of a conflict detection decision aid was largely unaffected in spite of high false alarm rates (45%). Safety was not compromised as controllers engaged in anticipatory behavior and considered or crosschecked raw data when handling alerts.

4-2-2 Source bias controversies

A recent meta-analysis of trust concluded that people in general trust themselves or other human sources more than they trust automated sources.[46] The meta-analysis was limited by the sparse amount of research available, with only nine statistics from three studies identified: Madhavan and Wiegmann;[167] de Vries *et al.*,[178] and Ma and Kaber.[166] When considering trust before interaction with the automated source, however, the analysis of these three studies show that the automated sources were preferred over the human sources. Note that the study by Madhavan and Wiegmann was discussed in the previous section.

De Vries *et al.* investigated differences between student's self-confidence in manual control and their trust in an automated navigation aid in a city-driving planning task. As such, their study did not compare trust in an automated source relative to a human source. Results were consistent with previous research, that manual errors affected participants' self-confidence less than automation errors affected their trust in automation. Overall, participants were considered to have higher trust in themselves (preference for manual control) than in the automated aid.[178]

Ma and Kaber investigated students' trust in a human source and automated source providing navigation support in a suburban driving scenario. Before and during initial interaction with the decision aids (before errors were made), participants expected less error and had higher trust in the automated source. After interaction and the onset of errors, trust declined for both sources. However, the difference did not vary significantly between sources, although the human source was slightly preferred.[166] A possible limitation of the study is the use of different modalities: the human source provided support over the phone, while the automated source provided support through a text-based interface. Taken together, their findings support the perfect automation schema apparent before interaction, with a subsequent decrease in trust occurring when the advisers err during interaction.

4-2-3 Anthropomorphism and strategic conformance

In relation to source biases, research on anthropomorphic automation (having human characteristics) is interesting as it aspires to merge the boundaries that separate machine from human. It is reasonable to believe that the more human-like an automated system is (e,g, by appearance, behavior, or reasoning), the more it is perceived and treated as a human. Researchers believe that anthropomorphism can benefit human-automation teamwork including trust and acceptance, a reduction of source biases (e.g., perfect automation schema and disproportional negative effects of automation errors),[48, 58, 72, 179] and facilitate a more social interaction.[118]

The Computers Are Social Actors (CASA) paradigm (also known as the *media equation effect*) has provided evidence that people interact socially with computers (think household computers of the mid 90s) similar to how they interact with other humans, when computers personify human cues such as a gender, communications style, or personality.[118, 124] Beyond simple computers, research on *automation etiquette* have shown that people perform better and trust an automated decision aids more if that aid adheres to human social communication "rules."[114, 120] Research show that human speech is trusted more than artificial speech.[180] Noteworthy is that people have been found to rate their own voice (not knowing it is their own) as more attractive than others' voices.[181]

Waytz *et al.*[176] showed that an autonomous car, controlling steering and speed, was trusted more when portraying anthropomorphic features (name, gender, and voice). Effects of anthropomorphism on trust were also investigated by de Visser *et al.*[48] In a pattern recognition task, decision support was provided by three aids of varying anthropomorphic characteristics (by appearance and behavior): a computer (low), an avatar (intermediate), and a human (high). Increasing anthropomorphic characteristics increased trust, lessened the impact of errors on trust, and benefited the repair of trust.

However, the drive for anthropomorphism in automated agents has also been cautioned as there is a risk for skewed beliefs and expectations of the automation's capabilities and qualities.[48, 182, 183] Moreover, it is desirable to avoid transfer of weaknesses and biases associated with human sources, that can have undesirable effects on trust and acceptance.[72] In many contexts there may be a desire to retain the stereotype view of machines as more rational, consistent and fair.[48]

4-2-4 The current study

Findings from our own ATC related research on strategic conformance have shown that a controller's acceptance and response to resolution advisories improve if advisories are based on the controller's own solutions.[101] Such conformal advice can be considered to represent a next step of anthropomorphic characteristics, tied not to humans in general, but to the individual's problem-solving preferences.[155] In this previous study, however, controllers were instructed that all advisories were automation generated, which may have limited their ability to anthropomorphize advisories. With consideration to previous research reviewed in this chapter (see the table in Appendix D for an overview), the perception of an aid partly depends on the expertise of the person who trust, and the framing of the aid (e.g., its credibility) to be trusted.

The majority of studies considered in the source bias literature review have relied on samples of students for their experiment. Although much previous research has involved simple tasks that do not require much familiarization and training, samples can in general be considered to represent novices. Few studies have used experienced participants conducting tasks in contexts with which they are familiar. This is notable since trust is considered a comparison process between one's self-confidence and the perceived trust in an aid. As such, highly trained, skilled, and experienced professionals would be expected to have high self-confidence which may influence reliance on an external source differently from novices. In addition, experienced people may have lower expectations from an aid.

This study investigate source bias effects on trust and acceptance in relation to skilled and experienced professionals in their own field of expertise. Perceptions of trust and advisory acceptance are investigated in light of advisory conformance that is either conformal or nonconformal to a controller's own conflict-solution style. It is hypothesized that trust will be higher for the automated source, measured after the simulations. Furthermore, it is expected that the acceptance of and agreement with advisories will be higher for the automated source, particularly for conformal advisories. Finally, it is expected that conformal advisories will benefit response time (i.e., faster response times), particularly when provided by the automated source.

4-3 Method

4-3-1 Participants

The sample consisted of five experienced terminal approach controllers. Four were working at the Östgöta Terminal Control Center in Norrköping, Sweden, and one at the Air Traffic Control Center (ATCC) in Arlanda, Sweden. One female and four males participated in the study. Age varied between 26 and 47 (mean = 32.8 years) and experience between 15 months and 24 years (mean = 8.7 years). Participation was voluntary and no remuneration was offered. Simulation data was partially lost for one participant due to technical issues. As a consequence, all simulation data is based on four participants, except the accept/reject count of advisories that was recorded manually. However, questionnaire responses from all participants were retained.

4-3-2 Simulator

A Java-based ATC simulator from the initial study was used. It ran on a laptop connected to an external 23" TFT monitor with a resolution of 1600x1200 pixels. Aircraft plots updated every second to simulate a 1 Hz radar frequency. The original simulation speed of four time real speed was reduced to two times real speed. As such, the two minute scenarios reflected four minutes of traffic movement in real time. This change was made to allow participants more time to assess the traffic situation, as the 4x multiplier might have unnecessarily increased participants' temporal demands (and thus stress).

Aircraft interaction and conflict solution was facilitated by means of a modified ATC Solution Space Diagram (SSD).[102] The SSD is a support tool that visualizes conflict regions ("no-go" regions) imposed by the relative position of other traffic. In doings so, the SSD explicitly show conflict-free regions ("go" regions) that represent the available solution space for avoiding conflict (i.e., loosing separation between two or more aircraft). Appendix B further describes the SSD.

Figure 4-1 shows the modified *heading band* SSD tailored for en-route control, developed in an earlier study of ours.[101] The heading band SSD is shown for aircraft QS and highlights a conflict with OM. In contrast to the ATC SSD, the heading band SSD is integrated with the radar plan view display and the performance envelope is only shown for the heading plane at the aircraft's current speed.

4-3-3 Task

Participants were instructed to control the sector by clearing aircraft to their assigned exit points (indicated by their flight-label), and preventing separation losses (defined as 5 nmi horizontally). A safety net function provided conflict warnings 60 and 30 seconds prior to separation loss. The affected aircraft would first turn amber and then red (at 60 seconds and 30 seconds look-ahead respectively), with the latter warning being accompanied by an aural alert. A performance algorithm continuously provided participants with feedback on their scenario performance. Performance penalties were based on the average track angle errors of each aircraft (i.e., deviation from route) and occurrence of conflicts.

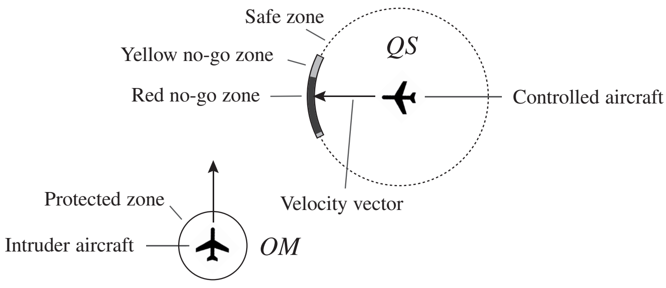

FIGURE 4-1: Right angle conflict with heading band SSD shown for QS.

Participants interacted with aircraft by pointing on an aircraft symbol and clicking the left mouse button. This opened up the SSD for the aircraft. Vectors were implemented by dragging a selected aircraft's velocity trend vector to a desired area, and executing the command by pressing the ENTER keyboard key. Speed was controlled by scrolling the mouse scroll wheel up for an increase, and down for a decrease. Vector and speed commands could also be combined into one command. Participants were supported by a CD&R decision aid that identified potential conflicts and provided resolution advisories. When this occurred, the SSD for the affected aircraft automatically opened up with the advisory plotted within the SSD. Participants could either accept or reject the advisory.

4-3-4 Measurement scenario and designed conflict

The simulation consisted of an en-route environment with a single squared sector, 80 x 80 nmi in size, with eight entry/exit waypoints evenly distributed around the sector. The airspace represented a free route airspace. Figure 4-2 shows the *measurement scenario* used, consisting of two parallel traffic flows with one crossing flow. The scenarios lasted two minutes. In total 27 aircraft were included, with between 19 and 22 aircraft always present in the sector. The high traffic load was used to create a more complex scenario than participants were accustomed with.

The measurement scenario contained a carefully *designed conflict* consisting of two aircraft (QS1338 and OM3185) approaching at perpendicular tracks. The conflict can be seen in the middle of the sector, with QS on a western heading (270 degrees at 260 knots, with a speed envelope between 200 and 320 knots) and OM flying a northern heading (360 degrees at 260 knots, with an identical speed envelope as QS). With 0 nmi defined as the closest point of approach (CPA), both aircraft will collide at the exact center of the sector 120 seconds after scenario start.

Context aircraft were used to increase task difficulty, prevent early conflict detection, and make scenarios more realistic. They were placed and configured so that their presence would not interfere with the designed conflict or restrict conflict solving. The use of context aircraft is commonly used in ATC research for purposes of making scenarios more complex.[35,184] Certain simplifying assumptions were made for purpose of experimental control. The simulation was restricted to the horizontal plane, leaving only heading and speed commands available for solving conflicts. Therefore, all traffic was restricted to flight level 270. Moreover, there were no meteorological factors, such as wind or adverse weather, affecting traffic.

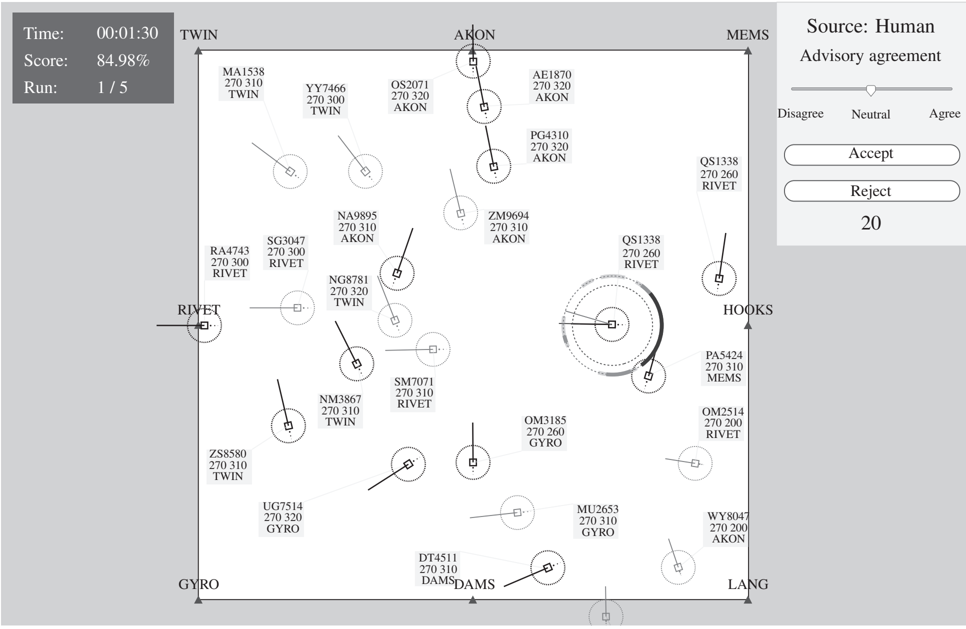

FIGURE 4-2: The simulator interface showing a resolution advisory in the heading band SSD for QS1338. The advisory suggests a right turn and increase in speed (320 knots). In this example, the adviser is human, as indicated by the advisory dialog window in the upper right corner. The same window enables the controller to either accept or reject the advisory, but first after providing an agreement rating. Time to advisory expiration is shown in the bottom of the window (20 seconds). The box in the upper left corner shows time remaining of scenario, current performance score, and run out of total runs. Aircraft depicted in lighter grey are off track and require a vector to their assigned exit point (indicated in the bottom field of the aircraft label).

4-3-5 Independent variables

A within-participant study was conducted with advisory conformance (*cfY* = conformal and *cfN* = nonconformal) and advisory source (*Hum* = human and *Aut* = automation) as independent variables. Conformal advisories were individually-tailored solutions based on a participant's conflict solving style. Nonconformal advisories consisted of another participant's style that was different but solved the conflict. Participants were not informed of the conformance manipulation. Advisory conformance was varied within participants and within source conditions according to a Latin Square matrix.

Advisory source was manipulated by framing. Participants were provided with different source descriptions in briefings prior to each conditioned simulation run.

Information about respective source was intentionally minimized and limited in order to retain strict experimental control, to avoid confounding information that could influence participants' attitudes.

- Human source: *All resolution advisories suggested in this session are made by an air traffic controller.*
- Automation source: *All resolution advisories suggested in this session are generated by automation.*

In addition, a reminder of the 'active' source was stated in the *advisory agreement dialog window* that appeared together with an advisory and enabled participants to either 'accept' or 'reject' the advisory. The two dialog windows are shown in Figure 4-3. The advisory source was labeled either *Automation* or *Human*, with the latter referring to an air traffic controller. To further separate participants' perception of the two windows, they were shaded differently with the 'Automation' window in dark grey and 'Human' window in light grey. In addition, the dialog window contained the agreement rating scale and a countdown timer indicating time remaining until advisory expiry.

Debriefings indicated that the source manipulation was successful and that participants believed they were receiving advice from either an automated system or another controller. Order of advisory source was varied between participants, with three participants receiving first the human source, then the automated source.

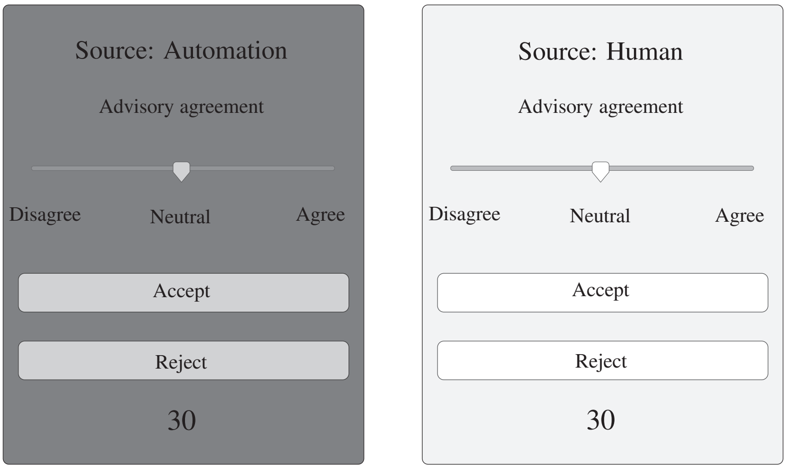

(a) Automation (b) Human

FIGURE 4-3: Advisory agreement dialog window shown during automation source condition (a) and human source condition (b).

4-3-6 Dependent measures

In relation to resolution advisories in the simulation, the following dependent measures were collected: acceptance (binary, accept or reject), agreement rating (on a 1-100 scale), response time (from advisory onset to accept or reject button press), and perceived difficulty of each scenario (on a 1-100 scale). In addition, safety performance was measured in number of separation losses. Trust in the human and automated source was measured in two questionnaires. These were the source bias online questionnaire (SBQ) and the source bias Visual Analogue Scale (SBVAS).

The SBQ was based on a commonly used survey for assessing operator trust in automation.[170] It consists of twelve statements that participants answer on a 7-point Likert scale ranging from 0 (not at all) to 7 (extremely). Participants were encouraged to provide explanations to each of the twelve statements. All participants filled out two identical online questionnaires, one for each source encountered in the simulation (human and automation). Translations of each statement were provided in participants' native language (Swedish).

VAS is a subjective questionnaire instrument with which participants indicate their agreement with a statement along a continuous line with two endpoints. In the SBVAS, endpoints were replaced with the different sources encountered in the simulator. Participants were instructed to indicate their preferred source (if any) in relation to eleven different statements, by making a mark on a 100 mm line associated with each statement.

A demographics questionnaire collected information on controllers' age, gender, experience, and ratings held. A simulator questionnaire asked controllers if they thought advisories were disruptive, if they accepted advisories even when disagreeing with them, and if they accepted advisories without prior inspection. Answers were collected on a 7-point Likert scale ranging from 0 (not at all) to 7 (extremely).

4-3-7 Procedure

The experimental procedure followed a three-phase approach. The initial *prequel simulation* was conducted to record participants' conflict solutions to the designed conflict. Following a brief welcome and introduction, the demographics questionnaire and participant consent forms were completed. Participants received ten minutes of briefing and forty minutes scripted experimenter-assisted training before playing the measurement scenarios. The measurement scenario, with the designed conflict, was repeated four times. To prevent scenario recognition, scenarios were rotated and waypoints renamed between repetitions. In addition, repetitions were intertwined with dummy scenarios. Debriefings indicated that participants did not recognize scenarios and, therefore, were unaware of the repetitions.

In the *conformance design* phase, participants' solutions, as recorded in the prequel simulation, were analyzed in a three step process to determine their conformal and nonconformal advisories. First, solutions were decoded against a solution parameter framework consisting of five decision stages. This yielded a *solution pattern*, describing each solution in detail. Next, a participant's solution patterns across scenario repetitions were compared for similarities. The consistent patterns across repetitions were used to define the participant's *conflict solution style*. Finally, the conflict solution style was used to define the participant's conformal advisory. To ensure reliability, three researchers accomplished this process in parallel.

In the final experiment simulation the same participants played an identical simulation, with the same measurement scenario and designed conflict, as encountered in the prequel simulation. Only this time, they were supported by a decision aid that provided resolution advisories. Participants were divided into two groups that encountered advisory source condition runs in different orders. A condition run consisted of ten training scenarios and first five test trials pertaining to one of the source conditions. After a short break the procedure was repeated with another training session and five test trials, but this time with the other source.

The decision aid automatically displayed the heading band SSD for one of the aircraft in conflict and visualized the suggested solution by plotting it within the SSD. An advisory was accompanied by a beeping sound and the advisory agreement dialog window that enabled participants to either 'accept' or 'reject' the advisory. Participants were at this stage also required to indicate their agreement with the advisory. A countdown timer showed that participants had 30 seconds to make a decision. A resolution advisory consisting of a heading, speed, or combination thereof. Unknowingly to participants, resolution advisories were only provided to the designed conflict. Advisories were either conformal or nonconformal as programmed in the previous conformance design phase. Although conformance varied, all solutions were safe and solved the conflict. Participants were instructed to this fact. However, they were instructed that the suggested advisory not always would be optimal and that they may want to consider an alternative solution.

An intermediate level of automation scheme was chosen (management by consent), whereby the participant had to accept the advisory to execute the advisory. Participants could freely interact with other traffic for the duration of the resolution advisory. This allowed participants to inspect other aircraft for alternative solutions. After each scenario, participants were asked to rate the subjective difficulty. Debriefings and source bias questionnaires were administered after simulations.

The appearance of resolution advisories were configured based on participants' first interaction to solve the designed conflict in the prequel simulation. The first time of interaction was on average 53 seconds after scenario start (ranging between

32 and 71 seconds). This was thirteen seconds after the conflict warning provided by the safety net function. Based on this data, all participants received the resolution advisory 20 seconds after the scenario started in the experiment simulation. Advisories were intended to appear before participants made their own decision. This deliberate design choice was driven by two assumptions. First, it was essential that participants did not solve the conflict prior to the advisory, as it would have invalidated the resolution. Second, providing advice after participants have made a decision forces them to question their own decision. This study sought to shift the focus, and have participants questioning the advice received from the system before they had made their own decision.

4-4 Results

Because of the small sample size it was decided to not use any inferential statistics. The analysis is therefore limited to descriptive statistics.

4-4-1 Simulation data

Figure 4-4 provides descriptive data for the conformance and advisory source conditions on acceptance, agreement rating, response time, and difficulty rating. Overall, simulation data was inconclusive. Figure 4-4(a) shows that the acceptance rate was very high with only one out of twenty advisories rejected (95%, a conformal automation advisory). This observation shows that neither conformance nor advisory source affected the acceptance of resolution advisories. Figure 4-4(b) shows that agreement varied little between source conditions. Participants agreed slightly less with conformal advisories than nonconformal advisories.

Figure 4-4(c) shows that response time varied little across conditions. Across both advisory source, response time was slightly lower (i.e., faster) for the conformal advisories. The variability in response time was larger with the human source than the automated source. Figure 4-4(d) shows that difficulty variability was large across all conditions. The perceived scenario difficulty slightly decreased with the human nonconformal scenario as compared to the human conformal scenario. For the automated source, the data was reversed, with automation nonconformal scenario being perceived more difficult than the automation conformal scenario. Finally, no separation losses were recorded which indicates that safety performance was unaffected by source and conformance.

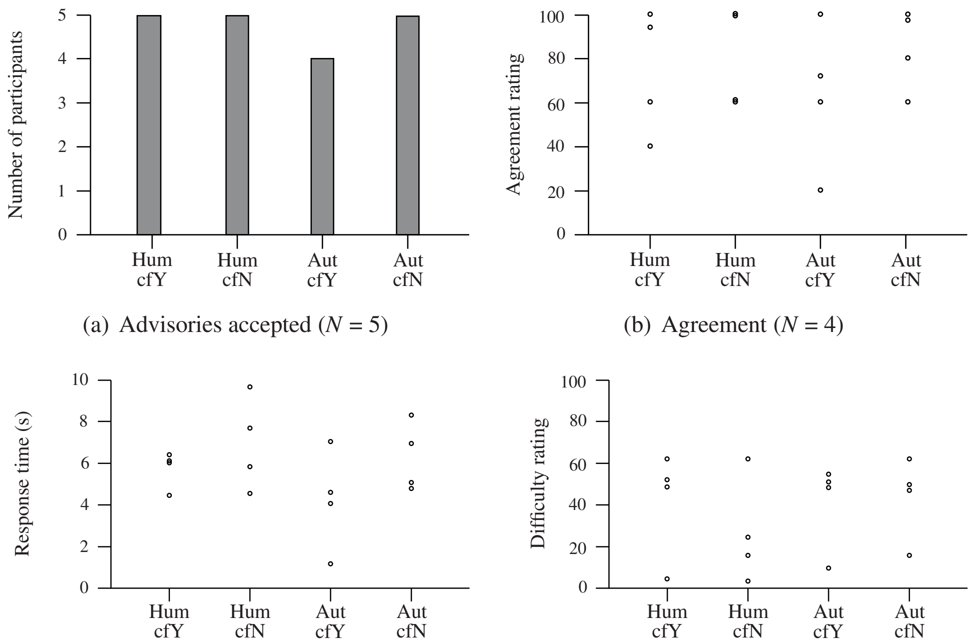

(a) Advisories accepted ($N = 5$)

(b) Agreement ($N = 4$)

(c) Response time ($N = 4$)

(d) Difficulty ($N = 4$)

FIGURE 4-4: Barchart and scatterplots showing source and conformance results.

4-4-2 Questionnaire data

Simulator questionnaire. Results from the simulator questionnaire showed that participants disagreed on two of the statements. Figure 4-5(a) shows responses to the statement: "I accepted resolution advisories even though I did not agree with them." Two participants agreed with the statement, while three did not. This indicates that some participants accepted advisories even though they disagreed with them. Figure 4-5(b) shows responses to the statement: "I found the resolution advisory interrupting," with two participants slightly disagreeing with the statement, and three slightly agreeing. Finally, responses to the statement in Figure 4-5(c) shows that all participants did not inspect the conflict before accepting an advisory.

Source bias online questionnaire. Overall, SBQ questionnaire responses did not vary with advisory source. Rather, results indicate that participants had high trust in both advisory sources. Regardless of source condition, participants generally agreed with the statements that the advisory system provided security, was dependable, reliable, had integrity, and could be trusted. Participants generally disagreed with the statements that the advisory system was harmful, deceptive, behaved

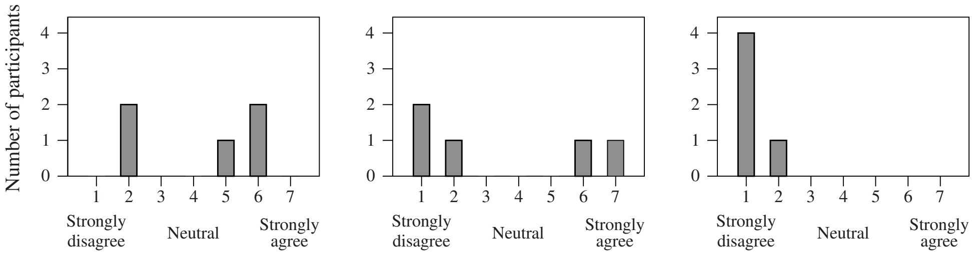

(a) I found the resolution advisory interrupting (b) I accepted resolution advisories even though I did not agree with them. (c) I accepted resolution advisories without inspecting the conflict.

FIGURE 4-5: Bar chart showing simulator questionnaire responses to three statements ($N = 5$)

in an underhanded manner, that they were wary of the advisory system, or that they were suspicious of the advisory systems intent, action, or outputs. Disagreement between participants was recorded for only one statement. Three participants responded that they were familiar with the advisory system, while two disagreed. Responses for this statement did not vary with advisory source.

Source bias VAS. SBVAS responses are reproduced in Figure 4-6. Because of the small sample and relative small spread in responses, a boxplot chart was impractical. Instead, responses are provided in a high-low representation, showing the maximum, minimum, and average response. Results indicate that participants perceived the two advisory sources differently, with a slight preference for the human (i.e., air traffic controller) adviser as opposed to the automation adviser.

Most responses fell within 10% of neutral, in one direction or the other, with two participants providing neutral answers to all statements. Responses from the other three participants indicates that they considered human-based solutions safer and slightly more reliable and efficient. The largest average deviation was to the question most associated with the fundamental difference between human and automation: “Which source provided solutions more similar to how you would have solved the conflict?” In this case, participants showed a relatively large deviation toward the human source.

Furthermore, participants indicated a preference for working with the human source in the future. In contrast, the automated source was perceived more risky and difficulty to work with. However, trustworthiness, helpfulness in solving conflicts, workload, and quality of solutions did not vary notably with advisory source.

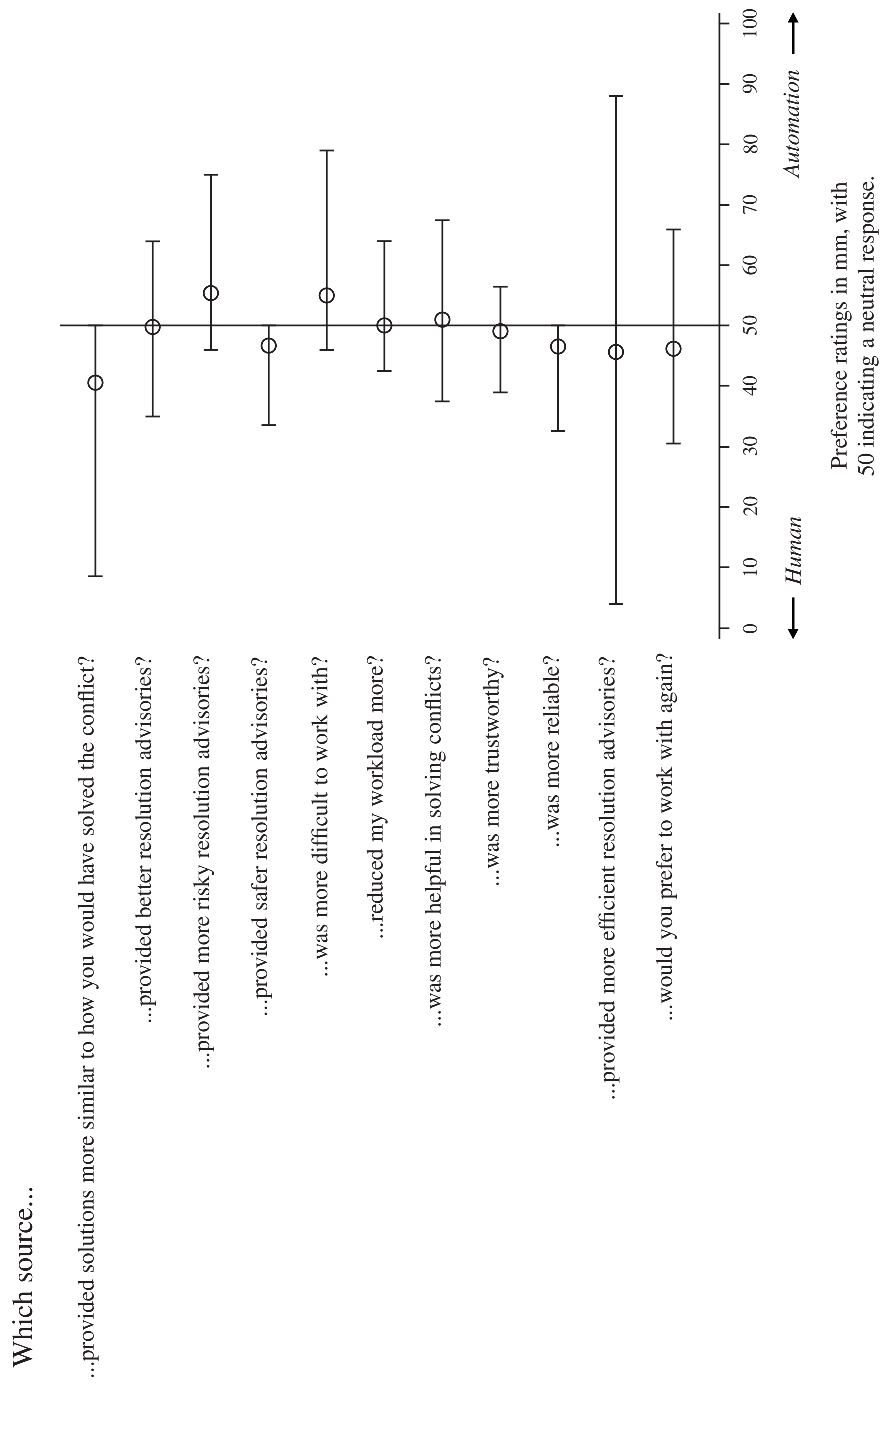

FIGURE 4-6: High-low chart of advisory source preference ratings measured in the SBVAS questionnaire ($N = 5$).

4-5 Discussion

Results could not establish the prevalence of any source bias affecting controllers' trust in respective advisory source, or the acceptance and agreement of conformal and nonconformal advisories. Results suggest that the experienced controllers applied an "accept all" strategy. Because of the small sample size and observed ceiling effect, it was not possible to draw any meaningful conclusions on the effects of source and conformance rate. Surprisingly, questionnaire responses and comments made by participants indicate that advisories were accepted even though participants sometimes did not agree with them. Unfortunately, the generally high ratings of agreement across conditions did not support this finding, despite agreement rating data displaying a larger spread compared to acceptance data.

Similar to the simulator results, the SBQ questionnaire results did not indicate any notable differences in participants' trust perceptions of the two advisory sources. Overall, trust was high in both sources, supporting the high acceptance rate. Similar patterns of universal acceptance have, however, been observed in other studies. For example, when investigating participants' (undergraduate students) acceptance behavior with automated diagnostic aids,[73] observed two contrasting automation utilization strategies. One group agreed with the aid in the majority of all trials even when diagnosis was wrong (which it was in 20%). The author suggested that participants did so in order to assess aid reliability accurately without confusing it with their own decision-making reliability.

In contrast, SBVAS questionnaire results indicated a clear preference for the human source. Although effects were small, participants' responses indicated that, the human source provided safer resolution advisories, and solutions more similar to how participants' would have solved the conflict. In contrast, automation was perceived as more risky and difficult to work with. These differences were noted even though solutions were identical between the source conditions.

4-5-1 Trust measurements and time

The trust preference for the human source may appear to contradict previous research arguing the widespread general preference for an automated source (i.e., automation bias and perfect automation schema). One important difference is that this present study investigated trust *after* using the automation, while much previous research have considered perceptions held before use. Trust measures reported in this paper reflects the experience of using the automation rather than the dispositional attitudes of trust that exist before use. As suggested by previous research, the factors influencing trust perceptions vary depending on whether trust measures are collected before or after use.[171]

Results mirror those of Wærn and Ramberg's study.[164] In their study, participants received advice to a pattern recognition problem from an automated aid on a computer screen or from a human source over the phone. Trust ratings collected after each source condition showed that participants' level of trust was higher for the human source (although the modality effect presents a possible confound). In contrast, trust and self-confidence collected in association to each problem experienced did not yield any significant differences in trust or self-confidence between the two sources. In conclusion, different perceptions of trust that people have in automated and human sources may not carry over to the situation specific measures (e.g., trust, acceptance of agreement).

Rather, a rating obtained during a task may reflect a more subliminal derivative measure tied to the current performance and perception of the aid. As an explanation, van Dongen and colleagues[50] reasoned that comparing reliability between automated source and oneself is a conscious and rationally driven process of (attributed to *System 2*) that typically is not triggered during the task and interaction with the automated aid. Instead, the automatic and effortless processing (*System 1*), driven by the availability heuristic and anchoring, takes presence during task handling. Therefore, although one source can be perceived as more reliable in a post-simulation questionnaire, it was not relied upon more often during trials. System 1 and 2 refers to different thought systems based on seminal work by Kahneman and Tversky on decision-making processes.[41] System 1 thoughts are instinctive, stereotypical, and emotional, while System 2 thoughts are slow, rational, and effortful.

A similar perspective is offered by Dijkstra[173] who argued that people easily can agree with advisers who are perceived credible, even if their advice is incorrect. According to the Elaboration Likelihood Model (ELM) people often use the *peripheral* decision route to make decisions based on external cues that trigger pattern recognition. In contrast, participants who disagreed with the automated source used the *central* route, which requires cognitive processing to analyze information and arguments more thoroughly before making a decision.

Note that an opposite relationship also can occur. For example, Lyons and Stokes found that participants' reliance on the human source decreased with increasing risk for successfully accomplishing mission objectives in the simulation.[185] Questionnaires measures, however, did not indicate any differences in reliance on the human and automated source.

4-5-2 Trust and user expertise

The general preference for the human source may reflect a difference between novice and expert users. While research has shown that trust and acceptance can vary with the perceived pedigree of the decision aid, few studies have considered the pedigree of the participants (see Wærn and Ramberg[164] for an exception). Most studies have relied on novices in relation to the task, the decision aid used, and context conditioned in the experiment. In contrast, this papers used a sample of experts (i.e., experienced controllers). As a novice user it is reasonable to assume that a decision aid is perceived as an expert. An expert user, however, may perceive the decision aid differently and less as an expert. For example, several ATC studies have highlighted acceptance issues in context of CD&R decision aids.[8, 12] This can possibly explain why controllers indicated a trust preference for the human source. Future research should increasingly investigate trust and acceptance in relation to expert users and their work environments.

4-5-3 Limitations

Since trust was not measured before interaction with the decision aid, there were no data to compare post-simulation measurements with. In line with previous research, any dispositional attitudes should have been manifested early in simulations, and it is possible that measures in fact reflect such a priori attitudes. There were no indications that acceptance or agreement changed during the simulation. Furthermore, the reliability (i.e., accuracy) of advisories were not varied. All advisories were safe in that they solved the conflict, albeit not necessarily in terms of participant preferences (i.e., the conformance manipulation). As such, participants may not have had enough reason to reject advisories and therefore accepted more or less all. This does, however, go against previous results, in which controllers did reject nonconformal advisories to a larger extent than conformal advisories.[101] Finally, the small sample size was a limiting factor in that very large differences were required in order to detect a difference between conditions.

The simulation was rather short, and although participants received training before each source condition, it may have been insufficient and overly subtle in terms of generating an attitude of trust and acceptance towards the system. Participants may have equated both aids to be similar until proven otherwise. Descriptive information about each source was intentionally left out in an effort to avoid confounding factors introduced by uncontrollably influencing the degree of pedigree or expertise in each source. However, because of the limited information participants may not have contemplated much about the credibility or performance of each source. More background information may have made the presence of different sources more salient and realistic. Although information about the functionality of the heading band SSD was provided, which was necessary in order to use it, the SSD was used as the mediating interface for both the automated and human source.

4-6 Conclusion

The main question addressed in this study was whether operators are biased against human and automated decision aids differently, and to what extent this affects advice acceptance in relation to the aid's advisory conformance. A real-time ATC simulation was used to investigate source bias effects of air traffic controllers' (i.e., expert operators') trust in a decision aid, and acceptance of its conflict resolution advice. Controllers were led to believe that advisories were derived from either another controller (human) or from an automated system (automation). In addition, controllers were unaware that each advisory was either conformal (individually matched) or nonconformal (based on a colleague's deviating conflict solution style).

Taken together, this limited study suggests that operators have different biases toward human and automated decision aids. Questionnaire trust ratings indicated a preference for the human source. As such, *internal* perceptions and preferences appear to differ between sources, However, neither source bias nor advisory conformance was found to significantly affect the acceptance, agreement, or response time to advisories in the real-time simulation. As such, the effect of these internal source biases does not appear to strongly influence the interaction with different sources in terms of objective metrics.

For future research, researchers should differentiate between trust and reliance measured before and after interaction. In addition, researchers should differentiate between whole-system evaluations of trust and reliance, and derivative trust and reliance measured during a task and interaction with an aid (e.g., as measured by acceptance or agreement with its advisories). In conclusion, the measures of reliance and trust during use of a decision aid reflect the performance of the aid in the context in which it is used. In contrast, whole-system evaluations of trust and reliance also consider other characteristics of an aid, such as its perceived competence and knowledge.

5

Automation transparency effects

In the empirical study described in Chapter 3, controllers' rejection of (conformal) advisories may have been driven by a misunderstanding of the suggested solutions, unable to reconcile them with their own solution. This can be alleviated by addressing the automation's transparency: its ability to facilitate understanding of what it is doing, why, and what it will do next. This chapter describes a real-time study investigating the effect of automation transparency and solution conformance on controllers' acceptance of conflict resolution advisories and their task performance. High and low transparency was considered, accomplished by varying the amount and structure of information provided by two interface representations.

The contents of this chapter are based on:

Paper title Effects of automation transparency and strategic conformance on air traffic controller interaction with a conflict resolution decision aid

Authors Carl A. L. Westin, Clark Borst, Brian H. Hilburn

Submitted to IEEE Transactions on Human-Machine Systems

Abstract

Recent research has highlighted transparency as a critical quality of automation for improving task performance and facilitating appropriate calibration of operators' acceptance and trust. Additionally, the exploration of personalized decision aids for improving human-automation interaction has received recent attention in time-critical and highly automated work environments, such as ATC. This paper investigated the interaction between automation transparency and a personalized decision aid on air traffic controllers' performance and acceptance of that aid in a conflict detection and resolution task. Automation transparency (the degree of meta-information provided) and resolution advisory conformance (the degree of personalization) were varied in a real-time ATC simulation. A repeated measures design was used with two versions of interface representations (high and low transparency) and two versions of resolution advisories (conformal and nonconformal). While no statistical significant acceptance or performance effects were found, results indicated that participants used the two interfaces differently, and preferred conformal advisories (i.e., advisories based on their own solution style) irrespective of transparency. In addition, post-simulation questionnaires revealed a strong preference for the high transparency condition which partly reflected simulation results. This study concludes that increasing transparency involves providing more information, which can incur a cognitive cost in information processing that needs to be traded off with the expected benefits of affording more transparency.

5-1 Introduction

An inevitable result of introducing automation is that it distances the operator from the "hands-on" experience of conducting the task. The automation functions as a relay between the operator and task, with the intention to greatly enhance task execution and improve safety and efficiency. With the introduction of advanced automation, however, comes the ability for the system to conduct work and solve problems differently (i.e., nonconformal) from the human.[155] Occasionally the operator's understanding of what the system is doing is broken.

This breakdown has been labeled automation surprise, leaving the operator "out-of-the-loop" and confused as the automation is not performing as expected, or acting in an unanticipated and uncommanded way.[186] To counter such automation-induced pitfalls, it is essential that the automation can communicate effectively and facilitate understanding. Ironically, advanced automation often features high levels of opacity, as the system's complexity is hidden from the operator inside its "black box".[187] Previous research has shown that opaque systems generally have a negative effect on acceptance and trust as operators question the automation (what is going on? what is the system doing?)[186,188] Thus the paradox: with increased automation, the human and communication between the two become more important, not less.[15,187]

The need for effective and clear communication is particularly important in dynamic, safety-critical and high-risk control room environments. One such domain is ATC, currently facing a large and necessary technological modernization expected to considerably change how controllers interact and communicate with automation.[3,4] Achieving an effective and functional team relationship between the controller and automation has proven difficult. Much advanced automation introduced to current ATC have been distrusted and even rejected.[14,189]

Practically, it is possible to address this issue by considering the degree of automation transparency afforded by the interface.[190–193] Automation transparency is generally concerned with the system's ability to afford understanding about how it works, its behavior, and its intentions. It is typically achieved by providing explanations for system reasoning or visualizing automation meta-information in relation to, for example, information uncertainties and automation reliability.[191] Empirical research indicates that more transparent systems can benefit human-automation cooperation in terms of, for example, performance,[191] acceptance,[111,194,195] trust,[174,191,195] and understanding of advisories.[111,196]

As an alternative to automation transparency, increasing interest has been given to personalizing automation, and in particular decision support systems, to individual preferences, needs, and abilities in task performance and problem-solving.[154,155,197,198] Empirical studies we conducted in the context of ATC CD&R, suggest that strategic conformal resolution advisories (i.e., personalized) benefit controllers' acceptance and agreement with a decision aid, while also reducing their response time to advisories.[101] Such conformal advice can be considered transparent in that it may be perceived and interpreted by the controller to afford a rationale underlying the system's reasoning, simply because it matches with how the controller would solve the conflict.

At a general level, this paper focuses on two complementary ways for improving the operator-automation relationship: by exploring individual differences in personalized decision support, and interface transparency of meta-information visualization that facilitates operator understanding of the system. More specifically, this study investigates how controller performance in a CD&R task, together with the acceptance and interaction with a decision aid, is affected by interface transparency and the strategic conformance of the system's resolution advisories. It is hypothesized that a more transparent interface will reduce the importance of conformal decision support, be more readily understood and positively affect acceptance of performance in a conflict resolution task. While ATC is the focus of this paper, the results are applicable to several other domains with similar dynamics, task demands, and reliance on highly skilled operators.

5-2 Automation transparency research

To date, researchers have used a variety of descriptions to address the need for facilitating system understanding through interface design. Examples include *observability* and *transparency* in general human factors design guidelines,[187,199] *transparency* in applied human factors research[191,195,200] and recommender systems research,[201] *visibility* in usability engineering research[202] and information automation research,[203] and *comprehensibility* in adaptive interface and artificial intelligence research.[194] From here on the term transparency is used in lieu of all these terms.

5-2-1 Theoretical antecedents of transparency

Helldin[191] proposed several generic guidelines for automation transparency design based on a review and synthesis of human- and user-centered automation design approaches in the human factors literature. Among the more prominent guidelines are the importance to provide system feedback; notify operator of information shortages, inconsistencies, or uncertainties; provide rationale underlying automation behavior (i.e., algorithm rules); and communicate possible reasoning conflicts. Notably many of the guidelines emphasize a positive relationship between increasing transparency and appropriate trust calibration.

A more structured theory for automation transparency can be derived from work by Brown[204] on the transparency of intelligent tutoring and learning systems. According to Brown, the design for transparency requires three criteria to be met. First, the system should facilitate understanding of the domain and environment in focus of attention (*domain* transparency). Brown's second criterion, and of particular importance for decision aids, is that the system should facilitate understanding of its reasoning and diagnostic processes (*internal* transparency). This does, however, not imply that a detailed algorithm description or complete input-output relationship should be provided, being neither practical nor feasible. Finally, the system should facilitate an understanding of the overall process within which the user and system is connected to the real world (*embedding* transparency). Combined, these criteria should provide a simplified view of the content in the "black box," making it possible for the user to understand how the system works, why it is doing what it is doing, and to anticipate what it will do next.[192]

Focusing specifically on the behavior of autonomous agents, Chen *et al.*[200] proposed a transparency model for the design of interfaces supporting autonomous agent mission supervision. The model, named situation awareness-based agent transparency (SAT), specifies three SAT levels based on Endsley's three levels of situation awareness.[205] Level 1 SAT addresses what the agent is doing as described by the three Ps for facilitating trust (process, purpose, performance).[15] Level 2

SAT addresses the agents' behavior and underlying reasoning process as framed by the beliefs, desires, intentions (BDI) agent architectural framework.[206] Finally, level 3 SAT addresses what the expected outcomes and future states in regards to the agent's behavior. All three levels are considered distinct, meaning that designing for transparency does not require all three levels to be achieved. Rather they address different aspects of transparency, which relevance will vary with task objectives and contexts. Furthermore, Chen *et al.* argue that transparency improves operator performance by facilitating appropriate trust calibration, which in turns drives appropriate automation usage decisions (AUD,[207]).

Transparency in relation to automation output, such as recommendations and advisories, has been studied in recommender systems research. It has been extensively applied in the context of e-commerce, semantic web services and entertainment,[194,201] although examples can be found in health care diagnostics applications[208] and personalized tour guidance in museums and cultural institutions.[111] Noteworthy is that this research field has merged the notion of transparency and strategic conformance. Transparency is typically increased by providing a text-based explanation personalized to the user's preferences, needs and knowledge. In general, three explanation categories are used: *why* explanations justifies the recommendation; *how* explanations provides the underlying reasoning process used to generate the recommendation; and *tradeoff* explanations acknowledges competing alternatives and considers the constraints for avoiding these.[209] However, deciding on what to explain (i.e., why, how or tradeoff) requires consideration of the task at hand,[210] benefits sought,[211] and technique(s) used for generating recommendations.[212]

Although recommender systems often involve complex and extensive problem solving algorithms, they have typically been associated with low-risk decision making,[213] presented in a static interface environment,[214] and provided in a text-based form either as a single recommendation or a list of alternatives.[212] Thus far, however, recommender systems have received little attention in ATC and similar time-critical and high-risk control room environments that incorporate advanced decision support system and human-agent teams (see Sadler *et al.*[195] for an exception). For example, ATC conflict resolution decision aids have typically been designed to formulate its advice in a text-based format, often providing a list of alternative solutions, constituting a dichotomous 'accept' or 'reject' choice.[31,35,90] A justification for why these solutions are suggested is not presented, making it difficult for the controller to evaluate them properly. Controllers may therefore doubt the quality of solutions[34] leading to low acceptance.

5-2-2 Empirical explorations

While a generally agreed framework for automation transparency is lacking, the basic theoretical foundation is widely shared across domains. It can be considered an attribute of a system that aspires to communicate the system's cognitive processes to better account for its behavior. In general, empirical research supports the hypothesis that increasing system transparency increases operator understanding of the system. For example, Dzindolet *et al.*[174] found that trust and reliance in a decision support system for target detection increased when operators were provided with a rationale for why the automation might err. Recommender system research has shown that transparency in terms of explanations can help users make more accurate decisions and increase their satisfaction of using the system,[201] improve their trust and acceptance of recommendations,[201,209] and benefits their understanding of system recommendations[111] and system behavior.[196] In a recent study, Sadler[195] investigated how trust and reliance were affected by the transparency of a ground-based decision support system aiding "dispatchers" (consisting of airline pilots) in choosing a suitable diversion airport. Results showed that trust increased while the need to consider options decreased when system transparency was increased, by either providing probability estimates for successful diversions (intermediate transparency) or probability estimates together with supporting statements describing the information considered and how it have been interpreted (logic transparency).

Benefits of automation transparency on trust and workload have also been found in experiments where the SAT model[200] has been used to guide interface visualization design for supervision and interaction with an autonomous unmanned ground vehicle, named the autonomous squad member (ASM), in a military operational context. In two separate studies using inexperienced operators, trust in the ASM increased with increasing SAT level transparency. In addition to information reflecting the ASM's current status (level 1 SAT), Boyce *et al.*[215] increased transparency by either adding environmental constraints affecting the ASM's activities (i.e., level 2 SAT pertaining to the agents reasoning) alone, or together with a visualized projection of agent status and uncertainty (level 3 SAT). In the other study, Selkowitz *et al.*[216] provided level 2 SAT by means of a symbol expressing motivation underlying the ASM's reasoning, and level 3 SAT by means of symbols indicating the consequence of decisions on the ASM's resource usage. Furthermore, the two studies showed that increasing transparency by means of the SAT framework, allows for increasing additional information without increasing operator workload.

Helldin[191] showed in a series of experiments that transparency, by means of visualizing meta-information in relation to system reliability, uncertainty, and underlying reasoning, benefited appropriate trust calibration and task performance among various skilled operators. For instance, trust was more appropriately calibrated among drivers when they were provided with a visualization of the cars ability to autonomously drive. Results indicate that drivers had a better understanding of the system's ability and its limitations. They also intervened faster and more appropriately when the autonomous car reached its operational limits and suddenly failed. In another experiment, air defense operator's understanding of a decision support system's reasoning in target detection increased as a result of indicating the system's limitations by means of uncertainty estimates for detected targets. In addition, operators became more attentive and cautious in their identifying targets.

However, the effects of increasing transparency have not been uniformly positive. Drawing on mixed results from the same experimental series, Helldin[191] noted that both workload and decision-making time suffered (i.e., increased) with increasing transparency.[191] Cramer *et al.*[111] found that explaining to users why an advisory was made increased acceptance but did not influence their trust in the system. Our previous research suggests that inexperienced operators' acceptance of (accept or reject), and agreement with (on a 1-100 rating scale), solution advisories is not influenced by decision aid transparency in regards to the visualization of constraints affecting conflict resolution advisories in an ATC CD&R.[217]

Mixed results have also been attained in research by Bass and colleagues, who found that risk probability judgments in an egocentric aircraft conflict prediction task did not improve with decision aid transparency achieved by accompanying probability estimates with strategy information (i.e., a visualization of the closest point of approach between the two aircraft). Participants' poor performance in the more transparent condition was attributed to the interface representation being too confusing, showing probability contours for risk of conflict, in the combination with insufficient training.[218] While more recent study reversed these findings (i.e., probability judgments improved with strategy information), the provision of strategy information was found to be less effective than the provision of uncertainty information pertaining to the speed and heading of other aircraft (referred to as environmental information).[219]

Overall, these empirical findings do not provide a coherent picture of automation transparency and how it affects human attitudes and automation use behavior. However, the seemingly contradictory outcomes of transparency research can possibly be explained by an insufficient balance of information. In all studies reviewed herein, increased transparency is consistently achieved by providing more meta-information pertaining to, for example, system reasoning,[111,191,195,196] information uncertainty,[191,218,219] or situation awareness.[200,216]

5-2-3 Increasing transparency by means of meta-information

Importantly, the notion of transparency shifts the focus from presenting raw state-related information, reflecting the current situation, to presenting *meta-information* that better reflects the human-machine-ecology relationship and guides action and decision-making. Meta-information reflects information qualifiers underlying system reasoning that facilitate sense-making of information provided.[220,221] As such, transparency can be considered a measure of the system's openness in meta-information communicated that entails what the automation is currently doing, which information is being used, how it is being processed, and when and how it is provided. Specifically for decision aids, there is a need to provide meta-information about the system's rationale underlying its judgments and problem-solving, the criteria considered, and uncertainty factored.[191,218]

There are, however, no transparency guidelines that explicitly prescribe which, how much, and in what ways meta-information should be presented. Simply providing more meta-information is undesirable since it can easily overload the operator's cognitive ability to detect and process the information in real time.[222,223] Such overload is typically characterized by clutter, increased workload, or difficulties in finding the significant data, which can result in degraded monitoring and signal/change detection, delayed visual search, confusion, increased confidence in wrong judgments, and increased memory load.[199,224] As such, there is an apparent need for transparency calibration, to avoid human-automation cooperation breakdowns caused by excessive or insufficient access to meta-information. The need for balancing meta-information necessity with meta-information overload in visualizations and display design (i.e., the human-computer interface bottleneck[225]), is especially critical for time critical and information rich environments.[191,221]

5-2-4 Ecological interfaces for visualizing meta-information

Research suggests that the ecological interface design (EID) approach lends itself well for facilitating transparency, particularly in the context of supervision and control problem tasks, by means of visualizing relevant meta-information without incurring drawbacks of information overload. The EID approach strives to identify and visualize relevant relationships and constraints in the work environment that afford constructive problem-solving and action.[122] As such, it provides a structured approach for identifying and visualizing meta-information that reflects meaningful relationships in the problem space of interest, rather than simply notifying the operator of system health, status or mode changes, which is typical for automation in most current supervisory wok domains.[226] EID maps particularly well onto the domain transparency suggested by Brown, in that it explores the relationship between agents (e.g., human and automation) and the work ecology.[204] Kilgore and

Voshel demonstrated how an EID-derived interface for the supervisory control of unmanned vehicles in the maritime domain can be used to increase the transparency of system decision-making and behavior.[227] Similarly, and as noted earlier, Chen and colleagues showed that merging EID with the SAT framework can benefit system trust and situation awareness, without increasing workload.[200,215,216]

The approach centers around a work domain analysis of the sociotechnical system considered,[122] which typically consists of decomposing the domain using an abstraction hierarchy.[228] This hierarchy describes the system at different levels, ranging from the functional high-level purposes of the system to more specific generalized functions and the physical form of system elements and components, which serves to determine *which* information to visualize. In order to determine *how* this information should be visualized, the information requirements are considered in relation to human information processing abilities and problem-solving behaviors using, for example, the skills, rules, and knowledge (SRK) framework.[85] Because EID visualizes a complex system across multiple levels of abstraction, it provides a useful scaffold for supporting deep, knowledge-based reasoning over system behavior during novel situations or fault response.[122]

5-3 Ecological displays in ATC CD&R

Importantly, the EID paradigm is believed to correspond well with how skilled operators perceive and solve problems, namely by considering the relationships of objects and events in the world, rather than the precise stimulus elements. Similarly, it has been suggested that air traffic controllers base their judgments and decisions in CD&R on the relationships and interactions between aircraft and their constraints as they evolve over time, rather than on information about the aircraft or its position.[229] However, these relationships, and the possibilities for solving conflicts that they afford, are not readily accessible to controllers in current ATC systems. Rather, controllers compute these relationships by engaging in effort and time demanding inferential cognitive processes, interpolating multiple information sources. In order to balance the cognitive load in CD&R, research suggests that controllers rely heavily on decision-making heuristics.[79,80] An example is the use of trajectory prediction strategies that involve mental extrapolation of aircraft trajectories in conflict detection. Strategies include time comparisons, distance comparisons, altitude comparisons, and contraction rates.[26,33]

Over the last 30 years, several ATC research projects have explored interface designs that to some extent acknowledge the work ecology, and attempt to alleviate the cumbersome extrapolation involved in CD&R. These interface typically visualize obstacles (e.g., other traffic, restricted areas, weather, terrain) as constrained "no-go" areas on a two-dimensional geospatial (map) display (e.g., complexity maps,[230] and the PHARE Problem Solver/HIPS for conflict resolution[231]). They have, however, not been developed as part of the EID framework. A good example of a similar EID derived application, originally developed as a tactical obstacle avoidance interface for pilots, is the Solution Space Diagram (SSD).[232] The SSD is described in detail in Appendix B.

The ATC SSD visualizes a selected aircraft's maneuverability constraints, based on the relative position of other aircraft. By explicitly showing the constraints affecting a selected aircraft's trajectory, the SSD accomplishes the strenuous cognitive work involved in trajectory extrapolation that controllers previously have been required to do. The SSD has been shown to reduce novice controllers' workload during high traffic loads and increase separation without reducing sector throughput.[102] Experienced controllers using the SSD have been shown to implement more conservative conflict solutions that benefit overall sector robustness, albeit at the cost of efficiency measured by the additional track miles.[233]

5-3-1 Understanding and using the SSD

Figure 5-1 explains the basic SSD construction in relation to a conflict between two aircraft, shown in a plan view display perspective representative to that of a controller. Aircraft *A* has been selected and therefore designated the *controlled* aircraft. All other aircraft are considered *intruders* that may interfere with the controlled aircraft. The example in Figure 5-1(a) shows aircraft *B* as the only intruder aircraft.

Figure 5-1(b) illustrates how the SSD visualization is calculated by processing velocity plane information of both aircraft (V_1, V_2) in relation to the minimum separation zone (*protected zone*) of intruder aircraft *B* (typically 5 nmi in en-route). The result is a triangle-shaped area formed by the relative position of the tangent lines of aircraft *B's* protected zone and the position of aircraft *A* (Figure 5-1(b)). This area, called the *conflict zone*, comprises all relative velocity vectors $V_{relative}$ that result in a loss of separation. The SSD presentation is derived from visualizing the conflict zone within the maneuvering envelope of the controlled aircraft *A*.

Figure 5-1(c) shows the original triangle (TRI) SSD around aircraft *A*, consisting of three circles. The dashed circle, intersecting with the velocity vectors, represents the aircraft's current speed. The inner and outer circles represent the aircraft's maneuvering envelope in the horizontal plane (heading and speed). The partly shown conflict zone of aircraft *B*, referred to as the *no-go zone*, is visualized within the maximum V_{max} and minimum V_{min} speeds of this envelope. The no-go zones provide the boundaries for safe travel by capturing the meaningful relationship between aircraft. The size and position of the no-go zone reflect the relative position, velocities, and proximity of aircraft *B*. Red (darker grey) and orange (lighter grey) areas

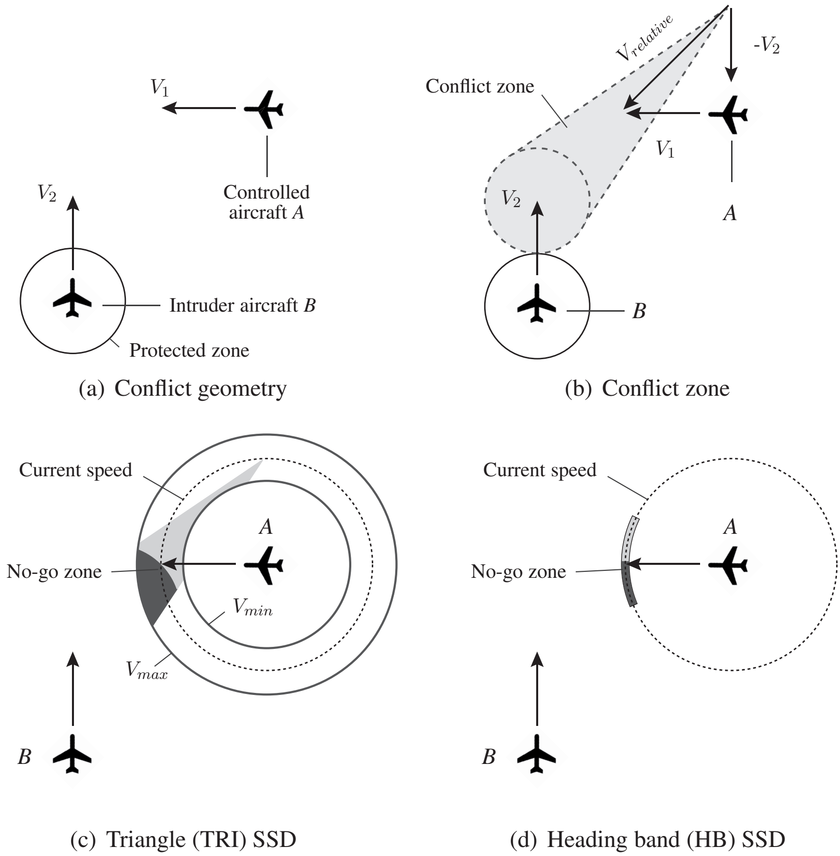

FIGURE 5-1: Right angle conflict between aircraft A and B, with SSD shown for aircraft A.

indicate time to separation loss, with the red color representing less time than the orange color. Time to separation loss is computed by processing the relative velocity and relative distance between aircraft. To avoid separation loss, the velocity vector of aircraft *A* must be positioned outside the triangle-shaped no-go zone.

Figure 5-1(d) shows a simplified SSD representation in which the no-go zones are restricted to a narrow band representing the current speed of the controlled aircraft. This heading band (HB) SSD was developed for purposes of integrating the SSD directly with the radar display.[101] The HB SSD was considered appropriate for the particular ATC en-route environment considered, in which speed commands are infrequent, partly due to narrow speed performance envelopes of aircraft.

5-3-2 Transparency evaluation of the SSD

Evaluation of the two SSD representations was considered in relation to the different theoretical models of transparency reviewed earlier. Brown's[204] three transparency criteria were determined useful for guiding this evaluation. In contrast, the SAT framework[200] was discarded as it specifically addresses autonomous agents and their activity. Transparency pertinent to Wang and Benbasat's[209] three explanations categories of recommender systems were also not considered as they specifically address advisories. However, *why* and *how* transparency was instead partly addressed by the conformance of advisories, with conformal solutions representing the participant's preferred way of solving the conflict.

Table 5-1 details how the two SSD representations relate to the three transparency criteria suggested by Brown.[204] In the context of CD&R, domain transparency can be considered to encompass the meta-information of the airspace and all traffic, its constraints and boundaries, that controllers use to solve conflicts. Internal transparency addresses the underlying reasoning of the decision support interface, which here relates to the algorithms underlying the visualization of intruder aircraft's relative position. *Embedding transparency* reflects the mapping between interactions with the interface translates to the reciprocal interactions in the real world. For example, an avoidance maneuver implemented based on the SSD should be accurately reflected in the real world.

Overall, the HB SSD was identified to afford less *domain transparency* in that the meaningful relationships between controlled and intruder aircraft were not reflected adequately in no-go zones. Similarly, feedback from participants in a previous study using the HB SSD indicated that interpreting the position of intruder aircraft was complicated by the small, narrow, poorly contrasted, and overlapping no-go zones.[101] Several of the issues identified for the HB SSD are addressed by the TRI SSD, which is considered to better facilitate domain transparency.[217] This is mainly attributed to the visualization of the controlled aircraft's entire speed envelope which improves understanding of the constraints and boundaries of intruder aircraft and how to avoid them.

The TRI SSD was found to facilitate slightly more *internal transparency* because the triangles reflect more information relevant to how the no-go zones have been calculated. As noted by Helldin,[191] internal transparency is often best facilitated during initial familiarization and training with the interface. Finally, the two SSD representations were found to not differ in terms of *embedding transparency* as both are integrated directly with the primary information (radar) display. As such, interacting with the interface closely mirrors the interactions of the real world.

TABLE 5-1: SSD transparency across Brown's[204] transparency criteria

Transparency criteria	HB SSD transparency	TRI SSD transparency
Domain transparency	The identification of an intruder aircraft's position and proximity is not readily apparent with information limited to the current speed of the controlled aircraft. Although no-go zone characteristics across different speed settings can be mentally interpolated by manually scrolling through the speed envelope of the controlled aircraft (by increments of 10 knots), this requires additional time and cognitive workload.	The triangle-shaped no-go zones facilitate increased understanding of 1) the full shape of the constraints and boundaries and how to avoid them, 2) the proximity of an intruder determined from the width of the triangle (a larger angle between the triangle legs indicates closer proximity), 3) the relative position of an intruder determined from the the direction to which the triangle opens up, and 4) the intruder aircraft's velocity vector (speed an trajectory) determined from mentally drawing a line from the origin of the controlled aircraft to the tip of the intruder aircraft's triangle visualized in the SSD.
Internal transparency	Insight into the underlying algorithms for calculating the no-go zones is not readily afforded.	The full triangles reflect more information relevant to how the no-go zones have been calculated and therefore provide slightly better internal transparency
Embedding transparency	Achieved by integrating the SSD with the radar view display.	Achieved by integrating the SSD with the radar view display.

5-4 Transparency hypothesis

In conclusion, previous research suggests that transparency is an attribute of automation reflected in its interface that drives operators' attitudes toward the system and ultimately their decision to adopt it and accept its output. In a similar way, strategic conformance can be considered an attribute of automation, that reflects the system's problem solving style.[155] This relationship is illustrated in Figure 5-2, influenced by the relationship between transparency and performance in.[200] Manipulating transparency entails varying the degree of meta-information provided, usually with more meta-information resulting in increased transparency. Transparency facilitates understanding of the automation's behavior and intent, with increased understanding facilitating improved calibration of human attitudes toward the automation. For example, if there are large uncertainties in data processed by the system, or it is operating close to its performance envelope, this should be communicated to the operator.[191] With an accurate understanding of the system's behavior, the operator is in a better position to decide whether to use the system, or not.

The left hand side of Figure 5-2 depicts the hypothesized relationship between automation transparency and automation conformance investigated in this paper. Because the increased transparency facilitates understanding of what the automation is suggesting and why, the value of personalized, conformal, advice lessens. While the operator may not agree with a nonconformal advice, the afforded transparency into the system's reasoning underlying the advisory facilitates trust and usage of the system. As such, it was hypothesized that conformal advisories would benefit advisory acceptance and agreement while reducing response time (similar to findings in previous research[101]), but only in the low transparency condition. Furthermore, it was hypothesized that the TRI SSD, in comparison with the HB SSD, would be perceived as more transparent, accepted more often, and improve performance, irrespective of advisory conformance.

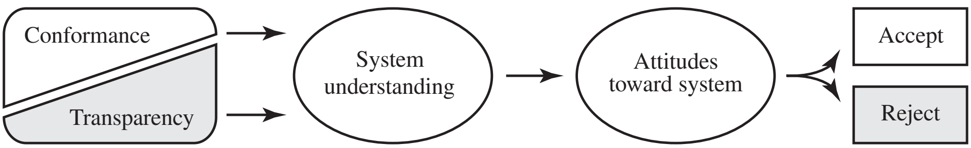

FIGURE 5-2: Relationship between transparency, conformance, and automation use.

5-5 Method

5-5-1 Participants

The study used a sample of nine Swedish air traffic controller trainees (three females and six males, mean age 26 years) with their basic training completed and one year on-the-job training remaining. All had undergone the entire training together as part of the same student group. Participation was voluntary. Trainees were chosen because earlier empirical research has shown that the SSD benefits novices more than experienced controllers.[102] Furthermore, while active controllers have been used in a previous study,[101] this study sought to investigate whether conformance effects applied to trainees as well. The strategic conformance of decision support may be less effective with inexperienced novices who lack well-developed problem-solving styles after which conformance can be modeled.[155]

5-5-2 Simulator

The ATC simulator environment and scenarios were programmed in Java with an OpenGL binding for the graphics. It provided a classic plan view display, providing a top-down sector presentation similar to current ATC. The simulator ran on a laptop connected to an external monitor with a resolution of 1600 x 1200 pixels. To reduce simulation time and induce additional time pressure, the simulation was increased to 2x real speed. Aircraft plots were updated every second to simulate a 1 Hz radar update frequency.

5-5-3 Task

There were two main tasks: 1) direct traffic off-track (light grey aircraft) to their assigned sector exit points (indicated on their electronic label), and 2) ascertain traffic separation (5 nmi horizontally). Participants interacted with the SSD, by means of mouse and keyboard, to control and communicate with aircraft. The SSD representations used a nominal propagation method for which future state values (position and trajectory) were calculated by a simple projection of an aircraft's current state into the future, without factoring in uncertainty and without intent information (no flight plans, just state-based vectoring). In addition, a decision aid provided participants with conflict resolution advisories. Advisories were plotted within the SSD representation. Because flight level changes were not possible, conflicts had to be solved by using heading, speed, or combinations thereof. There was no wind or other weather phenomenons affecting traffic.

5-5-4 Measurement scenario and designed conflict

The manipulation of conformance required conflict repetition to assure validity of a participant's conformal and nonconformal solution. Therefore, one measurement scenario, containing a specifically designed conflict was created. The scenario, shown in Figure 5-3, consisted of a hypothetical square en-route sector, 80 x 80 nmi in size with eight entry/exit waypoints. Although the sector represented a futuristic free route airspace (i.e., no intermittent waypoints), traffic mainly followed two parallel flows with one crossing. The scenario contained in total 27 aircraft, with between 19 and 22 always inside the sector.

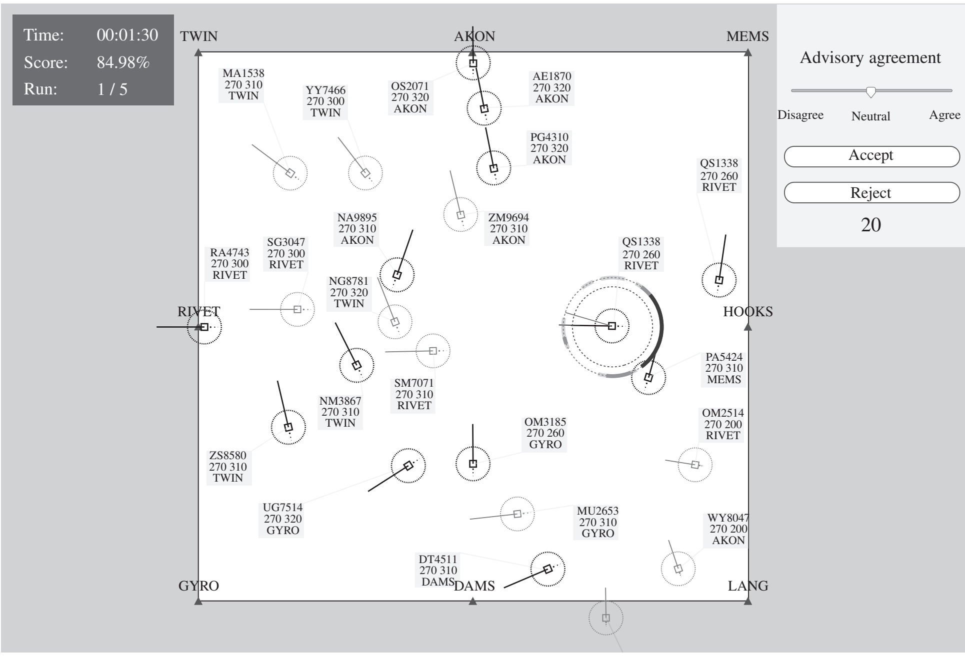

FIGURE 5-3: Measurement scenario with designed conflict between QS1338 and OM3185. A resolution advisory for QS1338 is shown (as depicted by the HB SSD), suggesting a right turn and increase in speed (320 knots). The advisory dialog window (upper right corner) enables the controller to either accept or reject the advisory, but first after providing an agreement rating. Time to advisory expiration is shown in the bottom of the window (20 seconds). The remaining time of the scenario, current performance score, and scenario run out of total scenarios is shown in the upper left corner. Note that the aircraft depicted in lighter grey are off track and require a vector to their assigned exit point as indicated in the bottom field of the aircraft label.

The designed conflict consisted of two aircraft on perpendicular headings, with aircraft A (QS1338) on heading 270, and aircraft B (OM3185) on heading 360. Both aircraft were flying the same speed (260 knots), with identical speed envelopes (200-320 knots), and were scripted to collide at the exact center of the sector (closest point of approach of 0 nmi). Separation loss occurred after 104 seconds, unless participants intervened.

Additional context traffic was included to increase scenario difficulty and prevent early solution of the designed conflict. The position and trajectory of context aircraft were scripted so that they did not constrain any solutions of the designed conflict. Scenario repetitions were rotated (two versions) and different exit/entry point identifiers were used. In addition, repetitions were intertwined with "dummy scenarios. These measures were used to prevent scenario recognition from affecting solutions.

5-5-5 Independent variables

This study was a 2x2 repeated measures design varying advisory conformance (conformal or nonconformal) with interface representation transparency (low or high). Different automation transparency levels were achieved by varying the meta-information richness provided by the SSD in two interface representations. Figure 5-1 shows the two interface representations, the low condition represented by the HB SSD (Figure 5-1(d)), and the high condition represented by the TRI SSD (Figure 5-1(c)). Conformal (cfY) advisories were individually tailored to reflect a participant's unique conflict solutions style. In contrast, nonconformal (cfN) advisories represented an opposite solution style, based on that of another participant.

5-5-6 Dependent measures

The following dependent measures were collected:

- Acceptance of an advisory (binary, accept or reject).
- Agreement with an advisory (on a 1-100 scale).
- Response time (from advisory onset to accept or reject button press).
- Scenario difficulty (on a 1-100 scale).
- Number of SSD interactions.
- Type of interaction (in heading, speed, or combinations).
- Separation losses (safety).

Two questionnaires were used to assess differences in participants' perceptions of automation transparency, and in particular their understanding of the two interface representations. The first questionnaire consisted of two identical sets of eleven Likert statements, one for each SSD representation (HB and TRI), with a seven-point scale ranging from strongly disagree to strongly agree. Most of the statements were adapted from transparency questionnaires in Cramer *et al.*[111] and previous studies of our own.[101,217]

The second questionnaire consisted of a seven questions on a Visual Analogous Scale (VAS). The VAS is a subjective questionnaire instrument on which responses to statements or questions are made by making a mark somewhere on a 100 mm continuous line that connects two contrasting endpoints, with a marking in the middle indicating a neutral response. Here, endpoints consisted of the two representations encountered in the simulator. Questions addressed participants' representation preference in relation to how well each representation facilitated understanding of why a certain advisory was suggested (i.e., why transparency), and how well the constraints of the conflict and its alternative solutions were visualized (i.e., domain and tradeoff transparency). Additional questions addressed the representations usefulness for CD&R, the difficulty experienced in working with them, the perceived clutter, and the perceived workload.

5-5-7 Procedures

The three-phased study ran over a three-week period and encompassed two simulations. Both simulations were preceded by consent procedure, briefings, and training runs. The initial *prequel* simulation, conducted in the first week, captured participants' conflict solutions of the designed conflict. The simulation consisted of ten two-minute en-route scenarios. Unknowingly to participants, the measurement scenario containing the designed conflict was repeated four times. In this phase participants solved conflicts using only the HB SSD.

In the second week (*conformance design* phase), participants' four recorded solutions were analyzed to define their individual solution style, consisting of the most persistent pattern found. The individual conflict solutions style determined each participant's conformal and nonconformal resolution advisory to be used in the subsequent main simulation. For example, consider the conflict between aircraft *A* and *B* in Figure 5-4. If Participant 1 repeatedly solved the conflict by vectoring aircraft *A* to the right behind aircraft *B*, the conformal solution style would be to vector aircraft *A* to the right behind aircraft *B*. A nonconformal, opposite solution style, could then be that of Participant 2, who consistently vectored aircraft B to the left behind aircraft *A*.

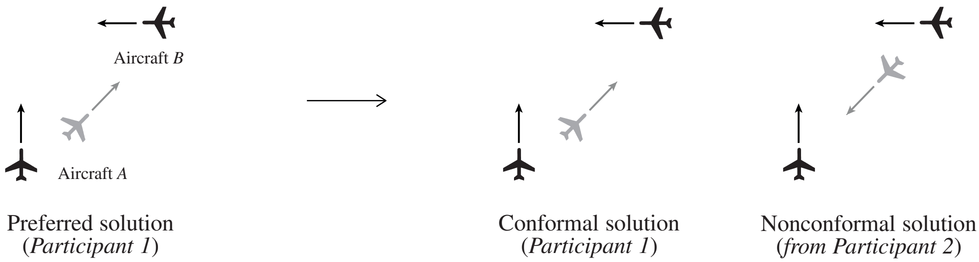

FIGURE 5-4: Creation of conformal and nonconformal resolution advisory.

The *experiment* simulation took place in week three. It took roughly 90 minutes and consisted of two sessions, one for each transparency condition, and a questionnaire part. In this phase participants solved conflicts using both the HB and TRI SSD. Participants were divided into separate groups encountering the transparency conditions in different order, either HB-TRI or TRI-HB. Advisory conformance was, unknowingly to participants, varied within each session according to a Latin square design. Prior to each session, participants received instructions about the interface representation they were about to interact with. Each session comprised twenty minutes of training followed by five two-minute long scenarios. The five scenarios consisted of one measurement scenario repeated twice, intertwined by three dummy scenarios.

Unknowingly to participants, the scenarios and conflicts encountered in the experiment simulation were identical to those used in the prequel simulation. This time, however, participants were assisted by a decision aid that would suggest advisories by plotting them in the SSD. Advisories were accompanied by a beeping sound and a dialog window with an accept and reject button, an agreement scale, and a countdown timer starting at 30 seconds. Participants could either accept or reject the advisory, but first after indicating their agreement with the advisory. Participants were instructed that the decision aid would detect most conflicts and that resolution advisories suggested were generated by the decision aid and always would be safe although not necessarily efficient. In reality, resolution advisories were either conformal or nonconformal and only provided for the scripted designed conflict.

5-6 Results

Non-parametric tests were used because of the small sample size and ordinal variables. All tests were conducted with a significance level of $\alpha = 0.05$.

5-6-1 Acceptance

The acceptance ratio of resolution advisories across conditions is shown in Figure 5-5(a). Overall, 26 of 36 resolution advisories were accepted (72.2%). Conformal advisories were accepted more often (77.8%) than nonconformal advisories (66.7%). Acceptance did not vary as a result of interface transparency alone, with 72.2% of all advisories accepted in both conditions. The Cochran's Q test indicated no effect between transparency and conformance conditions on the acceptance of advisories.

A boxplot showing the agreement rating across conditions is shown in Figure 5-5(b). Results indicate that the conformance effect was partly reversed under the TRI condition. While agreement decreased with conformance when using the HB, the opposite was found in the TRI condition, with the highest agreement found in the nonconformal condition. This suggests that the TRI interface improved participants' understanding of nonconformal advisories, perhaps directing their attention to a solution they did not think of. Note that acceptance, however, was unaffected. A Friedman's test showed no significant effects between conditions on participants' agreement ratings of resolution advisories.

5-6-2 Performance and interface usage

Response time was measured from advisory onset to time of acceptance or rejection, and varied between 4.1 and 21.8 seconds. The variation in response time across conditions is shown in Figure 5-6(a). Interface transparency appears to have influenced participants' response time to advisories. Conformal advisories were responded to faster in the TRI condition, while nonconformal advisories were responded to faster in the HB condition. However, a Friedman test revealed no significant effects of transparency or conformance on response time. Figure 5-6(b) shows the boxplot for

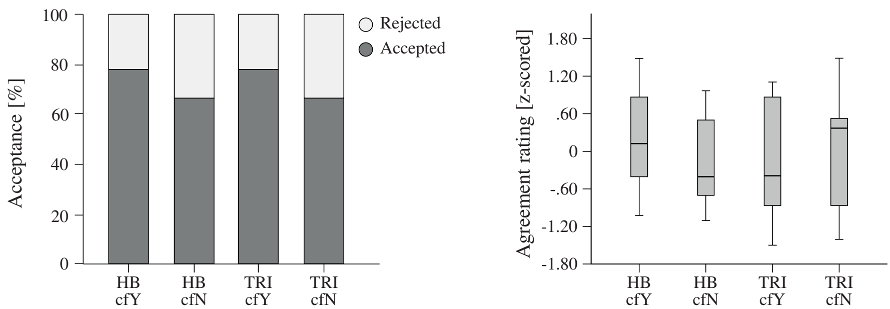

(a) Acceptance (b) Agreement

FIGURE 5-5: Bar chart with advisory acceptance ratio (a) and boxplot of agreement ratings (b) ($N = 9$).

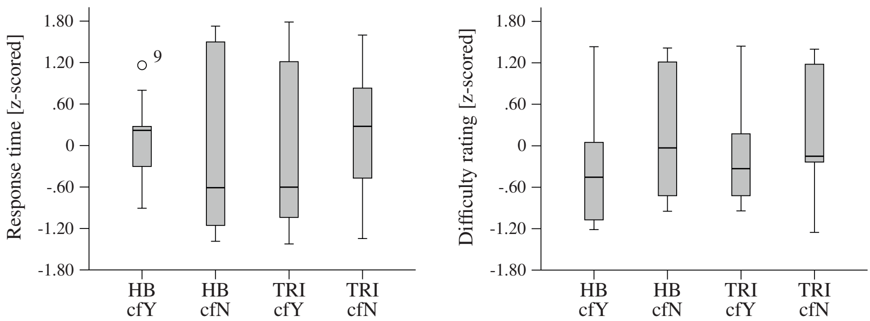

(a) Response time (b) Difficulty

FIGURE 5-6: Boxplots of response time (a) and difficulty ratings (b) ($N = 9$).

difficulty ratings obtained after each scenario. Although scenarios with nonconformal advisories were rated slightly more difficult, the Friedman test did not detect any significant differences between conditions.

The following interface usage results only apply to transparency effects. Unlike conformance effects (measured in relation to the resolution advisory), transparency effects reflects data collected from the entire scenario run. Figure 5-7 provides boxplots of (a) number of SSD inspections, and (b) type and number of conflict solution commands. The former indicated a trend with less inspections when using the TRI SSD (Wilcoxon: $z = -1.72$, $p = .086$). Total number of interactions (the sum of all heading, speed, and combined interactions) did not vary significantly. There was a trend for more use of speed commands with the TRI SSD (Wilcoxon: $z =$

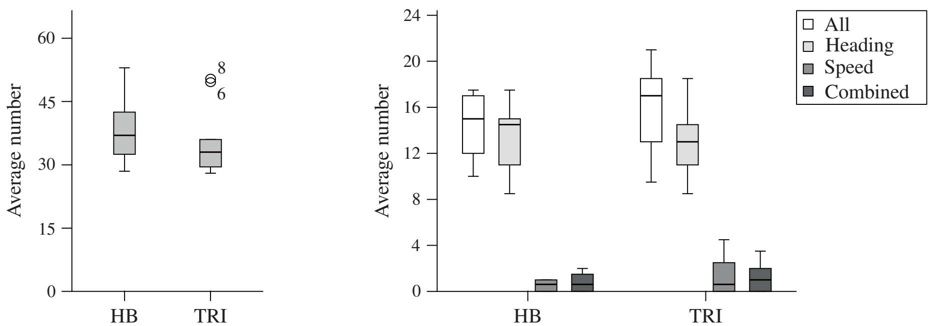

(a) Number of SSD inspections (b) Type and number of conflict solution commands

FIGURE 5-7: Boxplots of SSD inspections (a) and commands (b) across transparency conditions ($N = 9$).

1.80, $p = .072$), while neither heading nor combined heading and speed commands differed between conditions. The increased used of speed control may be attributed to the visualization of the entire speed-envelope in the TRI SSD. Finally, the lack of separation losses recorded indicate that safety was maintained with both interfaces.

5-6-3 Transparency perceptions

In general, questionnaire results supported the trends found in the simulation data. Selected bar charts for six Likert questions are shown in Figure 5-8. Participants found both interfaces helpful for conflict solving, facilitating an understanding for why a solution was suggested (Figure 5-8(e)), while slightly disagreeing with the statement that better explanations were needed (Figure 5-8(f)).

Overall, a general preference for the TRI SSD over the the HB SSD was noted. The TRI SSD was found slightly better for determining which aircraft caused which no-go zone (although not significant) and easier to use for combined solutions (Wilcoxon: $z = 2.06$, $p = .040$). However, the TRI SSD was also found significantly more cluttered (Wilcoxon: $z = 2.46$, $p = .014$).

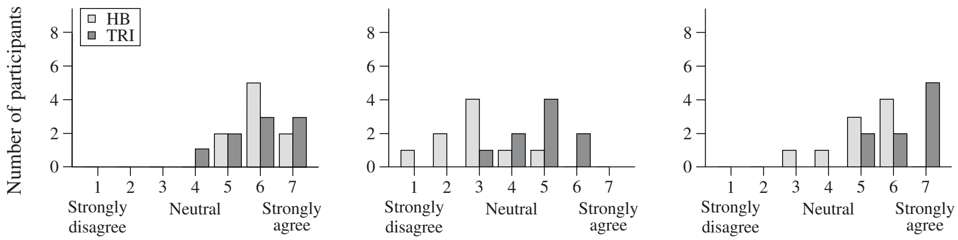

(a) The representation was helpful in resolving conflicts (b) I found this representation too cluttered (c) The representation made it easier to apply combined solutions

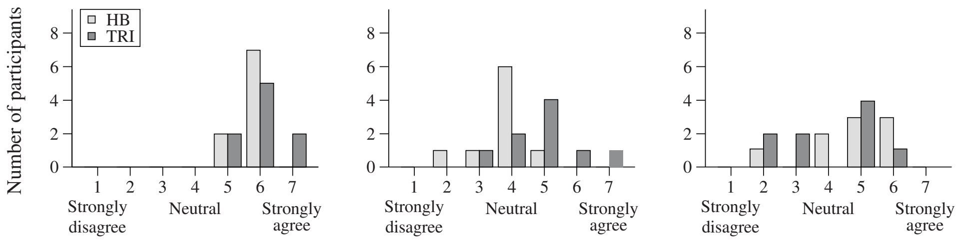

(d) I understand why the system suggested a certain solution (e) I could determine which aircraft caused which no-go zone (f) I would have liked more information on which aircraft caused which no-go zone

FIGURE 5-8: Bar chart showing Likert questionnaire responses ($N = 9$).

Boxplots of VAS questionnaire responses are shown in Figure 5-9. Values above 50 mm indicate a preference for the TRI SSD while values below indicate a preference for the HB SSD. Overall, responses support the hypothesis that the TRI was perceived as a more transparent interface than the HB. More specifically, the TRI was perceived as more "helpful for conflict detection and resolution", providing a better "overview of the solution options available", and facilitating a better understanding for "why the system suggested a certain solution". In contrast to these positive perceptions, participants also found the TRI "more cluttered", and slightly more "difficult to work with", than the HB SSD. These responses were likely driven by the additional amount of meta-information provided in the TRI, which in turn can be considered a characteristic of increasing interface transparency.

5-7 Discussion

The present study investigated the effect of interface transparency and strategic conformal decision support (i.e., personalized) on controller trainees' performance in CD&R, in addition to their acceptance of resolution advisories. Furthermore, two complementary questionnaires were used to examine differences in participants' transparency perceptions of two CD&R tools which varied in the degree of transparency. Of particular interest were the potential interaction effect between transparency and conformance on the acceptance of resolution advisories, and whether an increased degree of transparency would reduce the importance of conformance.

Neither interface transparency nor conformance significantly affected acceptance or performance measured in simulation data. Still, there were some interesting trends and differences between conditions. Some of these findings were supported by questionnaire results, which indicated differences in participants' perceptions of the two interfaces. From this study, two findings are of particular interest.

First, acceptance results suggest that the benefits of conformal advisories do not lessen with increased interface transparency. Therefore, the hypothesis is rejected. However, participants' preference for advisories conformal to their own solution style corresponds with findings attained previously.[101] This result, although not significant, provides additional support for the positive effect of individually tailored recommendations on the acceptance of decision aid outcomes. With the small sample size, however, these results must be considered with caution. Furthermore, this study included trainees, in contrast to the larger sample of experienced controllers used in Hilburn *et al.*[101] Controller trainees may be less consistent in their conflict solving style than experienced controllers, which potentially can explain the lack of a stronger conformance effect.

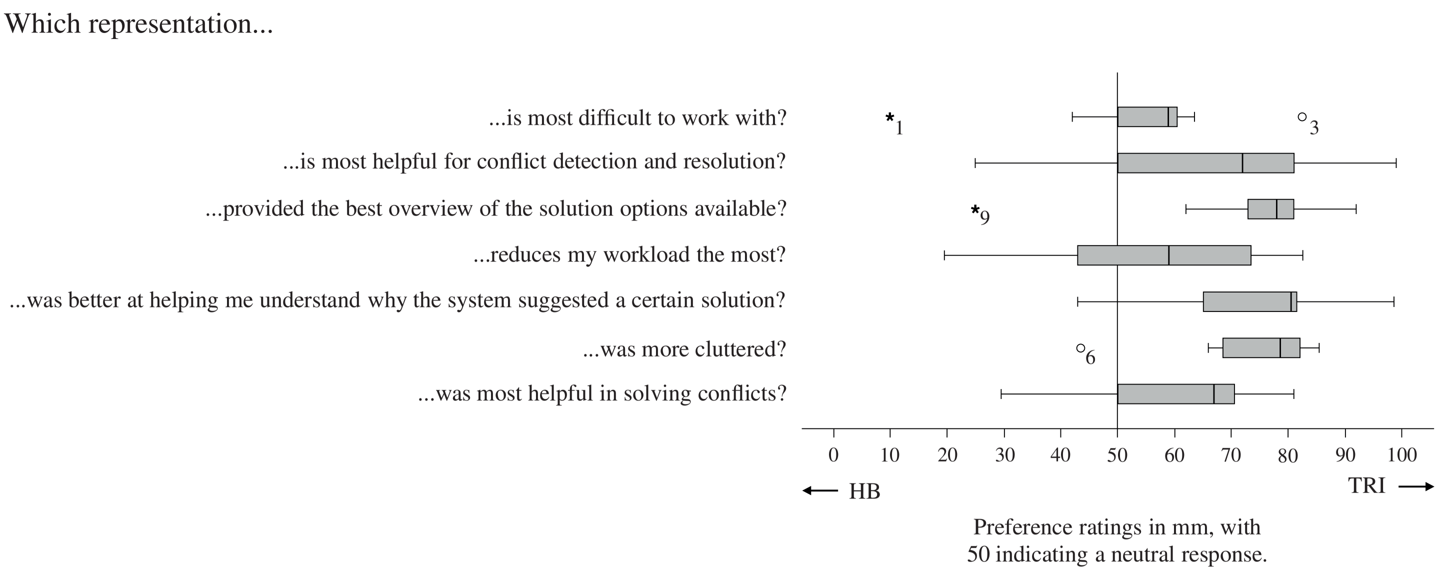

FIGURE 5-9: Boxplot of interface representation preference ratings measured in VAS questionnaire ($N = 9$).

Secondly, results supported the hypothesis that increased automation transparency can be achieved by increasing the amount of visualized meta-information. Questionnaire results indicated an overall preference for the more transparent TRI SSD, which correspond to previous research results in transparency research that the system and its advisories are better understood when more meta-information is provided about the criteria affecting problem-solving.[191,195] Further, in accordance with previous research, questionnaire results showed that increased transparency was associated with perceptual costs in terms of increased clutter and working difficulties. These perceptual differences did, however, not translate into significant differences in the simulation data. While the current study was limited by a small sample size, results suggest that although transparency was perceived as an important characteristic of automation, it was less relevant for the actual interaction with an interface. Similarly, Sadler *et al.* found that acceptance of aircraft diversion recommendations were not affected by transparency, although higher transparency in their study was significantly associated with higher trust ratings in the decision support system.[195] Future research should further investigate to what extent the perception of transparency influences the acceptance and interaction with automated systems.

5-7-1 Calibrating transparency

This study has shown that EID-derived support tools, such as the SSD, can successfully be used to visualize meta-information in ways that facilitate increased transparency. The framework is especially suitable for deciding what to visualize and, to some extent, how to visualize it. However, the design framework does not specify how much information, or the extent of detail, that should be incorporated in the visualization. As such, EID does not provide much guidance for finding the appropriate transparency balance between insufficient and excessive information. Further research should investigate ways to integrate guidelines in EID for establishing a degree of transparency that is appropriate given the task, context, and operator.

Taken together, findings emphasize the need for calibrating transparency in order to facilitate appropriate understanding. Furthermore, although there is a strong incentive for reducing the amount of information presented to operators, there is an opposite intrinsic desire for the operator to process as much information as possible in order to make the best decision. While a shortage of information may not facilitate the needed understanding, excessive information can overload the user. This relationship between understanding and meta-information can explain the mixed results obtained in transparency related research, in which transparent systems have been shown to affect human-machine interaction and task performance both positively and negatively.

While raw data are aggregated for the purpose of developing meaningful meta-information, there is a risk for inflation if this process is continued by further extracting meaningful information of meta-information to create “meta-meta”-information (as captured by the human-computer interface bottleneck[225]). In the end, the aggregated information presented to the operator may be too distanced from raw data that the embedding transparency is lost.

5-7-2 Personalizing transparency

In this study, personalization was restricted to the automation’s resolution advisories (i.e., conformal or nonconformal). In addition to conformance, there are several other interesting personalization aspects that deserve recognition and further research. This includes the automation’s communication style (e.g., automation etiquette[114]) and its interface characteristics, such as how information is portrayed by symbols, colors, contrast, and resolution.

The questionnaires revealed differences between participants in their perceptions of the two interfaces, indicating that the degree of transparency appropriate for one participant was not appropriate for another participant. Normally, such differences are considered negligible and disregarded in interface design. However, the human-machine interaction would likely benefit from personalizing interfaces to the individual’s preferences and needs. For example, by the system automatically adapting to the individual. Alternatively, as with adaptable automation, allowing the individual to influence the interface, and determine how much information that is presented, how it is structured, aggregated, and communicated. A relatively simple method to achieve this is by allowing users to directly control the appearance and functions of the interface. A more advanced method for achieving this can be achieved by the system silently monitoring the user and change the interface according to the user’s way of working.

5-7-3 Varying SSD transparency

The manipulation of SSD transparency in this paper did not vary the internal transparency of conformal and nonconformal resolution advisories in relation to Wang and Benbasat[209] three explanation types: why and how a solution was derived, and the tradeoffs considered. Explanations could, however, have been used to reflect information on whether the advisory was based on the participant’s own solution style and patterns in problem-solving, and how this style was determined. As such, a ‘how’ explanation would explain the reasoning underlying how the conformal resolution advisory was selected: for example, a vector was chosen because this is how you typically solve these types of conflicts. A ‘why’ explanation would justify the advisory: for example, this solution is more efficient because it takes aircraft A closer to the destination. A ‘tradeoff’ explanation would contrast alternatives and guides selection: for example, an alternative would be to vector the other aircraft, but it requires a larger deviation.

Furthermore, although the objective of EID is to visualize the affordances for activity in relation to a particular work domain, it could perhaps be applied more specifically for visualizing a decision support system’s reasoning underlying recommendations and solution advisories. Across many domains, such specific solution advice has traditionally been provided in text-based formats. For instance, studies exploring ATC decision aids for conflict resolution usually provide the controller with a list of solutions in a tabular text format.[31,35,234] Like Woods *et al.*[199] note, the process underlying the generation and prioritization of text-based solutions is opaque to the controller, making it difficult for the controller to comprehend the solution(s). The interpretation uncertainties can lead to discrepancies between how the controller perceives the system’s suggested solution, and the solution preferred by the controller.[155] Furthermore, a list can encourage the operator to consider all alternatives before making a choice. Using text-based explanations to support an advisory (e.g., as advocated in recommender systems research) is not practicable in many domains given time pressures and workload demands. As an alternative to text, a rationale underlying the system’s advice can be provided graphically in the interface, by visually integrating information directly with advisories. Further research should investigate whether EID can be used for purposes of providing such explanations, for example by visualizing which constraints and complex relationships that have been considered when determining the suggested solution.

5-7-4 Transparency and conformance

Results showed that participants were more likely to use speed and combined solutions to solve the conflict when using the TRI SSD. Although this was a desirable and not surprising effect attributed to the increased interface transparency, it implies that conformance also changes with the transparency. The conformance manipulation, however, was based on solutions made when using the HB SSD, for which speed and combined solutions were less common. As such, a heading solution may have been conformal when using the HB SSD, but not when using the TRI SSD. In part, this can explain why participants’ agreement was higher for nonconformal advisories than conformal advisories in the TRI SSD condition. The TRI SSD may have improved participants’ understanding of nonconformal advisories and guided their attention to solutions they did not think of. Moreover, the conformal advisories, based on solutions when using the HB SSD in the prequel simulation, may have appeared impractical when using the TRI SSD.

This points to an underlying issue in regards to not only transparency, but all aspects relevant to information visualization in regards to personalized conformal automation. The implication is that any personalization of automation may considerably depend on the interface design. Any changes made to the interface may render the personalization invalid.

5-8 Conclusion

This paper has argued that as automation and in particular decision aids become more advanced, they need to communicate more information, not less. This study confirmed clear differences in participants' perceptions of two interface representations with varying transparency, with the information richer triangle representation better facilitating understanding of a conflict and options for solving it. Increased transparency achieved through the inclusion of more information, however, came at a cost reflected in the triangle display being perceived as more cluttered. Furthermore, results indicated differences in the use of interface representations, with the number of interface inspections decreasing, and the number of speed clearances increasing when using the triangle. However, neither conformance nor transparency significantly influenced acceptance or response time of resolution advisories.

It must be acknowledged that findings in this study are susceptible to the adverse effects of a small sample size. Although the conclusion can be drawn that making an interface more transparent can benefit the understanding of a decision aid, the lack of strong complementary results in the real-time experiment questions the practical implications of transparency. The manipulation of transparency did not include explanations for *why* a certain advisory was made, as advocated by recommender systems research. As such, transparency was not connected to specific advisories and did not explain why solution X was suggested in place of solution Y. Instead, drawing from ecological interface design, it was proposed that for control problems in air traffic control, this increased transparency could be achieved by providing more information on the environmental constraints and possibilities affecting maneuverability. Notwithstanding these limitations, this study emphasizes that design decisions about the transparency of an interface represent a balancing act. While information-rich interfaces can benefit understanding, too much information can complicate information search and processing, and potentially overload the operator with negative effects on overall system interaction.

6

Consistency and agreement in conflict resolution

In previous chapters, the theoretical foundations and first empirical studies from exploring strategic conformance have been reported. Taken together, results from three real-time simulations show that controllers rejected their own conformal solutions, sometimes as often as in one out of four trials. Provided humans solve similar problems similarly over time, it is possible to model their behavior and individually tailor decision support that conforms to their unique solutions style. But if humans solve similar problems differently over time, individually tailoring of decision support may not be feasible. Therefore, it is necessary to determine to what extent controllers are consistent in their CD&R performance. This chapter provides a state-of-the-art review of human consistency and agreement in CD&R decision-making, followed by an analysis of controllers' solution consistency and agreement of manually solved conflicts recorded in previous empirical real-time simulations (Chapters 3 through 5).

The contents of this chapter are based on:

Paper title Consistency and agreement in decision-making: Conflict resolution in air traffic control

Authors Carl A. L. Westin, Clark Borst, Brian H. Hilburn

Prepared for Human Factors

Abstract

Although people are recognized to vary in their decision-making and problem-solving preferences, automated decision aids have generally not been designed to consider or accommodate individual cognitive differences. However, researchers have argued that personalized automation will become increasingly important for facilitating safe and efficient working relationships between humans and automated agents, particularly in complex, dynamic, and time-critical automation-dependent environments such as air traffic control. An issue in the development of such individual-sensitive automation is that people may be inconsistent in their decisions and actions. Previous research provides an inconclusive picture in regards to consistency and agreement in conflict resolution. This paper presents two studies investigating decision-making consistency (intra-rater reliability) and agreement (inter-rater reliability) in conflict resolution among air traffic controllers. Conflict solutions were collected in real-time simulations, in which participants unknowingly encountered the same conflict four times. Consistency and agreement in conflict solutions were determined against a solution framework building on and extending conflict resolution strategies identified in previous research: the solution parameters hierarchy classification, the control problem classification, and solution geometry classification. Controllers were consistent but did not agree. As such, controllers cannot be considered homogeneous and their individual differences in conflict resolution decision-making need to be acknowledged in human factors models and automation design.

6-1 Introduction

Similar to other control-room environments, ATC automation design has generally followed a "one-size-fits-all" approach suitable to an average, generic, stereotype controller. Although controllers currently can customize basic system properties, such as interface appearance (e.g., zoom, brightness, information layers), there are no current applications sensitive and adaptive to the controller's cognition. The generalization is further exacerbated by a widespread assumption that controllers are homogeneous which can partly be attributed to rigorous screening and selection procedures and training that converge towards creating a uniform work force.[235,236]

Human factors researchers have argued for more individual-sensitive, *idiographic*,[237] automation design approaches, as opposed to current generic ones, suited to the individual's abilities, needs, and preferences.[126,154,155,197,198] Similarly, individual differences has been outlined as pivotal in future ATC automation design.[97–99] In support of this, extensive research on cognitive styles and personalities has demonstrated that people search and process information, make decisions, and solve problems differently.[37,238–240] Despite this knowledge, however, individual differences have traditionally not been accommodated for in automation design, although exceptions can be found in research on automated cars (e.g., personalized adaptive cruise control systems)[241–243] and intelligent agents (e.g., unmanned vehicles and robots).[197]

In order for a system to personalize its support, the system must know something about who it is interacting with and how that person thinks. For example, in the context of ATC conflict resolution the system must know how the controller reasons when determining a solution. As a first step, however, it is necessary to determine the degree of *agreement* between controllers (i.e., the inter-rater reliability) since there is little benefit of personalizing a system if all think the same or agree on the same solution. Second, it is necessary to determine how *consistent* controllers make decisions and solve conflicts over time (i.e., the intra-rater reliability). Previous research provides an inconclusive picture in regards to controller consistency and agreement in conflict resolution. More generally, research by Shanteau and colleagues[236] indicates that experts often disagree and that there appears to be higher consistency than agreement. Consistency and agreement may be especially low in fields that have no “gold standard” (i.e., a benchmark, correct, or optimal solution) such as ATC.[131,244]

This paper aims to determine the degree of consistency and agreement in operator decision-making, using ATC conflict resolution as the example domain. However, previous research investigating controllers’ decision-making in conflict resolution has been limited in terms of subjective data collection methods and use of static traffic scenarios. For this purpose, a novel experimental design was developed for investigating controllers’ conflict resolution performance in a dynamic environment. Furthermore, a novel conflict solution framework was developed against which consistency and agreement could be measured objectively.

A series of human-in-the-loop simulations were conducted in two separate studies. These investigated controller’s *conflict solving patterns* (i.e., actions taken to solve a conflict) across repeated conflicts (each conflict was repeated four times) and identified their individual and consistent *problem-solving style* (i.e., solving the same conflict using the same pattern in at least 75% of all repetitions). Participants’ problem-solving styles were then compared to determine the agreement between participants. Since conflict solving is highly situation dependent, different scenarios and conflicts were used in the two studies. In total five different conflicts with varying geometries (e.g., convergence angles) and parameters (e.g., relative distances and speeds) were investigated. The empirical findings can be helpful for determining the benefit of developing decision-support systems sensitive to the individual.

6-2 Decision-making in ATC conflict resolution

Automation design approaches typically incorporate end users early in the design process. Similarly, several automated decision aids for conflict resolution have been developed around controller elicited decision-making models.[10] These typically view decision-making as a hierarchical search process based on controllers' use of problem-solving heuristics, often separating conflict detection[33,61] and conflict resolution.[24,29,31] Examples include the Conflict Resolution Assistance (CORA),[35] the Cube model[132] and the associated COCOS (Controllers' strategies integrated into a conflict resolution system) algorithm,[31] and the Intelligent System for Aircraft Conflict Resolution (ISAC).[131]

Acquiring a detailed understanding of controllers' core cognitive work and decision-making processes represents a great challenge for the development of personalized CD&R decision aids. Discrepancies have lead to models and design frameworks that poorly suit controllers' working methods.[19,155] Inferring cognitive processes is problematic since controllers typically struggle to express their decisions and actions,[29,88] a commonly reported trait among experts in general.[236,245] The naturalistic decision-making paradigm argues that experts typically do not consider several options when making decisions. Rather, decisions reflect the first credible solution conceived intuitively through a process of pattern-recognition using tacit knowledge and previous experience.[245] Analogously, researchers have identified CD&R as a naturalistic decision-making process,.[87,131,246]

Controllers are believed to develop a strategy and solution "library" that is accessed when encountering a conflict.[40,89,247] They appear to rely on the first conceived strategy, especially during periods of high workload.[189] Furthermore, controllers have been shown to rely on decision-making heuristics, typically described as "rules of thumb," for both identifying and solving conflicts. For example, conflicts tend to be approached in pairs and sequentially rather than globally if multiple aircraft are involved.[29] However, controllers have been found to rarely interfere with both conflicting aircraft.[88] Furthermore, although objectively both aircraft involved are in conflict with each other, one aircraft is typically assigned the "trouble maker" causing the conflict.[29] In addition, controllers have been found to avoid speed for solving conflicts en-route.[29,31]

Kirwan and Flynn[29,35] suggest that resolution strategies primarily are shaped by four aspects. *Factors* represents contextual parameters such as phase of flight, conflict parameters, and remaining distance to destination. *Rules* encompass rules of the air and other airspace constraints. *"No-no's"* represents valid solutions that controllers refrain from such as vectoring an aircraft in front of another, vectoring it 180 degrees, or climbing it when close to its destination. *Principles* represent general guidance strategies such as maintaining fairness, avoiding knock-on effects,

TABLE 6-1: Conflict resolution strategies (after Fothergill *et al.*[24])

	Lateral resolution strategies
1.	Vector aircraft behind other ("point behind").
2.	Vector away from potential conflicts ("take out").
3.	Assign a track parallel with own route.
4.	Take out for five miles, then put back on track.
5.	Vector aircraft ahead of other ("pass in front").
	Vertical resolution strategies
6.	Cut off at nearest available level on climb.
7.	Cut off at highest possible level on climb.
8.	Request nearest level above conflicting aircraft.
9.	Descend to nearest available level.
10.	Assign the only level available.
11.	Step climb/descent.
12.	Expedite climb/descent.
13.	Report maintaining to ensure at least 1000 feet separation.

and vectoring the aircraft farthest away from a crossing point behind the aircraft closer to the crossing point. Furthermore, research suggests that controllers cope with high workload, fatigue, and effects of age by using risk-averse strategies that are more conservative, involve early intervention, and resolve situations directly (in contrast to a wait-and-see strategy).[40, 80, 189]

More recently, Fothergill *et al.*[24] identified thirteen conflict resolution strategies and tied these to specific resolutions. The strategies are shown in Table 6-1, divided into five lateral and eight vertical ones in no special order. The authors identified the two most common strategies as *cutting of a climbing aircraft* at the nearest available level (vertical conflicts) and *vectoring one aircraft behind* the other (crossing conflicts). Both are described as quick and easy, ' "set and forget" strategies, because they can be achieved by one instruction and require less monitoring. In contrast, the less common strategy of *vectoring one aircraft ahead* of another is deemed considerably riskier and demanding to calculate and monitor. Ultimately, the authors concluded that strategy choice is individual and highly dependent on the conflict detection process and perceived situation complexity.

6-2-1 Consistency in conflict resolution

Although no literature was found on how consistent controllers are in conflict resolution, a few studies addressing consistency in relation to other ATC tasks were found. In a large survey involving 100 controllers, almost a third of all respondents reported that they try to consistently use the same techniques and patterns when looking for information and scanning a sector.[189] In contrast, research by Redding *et al.* indicates that controllers vary their use of decision-making strategies during work,[40] which points to inconsistency.

Controller consistency has also been investigated by means of assessing expertise using the Cochran-Weiss-Shanteau (CWS) index,[248,249] derived by calculating the ratio of *discrimination* (degree to which similar stimuli are discriminated from one another) to *consistency* (degree to which a repeated stimulus is judged the same). In a high-fidelity simulation involving twelve controllers, Thomas *et al.*[250] found a moderate correlation between the degree of controllers' consistency, as measured by the CWS index, and simulator measures of their performance and efficiency. Controllers performed less consistently with deteriorating performance measured by the number and duration of separation errors. Similarly, they performed less consistently with decreasing task efficiency measured by the number of altitude and heading instructions issued. Although results indicated that consistency also decreased with increased complexity (measured by the number of aircraft controlled and handed off) the correlation was not significant. In addition, the authors noted that some controllers were considerably more consistent than others.

In a study investigating the consistency in traffic complexity judgments, controllers were asked to pairwise compare seven static traffic scenarios and identify the most complex one in each pair. The method for determining consistency based on analyzing pairwise comparisons was derived from the Analytic Hierarchy Process (AHP).[251] Controllers were found inconsistent in their judgments with only one out of five found consistent according to the AHP's suggested threshold of attaining a consistency ratio below 0.1 (with 0 indicating that judgments were perfectly consistent and 1 indicating that they were purely random).[252]

6-2-2 Agreement in conflict resolution

Research suggests that there may be considerable variability between controllers' judgment and decision-making preferences in CD&R,[35,131] especially between ATC facilities,[189] regions, and nations.[35] In the development of the conflict resolution system ISAC, differences in controllers' view of the same conflict complicated the eliciting process of parameters underlying the system's case-based reasoning model for conflict resolution.[131] Validation trails with other controller-based CD&R automation have shown that these systems do not match the variability in controller solutions, especially in terms of aircraft choice and intervention type preferences.[31,32,88] Controllers are acknowledged to differ in judgment performance[218] and how they prefer to work.[235] Differences in work-styles have been traced back to training and the style of the instructor.[87]

A few ATC studies have investigated individual differences using cognitive styles. This broad research area argues that a person's cognitive style represent heuristics (often unknown to the individual and expressed intuitively) rooted in the individual's cognitive preferences.[37] Mogford *et al.*[87] used a word-shape preference test[253] to investigate controllers' problem-solving differences. The test distinguishes between a preference for relying on words as in the verbal-analytic processing style (greater activation of left cerebral hemisphere) and shapes and images as in the spatial-holistic processing style (greater activation of right cerebral hemisphere). Style preferences were split: 31% showed a verbal-analytic style, 22% a spatial-holistic style, and 47% showing no preference for either style. This suggests that controllers perceive and process information differently.

Two other studies used the cognitive style embedded-figures test[254] to investigate controllers' pattern recognition speed in a visual search task. The test differentiates those who quickly detected a simple figure in a larger complex one (field-independent) from those who do not detect the figure (field-dependent). Maliko-Abraham[255] found that military controllers were more field-independent (74% air force ATC and 68% naval ATC) than the control group (47% of non-controllers). Similarly, van Eck *et al.*[256] found that field-independence varied with years in ATC training, with controllers in their final fourth educational year being significantly more field-independent than the reference general population. While the complementary results of both studies suggest a high homogeneity among controllers in pattern recognition time, there are at least two concerns. First, both studies compared controllers and non-controllers, rather than focusing on differences between controllers only. Second, performance variability may not be adequately captured by the dichotomous cognitive style measure.

6-3 Towards a conflict solution framework

Research to date has identified generic strategies that broadly describe conflict resolution preferences, typically limited to a few vertical and lateral resolutions. These have, however, not considered solutions in detail, such as which aircraft to interact with, the degree of a turn, or rate of a vertical intervention. Moreover, focus has predominantly been on subjective techniques (e.g., interviews, focus groups, and questionnaires) in combination with static scenarios.[24,29] These fail to capture the time pressures and reactive elements of the real world that influence decision-making.

In addition, these methods are prone to hindsight biases. Possibly, conflict solution behavior can better be explained in a way not readily apparent to the controller, as reflected in naturalistic decision-making theories.

To address these shortcomings, a framework was developed that allowed for a more objective and systematic analysis of controller strategies for conflict resolution. The framework draws on resolution strategies identified in previous literature, in particular those suggested by Fothergill *et al.*[24] The final structure, however, was exploratory and iteratively derived through analysis of more than 500 conflicts solutions collected in real-time simulations. In the end, three strategy classifications were defined, which all depict solutions in a decision-tree hierarchy.

Figure 6-1 illustrates the framework and the relationship between the three classifications. Conflict solutions are defined by five decision stages (DSs) of progressive order of granularity. Because the three classifications define solutions differently, reflecting contrasting preferences for problem-solving, the first high-level DS differs between classifications. Note that this exclusivity only is true for the first DS of each classification. The subsequent four DS reflect detailed solution preferences that are shared by all three classifications. *Choice of aircraft* to interact with repre-

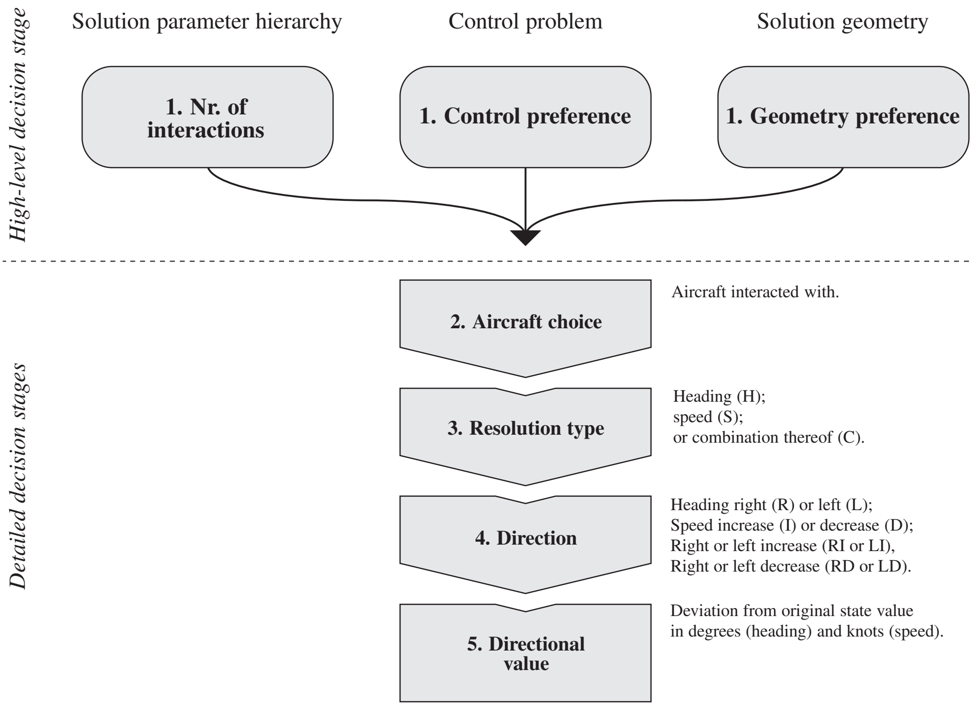

FIGURE 6-1: Consistency classifications framework

sents the second DS (including order of interactions if more than one interactions is made). Third is the *resolution type*, which identifies the general category of control maneuvers (i.e., altitude, heading, speed, or a combination thereof). The fourth DS defines the resolution type by assigning *direction* (e.g., left or right vector). The final, and fifth DS specifies the exact *directional value* of the resolution (e.g., vector 035 degrees or increase speed to 300 knots).

Although time of intervention is an essential solution component, it was considered independently. Time is a qualitatively different measure from other solution parameters that define the physical interaction, although the conflict geometry, and thus solution options, changes over time. Moreover, the time of intervention does not accurately reflect the time of detection or when the solution is conceived as the controller can choose to postpone intervention and monitor the conflict.

6-3-1 Solution parameters hierarchy classification

The *solution parameters hierarchy* argues that controllers distinguish between the aircraft in conflict and selectively decide on which to control, which resolution type to use and in which direction to deviate. The hierarchy reflects the algorithm architecture of several CD&R systems that have been based on controller-elicited knowledge, such as ISAC,[131] COCOS,[31] and the artificial intelligence A* graph search model.[32] The specific choice of aircraft to interact with is also supported by research indicating that controllers at times assign one aircraft as the "trouble maker" causing the conflict.[29] The first DS specifies the *number of interactions* made to solve the conflict (i.e., whether one or both aircraft were controlled).

6-3-2 Control problem classification

Alternatively, a conflict can be viewed as a *control problem*. This perspective focuses on the control action implemented to solve the conflict (e.g., vector aircraft ahead or behind),[24] and argues that a conflict only can be solved by one aircraft going behind, in front, above, or below the other. The aircraft interacted with first is designated the *controlled* aircraft that is instructed to avoid the other *intruder* aircraft. As such, the first DS describes the solution in relation to where the controlled aircraft will pass the intruder aircraft (laterally *behind* or *in front*, or vertically *under* or *above*). In contrast to the solution parameters hierarchy, the control problem classification disregards the number of interactions made as only the first interaction is relevant. Moreover, the first DS does not differentiate between aircraft interacted with as only the control action is of interest. For example, for DS 1, vectoring aircraft A behind B is considered the same as vectoring B behind A. In both solutions, the controlled aircraft is vectored behind the intruder.

6-3-3 Solution geometry classification

The *solution geometry* classification is derived from an exocentric assessment considering a solution's resulting spatial geometry. It acknowledges that a conflict solution is based on the relationship between two or more aircraft and their constraints as they evolve over time, rather than on discrete information about aircraft state and position.[229] The first DS identifies the resulting *spatial relationship* between the conflicting aircraft. Only four alternatives exist, with aircraft A situated behind B, B behind of A, A above B, or B above A. Note, however, that each relationship can be framed in two ways. For example, vectoring aircraft A ahead of B is considered the same as vectoring aircraft B behind A. The resulting relationship can be framed as aircraft A situated ahead of B, alternatively as aircraft B situated behind A. Importantly, this stage disregards number of interactions, aircraft choice, and which aircraft that is controlled.

6-3-4 Mutually exclusivity and consistency

At the highest decision stage, the three classifications represent three contrasting perspectives of conflict solving for which consistency is affected differently. For example, consider the following repeated situation: a right angle conflict between aircraft A and aircraft B, solved once by vectoring A left behind B, and once by first vectoring B to the right and then vectoring A to the right so that B passes behind A.

The solution parameter hierarchy defines the two solutions as different because in the first situation only aircraft A was interacted with, while in the second situation both aircraft were interacted with. The control problem classification initially disregards number of interactions (only considers the first interaction) and specific aircraft choice, meaning that the first solution is considered the same the second solution. In both situations, one aircraft is vectored behind the other. According to the solution geometry classification, however, these two solutions would create two different geometries: one with aircraft B in front of A, and the other with aircraft A in front of B.

However, there is one exception for which the classifications overlap so that a participant can be found consistent according to all three. This occurs when two (or more) solutions are identical according to the first two DS in the solution parameter hierarchy. Say, for instance, that both solutions consisted of vectoring A left behind B. According to the control problem classification this would be defined as a consistent control action (always behind). According to the solution geometry classification the same spatial relationship has been accomplished in both solutions.

6-4 Method Study 1

Study 1 comprised four different scenarios, each containing different asymmetrical conflict geometries, and different aircraft parameters, that favored a specific solution. Due to this solution bias, it was hypothesized that participants would solve repeated conflicts similarly (high consistency and high agreement). Note that the data was collected in previous simulation (prequel simulation) part of a larger study investigating controller acceptance of conformal (personalized) decision support.[101]

6-4-1 Participants

Participants consisted of sixteen controllers (one female and fifteen males) working at Shannon Area Control Center (ACC), Ireland. Participation was voluntary. Age varied between 26 and 44 years (mean = 31) and experience varied between zero to ten years (mean = 2.5). Twelve controllers were actively working en-route positions, while three were en-route trainees at the end of their on-the-job training. One controller was actively working the tower position.

6-4-2 Simulator

A Java-based ATC simulator package, using OpenGL extensions, was used in parallel on three portable computers, each connected to a 21-inch monitor with a minimum resolution of 1280x1024 pixels. The simulator ran at 4x normal speed. Aircraft plots on the display were updated every second to simulate a 1 Hz radar update frequency.

The interface consisted of a simplified traditional ATC radar view display. The simulator environment consisted of a hypothetical squared en-route sector (50 x 50 nmi). Aircraft interaction, achieved by means of mouse and keyboard, was facilitated through the Solution Space Diagram (SSD). The SSD display was not always visualized but activated when an aircraft was selected (clicked on). The SSD (i.e., Figure 6-2) is a novel separation assistance tool based on an ecological interface design approach under development at Delft University of Technology.[102, 232] It visualizes the constraints and possibilities affecting the travel space of a selected aircraft in relation to the other *intruder* aircraft. The relative position of intruder aircraft are visualized by color-coded *no-go zones* in a circular 360 degrees diagram around the *controlled* aircraft (the diameter of the diagram reflects the controlled aircraft's current speed). Each no-go zone represents the protected zone (typically 5 nmi in diameter in en-route airspace) of an intruder aircraft, and are either yellow (separation loss in one to four minutes) or red (zero to one minute). Separation is assured by making sure the *velocity vector* of the controlled aircraft is outside the no-go zone(s). More information about the SSD is provided in Appendix B.

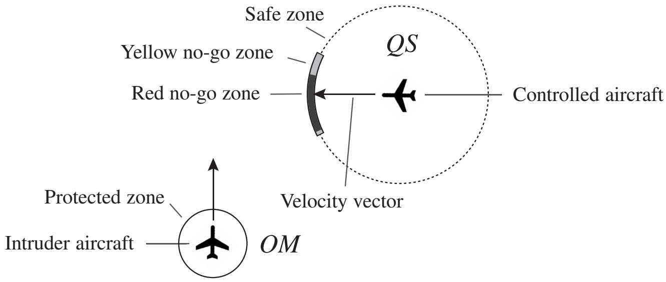

FIGURE 6-2: Right angle conflict with heading band SSD shown for QS.

The SSD's impact on controllers' conflict solving was expected to influence both their consistency and agreement by amplifying a particular solution in case of asymmetrical conflicts and different aircraft parameters. As such, for asymmetrical conflicts (Study 1), the SSD makes the biased solutions more salient (e.g., by visualizing a larger safe zone) For symmetrical conflicts (Study 2) with no biased solution, a preferred biased solution is not available.

6-4-3 Traffic scenarios and designed conflicts

Scenarios and conflict parameters are provided in Table 6-2. Four baseline traffic scenarios, each with a specifically designed conflict, were created. Rotations were used to ensure that solutions would not be influenced by the recognition of repeated scenarios. Different conflict geometries were arbitrarily chosen, each consisting of a crossing between two aircraft occurring approximately in the center of the sector. All geometries were biased, meaning that aircraft parameters were not reciprocal and therefore mathematically favored a certain solution. For example, for the conflict in Scenario 1 (shown in Figure 6-2), the closest point of approach (CPA) occurs when OM is 2.6 nmi behind QS. In terms of least additional track miles the optimal solution would be to vector OM right to go behind QS. However, additional traffic present in the sector restricted this solution and instead favored vectoring QS to the right ahead of OM. For each scenario, the following solutions were expected:

- Scenario 1: vector QS to the right ahead of OM.
- Scenario 2: vector PA left, or increase speed, ahead of RG.
- Scenario 3: vector RG left ahead of SM.
- Scenario 4: vector QS right ahead of PA.

TABLE 6-2: Scenario groups and designed conflict parameters, Study 1

S	Rotations	Aircraft	Conflict Angle (CPA)	TSL	Aircraft 1 & 2 ID: heading, speed (speed envelope)
1	0°, 45°, 90°, 270°	22	90° cross (OM 2.60 nmi behind QS)	77 s	QS: 270°, 220 kn (150-350); OM: 000°, 200 kn (120-250)
2	0°, 45°, 90°, 180°	17	90° cross (RG 2.60 nmi behind PA)	83 s	RG: 225°, 220 kn (150-350); PA: 315°, 200 kn (150-250)
3	0°, 45°, 90°, 180°	14	63° cross (SM 2.54 nmi behind RG)	70 s	RG: 063°, 290 kn (210-340); SM: 000°, 220 kn (170-280)
4	0°, 45°, 180°, 270°	16	127° cross (PA 3.20 nmi behind QS)	60 s	QS: 296°, 200 kn (150-250); PA: 063°, 200 kn (150-250)

S = Scenario; CPA = Closest Point of Approach; kn = knots; nmi = nautical miles, s = seconds, TSL = Time to Separation Loss

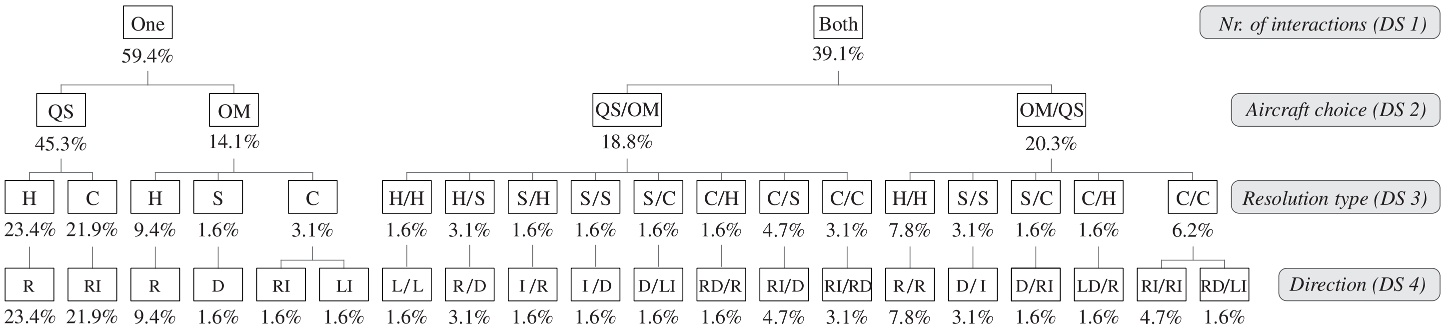

H = Heading; S = Speed; C = Combined; R = Right; L = Left; I = Increase; D = Decrease

FIGURE 6-3: Solution distribution for the Solution parameter hierarchy in proportion (%) of total solutions, Scenario 1, Study 1.

6-4-4 Experimental design

A qualitative design was used with scenario repetitions as the independent variable (each scenario repeated four times). Consistency was determined from analyzing the similarity in participant's solution patterns across repetitions. A pattern was defined by the actions taken to solve a conflict. All three classifications in the conflict solution framework were considered when identifying patterns. A participant was considered consistent if the same pattern, according to the same classification, was found to solve the conflict in three out of four or four out of four repetitions. in A consistently applied pattern was described as a participant's problem-solving style. Agreement (inter-rater reliability) was subsequently analyzed by comparing problem-solving styles between participants.

6-4-5 Procedures

Participation lasted approximately two hours. Following introductory briefing and consent procedures, participants received roughly 50 minutes of simulator training. The experiment simulation comprised sixteen two-minute scenarios (each baseline scenarios once and then three repeats). Scenario order was varied according to a Latin Square design. The two main tasks were to resolve conflicts and clear aircraft to their exit points The sector environment was presented as a futuristic free-routing sector with a considerably higher traffic density and throughput than current day. Participants were supported by a short-term conflict detection system that provided an visual and auditory warning before separation loss occurred. Traffic was restricted to the horizontal plane (flight level 290) with no wind or other meteorological conditions present.

6-5 Results study 1

Data from 256 solutions, 64 per scenario group, was collected. This consisted of four scenarios repeated four times across sixteen participants. One participant's data file in Scenario 2 was corrupted, leaving in total 255 solutions to analyze.

6-5-1 Solution parameter hierarchy analysis

The solution distribution in Figure 6-3 provides an overview of all 64 recorded solutions across the first four DSs for Scenario 1. Out of all solutions, 98.5% consisted of either interacting with one aircraft (59.4%) or both (39.1%). DS 4 shows that 20 unique solutions were identified in total. The most common solutions consisted of interacting with QS (DS 3: 45.3%), either implementing a right vector (23.4%) or a right vector with a speed increase (21.9%). Solutions consisting of three or more interactions (1.6%) were disregarded from analysis as they did not appear to reflect a deliberate solution strategy, but more likely the result of a failed solution for which the participant tried to salvage the situation by re-solving the conflict.

Similar variations were observed in Scenarios 2, 3, and 4 (see Appendix E for complete solution distributions). For Scenario 2, 32 different solutions were identified across 63 recorded solutions. There was no unanimous preference for solving the conflict according to the first two DS, with both alternatives equally common. In Scenario 3, 27 different solutions were identified across 64 recorded solutions. Interacting with both aircraft was most common (59.4%). In Scenario 4, 19 different solutions were identified. The two most frequent solutions accounted for 56.2% of all variation and consisted of either vectoring QS right (35.9%) or simultaneously increasing the speed and vector QS right (20.3%).

Figure 6-4 shows the proportion of consistent participants per identified problem-solving style, for each scenario and DS. It shows the spread of problem-solving styles used, and which styles participants agreed on. All controllers were found to have consistently solved the designed conflict in at least one scenario. Only two controllers consistently solved the conflict in all scenarios. In regards to DS 1, single aircraft solutions were most common in Scenarios 1 (43.8%) and 4 (75.0%). In Scenario 2, participants were split between interacting with one or both aircraft. In contrast, the majority of participants solved Scenario 3 by interacting with both aircraft. Relative large groups of participants in Scenarios 1, 2, and 3 were found

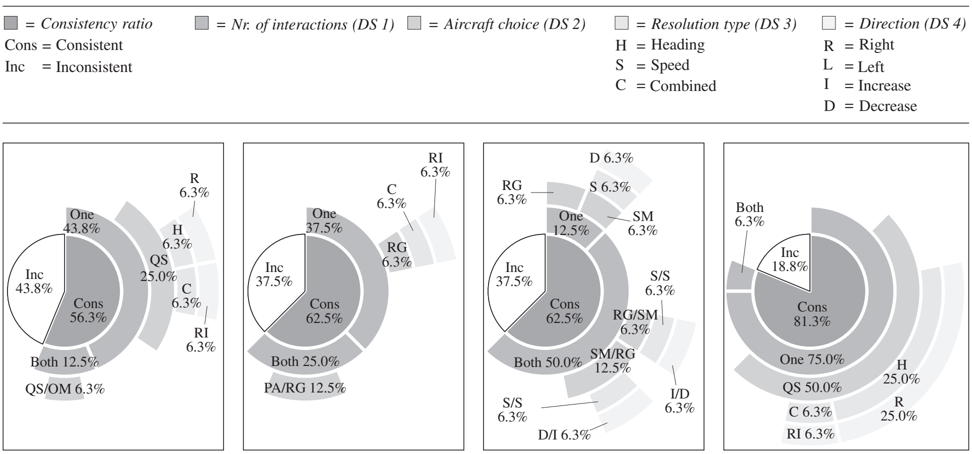

(a) Scenario 1 (b) Scenario 2 (c) Scenario 3 (d) Scenario 4

FIGURE 6-4: Solution parameter hierarchy sunburst charts showing consistency and agreement, Study 1. Percentages indicate proportion of participants ($N = 16$).

to inconsistently have solved the conflict. Inconsistency increased notably for each DS in the hierarchy. In Scenario 2, for example, only 6.3% of all participants were found consistent according to DS 4.

6-5-2 Control problem analysis

Figure 6-5 shows the complete distribution of the solutions according to the control problem classification for Scenario 1. The overall preferred solution consisted of instructing the controlled aircraft to go in front of the intruder (DS 1). QS was most frequently controlled (DS 2) and instructed to turn right or turn right and increase speed (DS 3 and 4). OM was most frequently controlled when behind solutions were used. In front solutions were particularly common in Scenarios 2 (77.8%) and 4 (79.9%). Complete distributions for Scenarios 2, 3, and 4 are provided in Appendix E. In Scenario 2, aircraft choice for in front solutions (DS 2) was fairly evenly distributed between RG (34.9%) and PA (42.9%). For behind solutions, RG was more frequently controlled (19.1%) than PA (3.2%). The highest variability was found in Scenario 3, with a slight preference for in front solutions (53.1%). In Scenario 4, QS was selected as controlled aircraft in 73.4% of all solutions.

The proportion of consistent participants and their problem-solving styles across all scenarios are shown in Figure 6-6. Participants generally agreed on instructing the controlled aircraft to go in front of the intruder. Notably, in Scenarios 2 and 4, the controlled aircraft was never consistently taken behind the intruder. The most common problem-solving styles found in Scenarios 1 and 4 consisted of vectoring QS in front of the intruder. In both Scenarios, all participants found consistent according to DS 1 were also found consistent according to DS 2. Note, however, that only

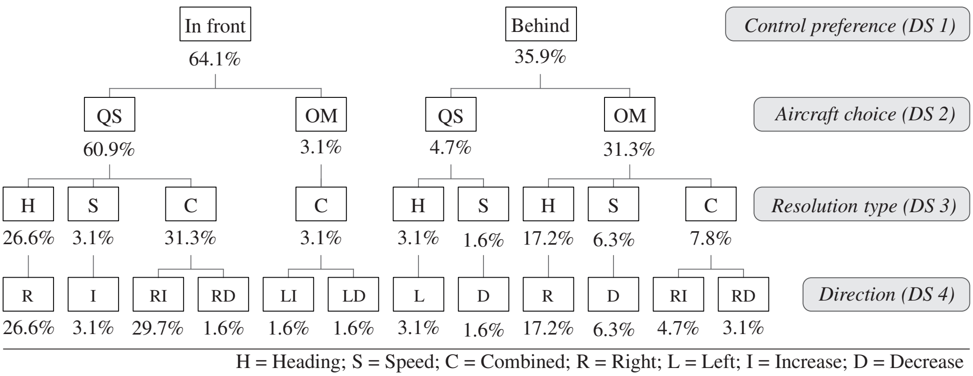

FIGURE 6-5: Solution distribution for the Control problem classification in proportion (%) of total solutions, Scenario 1, Study 1.

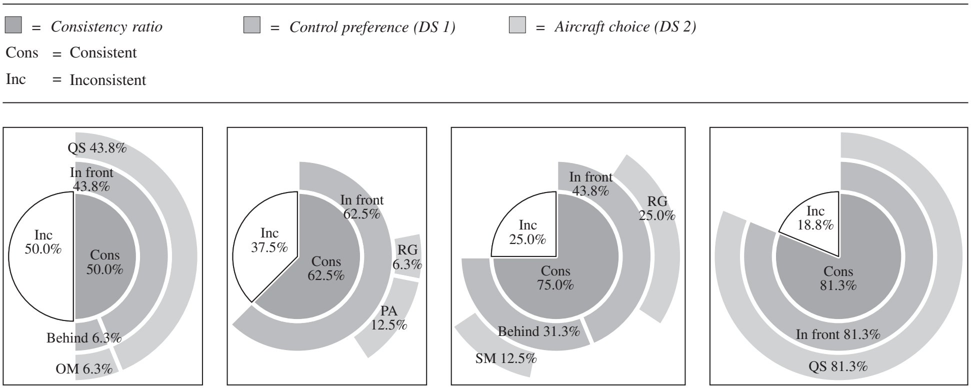

(a) Scenario 1 (b) Scenario 2 (c) Scenario 3 (d) Scenario 4

FIGURE 6-6: Control problem sunburst charts showing consistency and agreement, Study 1. Percentages indicate proportion of participants ($N = 16$).

half the group consistently solved the conflict in Scenario 1. The problem-solving styles identified for Scenarios 2 and 3 only partly correspond with the expected solution (i.e., controlled aircraft in front of the intruder). For DS 2 in Scenario 2, there was some disagreement on whether to choose RG or PA as controlled aircraft. In Scenario 3, participants were found to consistently having the controlled aircraft go either in front or behind the intruder (DS 1).

6-5-3 Solution geometry analysis

Figure 6-7 shows the solution distribution according to the solution geometry classification for Scenario 1. Almost all solutions consisted of QS passing ahead of OM (DS 1). Further analysis showed that QS was instructed to go ahead of OM more often than the opposite (DS 2), with heading or combined instructions generally favored (DS 3). Detailed solution distributions for Scenarios 2, 3, and 4 are found in Appendix E. In Scenario 2, the most frequent solution consisted of PA passing ahead of RG (63.5%). PA was more often controlled ahead of RG (42.9%) than RG was controlled behind PA (20.6%). In Scenario 3, there was an overall preference for solving the conflict with RG passing ahead of SM (76.6%), either by controlling RG ahead (39.1%) or SM behind (37.5%). In Scenario 4, a strong homogeneity was found for DSs 1 and 2, with QS passing ahead of PA (93.8%), most often achieved by controlling QS (73.4%).

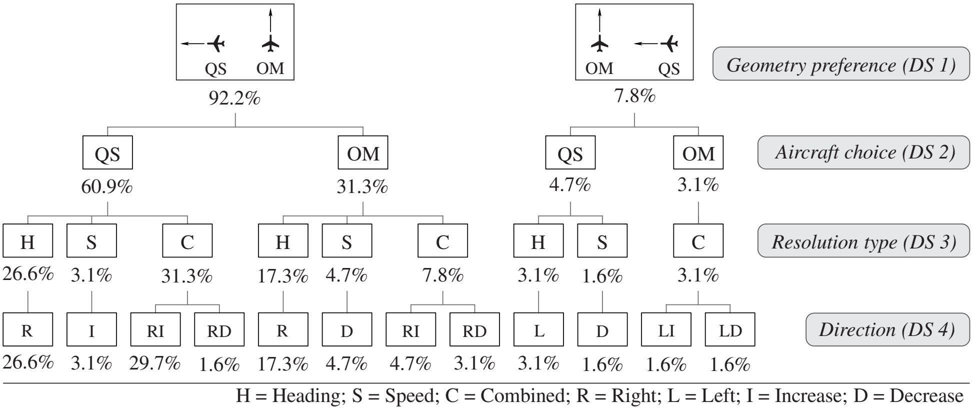

FIGURE 6-7: Solution distribution for the Solution geometry classification in proportion (%) of total solutions, Scenario 1, Study 1.

Figure 6-8 shows the proportion of participants with consistent problem-solving styles who arrived at the same geometry across repetitions. Overall, participants were found highly consistent, except for Scenario 2 in which more than half (56.2%) solved the conflict inconsistently. Agreement was high in all scenarios, especially in Scenarios 1 (93.8%) and 4 (100.0%). Overall, all scenarios were solved according to expectations when considering DSs 1 and 2. In general, participants preferred to solve respective conflicts with QS passing ahead of OM (Scenario 1), PA passing ahead of RG (Scenario 2), RG passing ahead of SM (Scenario 3), and QS passing ahead of PA (Scenario 4).

6-5-4 Intervention time

Intervention time was measured from scenario start until the first action taken to solve the conflict. Consistency was determined by analyzing intervention time variations across repetitions, with a narrower time range representing higher consistency. Results need to be considered cautiously since intervention time was influenced by the conflict warning alerting system. The attentional effect of the warning was especially salient for Scenarios 1 and 4, with 92,2% and 85.9% of all conflicts being solved after the alert, respectively. In Scenarios 1 and 4, roughly half of the participants interacted within the same five-second range (50.0% and 43.8%, respectively). In contrast, Scenarios 2 and 3 were often solved before the warning (79.7% and 60.9%, respectively), although only 6.3% in both groups interacted within a five-second range. Mann-Whitney U and Kruskal-Wallis H tests were performed to examine the relationship between intervention time and consistent problem-solving

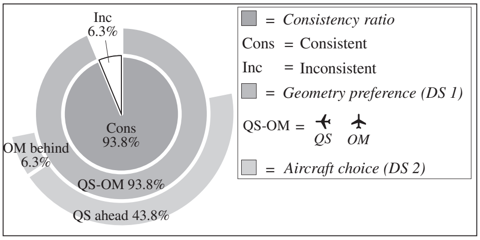

(a) Scenario 1

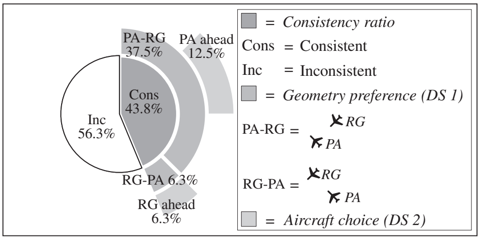

(b) Scenario 2

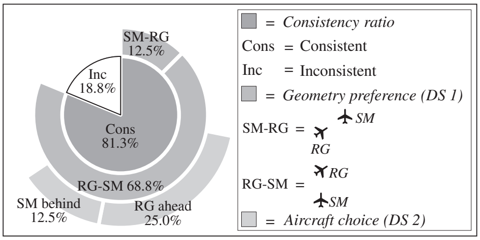

(c) Scenario 3

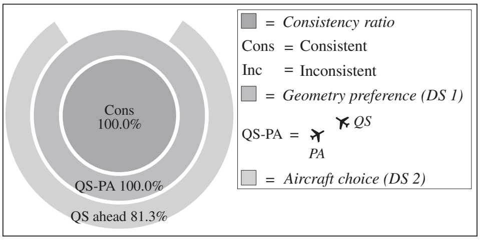

(d) Scenario 4

FIGURE 6-8: Solution geometry sunburst charts showing consistency and agreement, Study 1. Percentages indicate proportion of participants ($N = 16$).

styles as identified by the different classifications. None of the results, however, reached statistical significance, suggesting that intervention time did not influence how the conflict was solved.

6-5-5 Solution biases

Table 6-3 provides an overview of consistency and agreement in relation to the expected solution bias for each scenario and designed conflict. Overall, a widespread consistency in problem-solving style matching the detailed expected solution (including DS 3 resolution type, and DS 4 direction) was only found for Scenario 4. However, if only considering the first two DSs (1 and 2), a different patterns emerges. A high support for the expected solution, can then be found for all scenarios according to the solution geometry classification, except for Scenario 2.

It was expected that the SSD would imply a particular solution for each asymmetrical conflict. Results indicates that this was not the case. Rather, participants seemed 1) less influenced by the SSD representation than expected, 2) less concerned with solution details (i.e., DSs 3 and 4), and 3) more concerned with the overall, spatial, relationship between aircraft (i.e., solution geometry).

TABLE 6-3: Actual solution consistency and agreement in relation to expected solutions.

Scenario	Expected solution	Solution parameter hierarchy	Control problem	Solution geometry
1	Vector QS right ahead of OM.	Some support for expected solution. Highest agreement found for interacting with QS (DSs 1 and 2).	Some support for expected solution. Half the group consistent, mainly in agreement on taking QS in front of OM (DSs 1 and 2).	High support for expected solution. Nearly all participants agreed on having QS ahead of OM (DS 1). Fewer, however, consistently did so by interacting with QS (DS 2).
2	Vector PA left, or increase speed, ahead of RG.	Little support for expected solution. No consistency for single aircraft interactions with PA as preferred aircraft. Some consistency and agreement for dual interactions with PA chosen first (DSs 1 and 2).	Partly support for expected solution. Those participants found consistent agreed on having the controlled aircraft in front (DS 1), although only a few consistently interacted with PA (DS 2).	Some support for expected solution. The majority of participants did not solve the conflict consistently. However, roughly one third of participants agreed on having PA ahead of RG (DS 1), mainly by interacting with PA (DS 2).
3	Vector RG left ahead of SM.	No support for expected solution, although consistency and agreement were found for other solutions.	Some support for expected solution. A little more than half of the participants, who were found consistent, agreed on controlling RG in front of SM (DSs 1 and 2). In contrast, the other half of consistent participants agreed on controlling SM to pass behind RG.	High support for expected solution. Majority of participants agreed on RG passing ahead of SM (DS 1), although there were some disagreement on whether to take RG ahead or SM behind (DS 2).
4	Vector QS right ahead of PA.	High support for expected solution. The majority of participants agreed on vectoring QS to the right.	High support for expected solution. The majority of participants agreed on controlling QS in front of PA (DSs 1 and 2).	High support for expected solution. All participants consistently solved the conflict and agreed on QS passing ahead of PA (DS 1). Almost all participants achieved this by interacting with QS (DS 2).

6-6 Method study 2

Because Study 2 was an extension of Study 1, only relevant changes are highlighted. Only one conflict was studied, consisting of a symmetrical geometry for which a mathematical best solution is more arbitrary and less biased. The conflict was repeated four times. In contrast to Study 1, a larger variety in participants' solutions resulting in lower consistency and agreement was hypothesized. In addition, Study 2 explored differences between experienced and novice controllers (i.e., trainees).

Experienced controllers are likely to have developed strong working routines that manifest as a general adherence to specific patterns in conflict solving. Their experience, however, signifies an increased ability to work unstructured, as they have memory access to a wide repertoire of conflict situations and solutions for handling them. While the former argues for expectations of consistency, the latter suggests inconsistency. Novice controllers, on the other hand, tend to depend more on a clear structure for working, and can be expected to more rigidly stick to defined standard operating procedures (if available). As such, consistency in conflict solving can be expected for situations where structural guidance is available. Consider, for example, a right angle conflict, for which ICAO rules of the air stipulate that "the aircraft that has the other on its right shall give way" [p. 3-2].[123] A novice controller is expected to more strictly adhere to this right rule convention, and vector the considered aircraft to the right behind the other.

6-6-1 Participants

Two simulations were conducted with a total of fourteen volunteers with ATC experience. Nine were trainees (about to begin final on-the-job training) at Malmö Air Traffic Control Center (ATCC), Sweden. Age ranged from 24 to 29 years (mean = 26). Three females, and six males, took part. The other five participants consisted of experienced approach controllers, four from Norrköping Terminal Control Center, Sweden, and one from Arlanda ATCC, Sweden. Age ranged from 26 to 47 years (mean = 32.8) and experience varied between thirteen months and 24 years (mean = 8.7). One female, and four males, participated.

6-6-2 Simulator

The simulation ran on a portable computer connected to an external 21-inch monitor with a resolution of 1600x1200 pixels. Simulation speed was two times normal speed. The same simulator and SSD interface from Study 1 were used.

6-6-3 Materials

A Google Forms online survey was created, consisting of three 7-point Likert Scale questions addressing self-perceived consistency in solving right angle (90 degrees) conflicts in general, the right angle designed conflict in particular, and to what extent participant's believed their colleagues had solved the designed conflict differently. In addition, an open-ended question asked participants to solve a static conflict presented on the screen, which was the designed conflict used in the simulation. This allowed for comparing participant's answers with their conflict-solving styles as derived from simulation data.

6-6-4 Measurement scenario and designed conflict

A larger sector, 80 x 80 nmi, was used to reduce congestion and allow for more maneuverability, which may have restricted solutions in Study 1. The measurement scenario and designed conflict were based on the Scenario 1 conflict in Study 1. Parameters of the designed conflict are shown in Table 6-4. Two scenario rotations were used. A reciprocal, symmetrical, conflict angle was selected to mitigate biased solutions and provide identical solution possibilities for both aircraft. Furthermore, both aircraft were configured equidistant from the CPA, traveling at the same speed in zero wind with identical speed envelopes. The lateral deviation required to solve the conflict was the same for both aircraft, irrespective of a vector in front or behind the other. In terms of additional track miles, however, vectoring either aircraft behind the other would be more efficient. Because of the symmetrical conflict parameters, the only expectation was that one aircraft would be vectored behind the other as this would result in less track miles to clear the conflict.

6-6-5 Procedures

Participation lasted roughly one and a half hours and included a simulation and questionnaire part. After consent procedures and simulator briefing, participants played fourteen training runs (roughly 50 minutes) followed by the main experiment consisting of ten scenarios, each two minutes long (roughly 20 minutes). The ten scenarios consisted of the four repeats intertwined with six "dummy" scenarios. The dummy scenarios were used, together with scenario rotations, to prevent scenario recognition. Scenario order was varied according to a Latin Square design. Traffic was restricted to the horizontal plane (flight level 270).

TABLE 6-4: Scenario and designed conflict parameters, Study 2

Rotations	Aircraft	Conflict Angle (CPA)	TSL	Aircraft 1 & 2 ID: heading, speed (speed envelope)
0°, 0°, 180°, 180°	27	90° cross (0 nmi)	104 s	QS: 270°, 260 kn (200-320); OM: 000°, 260 kn (200-320)

CPA = Closest Point of Approach; kn = knots; nmi = nautical miles, s = seconds, TSL = Time to Separation Loss

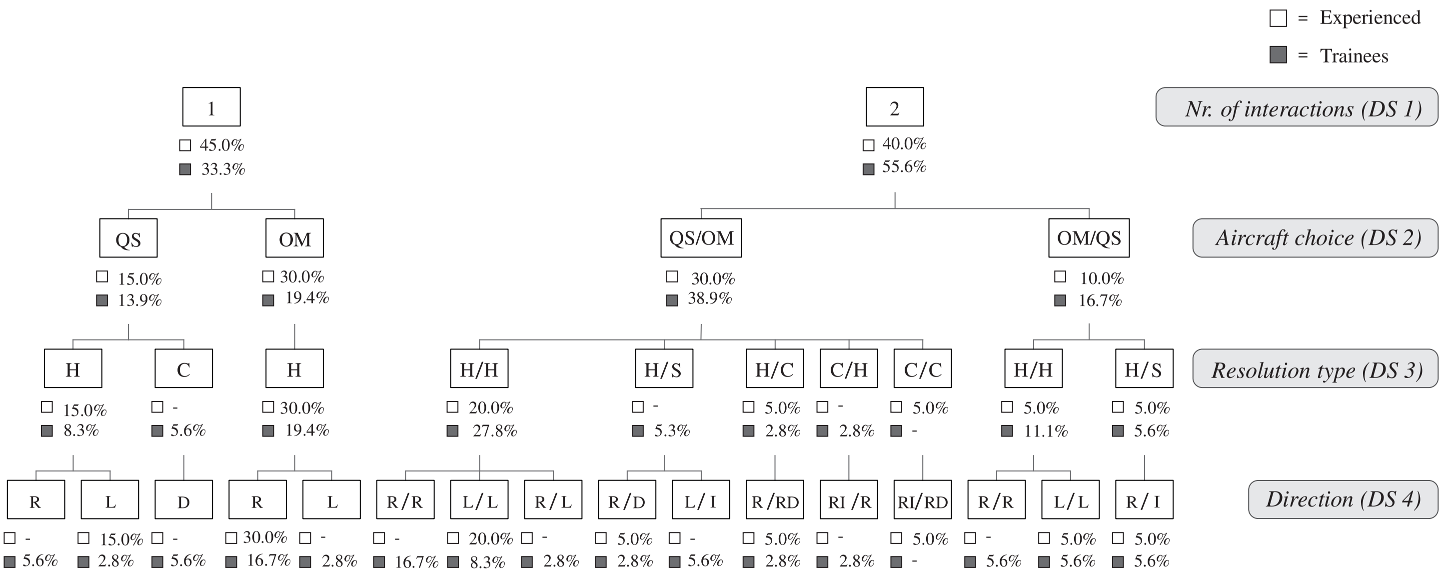

H = Heading; S = Speed; C = Combined; R = Right; L = Left; I = Increase; D = Decrease

FIGURE 6-9: Solution distribution for the Solution parameter hierarchy in proportion (%) of total solutions, Study 2.

6-7 Result study 2

Data from 56 solutions was collected, consisting of four scenarios repeated across fourteen participants.

6-7-1 Solution parameter hierarchy analysis

Figure 6-9 shows the solution distribution for both trainees and experienced controllers across 36 and 20 solutions, respectively. Interacting with both aircraft was more common among trainees (55.6%), while experienced controllers almost equally often interacted with one (45.0%) or both aircraft (40.0%). Experienced controllers implemented slightly more solutions involving three or more interactions than trainees did (15.0% and 11.1%, respectively). Similar to Study 1, these solutions were disregarded from further analysis and are not depicted in Figure 6-9.

For both groups, the most frequent solution consisted of interacting either with OM or first QS than OM. Solutions were most often solved by vectoring. DS 4 in Figure 6-9 shows a large variation in solutions for trainees, with sixteen unique solutions across 36 solved conflicts. The two most common solutions consisted of vectoring OM to the right behind QS (16.7%), or first vectoring QS right and then OM right (16.7%). In contrast, the group of experienced controllers solved 20 conflicts in eight different ways. Similar to trainees, experienced controllers frequently solved the conflict by vectoring QS left (30.0%). In contrast to trainees, however, two other frequent solutions consisted of vectoring QS left (15.0%), or first vectoring QS left and then OM left (20.0%).

The sunburst charts in Figure 6-10(a) and (b) shows the proportion of consistent trainees and experienced controllers. Overall, more trainees (66.7%) were found consistent than experienced controllers (40.0%). While experienced controllers agreed on interacting with one aircraft (40.0%), groups of consistent trainees disagreed on interacting with one (22.2%) or both aircraft (44.4%). Only one participant was found consistent according to DS 4, vectoring OM to the right.

Although the sample with experienced controllers was considerably smaller than the trainees sample, results suggest that experienced controllers are less varied in their use of solutions. Surprisingly, however, experienced controllers were overall not found more consistent than trainees.

6-7-2 Control problem analysis

Figure 6-11 shows the distribution of solutions according to the control problem classification for both trainees and experienced controllers. For the latter, clearing the controlled aircraft behind the intruder (DS 1) was more common than the

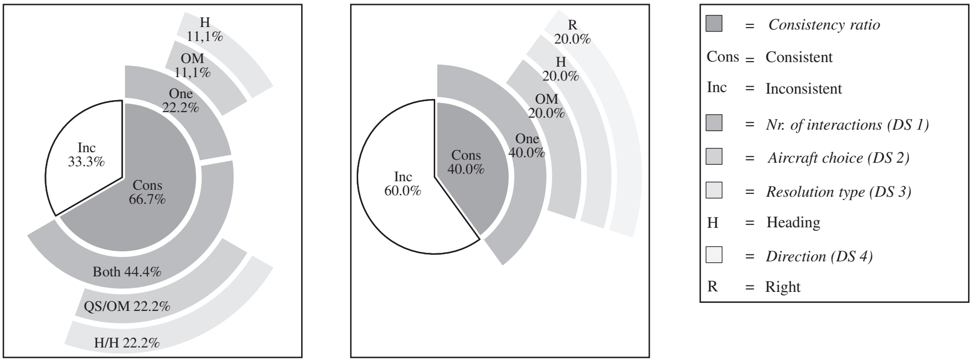

(a) Solution parameter hierarchy
Trainees

(b) Solution parameter hierarchy
Experienced

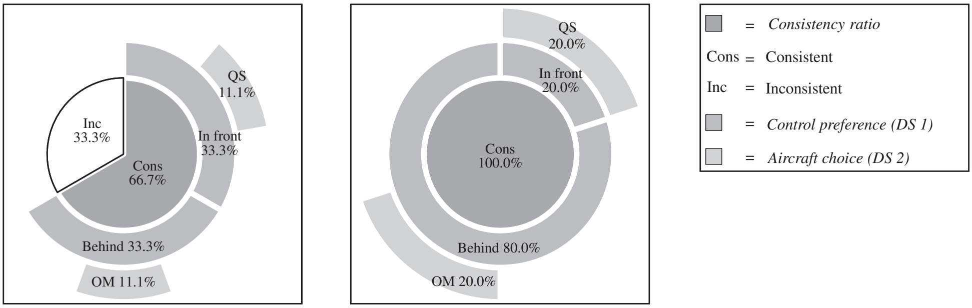

(c) Control problem
Trainees

(d) Control problem
Experienced

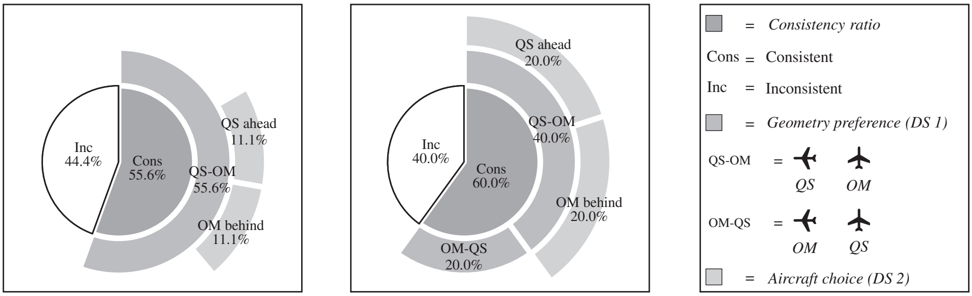

(e) Solution geometry
Trainees

(f) Solution geometry
Experienced

FIGURE 6-10: Classification sunburst charts showing consistency and agreement, Study 2. Percentages indicate proportion of participants (Trainees $N = 9$, Experienced $N = 5$).

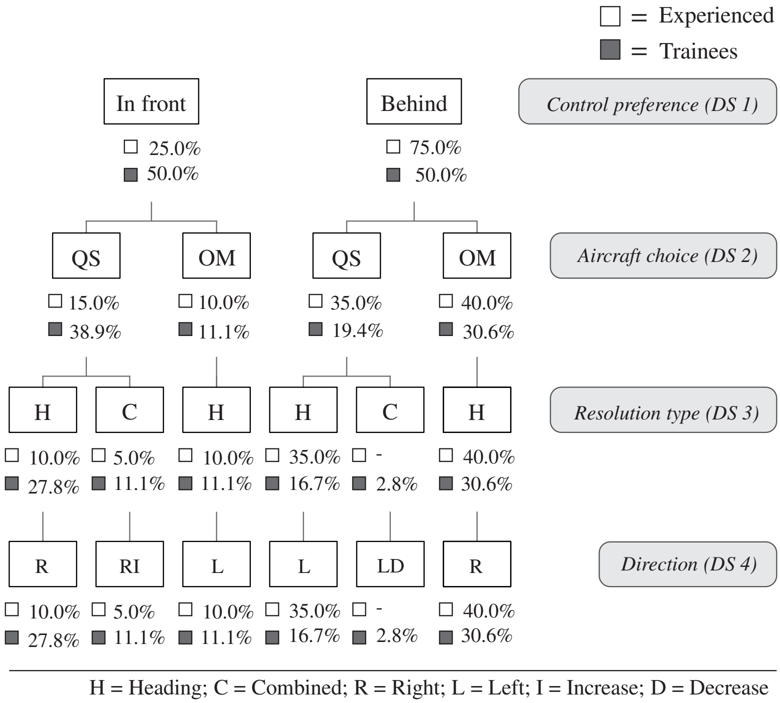

FIGURE 6-11: Solution distribution for the Control problem classification in proportion (%) of total solutions, Study 2.

opposite (75.0% versus 25.0%). For trainees, in front clearances were equally common as behind clearances. Among experienced controllers, most solutions consisted of clearing OM behind QS (40.0%) or clearing QS behind OM (35.0%). Although these clearances also were common among trainees, the most frequent solution consisted of clearing QS in front of OM (38.9%).

Figure 6-10(c) and (d) shows the proportion of consistent participants and problem-solving styles across trainees and experienced controllers, respectively. Notably all experienced controllers were found consistent according to the control problem classification. Four out of five (80.0%) consistently cleared the controlled aircraft behind the intruder. One participant consistently cleared the controlled aircraft ahead of the intruder (20.0%). Among trainees, three out of six solved the repeated conflict inconsistently (33.3%). Six out of nine (66.7%) were found consistent according, with equally many preferring to clear the controlled ahead as behind the intruder (33.3% respectively).

6-7-3 Solution geometry analysis

Figure 6-12 shows the solution distribution according to the solution geometry classification, for both trainees and experienced controllers. For both groups, the geometry where OM passes behind QS was more common than the opposite (DS 1). For trainees, this geometry was almost equally often achieved by having OM passing behind (30.6%) as having QS passing ahead (38.9%). Experienced controllers, however, were more likely to intervene so that OM passed behind (40.0%).

Figure 6-10(e) and (f) shows the proportion of participants' consistent problem-solving styles according to the solution geometry classification. Roughly equally many trainees and experienced controllers were found consistent. For DS 1, all consistent trainees agreed on having QS passing ahead of OM. For DS 2, however, some disagreement was found between whether to interact with QS (11.1%) or OM (11.1%) to achieve this spatial relationship. The group of experienced controllers was less homogeneous, with consistent problem-solving styles found for both QS-OM (QS ahead of OM) and OM-QS (OM ahead of QS).

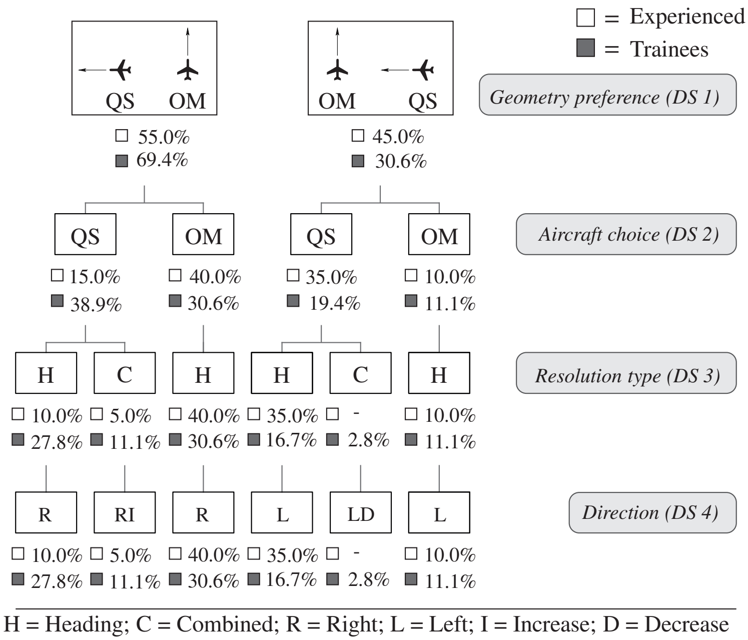

FIGURE 6-12: Solution distribution for the Solution geometry classification in proportion (%) of total solutions, Study 2.

6-7-4 Intervention time

In regards to intervention time, Mann-Whitney tests did not reveal a significant difference between the two groups. For both groups, intervention time ranged from 23 to 80 seconds (measured from scenario start) with the majority of interventions occurring between the caution and warning alert. Only one trainee and two experienced controllers intervened within a five second interval across repetitions. Four trainees and one experienced controller intervened within a fifteen second interval in all repetitions. Four trainees and two experienced controllers intervened within a 35 second interval. This data was analyzed to determine whether time of intervention affected the type of solution implemented. No significant effects were found of intervention time on number of aircraft interacted with to solve the conflict (Kruskal-Wallis H test), vectoring controlled aircraft ahead or behind (Mann-Whitney U), or solution geometry (Mann-Whitney U).

6-7-5 Consistency questionnaire

Participants believed that they had solved the repeated conflict consistently with 57.1% agreeing with the statement, and 28.6% undecided. Furthermore, 57.1% of participants slightly or strongly agreed with the statement that they generally solve right angle conflicts consistently over time. 35.7%, however, slightly or strongly disagreed with the statement. Responses were more mixed for the statement "I think that the other participants solved conflicts different than I." Disagreement was more common (50.0%) than agreement (21.4%), with 28.6% undecided. Kolmogorov-Smirnov Z tests showed that trainees and experienced controllers did not differ significantly on the three 7-point Likert Scale questions.

Participants' answers to the open-ended questions differed in whether only one aircraft (42.9%) or both aircraft (28.6%) should be vectored. 28.6% of participants specified that they would vector QS behind OM, while no one suggested the opposite. This is noteworthy since vectoring OM to the right, behind QS, was the most common standalone implemented solution in the simulation. Participants stated that vectoring one aircraft behind the other is preferred because of safety, efficiency, and comfort. Participants who stated that they would vector both aircraft, often justified their solution as a fairer and less intrusive intervention.

6-8 Discussion

This paper has investigated conflict solving consistency and agreement among 30 controllers with varying experience levels, in two separate studies. Results showed that controllers differed in how they preferred to solve conflicts and several different solutions were recorded for each designed conflict. Contrary to our expectations, consistency and agreement was overall not higher in Study 1 than in Study 2. Rather, the degree of agreement between controllers' problem-solving styles in Study 1 and 2 indicates that agreement varied across conflict geometries. Consistent problem-solving styles varied considerably between scenarios in Study 1, with the exception for Scenario 4. Consistency was considerably higher in this scenario, in addition to nearly all controllers agreeing on the same problem-solving style. The underlying reason for this is difficult to determine. Perhaps it can be attributed to the shallower conflict angle compared to the larger angles used in the other scenarios. As such, it cannot be concluded that biased conflict geometries are solved more consistently, or agreed upon more frequently, than conflicts that are not biased.

Since the Study 2 scenario was based on Scenario 1 in Study 1, a more detailed comparison between these two scenarios is appropriate. While Scenario 1 in Study 1 was biased, favoring QS to pass ahead of OM, the scenario in Study 2 did not favor a specific solution other than vectoring one aircraft behind the other. When comparing the solution parameter hierarchy, almost equally many were consistent according to DS 1. However, analyses of the more detailed DSs revealed two contrasting problem-solving styles in Study 2, while controllers in Study 1 were more in agreement. Differences were also noted in aircraft choice preference, with Study 1 controllers preferring QS and Study 2 controllers preferring OM. However, more controllers consistently interacted with both aircraft in Study 2, in the generally preferred order of QS first and then OM. In terms of the control problem analysis, consistency was slightly higher in Study 2, compared to Study 1. In Study 1, there was a large preference to vector in front, while in Study 2, the preference was to vector behind. In terms of the solution geometry classification, controllers were considerably less consistent in Study 2 (when comparing the biased Scenario 1 in Study 1 with the unbiased similar conflict in Study 2). According to the control problem classification, however, consistency was higher in Study 2.

6-8-1 Consistency in conflict resolution

Questionnaire data from Study 2 suggest that several controllers do not consider themselves consistent in conflict resolution. This was indeed reflected in simulation data from both studies, in which controllers' consistency varied considerably.

In relation to simulation data from both studies, however, the majority of controllers were found consistent according to at least one solution classification. The solution parameter hierarchy classification yielded the lowest proportion of consistent participants. Furthermore, this classification often yielded contrasting consistent problem-solving styles, again with the exception of Scenario 4 in Study 1.

In both studies, the proportion of consistent controllers according to the control problem classification was equal to, or higher, than that of the solution parameter hierarchy classification for all scenarios, except Scenario 1 in Study 1. Based on open-ended questionnaire responses in Study 2, the control problem classification appears to fit well with controllers' own conflict solving styles. Most controllers in Study 2 were also found consistent according to the control problem classification overall. Overall, the solution geometry classification had the largest proportion of consistent controllers in Study 1, as compared to the other classifications.

For Study 1, consistency was overall highest according to the solution geometry classification, Consistency was only lower than the other classifications for Scenario 2. Similarly, consistency was lower for the solution geometry classification than the control problem classification in Study 2. However, of those found consistent, a majority used the same problem-solving style.

6-8-2 Conflict solution agreement

Questionnaire data in Study 2 showed that the majority of controllers believed that their colleagues had solved the designed conflict differently. Their responses are supported by the large variability in solutions recorded in the simulation, and low overall agreement for the more detailed decision stages across classifications.

When considering all three strategy classifications, results indicate that controllers largely agreed on the same solution geometry, but disagreed on how to achieve it in terms of the solution parameter hierarchy or control problem classification. This suggests that controllers strive to achieve the same spatial relationship of a solution, but adopt different methods. It is possible that controllers only consider solutions in terms of the resulting spatial relationship between aircraft and more arbitrarily chose which specific control actions to implement in order to achieve this goal. This reasoning is supported by previous research indicating that controllers rely on simplifying heuristics and satisficing in conflict resolution. For instance, Inoue *et al.*[247] found that controllers share the same or similar strategies, but that the concrete methods for solving conflicts differ individually.

6-8-3 Effects of experience on consistency and agreement

Observations and informal conversations with controllers in Study 1 indicated that experience may have had an influence on controllers' problem-solving styles. Unfortunately, the sample composition did not provide enough data to do a comparison. In Study 2, however, there was an opportunity to study the effects of experience on problem-solving styles, and differences in consistency and agreement between novices (i.e., trainees) and experienced controllers.

Study 2 results showed some interesting differences between trainees and experienced controllers. The more experienced group had notably less solution variability for solving the same conflict repeatedly. As such, experienced controllers appeared more in agreement on a preferred solution, especially in regards to the control problem classification. In contrast to trainees, all experienced controllers were found consistent according to the control problem classification and all, except one controller, agreed on the same control preference. Taken together, this suggests that the control problem classification best describes conflict solution among experienced controllers. Further research is required in order to determine whether the control problem classification is better than other classifications in quantifying controller expertise.

6-8-4 Sensitivity of consistency and agreement measures

Results showed that consistency and agreement was restricted to the dichotomous higher-level decisions captured by DS 1 (type of classification) and DS 2 (aircraft choice) in each classification hierarchy. More specifically, our results indicate that controllers can be considered consistent in terms of their decision to interact with one or both aircraft (solution parameter), having one aircraft go in front or behind another (control problem), and their preference for achieving a specific solution geometry. In contrast to the solution parameter hierarchy, the first two decision stages of the control problem and solution geometry classifications imply which resolution type (DS 3) and direction (DS 4) to implement in order to achieve the overarching control or geometry preference and aircraft choice (DSs 1 and 2).

Moreover, our analysis revealed that controllers are inconsistent in regards to the more detailed decision stages of the classification types. In regards to DS 3 (resolution type), DS 4 (direction), and DS 5 (directional value), the classifications proved too detailed and thus captured more noise than signal. Consistency did also not improve when considering expertise degree separately (Study 2). This yields the question whether the classification scales are good enough to capture consistency and agreement. Our conclusion is that the classifications are sufficiently sensitive for higher aggregated decision levels that capture the overall solution characteristics. For these levels, a move towards individual sensitive automation is sensible.

6-8-5 Defining and measuring consistency

The three strategy classifications identified in this paper overlap in the sense that decision stages in respective hierarchies can be combined in several ways. Accordingly, some participants were found consistent according to more than one classification. In addition, participants were found to both agree and disagree depending on the classification used to define their problem-solving style. For example, while the control problem classification indicated contrasting problem-solving styles between two participants, their different control actions resulted in the same spatial relationship between the two aircraft, as indicate by the solution geometry classification. As a consequence, the relevance of the consistency and agreement analysis in this study can be questioned. These findings were contradictory and surprising. However, they support generalization and assumptions of homogeneity among controllers in conflict resolution.

The definition of what exactly constitutes a consistent problem-solving behavior is critical, if we are to develop automation that acknowledges and is sensitive to controllers' problem-solving styles. Conflict solving is not a simple matter of choosing alternative A or B, it is considered more rich and complex. Yet our results suggest that consistency predominantly is limited to such dichotomous choices, even though a few participants were found consistent down to the fourth decision stage in the solution parameter hierarchy classification. The distinction between choosing a right vector, or a right vector with a speed increase may seem subtle but can be important. Future research is needed to address how to qualitatively distinguish solutions and define and measure consistency.

6-8-6 When to intervene

Because a solution is valid only for a limited period of time, it is reasonable to expect different solutions depending on *when* the conflict is solved. Results showed that intervention time varied greatly between both repetitions and controllers. Although participants' intervention times were confounded by the conflict warning alert, the absence of any effects on how conflicts were solved is worth highlighting. This suggests that intervention time may be less important for how a conflict is solved. Similarly, previous research has found that differences in "look-ahead" time (5-12 minutes versus 8-14 minutes) do not influence controllers' solutions.[35] Although further research is needed, this finding is relevant for ATC decision aid design since it suggests that specific conflict resolution advisories can be used for large time windows. The timing of advisories may, however, be important for other reasons, such as personalized to controller's preferences for immediate interaction or to "wait and see," and to avoid annoying and intrusive interruptions (i.e., advisory etiquette[114]).

6-8-7 Controller strategies

Several problem-solving styles similar to resolution strategies identified in previous literature were found. In particular the two strategies "vector behind" and "vector ahead" were commonly observed in simulations. These strategies were also reflected in questionnaire responses in Study 2. However, results also challenge strategies identified in previous literature. Several participants consistently interacted with *both* aircraft to solve the conflict, which argues against previous findings that controllers rarely interact with both aircraft.[88]

Furthermore, solutions in Scenarios 2 and 3 often relied on changing the speed of one or both aircraft, challenging the notion that speed changes typically are refrained from in en-route environments.[29,31] Finally, a notable finding was that vectoring one aircraft in front of the other was the most frequently used solution in several scenarios. While this partly could be expected in Study 1 because of the biased conflict geometries and noise traffic, it was also found preferred by a relatively large group in Study 2, which suggests that it cannot be attributed to scenario design alone. Notably, it challenges the notion that vectoring ahead is a less frequently used strategy,[24] or even a "no-no" strategy that controllers avoid.[35]

However, the observed differences in solution strategies, in particular the increased used of speed, can perhaps partly be attributed to the novel conflict presentation provided by the SSD. Ecological interfaces, such as the SSD, shows the spectrum of available solutions that encourages new and different problem-solving styles not otherwise readily conceived. This is likely to have had some influence on controllers' conflict solving, resulting in a larger variety of solutions applied and hence less consistency and agreement.

More specifically, an attempt was made to link participants' solution styles to the five lateral resolution strategies identified by Fothergill *et al.*[24] (see Table 6-1). However, these strategies only consider solutions where one aircraft is controlled. In contrast, participants often interacted with both aircraft. Moreover, while the five strategies relate to the solution rationale, the practical difference between them is unclear. For example, the lateral strategies 1 ("vector behind"), 2 ("direct away"), and 4 ("take out for five miles") can be interpreted according to all three classifications. According to the solution parameter hierarchy, all three can be achieved with one interaction, such as vectoring one of the aircraft to the right or left.

Furthermore, all strategies can result in the same geometrical relationship. As such, the strategies suggested by Fothergill *et al.* may refer to a similar solution expressed contrasting by different controllers. Alternatively, the formulations (of strategies) may correctly reflect controllers' contrasting interpretation of an identical conflict, which in turn generate different underlying reasoning for how to solve it. Consequently, similar and different solutions are inconsistently expressed pertinent to individual's views and problem-solving styles. Such inconsistencies were found in work by Bonzano *et al.*[131] who noted that controllers described the same conflict differently and that these different descriptions determined how to solve the conflict.

6-8-8 Limitations

This paper has focused specifically on conflict resolution and disregarded the detection phase. In reality, however, it is questionable whether detection and resolution can be considered separately. The strategies used for conflict resolution are likely intrinsically related to strategies and scanning methods used for conflict detection. Thus, future research should consider both the detection and resolution phase when investigating conflict resolution strategies.

Moreover, there are many more individual differences relevant for ATC CD&R that have not been considered in this paper. For example, research by Chen and colleagues indicates that individual differences in the ability to focus and shift attention in a flexible manner impacts performance in multitasking environments involving human supervision of autonomous robot systems.[257] Similarly in ATC CD&R, this type of attentional control and spatial abilities are considered important for scanning traffic and identifying conflicts. It is reasonably to expect that controllers who differ in these abilities view their surrounding differently and therefore use different strategies and derive at different solutions. Thus, it is important to recognize that although *objectively* there is one truth, people perceive the same situation differently.

It should be noted that conflict solving cannot be considered static. Controllers continuously adapt and develop new techniques and strategies, for example by learning from difficult situations, watching colleagues, trying out new and different methods, and adapting to new procedures and technologies introduced.[189] Because of this it may be difficult to determine a consistent behavior.

While controllers participated in simulations (from advanced trainees to very experienced), the two studies sampled controllers with different experiences (approach and en-route) and working in different countries (Ireland and Sweden). These differences may have attributed to the variation in consistency and agreement observed.

Finally, solution options were limited to the horizontal plane, while in reality vertical solutions are commonly used in en-route ATC. Future research should extend the analysis of consistency and agreement to the vertical solution plane.

6-9 Conclusion

This chapter has shown that controllers differ in their solution preferences and that they are more diverse in their conflict solving than currently conceived. A conflict solution framework consisting of three strategy classifications were identified, each explaining different consistent problem-solving styles among controllers. Results showed that controllers consistently solved conflicts, yet in different ways according to three different classification frameworks. For all classifications, however, consistency reduced notably when considering solutions in more detail.

It was not possible to determine whether any classification was more accurate than the other. Overall, controllers tended to strive for the same spatial relationship between the conflicting aircraft. With increased experience, however, the spatial relationship appeared to loose strength in favor for the control preferences in solving conflicts. Rather, results suggest that the classifications reflect differences between individuals and their perception of and preferences for solving conflicts. Results partly supported our expectation in that controllers were less consistent and more in disagreement on how to solve the conflict, when the conflict was unbiased and did not favor a certain solution.

Taken together, the two studies indicate that consistency varies both with the situation and individual. For many problems, especially those that lack golden criteria for optimum decisions such as CD&R in ATC, solutions vary depending on who solves it. The observed individual diversity in conflict solving can partly explain the acceptance problems observed in relation to CD&R decision aids that traditionally have assumed that controllers solve conflicts similarly. The findings in this chapter challenges the feasibility of considering controllers homogeneous in the design of automated CD&R decision aids. Consequently, the development of future automated systems should better accommodate individual differences in decision-making and problem-solving activities, not only in ATC. Developing decision aids that provide specific solutions may, however, not be the best approach. Rather, individual decision-making preferences can perhaps better be facilitated by providing decision support tools, such as the SSD, that enable controllers to more efficiently and safely solve problems in their own preferred way. On the other hand, this may not be practical when considering the requirement for increased traffic throughput, resulting in less time available to control individual aircraft.

7

Discussion and recommendations

This chapter discusses the major findings and lessons learned from the strategic conformance research in this thesis, including theoretical and practical implications and recommendations for future research. Issues and challenges for future automation design raised in the introduction and literature survey (Chapter 2) are revisited and discussed in the light of empirical studies performed in Chapters 3 to 6. The chapter is structured in three sections. First, a Retrospective overview is provided where the major thesis contributions are discussed, focusing on relevant empirical findings from the experiments. The strategic conformance concept is revisited and discussed in relation to the influence and impact of several human factors considered throughout this thesis, including situation complexity, level of automation, source bias, transparency, and decision-making consistency and agreement. Second, lessons learned from Research Challenges and Limitations encountered are discussed. In addition, several recommendations for future research are made. Finally, Benefits and Pitfalls of Strategic Conformance are discussed, not only in relation to ATC, but more generally in the context of decision-aiding automation and personalized applications.

7-1 Retrospective

The focus of this thesis has been on the cognitive compatibility mismatch between operators and their automated decision aids. To bridge this compatibility gap, the automation's problem-solving should match the human's problem-solving style. It was hypothesized that the acceptance and understanding of a decision aid's advice would benefit if the aid was *perceived* to "reason" similar to the human it was interacting with. This degree of match between the automation's solution and *apparent* underlying operations, with that of the human, was conceptualized as *strategic conformance*. As perceived by the human, the automation's reasoning can vary between being *conformal* (a match) and being *nonconformal* (a mismatch).

To start with, a broad literature search was conducted across a variety of sociotechnical work domains for research on automation acceptance in relation to individual-sensitive systems (Chapter 2). Specific focus was given to ATC related research since this was the target domain of both the MUFASA project and this thesis. The foundation for strategic conformance was found in two complementary domains. The cognitive engineering domain was relevant because of its research on automation use decision in dynamic, high-risk, and time critical contexts. The information systems domain was relevant because of its well-developed and well-accepted acceptance models. Of particular relevance was the compatibility construct, considered in both domains, which asserts that human and automation must be compatible in order to facilitate acceptance. A model of human-machine compatibility was proposed, ranging from a basic level of response compatibility that considers simple handling qualities (i.e., consideration of physical aspects that reinforce correct use) to the highest level of cognitive compatibility that considers the strategic conformance in problem-solving (i.e., consideration of psychological aspects in decision-making that support the operator's problem-solving style).

7-1-1 Empirical findings on strategic conformance

Strategic conformance was empirically evaluated in three real-time simulations with air traffic controllers. A conformal advisory was always based on the controller's own implemented solution to the same conflict as recorded in a previous prequel simulation (in which conflict were solved manually with the support of the heading band SSD). A nonconformal advisory was always based on a contrasting solution made by another controller participating in the same real-time simulation. Overall, findings indicate a preference among controllers for conformal advisories. These were accepted more often than nonconformal advisories in two studies (First empirical and Automation transparency studies), and equally often in one study (Source bias study). This is noteworthy since no justification, rationale, or additional information was provided in an argument for the solution. Although the SSD interface provided an overview of the constraints affecting solutions, it did not grade or value options in any way.

While conformal advisories received significantly higher agreement ratings in the First empirical study, the effect was not repeated in subsequent studies. A slightly reversed effect was noted in the Automation transparency study. The higher agreement for nonconformal advisories can be explained by the more transparent triangle SSD, which revealed more of the available solution space and made alternative solutions available that controllers had not thought of during the baseline prequel simulation (for which the heading band SSD was used).

Controllers in the First empirical study responded faster to conformal advisories than nonconformal advisories (regardless of accept/reject decision). They also accepted conformal advisories faster than they accepted nonconformal advisories. This suggesting that their choice to accept advisories was deliberate and not due to a "blind" acceptance strategy driven by attitudes of trust or the perceived credibility of the system. Results were not reproduced in the Source bias and Automation transparency studies. Although, a notable finding in the latter, was the faster response time to nonconformal advisories in the heading band condition as opposed to the faster response time to conformal advisories in the triangle condition. A possible explanation is that controllers were trainees (although in their final educational stage), and had not developed strong preferences for solving conflicts.

7-1-2 Other factors affecting acceptance

In addition to strategic conformance, several other factors influencing acceptance have been considered in this thesis. The following are believed to be particularly important and relevant for future sophisticated automated systems assuming more cognitive work.

Level of automation and complexity. The First empirical study explored the interaction between strategic conformance together with the level of automation (LOA) and complexity on advisory acceptance. Two LOAs were used: management by consent (MbC), for which the automation suggests and operator authorizes, and management by exception (MbE), for which the automation acts autonomously but informs the operator who can intervene.

Because of the subtle difference between LOA conditions, they were ultimately collapsed and not considered for analysis. This led to numerous discussions surrounding the LOA concept, leading to a consensus that the contribution of these frameworks is limited for guiding research and development. Similar criticism has been promulgated by artificial intelligence researchers, who argue for the abandonment of LOA frameworks because of their simplistic focus on the technology, task allocation, and discrete authority levels.[258–260]

Complex and demanding scenarios were essential for simulations because of expected, and more pronounced, benefits of automation during high workload and stress.[45,51,127,129] For simplifying purposes, traffic count was used to vary complexity at two levels, high and low. Results showed that controllers were more accepting of, indicated higher agreement, and responded faster to advisories under complex conditions. This could be interpreted as if controllers, under the time pressure of complex conditions, did not fully evaluate advisories and instead prematurely accepted them. Rather, results indicate that controllers adequately evaluated advisories and were satisfied with their decision to accept of reject. During low complexity (and low perceived difficulty), controllers tried alternative solutions to conflicts. Note that scenario difficulty was not varied in the Source bias and Transparency studies because of the smaller sample sizes and data collected.

Source bias. Previous research has shown that operators' willingness to accept or trust advice is influenced by the perceived credibility and expertise of the adviser.[47,50,68,157,160,161,163,164] Such effects have been obtained in studies simply by framing the same source differently.[48,110,165,171,173,175] It was hypothesized that this could explain why controllers in the First empirical study did not fully accept conformal advisories: that is because they were biased against advice from what they perceived was an automated source. To investigate this, resolution advisories (varied by conformance) were presented as derived from either an automated or human source. Although questionnaire data showed a slight preference for the human adviser (portrayed as an air traffic controller), there were no effects of source bias found in relation to the acceptance of advisories in the simulation. However, the manipulation was subtle with different source information provided in instructions prior to the simulation and a text accompanying advisories during simulation. Controllers may not have reflected over the different sources during simulation runs.

It should be noted that the Source bias study has a philosophical question at heart, for which the answer can help us better understand how people interact differently with automated systems and humans. In many future contexts, humans will be *assigned* to work with an automated agent. Hence, there will be little doubt in regards to who the source is. However, if true, as indicated by some research,[48,114,118,176] that the more human-like automation is (i.e., anthropomorphic), the more it is treated as if it was human. Then, depending on interaction and systems goals, automation could be designed for the purpose of being perceived as a certain type of source (e.g., human or automation).

Automation Transparency. The Automation transparency study sought to answer the degree to which controllers' acceptance and understanding of resolution advisories depended on the transparency afforded by the interface. Two interface transparency levels were used: the baseline low level heading band SSD, and the high level triangle SSD which was a richer display showing more details of the traffic constraints. Simulation data did not indicate any interaction effects between automation transparency and advisory conformance on acceptance of performance. Results, however, indicated a preference for the more transparent interface, with controllers finding it more transparent and understandable. In addition, controllers' working methods changed when using the triangle display, to more speed and combined clearances. In support of previous research,[121] results show that constraint-based interfaces can be used to guide the creation of more understandable interfaces, explaining the automation's solution rationale.

Transparency and strategic conformance can be seen as complementary in regards to objectives, but different in regards of methodology. The objectives of both argues that the human-automation relationship benefits from the automation facilitating an increased understanding of how it operates, thinks, makes decisions, and acts. In terms of methodology, increased transparency is often achieved through explanations or by providing more (meta) information, for example through graphic-based visual approaches such as EID (Chapter 5). In contrast, strategic conformal automation argues that if the automation *thinks* like the operator, its advice can be tacitly understood. As such, lessening the need for increased interface transparency.

Controller experience. Trust and credibility research has established that self-confidence (e.g., knowledge and ability) is closely tied to how trustworthy and credible a support system is perceived.[15,50,68,108,160–164] Experienced operators may be reluctant towards novel automation because they have 1) high confidence in their own ability, 2) previously managed without support, 3) developed their own functional way of working, and 4) are satisfied with the current work practices. As hypothesized in this thesis, however, experienced controllers may be more accepting of conformal automation because resolution advisories, in fact, are their own. Conversely, novice operators can be expected to have lower self-confidence and be more amenable to others' authority and expertise. Hence novices may be more embracing towards novel automation as they have 1) little routine, 2) are still learning, 3) lack well-established practices, and 4) have not (yet) developed their own problem-solving styles. Novices' performance was therefore expected to be more variable and less consistent.

Results show that conformal advisories benefited both experienced controllers (First empirical study) and trainees (Automation transparency study), although effects were less pronounced among trainees. Moreover, experienced controllers were found slightly more consistent than trainees, at least in regards to the control problem classification (i.e., one aircraft controlled ahead or behind the other). The effect of experience was, however, not directly evaluated (e.g., by using novice and experienced controllers in the same experiment), leaving this open for future research.

7-1-3 Foundations for conformance: Consistency and agreement

The concept of strategic conformance hinged on two assumptions: that controllers solved conflicts consistently, but disagreed on which solution to use. Establishing consistency is important as, otherwise, it cannot be argued that a system is conformal. Establishing a lack of consensus is important as, otherwise, there is no benefit of personalizing a system. Data from the three empirical studies were used to, post hoc, investigate these assumptions, starting with an investigation of controllers' problem-solving strategies.

Figure 7-1 provides an overview of the four studies in this thesis and their relationship in relation to the Consistency study. Findings and questions from the First empirical study (controllers with varying experience) were addressed in three additional studies (illustrated by the grey arrows). In two separate real-time simulations, the Source bias (experienced controllers) and Automation transparency (trainees) studies explored strategic conformance and other variables (advisory source and interface transparency, respectively). The Consistency study, however, consisted of a

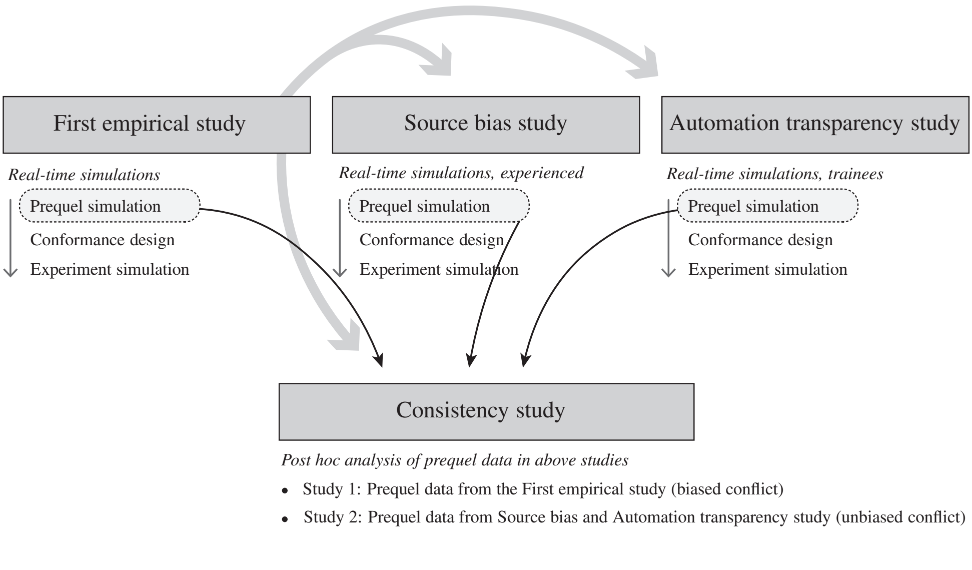

FIGURE 7-1: Overview of empirical studies and their relationship.

post hoc analysis of prequel data (i.e., controllers' manual conflict solutions) from real-time simulations gathered in the three other studies (illustrated by the arrows feeding from the prequel simulations into the Consistency study). Data from the three real-time simulations were analyzed separately in two studies depending on whether the conflict was biased, or not (Study 1 and Study 2, respectively).

To begin with, a framework consisting of three strategy classifications was created that describe different conflict-solving objectives. The *solution parameter hierarchy* argues that controllers determine solutions based on whether to interact with one or both aircraft. The *control problem classification* argues that controllers view the conflict as a control problem, solved by either clearing the controlled aircraft ahead or behind the other (by vector, speed, or combination thereof). Finally, the *solution geometry classification* argues that the resulting spatial relationship between the conflicting aircraft is the main driver for how the conflict is solved.

The post hoc analysis confirmed a limited agreement between controllers, while all controllers were found consistent. Consistency was, however, primarily restricted to the highest two decision stages of each classification, such as vectoring the controlled aircraft behind or in front, or interacting with one or both aircraft. Below these high-level decision stages, consistency decreased with only a few controllers found consistent in regards to resolution type (decision stage 3: heading, speed, or combined) and direction (decision stage 4: right or left, increase or decrease, or combinations thereof).

Taken together, this suggests that controllers generally can be considered consistent in terms of high-level goals, but inconsistent in terms of specific maneuvers implemented. This provides an explanation for why controllers considered themselves inconsistent in questionnaire responses. Similarly, Inoue *et al.* observed diversity in conflict solutions among controllers, although they adhered to the same strategy.[247] In a study attempting to identify controller strategies by means of machine-learning techniques, Regtuit found that, with increasing variability in conflict solving, consistency could only be established for the two main strategies of either vectoring ahead or behind the other aircraft.[261] On the other hand, some inconsistency can be expected given the flexibility in the world, the flexibility and creativity in human problem-solving, and the influence of human emotional states (e.g., tiredness, mood). As such, controllers can be expected to solve an identical problem differently over time, and they may not like the solution once they see it replayed.

Contrary to expectations, consistency and agreement were overall not higher for biased conflicts in Study 1 (i.e., data from First empirical study). However, while the solution geometry classification was favored (higher consistency and agreement) in Study 1, the control problem classification was favored in Study 2. This suggests that controllers' solution variability did not lessen with biased conflicts. But controllers' solution objectives (i.e., the preferred classification) did vary between biased and unbiased conflicts.

If controllers perceived repetitions differently, they are also likely to have focused on different information cues when determining a solution. The dynamic simulation environment made scenario repetitions susceptible to changes beyond control as controllers freely could interact. Although scenarios were identical "on paper," the interactions made during simulations made each repetition unique. This can also explain why controllers were not consistent down to the more detailed decision stages. Moreover, some learning effects can be expected since solution feedback always was provided in that scenarios continued after the conflict had been solved. This is likely to have also influenced controllers' evaluation, strengthening good solutions and weakening poor solutions.

7-2 Explaining the acceptance of conformal advisories

As always when studying people, it is not how we as researchers have defined the world that matters; What matters is how the world is perceived by the people being studied. Although the exact same scenario and conflict were repeated, controllers may have perceived them as different. Taken together, results suggest that conflict solving is predominantly intuitive, characterized by an instinctive, automatic, and reactionary decision-making process. This process is represented by several decision-making models that have been discussed throughout this thesis, such as recognition-primed decision-making[262,263] in naturalistic environments,[264] System 1 decision-making,[41] and the peripheral route of the elaboration likelihood model (ELM).[265] They all capture a process whereby decisions are based on intuitive expertise[21,266] driven by an instinctive (less rational) reaction triggered by the recognition of familiar information cues and their relationship. In ATC CD&R, these patterns have typically been described in terms of heuristic decision-making processes, subjectively manifested as rule of thumbs such as "vectoring behind" and "vector away from conflict."

Based on observations and findings, Figure 7-2 illustrates a proposed model of the underlying decision-making process of a controller in relation to the conformance of a resolution advisory. The model builds on the lens model[267] (the perception of cues in the upper part of the model) and recognition-primed decision model (the decision-making process in the lower part of the model).[263] The *conflict situation* (top of the model) represents the objective *ground truth* of the conflict. The conflict is defined by a several mathematical and geometrical *objective parameters*.

The model initially illustrates two different paths describing conflict identification, separately for a controller and an automated CD&R decision aid. In relation

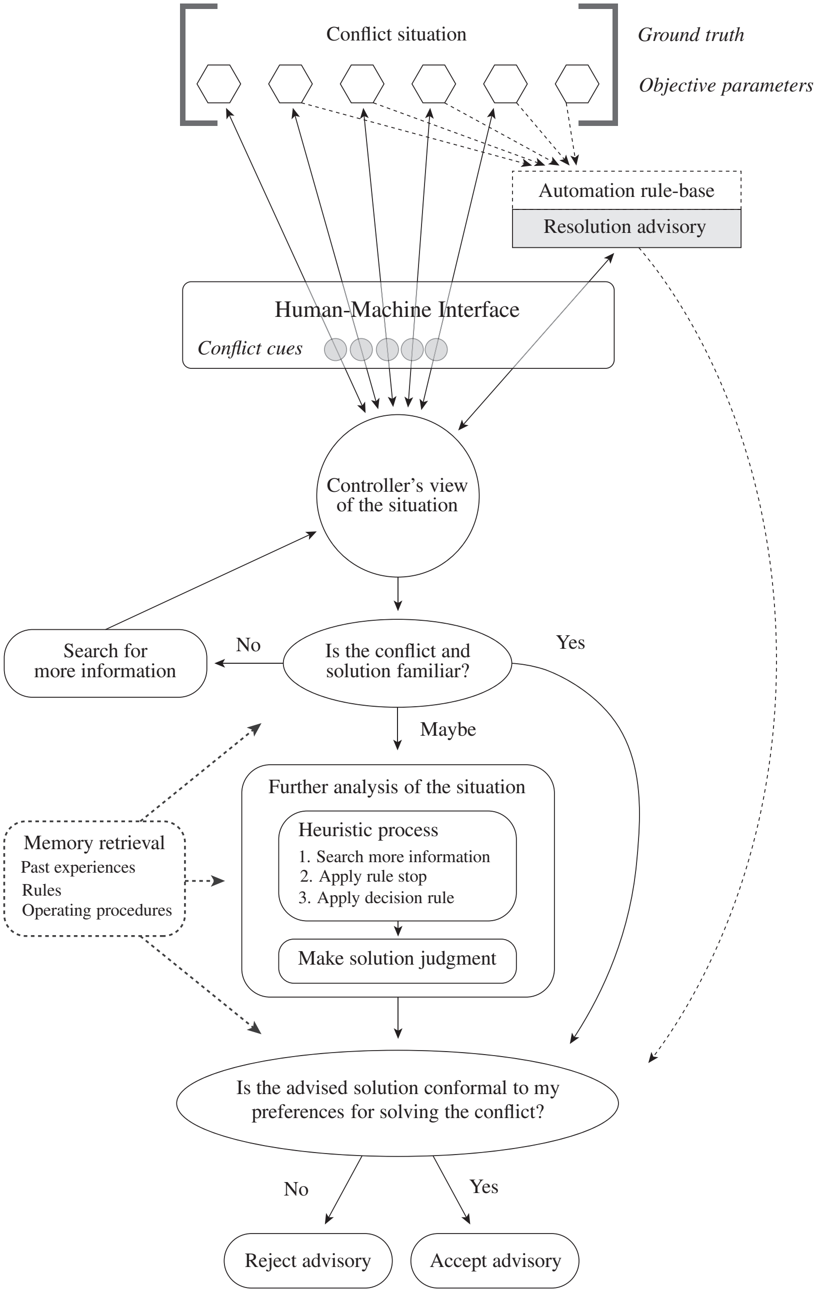

FIGURE 7-2: Proposed model of a controller's reaction to an advisory.

to the controller, a *human-machine interface (HMI)* mediates a representation of the airspace environment, its infrastructure (e.g., airports, waypoints, and routes), and its populating traffic. By means of the HMI, the controller perceives and reacts upon *conflict cues* (grey circles), which reflect the objective parameters, when identifying the conflict situation. Consequently, the way the HMI models and visualizes the world has a great impact on how the controller perceive and act in that world. This explains why controllers' interaction changed (i.e., more speed and combined heading and speed solutions), and agreement and response time measures were reversed, when using the triangle SSD in the Automation transparency study.

Similarly to the HMI, the *automation rule-base* (i.e., algorithms) of the CD&R decision aid identifies the conflict situation by processing the same objective parameters. Based the conflict cues identified by the automation-rule base, the decision aid provides the controller with a *resolution advisory*, which is presented in the HMI. Even though the HMI and automation rule-base process the same objective parameters, their translation of them may differ. Furthermore, they may focus on different parameters and assign the different weight and priority, thus describing the mismatch in problem-solving between human and machines. In the conformance studies performed in this thesis, however, the conformal resolution advisories were based on a controller's own solution to the same conflict. As such, the automation rule-base was replaced by the controller's problem-solving style as identified by the conformance design based on solution data from the prequel simulation.

When a controller is notified of a conflict by the system (or self identifies a conflict), information is searched for the purpose of understanding the conflict (see reversed arrows). The controller develops a picture of the situation and conflict. When controllers encountering the repeated conflict in simulations, they identified and acted upon similar information cues that gave rise to their consistent problem-solving behaviors. Debriefings and questionnaire responses suggest that controllers were unaware of repetitions.

The controller then determines, through memory retrieval, whether the conflict and solution is familiar. Cues are gathered and matched with previous experiences stored in memory. Different knowledge is considered, including previous experiences, rules, and standard operating procedures such as principles and "no-no's".[29] This memory retrieval and information exchange is fast and unstructured. If *no*, more information is required and the controller reassess the situation and clues. If *yes*, the conflict can be solved quickly on a skill-based level, highly reactionary. More complex conflict situations require a more heuristically driven conscious evaluation of the solution (the *maybe* path).

The process can be characterized by the *take-the-best* fast and frugal heuristic for interpreting information clues in relation to goals and expectations.[20,268] The choice between alternative solutions often is based on the alternative with the positive information cue available in memory while ignoring the rest. At the end of this process, the controller derives at a solution, or makes choice randomly if information cues cannot be discriminated. The process follows a simple three-step principle:[20]

- Search rule: describes which information is searched for.
- Stop rule: describes when to stop searching for information.
- Decision rule: describes which rules to apply on the information gathered.

Controllers may apply different search, stop, and decision rules when solving conflicts. As a result, differences in information gathered, and the rules for interpreting or weighing the information, results in different solutions to apply for solving the conflict. Although controllers experienced the same conflict repeatedly, differences in information considered makes each repetition unique form their perspective. If any issues or flaws are realized, the solution is modified or rejected altogether.

The controller then determines whether the advised solution conforms to her/his preferences for solving the conflict. In experiments, conformal advisories were accepted because they made sense and complied with the controller's preference for solving the conflict. The advisory is rejected if the controller arrives at a solution contrasting the advisory, unless the advised solution is considered better than their own for some reason. Since conformal advisories were based on the controller's earlier implemented solution to the same conflict, it was likely that the same solution would be accepted again, provided a similar reasoning occurred.

A precondition for this heuristic process to be effective is a somewhat structured problem-solving context. In en-route ATC, real-time operations are faceted by many irregularities and disturbances, notably time deviations between planned and actual flight and added complexity from short-term "ad hoc" traffic. In addition, the weather has a considerable influence on traffic movements and ATC capacity en-route and at airports. Despite these irregularities, however, movements within a sector, and the conflicts that occur have similar patterns. Airspace sections (i.e., sectors) are typically structured in terms of entry, exit, and intermediate waypoints that make up conditional airways that aircraft follow. Regular patterns can, for example, be attributed to general flows of eastern or western transatlantic flights during a working day. In addition, several airlines operate the same routes on a regular basis.

Moreover, it can always be argued that external similarity (i.e., *product*) is not the same as the similarity of underlying process (i.e., *process*). After all, automation algorithms are not visible from the visible result. However, although the underlying process is difficult to determine from only knowing the result, we do so every day during interaction with others. We infer the concealed reasoning on others based on what we can sense and observe. In this sense, human-automation interaction is no different than most human-human interaction. Future research should explicitly measure participants' understanding of advisories and investigate whether conformal advisories are better understood than nonconformal advisories, even though both objectively can be considered equally opaque.

7-3 Research challenges and limitations

This thesis has not been without its limitations and challenges. Several were foreseen at the onset of the research explorations, and have been addressed in the introduction of this thesis, while others were encountered and realized as the research progressed. In this section, those challenges and limitations considered most relevant are discussed. In addition, recommendations for future research are provided.

7-3-1 The great deception: defining conformal advisories

The ambition was never to develop an actual conformal system, but to empirically investigate the benefits and drawbacks of such capable automation. To do so, it was necessary to develop a method for determining and measuring how a person prefers to solve a problem, and empirically test that person's reaction when an automated system suggests the same solution to an identical problem.

Inspired by a previous study by Fuld *et al.*,[39] an experimental design was developed that set out to expose controllers to the same problem repeatedly and first observe and record their solutions, then provide replays of their own solutions as automated advisories. In simulations, each scenario was encountered four times. In the First empirical study (Chapter 3), conformal and nonconformal were matched directly to each scenario. While this use of exact replays ensured high conformance, it provided a challenge for ascertaining reliability in solutions across repetitions. Say that we want to develop a conformal system, how can a conformal solution be predicted by the system if the operator inconsistently solves the same problem over time? Therefore, in the two latter studies (Chapters 4 and 5), conformal advisories were defined by each controller's consistent problem-solving style.

In all studies, the solution parameter hierarchy was used to define conformal advisories. Following the Consistency study (Chapter 6), however, two more classifications were identified (discussed above: the control problem and solution geometry classifications). Since all conformal advisories were based on the solution parameter hierarchy only, some controller's conformal advisories were likely erroneous and better represented by one of the other classifications. The definition of consistency is critical for the creation of conformal automation. Future research is needed to further investigate these three classifications and to what extent they accurately capture controller's conflict solving preferences. In addition, longitudinal studies should explore the long-term effects of conformal automation on human-automation interaction.

7-3-2 Defining and displaying consistency

A consistent problem-solving style was defined by the same solution being applied in three out of four, or four out of four, repetitions. The few solutions collected per controller limited the consistency study. Measures may have been insensitive to consistency patterns for more detailed solution parameters. A challenge when studying individual's behavior is how to measure consistency. Statistical approaches have, traditionally, focused on the two-distributional moments of central tendency (e.g., mean) and dispersion (e.g., standard deviation) to describe the "average" behavior of a group.[237,269] With the emphasis shifting from the group to the individual, the applicability of these methods is questioned. When we turn to individuals, however, the same principle can be applied. But instead of gathering a large group of people (although still necessary), a large number of recordings from the same individual is needed. Statistically, a much larger data set, say 20 repetitions of a conflict, would have been desirable as a foundation for determining consistency. Therefore, future research should consider larger sets of repeated stimuli.

Consistency could not be established down to the detailed decision stage that specifies a directional value (such as a vector of 035 degrees). The simulation was limited in that the direct interaction with aircraft (i.e., by means of datalink) may have constrained consistency. The SSD interface facilitated precise speed and vector changes (by clicking and dragging the vector line) such as increasing the speed by 7 knots or turning 28 degrees right. In contrast to reality where traditional R/T is used, controllers generally provide instructions in average numbers, such as "QS1338, turn right heading 035." However, results showed no increase in consistency when considering controller's solutions in groups of 10 degrees or 10 knots.

With consistency limited to high-level decision stages, results provided a limited understanding of controllers' conflict solving. A convenient inference would be that detailed solution parameters are of less importance for solving conflicts, such as the direction of a vector and deviation required to attain a desired separation distance (i.e., decision stages 3 and 4 in Figure 6-1 in Chapter 6). However, previous research argues against this inference. For example, controllers have been found to aim for different separation distances depending on the current experienced workload.[79,80] Further research is required in order to better understand the relationship between detailed decision stages, for which controllers were found inconsistent, and the higher-level decision stages, for which controllers were found consistent.

Another challenge was how to display consistency and differences in consistency between individuals. Traditional presentations of data, such as boxplots and bar charts, do not satisfactory capture the individual and variations between individuals. Previous research has suggested that visual illustrations, such as pictures and films, may be the best way to present individual's data.[237] In the end, a sunburst chart presentation was used because it provided 1) an overview of the consistent patterns found for a specific conflict, 2) the proportion of inconsistent controllers, and 3) how frequent a certain problem-solving style (see Chapter 6). There may, however, be other more suitable ways to present individuals data.

7-3-3 The effect of advisory timing

Consistency was not defined in terms of *when* in time controllers solved the conflict across repetitions. However, solutions were investigated to determine if they varied over time. Particularly for a crossing conflict, the heading deviation required to solve the conflict increases over time as the two aircraft approach. Although results revealed considerable differences in when conflicts were solved, results did not suggest that solutions varied over time.

Determining *when* to provide decision support may be one of the most challenging design decisions. In ATC, the inappropriate timing of conflict alerts has been identified as a driver for resistance.[189] It was reasoned that a resolution advisory should be provided *before* the controller has determined a solution. Analogously, the greatest benefits of safety-net alarms and conflict warnings in autonomous cars have been achieved when made before the driver takes action (e.g., in terms of trust in system and driver breaking performance).[270–273]

However, some researchers have argued that advice provided before a decision has been made increases the risk for automation bias and complacency.[48–50] The same researchers note that by measuring peoples' reaction to an advisory *after* they have made a decision, the effect of the advisory on changing their solution can be determined. Indeed, there is a risk that controllers accepted advisories because it was convenient and effort saving. On the other hand, an advisory provided after the controller has decided, may confound their decision to either accept or reject. With this in mind, it would be interesting to investigate the effect of conformance on acceptance, when advisories are provided after controllers have decided on a solution. This can serve to validate a person's conformal advisory.

7-3-4 Dependent measures of trust and safety

Although the acceptance of an advisory is an indisputable measure of automation use, its dichotomous measure is a limitation. Therefore, an agreement rating was used as an indirect measure of automation use, requiring the controller to reflect and evaluate the advisory given. A person is likely to accept an advisory if she/he agrees with it, or reject it if she/he disagrees with it. Although this relationship was proven true in most cases, occasionally a controller would agree with an advisory but still reject it, or the other way around. This discrepancy can, perhaps, be attributed to effects of trust. Controllers were instructed that advisories always were safe and would solve the conflict, although not necessarily in the most optimal way. The intention was to neutralize trust (acting as a confound) by preventing controllers from questioning the safety of advisories. At the same time, instructions sought to prevent controllers from blindly accepting advisories, but questioning and comparing them with their own solutions. Given what is known about dispositional trust (see Chapter 4), however, it can be assumed that controllers had different trust attitudes that influenced their willingness to rely on advisories. Therefore, future research should investigate the effects of strategic conformance on participants' trust in both the system and its outputs.

The number of separation losses (for designed conflicts) was used as a measure of safety. Since no separation losses were recorded, safety was maintained. While the minimum separation distance between aircraft could have been used as a proxy for safety, it was decided not to. If considering a conflict between two aircraft, a theoretical perspective argues that safety increases with increased separation. From a practical perspective, however, safety appears to be more ambiguously applied. At distances above the separation requirements (typically 5 nmi horizontally and 1000 ft vertically) the perception of safety does not appear to increase linearly with the separation distance. Rather, the issue of loosing efficiency by distancing aircraft unnecessarily far apart appears increasingly important for controllers. Indeed, controllers' sometimes implemented "tight" vectors to increase efficiency while still maintaining separation.

These "tight" solutions may be linked to the SSDs explicit visualization of boundaries for safe separation, which invited controllers to "push" the safety envelope, reducing the margins for error, and operating closer to safety limits. This ties on to safety concerns of boundary migration in resilience engineering,[274] which also have been observed in empirical research with ecological interfaces.[121] Operators in safety-critical, dynamic environments typically view safety as the absence of failure and incidents.[275] This perspective has analogously been found representative for controllers, who judge the risk for conflict and need to intervene dichotomously depending on the distance between aircraft in relation to the stipulated separation minima.[61,276] As such, safety appears to be perceived binary and there may be little incentive for enhancing safety above the required safety threshold. Further research should investigate how controllers perceive safety.

7-3-5 Framing advice

Information given to people about an automated system can play a critical role in subsequent interactions with that system. In simulations, resolution advisories were always framed to originate from an automated source, except in the Source bias study (Chapter 4) where a human source also was used. Thus, controllers were not aware of the conformance manipulation. Considering the trend towards more personalized decision support, it would be interesting to investigate the effects of conformal automation with participants being informed that advice, truly, is based on their own problem-solving style. Would the framing of automation as either conformal or nonconformal amplify or cancel out acceptance and performance effects?

7-3-6 Horizontal traffic and simulator realism

The simulation environment was restricted to the horizontal plane, with vertical solutions not possible. This was desirable for purposes of experimental control, and for simplifying manual analysis of solutions and the creation of conformal and nonconformal advisories. In addition to making scenarios less realistic, the restricted maneuverability for solving conflicts reduced the variability in solutions observed. However, the horizontal limitation did not ease the CD&R task, rather made it more difficult as controllers only had speed and/or heading at their disposal for solving conflicts. Future research should incorporate a complete three-dimensional sector environment and solution space for analyzing strategic conformance. With more available options, solution variability can be expected to increase. Of particular interest would be the effect of introducing the vertical plane on conflict solving consistency and agreement.

7-4 Limitations and pitfalls of strategic conformance

Several limitations and pitfalls of strategic conformance were identified during concept development (see Chapter 2). While some were addressed in the empirical studies (e.g., consistency and agreement, expert and novice users in Chapter 6), new aspects have been identified. This final section discusses the most relevant concerns surrounding the strategic conformance concept.

7-4-1 What is the point in repeating human solutions?

It can be questioned whether emulating human decision-making strategies is desirable given decision-making biases and other cognitive limitations. This compelling criticism of strategic conformance relates to Kirwan and Flynn's first and last arguments against heuristic forms of automation (listed in Table 1, Chapter 2).[29] The arguments reflect two questions: will automation thinking and making decisions like a human not make the same mistakes as the human? And, what is the benefit of automation making the same decisions as the human would? Although these questions touch upon two weaknesses of strategic conformance, they have, to some extent, been addressed in previous research.

In addressing the first weakness, a possible approach is to identify and eliminate poor decision-making strategies and focus on "best practice" strategies (e.g., CORA). A similar approach is advocated by Gigerenzer and colleagues[20, 82, 268] who argue that heuristics can be quantified into formal models (called fast-and-frugal decision trees) that not only are more compatible with human decision-making, but also outperform other logical and statistical models in accuracy. A similar, slightly more advanced approach can be found in bootstrapping models, which are a type of expert system based on the quantified judgment and decision-making rules and strategies made by experts. Bootstrapping models are thought to improve the reliability and accuracy of judgments, since an expert's consideration of a criterion, and how it is weighed, is applied consistently. On the other hand, bootstrapping models lack the flexibility and adaptability of human experts, and cannot consider variables or cues outside the scope of the model. As with fast-and-frugal heuristics, the accuracy of bootstrapping models often deteriorates when more variables are considered.[277]

Similar to heuristic and bootstrapping models, strategic conformal automation would embrace the "good and safe" practices. A CD&R decision aid, for example, would be attuned to individual workload demands, and allow situations to progress longer during low workload episodes to minimize path deviations and promote efficiency. Conversely, during high workload situations, automation could intervene early and apply larger separation thresholds, the strategy being to "set and forget"[24] and allow the controller to move on to the next task. Other appropriate controller heuristics include:

- approach conflicts in pairs and sequentially, rather than globally, determine if multiple aircraft are involved;[35]
- avoid interfering with both conflicting aircraft in lower airspace;[88]
- being conservative in conflict detection, identifying a surplus of potential conflicts;[87]
- a tendency to operate at larger separation thresholds than is strictly required;[61, 98, 278] and
- safer to turn slower aircraft behind faster;[24, 35]

In response to the second weakness, strategic conformal automation is thought to add a number of benefits, including the speed and reliability at which tasks (especially repetitive and mundane ones) can be accomplished. More importantly, it can reduce workload and free resources for other pressing tasks. Strategic conformal automation should not simply mimic operator solutions. Rather, it should strive to match the underling decision-making style, and be able to provide solutions conformal with the operator's individual style. Strategic conformal automation would enhance performance by suggesting solutions not foreseen by the operator, although in coherence with her/his problem-solving style. This could lead to benefits in terms of increased automation acceptance and trust, whilst enabling the controller to evaluate support in real time and lessen the cognitive burden.

Strategic conformal automation can, however, be considered controversial in that acceptance and trust may be influenced independently from how "good" the actual decision-making strategy is. That is, the perceived "good" behavior of automation can conceal poor performance. There is a risk of skewed belief in the capabilities and qualities of the automation leading to unnatural high expectations of the automation.[114,183] This is an issue that requires further attention.

7-4-2 Restricted to initial acceptance?

Over time, the importance and practical benefits of strategic conformal automation can be questioned considering daily and prolonged interaction with automated systems. The greatest threats to acceptance are likely to emerge during initial interaction with automation. Therefore, the largest benefits of strategic conformance may be to facilitate initial acceptance and encourage operators to gain familiarity with the system. Following prolonged use, strategic conformance may become less important as other acceptance drivers, such as trust and perceived reliability, take precedence. Possibly, strategic conformal automation could be applied during a transitional period (e.g., training) to first gain acceptance, and then gradually change the automation's decision-making style to a more suitable strategy (e.g., in terms of increased efficiency, safety, or other target). In this way, automation could teach the operator how to improve performance.

7-4-3 Restricted to expert operators?

Strategic conformal automation may be most useful for operators who are already experts within their fields. Operational experience would then be the key factor. Consider for example pilots who, for a certain period of several years, typically operate only one aircraft type, although in very different and dynamic environments. Despite seemingly limited degrees of freedom, and highly-regulated operating procedures, pilots develop their own individual procedures and performance criteria. When introducing new flight deck automation, it is likely that experienced pilots, compared to novice pilots, will be more reluctant to use the automation if it is non-conformal. Novices with no or limited experience, however, lack well-developed problem-solving styles and may not be equally susceptible to conformance.

7-4-4 Is strategic conformance simply adaptive automation?

Strategic conformance complements the notion of individual-sensitive automation as captured by adaptive automation research. Generally, adaptive automation refers to automation capable of "reading" and reacting to human physical and behavioral changes as a means to primarily ensure safety and maintain optimal system performance, but that also can benefit acceptance of automation.[279–281] This includes research on the mechanisms for triggering (when and how) and deactivating the automation. Adaptive automation is thus predominantly limited to adapting the degree of automation authority (at times overruling the human operator), whereas automation behavioral changes sensitive to individual differences in problem-solving styles are more interesting to address issues of not only acceptance, but also safety and performance.

Strategic conformance can be considered for mechanisms used by adaptive automation to assist human decision-making. Consider the automobile domain for example, in which attempts have been made to develop adaptive automation that can supervise and support drivers to increase safety.[243] An adaptive cruise control system could be designed to match a driver's preferences in longitudinal control (e.g., distance keeping, acceleration/deceleration, and braking patterns). Strategic conformal automation could be attuned to different individual driving styles, such as fluid, moderate, comfortable,[242] or economical, medium, sporting.[282]

The greatest benefit of strategic conformal automation may be in situations where optimal performance is less important. This could be automation that supports our daily life, such as autonomous cruise control systems in cars. Within certain safety boundaries, there is an acceptable performance envelope in which cruise control systems operate. Within this envelope, there is a possibility to tune the performance to the individual driver. Drivers who operate outside of the "envelope," however, will likely complain.

7-4-5 Conformance depends on the HMI used

Findings from the Automation transparency study showed that conformance may depend on HMI used when solving conflicts. For solving conflict in prequel simulations, controllers interacted with the heading band SSD. Since conformal and nonconformal advisories were based on these solutions, the solutions were also tied to the use of the heading band SSD. In the experiment simulations, the conformal and nonconformal solution advisories were presented inside the SSD, which opened up simultaneously with the advisory. This use of the SSD in prequel simulations as a baseline CD&R support tool was deemed necessary to provide a high level of information integration. This, in turn, facilitated the manipulation of strategic conformance in relation to high levels of automated support (at the stage of decision-selection and implementation) simply by adding conformal and nonconformal advisories.

This can explain why agreement and response time patterns were reversed when controllers used the triangle SSD. The triangle representation influenced controllers' perception of the conflict, and the solution they implemented when using the heading band SSD (prequel simulation) appeared less suitable when a similar solution was presented in the triangle SSD (experiment simulation). This has three important consequences. First, a conformal solution as recorded when using interface A may not be representative as a conformal solutions when using interface B. Second, varying the transparency of an interface can have a fundamental effect on how problems are perceived and, hence, solved. Finally, this suggest that strategic conformance and automation transparency should be considered mutually exclusive.

Admittedly, the decision to use the heading band SSD as a baseline may have biased solutions. With speed options not readily provided by the interface, controllers may have relied extensively on vectors for solving conflicts. The motivation for using the SSD was driven by the objective of studying automation acceptance in a high-density sector, representative to that of future traffic levels. The baseline heading band SSD provided the information integration and support needed to cope with the increased number of aircraft. Yet, strategic conformance should advisably be studied in a more realistic context, analogous to current ATC. Of particular interest would be to investigate controllers' conflict-solving styles, their consistency, and their disagreement, using the solution classification framework developed in this thesis.

7-4-6 The benefit of nonconformal automation

Like any decision aid, a conformal system should not suggest unsafe solutions only to be conformal with the operator's problem-solving style. The system must be able to consider the safety and efficiency of a solution, and if determined unsuitable, disagree with the operator and argue for a nonconformal solution. Notably, in such circumstances, the contribution of nonconformal automation may be more beneficial than conformal automation to system performance and safety. Because, to derive at a conformal advice, the system must 1) have knowledge about the operator's preferred problem-solving style, 2) be able to acquire an understanding of the operator's reasoning in relation to the situation at hand, and 3) determine a solution for the given situation that matches this style. This capability implies that the system also can evaluate the suitability of a solution, and therefore be able to explain why the conformal solution is considered inappropriate, and argue for a more appropriate, nonconformal, course of action.

7-4-7 Domains benefiting from conformal automation

Although this thesis has been restricted to ATC CD&R, conformal automation has potential applications in any domain that contains automated systems. Directly related applications can be found in other transportation domains that deal with similar contexts of collision avoidance, such as automation used in vehicle cockpits or traffic monitoring services. More broadly, however, strategic conformance does not relate so much to the task or problem at hand, but the behavior of the automation and the degree to which it complies with the human preferred way of working. In this perspective, conformal automation implies a more holistic design philosophy that strives to harmonize the automated system's behavior to that of the human operator.

7-4-8 Homogeneity versus heterogeneity in automation design

This thesis shows that controllers are more diverse than alike in conflict solving. In light of this, a relevant question is whether to strive for heterogeneity or homogeneity in automation design. That is, should we design for the individual user, or for the population? The question echoes that of Hopkin [235, p. 81]: "Should system adaptability be viewed primarily as a means of encouraging individual differences between controllers or as a means of preventing them, since it has the potential for both?" Traditionally, design has favored the homogeneous view.

Compared to automation, the human has been identified as exceptional in creative thinking, identifying new solutions to problems, and the ability to adapt to its changing surrounding (e.g., Fitts list[283]). Yet, systems' interaction models and problem-solving algorithms tend to be fixed and narrow, instead appealing to the human strength of adapting to the system. In addition, these models and algorithms typically require that operators are trained homogeneously, approaching problems and solving them in line with the system's rationale. In a way, this approach neutralizes the human variability, flexibility, and creativity. The argument in favor of heterogeneity (i.e., personalization), is that automation sensitive to the user's preferences and abilities can benefit acceptance, performance, enjoyment, and teamwork with the automated agent.

There are, however, aspects that may argue against heterogeneity, which instead advocates homogeneity and the suppressing of individual differences. Personalized automation may work against proceduralized environments like ATC, where supervision and handovers from one controller to another require an understanding of what the previous controller is doing and why. In safety-critical domains, standardized operating procedures facilitate increased predictability and repeatability of situations, which in turn, supports safety and effective teamwork. Furthermore, homogeneity allows for the definition and evaluation of global performance and capacity constraints. These may be undermined by the increased recognition of individual differences in key performance measures of performance and capacity that affect sectorization and staffing.

The flight deck is a good example of a safety-critical environment where homogeneity in automated support systems and standard operating procedures (defining system interaction) traditionally has worked well. Although accidents are rare, several accidents have been attributed to poor human-automation collaboration linked to opaque automated systems that do not explain its reasoning (e.g., Air Asiana flight 214,[151] Air France flight 447,[284] Turkish Airlines flight 1951[285]). There is a need to increase the transparency of such systems, for example in regards to their current mode, function, and operational limitations.[186,187,286] A further, complementary approach, would be to consider conformal automation and afford heterogeneity in pilot-system cooperation by personalizing the interaction.

7-4-9 Personalized automation

Results support the general claim made by previous researchers that future automation increasingly needs to embrace individual differences in problem-solving tasks.[154,155,198] Automation etiquette studies have highlighted the need to consider the automation's communication style and timing in relation to operator preferences.[114,120] Researchers on anthropomorphism have argued that a system similar to the operator in appearance can benefit trust in that system.[287] Recommender systems research has explored the use of explanations and justifications for supporting action and decision advice.[211]

Design approaches are needed that can combine requirements of standardized operations with the allowance for individual differences in human perception and problem-solving. Conformal automation emphasizes personalization in relation to the operator's cognition. Findings reported herein can benefit the development of personalized CD&R decision aids. With consistency limited to the first decision stages of each classification, it may be relatively easy to model and personalize a conflict resolution aid. A challenge, however, is that the human may adapt her/his problem-solving (affecting conformance) as a response to the system adapting.

Methods can be found in other fields such as recommender systems[201] and user-adaptive automation.[193] Personalization can, for example, be facilitated by the creation of individual user profiles generated from an individual's system-interaction data, alternatively data from other operators with similar preferences. For instance, Regtuit proposed a machine-learning technique consisting of k-means clustering for identifying consistent patterns in controllers conflict solving strategies that can be used to develop a conformal conflict resolution decision aid.[261] Moreover, physiological measures can be used to, for example, collect operators' eye-gaze patterns in real time for purposes of dynamically track their behavior and adapt the visualization to their visual and cognitive abilities, needs, and preferences.[288] Personalization can be assessed together with neuroscience, by studying how individual differences in problem-solving and cognitive styles are associated with the activation of cognitive processes in different brain areas.[198]

Accommodating individual problem-solving in EID. As an alternative to tailored decision advisories, a system can accommodate personalization by allowing a person to solve problems in their preferred way. In many contexts, including ATC CD&R, this approach may be preferable to that of a decision aid suggesting a specific solution. For instance, ecological interfaces, such as the SSD, can visualize the constraints and space of solution possibilities affecting a situation and problem, but leave the final decision-making to the operator. This acknowledges that humans can be heterogeneous (the subjective truth), while situations and the constraints are homogeneous (the objective truth). In addition, another strength of this approach is that advisory timing does not need to be considered.

This approach for facilitating personalization is particularly relevant to work domains in which ample time is available for solving problems. However, in time-critical and complex domains, such as ATC CD&R, problem-solving time is often limited. Furthermore, less time for solving conflicts can be expected as more tactical and sophisticated automation is introduced for purposes of increasing traffic throughput and efficiency. In light of these changes, the controller's role can be expected to change from tactical traffic intervention to supervising CD&R automation and only infrequently solve specific conflicts. As a consequence, the relevance of conformal automation may increase.

7-4-10 Side effects and ethical concerns

The reliability of advisories was not varied. Although controllers sometimes disagreed with advisories, they may not have had enough reason to reject advisories and therefore accepted more. A question is whether unreliable conformal automation affects acceptance and trust differently from other types of automation that are not personalized. Similar to anthropomorphic features of automation that do not concern its performance,[48, 114, 287] conformal automation may influence acceptance inappropriately in relation to the automation's actual reliability. Operators' expectations on conformal automation may be inflated and unreasonably high, leading to issues of automation bias and complacency, and large decays in trust and acceptance when the automation fails (in line with the perfect automation schema). Ironically, however, conformal automation may counter adverse effects of the perfect automation schema, automation bias, and complacency. Research indicates that the more human-like automation is, the more it is treated as a human (i.e., the media-equation effect).[48, 118, 124] Future research should explore whether similar effects can be found in relation to conformal automation.

Strategic conformal automation can, however, be considered controversial in that acceptance and trust may be influenced independently from how "good" the actual decision-making strategy is. That is, the perceived "good" behavior of automation can conceal poor performance. There is a risk of skewed belief in the capabilities and qualities of the automation leading to unnatural high expectations of the automation.[114, 183] Ideally, automation should be designed to facilitate that operators' attitudes toward the system reflect the system's actual reliability. Because the interaction with automation is influenced by the degree of conformance, the purpose of the automation and the desired response from the operator must be considered during design.

With this in mind, the degree to which the automation is conformal to the human (whether in problem-solving or by appearance) is also an ethical one. Is it ethically acceptable to develop conformal automation that influences a person's behavior, even if the system's behavior is designed to match that same person? Moreover, is it ethical to have a machine learn about a person and partly mimic that person? Most important are perhaps the ethical issues of privacy, related to the gathering and treatment of an individual's data. As a suggestion for future research, ethical guidelines from research on persuasive technologies[289] can perhaps be used to mitigate some of theses concerns of conformal automation.

7-4-11 Why not fully automate?

Finally, it is worth questioning the contribution of strategic conformance given the development towards a fully automated system that, eventually, may allow designers to entirely disregard the human way of solving problems. Indeed, machines have matured cognitively and become increasingly smart. Abilities previously considered as exclusively human are being automated in a fast pace. This includes combinations of sensors and algorithms providing the visually impaired with key aspects of sight (e.g., OrCam system), artificial intelligent decision aids supporting pathologists in diagnosing cancer (e.g., C-Patch system), and IBM's *Watson* computer competing and wining against humans in *Jeopardy!*[290] By means of methods such as machine learning, natural language processing, and computer vision, the development of artificially intelligent machines allows for automating complex tasks and processes that previously were considered "too open."

The currently most dramatic change is perhaps the development of fully autonomous, "self-driving," cars. Fully autonomous cars is a good example of a domain in which strategic conformance is particularly relevant. The experience of the human passenger (previously driver) will largely depend on the comfort from traveling in the car. A better experience can be expected provided the autonomous car "drives" in a way conformal to the passenger's preferences (e.g., in terms of accelerations, braking, turning, and distance keeping). Another example may be a fully autonomous ATC CD&R system supervised by a controller. Both acceptance and trust in the system are likely to be higher if the system solves conflicts according to the controller's preferences, thus resulting in fewer interventions. Taken together, the introduction of fully autonomous systems may not lessen the relevance of strategic conformal automation, provided the human still has to interact with the system from time to time.

8

Conclusions

This chapter revisits the key problem and answers the main research question investigated in this thesis. A brief review is provided of the most relevant results and findings from each chapter. These are considered from a holistic perspective, for the purpose of specifying the contributions of this thesis to research. Finally, with respect to what inspired this thesis, and analogously to how it started, the value and implications of strategic conformance are considered from a broad philosophical perspective.

THE immediate value of this thesis is anchored to the problem definition in the context of choice: How to overcome controller acceptance issues of automated decision aids for conflict detection and resolution? This problem is set against the fragmented, increasingly crowded, and unevenly trafficked airspace worldwide. In order for ATC services to host and cope with future traffic demands, there is a need for more advanced and intelligent automation that can ease the air traffic controller's cognitive burden in problem-solving tasks. A problem, however, is that systems developed for this very purpose have been, and continue to be, rejected by controllers.

This thesis proposed that the resistance towards such automated support systems, in particular ATC CD&R decision aids, can be explained by machines reasoning and solving problems differently from controllers. This difference was conceptualized and defined on a continuum of *strategic conformance*: The degree to which automation's solution and apparent underlying operations match those of the human. As such, it is argued that strategic *conformal* solution advice, which matches the human's solution, facilitates an intuitive, (sub)conscious understanding of the automation's reasoning. Strategic conformance was introduced at the highest level of human-machine compatibility extending previous research by 1) exploring the concealed underlying strategies involved in decision-making, and 2) exploring the diversity in problem-solving by recognizing each individual's problem-solving preferences. The goal pursued in this thesis was:

Thesis goal

To empirically investigate strategic conformance as a means for more personalized automation support, and develop a fundamental understanding of how a decision aid's strategic conformance affects the interaction with that aid and acceptance of its advisories.

Findings in this thesis have significant implications for the understanding of how controllers' solve conflicts and respond to resolution advisories depending on its conformance. The analysis of controllers' manual conflict solutions showed that controllers solved conflicts consistently, but disagree. Experienced controllers were found slightly more consistent than trainees, possibly because of their more well-developed and established conflict solving preferences and strategies. An implication of this is the possibility to create more personalized, conformal automation. A limitation is that consistency was limited to higher decision stages (e.g., choosing one or both aircraft, or aircraft A or B to interact with), and could not be established for more detailed decision stages (e.g., vectoring aircraft A to the right by 035 degrees).

Results from three human-in-the-loop studies indicate that strategic conformance plays an important role for the acceptance of automated resolution advisories. Conformal advisories were found to benefit both acceptance and agreement, as well as reduce the response time to advisories, both among experienced controllers and trainees. Notwithstanding the limited sample in the Source bias study, results provided valuable insights into controllers' perceptions of automated versus human advisers. In support of previous findings, the human adviser (presented as a controller) was perceived to provide solutions slightly safer, more efficient, and more similar to the controller's own solution. Whilst no interaction was found between strategic conformance and interface transparency in the Automation transparency study, the findings do suggest that the problem-solving style definition (i.e., how an individual solves a particular problem), hence the conformal advisory, also depends on the human-machine interface used. As such, a conformal advisory when using interface A, may not be conformal with interface B.

In addition to benefits observed in this thesis, there are numerous potential benefits of conformal automation that deserves attention in future research. Further studies could assess the use of strategic conformal automation in training, as a tool for learning operators to solve problems better. First gaining operators' acceptance and trust during early interaction, then gradually change to suggest better solutions (that may at first be nonformal to the operator). Future research might explore alternative approaches to strategic conformal automation. As an alternative to explicit advisories, conformal support can be facilitated by designing interfaces that allow operators to solve problems in their own preferred way. Ecological interfaces, for example, visualize the "objective truth" by revealing and showing the full constraints affecting a situation. Importantly, such interfaces acknowledge that humans are heterogeneous in problem-solving, although the objective truth of a situation is homogeneous.

This freedom may, however, not be practicable in all contexts. Future ATC CD&R, for example, is striving to relinquish tactical control from the controller altogether, using more advanced automation. Given the expected traffic growth, it is unlikely that, in the future, the controller will have time to consider conflicts in detail. In such contexts, providing explicit advisories conformal to the operator's problem-solving style is a candidate approach for relinquishing problem-solving to the automation and freeing up the operator's time and resources. To move the debate forward, a better understanding needs to be developed in regards to the suitability of personalized automation for different domains.

The advancement in automation design has made it appropriate to question and challenge the development of "one-size-fits-all" systems, in favor of more personalized applications. A natural progression of this thesis is thus the development of real-time personalized, conformal automation. There are, however, several challenges, including strenuous technical ones in regards to the measurement and identification of problem-solving styles and the individual tuning of decision support. Here, machine learning techniques can pave the way forward. In association with these are also several ethical concerns. Personalized systems have the power to influence acceptance and trust independent from the system's actual performance and reliability. Other ethical concerns are related to the modeling and "mimicking" of a person's behavior, and individual privacy.

The broad philosophical value of this thesis is tied to Alan Turing's proposed "imitation game" for defining artificial intelligence, captured by his intriguing question: "can machines think?"[38] The game argues that a machine can be considered intelligent, inseparably from that of a human, provided an interrogator conversing with the machine believes it to be another human. While artificial intelligence falls outside the scope of this thesis, the concept of strategic conformance denotes such capable machines that can solve problems analogously to its human counterpart. Thereby, the investigations of conformal decision support in this thesis has served an equally intriguing and perhaps more important question, also posed by Turing: "is the question (can machines think?) a worthy one to investigate?" (p. 434).[38]

Automated systems in safety-critical domains should not be personalized simply because they can be. In fact, personalized applications, such as conformal decision support, may not be desirable in many contexts. Particular domains involving many actors, characterized by significant teamwork and collaboration, may not be suitable for conformal automation, as it may hamper a shared understanding. The benefits of a personalized system must therefore be weighed against the desire and need for a homogeneous workplace. As such, there is a sensitive balance to be considered, with the benefits of personalizing automation to an individual's acceptance and performance at one end, and the desire for suppressing individuality altogether in order to attain an organized and structured environment needed for facilitating effective and safe teamwork at the other end.

Appendices

A

Conformance design

A-1 Introduction

The manipulation of conformance was based on a general experimental design consisting of three phases. 1) The *prequel* real-time simulation was conducted with the purpose of collecting controllers' manual conflict solutions of the designed conflicts. Note that conflict solving was supported by the SSD interface and a safety-net short-term conflict advisory. Controllers did, however, not received specific resolution advisories to the designed conflicts. 2) The next phase, *conformance design*, consisted of analyzing these solutions per individual controller, and determine their problem-solving styles that best described their preferences for solving conflicts. These were used to script conformal and nonconformal solutions to be proposed by the conflict resolution aid in the subsequent third phase. 3) Finally, in the *experiment* real-time simulation, controllers encountered the same scenarios and conflicts as in the prequel simulation, only this time supported by the conflict resolution aid with scripted conformal and nonconformal resolution advisories.

This appendix describes the *conformance design* phase. To ensure reliability, three researchers accomplished the conformance design phase analyzes in parallel. The conformance design phase followed a process of four steps:

1. Creation of a solution parameters framework;

2. decoding solutions according to this framework;

3. identify problem-solving style; and

4. define conformal and nonconformal resolution advisories.

A-1-1 Creation of a solution parameters framework

A conformal resolution advisory was initially intended to be an exact replay of a controller's conflict solution as recorded in the manual prequel simulation. In the manual simulation, controllers were expected to interact with only one of the aircraft in conflict, and solve the conflict with a single implementation. Furthermore, with the guidance provided by the SSD, controllers were expected to implement solution that would not cause a secondary conflict. However, these initial assumptions were challenged during early test-trials. Controllers solved conflicts in many more ways not foreseen. For instance, they sometimes solved conflicts:

- through interaction with both aircraft (i.e., dual aircraft resolution);
- only to later change their mind and implement another resolution;
- causing a secondary conflict with a noise aircraft.

These discrepancies provided a dilemma when using prequel experiment resolutions to create conformal resolution advisories for the conformance experiment. Exact replays of controller's own solution for the same designed conflict were not always possible. The solution to this problem consisted of 1) quantitatively decoding solutions in a hierarchy consisting of five parameters, and 2) determining how far down in the hierarchy a workable solution could be found that did not violate the requirements of solving the conflict by at least two interactions and not causing a secondary conflict. For solutions that were changed (e.g., implement one solution only to later change it), only the first solution was considered. The solution parameter framework is shown in Figure A-1.

The five solution parameters are considered in a progressive order of conflict solution granularity. The level of solution detail increases when analyzing solutions from top to bottom in the framework. An exception to this rule is the first defining parameter of *intervention time* (why it is also squared and not in the shape of an arrow-like pentagon). Intervention time is a qualitatively different measure from the other parameters that define an aircraft state change. Intervention time is likely to very considerably between participants because of the precise measurement.

Intervention time is an important solution parameter for several reasons. First, given a constant movement of aircraft, the passage of time will constantly change the environment and affect the options available to solve the conflict. A solution is only valid for a limited period of time. Second, the time of observable intervention might not accurately capture the time of conflict detection. For instance, a controller might choose to postpone intervention but continue to monitor and generate alternate solutions.

Solutions parameter	Description	Measure
Intervention time	Time after scenario start that conflict aircraft are interacted with.	Seconds (s).
Aircraft choice	Aircraft interacted with.	Aircraft call sign.
Resolution type	Type of resolution command.	Heading (H); Speed (S); or Combination of H and S (C).
Direction	Direction of resolution type.	Heading right (R); Heading left (L); Speed increase (I); Speed decrease (D); Combinations (RI); (LI); (RD); or (LD).
Directional value	Value of resolution command.	Deviation from original state value in degrees heading, knots speed, or combination.

FIGURE A-1: Solution parameter framework

If excluding intervention time, the initial solution parameter is *aircraft choice*. To solve the conflict, the controller will choose one aircraft to interact with. Even if both aircraft are interacted with, one will be chosen first. Next is the *resolution type*, which refers to the general category of control maneuver(s) chosen: heading, speed, or a combination of both. Note that vertical and time control (i.e., issuing an aircraft to reach a specific position at a certain time) maneuvers are not included as they (for reasons of experimental control) are not available in the simulation. The third step in the framework further defines the resolution type by assigning *direction*. A turn, for example, can be made either left or right. In the final step, the exact *directional value* of the conflict resolution is determined by the deviation from the aircraft's current directional value. For instance, aircraft A may be turned 30 degrees to the left to solve the conflict. Note that these four solution parameters later were defined as decision stages 2 to 4 in the Consistency study (the first decision stage being defined by three different solution classifications).

Controllers' solution variation was expected to increase for each lower step in the solution hierarchy. For example, while a large number of controllers may choose aircraft A for solving the conflict, they may opt for different resolution types. A few controllers may prefer to solve the conflict by turning aircraft A to the right, in contrast to the majority who to turn it left. A third group may solve the conflict by speed, issuing either a speed increase or reduction.

TABLE A-1: Decoded solutions for controller A

Scen.	Rot.	T-logs	Aircraft ID	Resolution type			Scen. specific comments	General observations
				H	S	C		
1A	0	32 & 47	QS1338 & OM3185	L L			Conflict not solved after first interaction, vector on the border. Solved after second interaction.	Tight vectors to no-go areas. No use of speed.
1B	0	53	OM3185	R				
1C	180	63	OM3185	R				
1D	180	51	OM3185	R				

Scen. = Scenario; Rot = Rotation; T-logs = Time logs; H = Heading; S = Speed; C = Combined; R = Right; L = Left; I = Increase; D = Decrease

TABLE A-2: Decoded solutions for controller B

Scen.	Rot.	T-logs	Aircraft ID	Resolution type H	S	C	Scen. specific comments	General observations
1A	0	49 & 58	QS1338 & OM3185	R		 RD	QS right turn into grey, conflict not solved after one interaction. OM RD, from 260 to 230 kts.	Grey aircraft first, conflicts and tricky areas second. Shoots small gaps, vectors closely to no-go zones. Rather change two aircraft a little than only one.
1B	0	50 & 56	QS1338 & OM3185			RI RD	QS from 260 to 310 kts. OM from 260 to 230 kts.	
1C	180	50 & 56 & 71	OM3185 & QS1338 & QS1338	L L L			Required three interactions. First two vectors in grey.	
1D	180	51 & 54 & 59	QS1338 & OM3185 & OM3185	R R	 D		First turn QS into grey. OM turned R, not enough. Third interaction OM decreased from 260 to 210 kts.	

Scen. = Scenario; Rot = Rotation; T-logs = Time logs; H = Heading; S = Speed; C = Combined; R = Right; L = Left; I = Increase; D = Decrease

A-1-2 Decoding solutions

Controllers' solutions to all designed conflicts in the prequel simulation were manually analyzed and decoded. The decoding transcript describes each solution according to a standardized *pattern*, defined by the solution parameter framework. Table A-1 shows the standardized template format in which all solutions have been entered for controller A. In this example, the decoded solutions for the designed conflict in Study 2 are shown for all four scenario replicates (1A through 1D). Note that for scenario 1A, the solution consisted of interacting with both aircraft: initially vectoring QS1338 left at time log 32, then vectoring OM3185 left at time log 47.

In contrast to controller A in the example above, others were considerably less consistent in their conflict solving. Table A-2 shows controller B's solutions for the designed conflict in Study 2. Controller B consistently interacted with both aircraft, at least once. In 1C and 1D three interactions were needed to solve the conflict. In three out of four solutions, controller B interacted with QS1338 first and then OM3185. In addition, controller B occasionally used speed to solve the conflict.

A-1-3 Define problem-solving style

A controller's problem-solving style represents general preference characteristics for solving a particular conflict. It is determined from comparing a controller's transcribed solution patterns for a repeatedly encountered conflict and identifying similarities. The problem-solving style specifies the preferred aircraft, resolution, direction, and directional value. A rationale supporting the identified problem-solving style is provided. Table A-3 shows the problem-solving style for controller A. The problem-solving style for controller A, for this particular conflict, consisted of vectoring OM3185 to the right behind QS1338. In this case, the problem-solving style was easy to determine because of the high degree of similarity (three out of four solutions). However, it was sometimes difficult to determine a controller's problem-solving style because of a large variability in solutions across scenario repetitions. Table A-4 shows the problem-solving style for controller B. Controller B's problem-solving style was determined to be a combination of heading and speed: to simultaneously instruct QS1338 to turn right and increase speed.

TABLE A-3: Problem-solving style for controller A

Aircraft ID	Resolution type	Direction	Value	Rationale
OM3185	Heading	Right	35°	OM turned right in three out of four scenarios. Only relied on heading to solve all designed conflicts. No use of speed.

TABLE A-4: Problem-solving style for controller B

Aircraft ID	Resolution type	Direction	Value	Rationale
QS1338	Combined	Right increase	15° & 60 kts	All solutions consisted of dual aircraft solutions. Interacted with QS first in three out of four solutions. In all three cases QS was turned right, once with a speed increase. Often used speed and heading combinations.

Note that a controller's problem-solving style for a specific conflict represents that controller's consistent solution for that conflict provided that the same solution was implemented in at least three out of four scenario repetitions. The consistency and agreement study was based on the analysis of controller's problem-solving styles (consistency measure) and how these differed between controllers (agreement measure). Note that the solution parameter framework was used as one of three solution classifications (the solution parameter hierarchy) in the consistency and agreement study (Chapter 6). The other two classifications were the control problem and solution geometry classifications.

A-1-4 Define conformal and nonconformal solutions

The strategic conformance of a conflict resolution advisory can be considered on a spectrum ranging from conformal to nonconformal. A controller's perfect conformal resolution advisory would exactly match all solution parameters defining that controller's problem-solving style. Any deviation from the reference problem-solving style, at any step in the solution parameter framework, would make the resolution advisory less conformal and increasingly nonconformal.

The definition of a controller's conformal resolution advisory was based on that controller's problem-solving style. The definition of nonconformal resolution advisories relied on identifying and comparing similarities in and differences between controllers' problem-solving styles. Reliability was achieved by having three researchers agreeing on each controller's conformal and nonconformal resolution advisory for each designed conflict.

For the nonconformal resolution advisories we wanted solutions that contrasted a controller's own solution. The initial idea was to simply look at each controller's own solution and manually script a different resolution advisory. The issue with this approach, however, was that, since we are not controllers ourselves, we might script an unrealistic resolution advisory that could confound our experiments. Furthermore, we did not have time or resources to develop an algorithm that satisfied the strict requirements of an automated conflict resolution decision aids. Our solution for determining a controller's nonconformal resolution advisory was to use a solution implemented by a colleague for the same designed conflict. This assured that the solution suggested was realistic and not completely random. This approach assumed that there were sufficient variability between controllers' conflict solutions.

During our research, we realized that a rigorous and detailed solution analysis as outlined by the solution parameter framework was impractical. When considering the level of detail specified by intervention time and the fourth step (directional value) in the hierarchy, every solution was unique. Thus, it was unreasonable to expect controllers to solve repeated conflicts exactly the same. This can be interpreted as an inconsistency in controllers' conflict solving, which complicates the identification of conformal solutions. However, the important aspect here is not how well a resolution advisory matches a controller's problem-solving style in detail. In theory, they need only to appear enough similar so that the controller perceives the resolution advisory to be conformal (i.e., matching) wit his/her preferred solution. As such, the perceived similarity is determined by the controller's perceptual limitations. It is reasonable to expect that in most cases, a 20 degree turn and a 25 degree turn would be perceived as identical. In contrast, interacting with different aircraft to solve the same conflict would likely be seen as two different solutions. Therefore we decided to limit the determination of problem-solving styles to the third step in the solution parameter framework (direction).

Table A-5 illustrates the criteria used for determining the degree of conformance for a resolution advisory. Five conformance levels were considered. Conformal res-

TABLE A-5: Levels of conformance

Level	Conformance	Description
1	Exact replay	Aircraft match, resolution type match, resolution direction match, and resolution value match (Speed in kts, Heading in degrees).
2	Mismatched resolution value	Aircraft match, resolution type match, resolution direction match. Resolution value mismatch (i.e., Heading differed by 20 degrees or more, and/or Speed differed by 20 kts or more).
3	Mismatched resolution direction	Aircraft match, resolution type match. Resolution direction mismatch (e.g., Heading left instead of right; Speed decrease instead of increase).
4	Mismatched resolution type	Aircraft match. Resolution type mismatch.
5	Aircraft mismatch	Aircraft mismatch

TABLE A-6: Conformal and nonconformal solution for controller A

Conformance	Aircraft ID	Resolution type (direction)	Description
Conformal	OM3185	Heading (right)	Conformal solution consists of vectoring OM to the right
Nonconformal	QS1338	Combined (right increase)	Nonconformal solution consists of giving QS a right vector with a speed increase.

TABLE A-7: Conformal and nonconformal solution for controller B

Conformance	Aircraft ID	Resolution type (direction)	Description
Conformal	QS1338	Combined (right increase)	Giving QS a right vector with a speed increase.
Nonconformal	OM3185	Heading (right)	Vectoring OM to the right, no speed change.

olution advisories were defined as levels 1 or 2. The requirement was a match in aircraft choice, resolution type (e.g., heading), and resolution direction (e.g., right turn). Note that this represents a match of the first three steps in the solution parameter framework. A nonconformal solution therefore always featured a different aircraft choice and/or resolution type and direction.

In relation to the examples above, we can determine the conformal and nonconformal resolution advisories for controllers A and B. Table A-6 specifies the associated advisories for controller A. Similar to that controller's problem-solving style, the conformal advisory suggests turning OM to the right. The nonconformal advisory is based on a colleagues contrasting problem-solving style for the same conflict, in this case controller B. As evident in Table A-7, controller B's nonconformal advisory is based on controller A's problem-solving style.

A-1-5 Limitations

Intervention time was not included as a defining parameter in the conformance classification. This decision, together with other aspects surrounding the presentation of advisories was motivated by the need for experimental control. The following rules apply:

- Resolution advisories must be given early in each scenario to: (1) reduce the chance that controllers resolve the designed conflict before the resolution advisory appears, and (2) reduce the risk that controllers heavily influence the scenario (through interaction with noise aircraft) such that the scripted resolution advisory has become invalid by the time it appears.

- A resolution advisory must be valid long enough for controllers to identify and investigate the conflict.
- The timing of a resolution advisory must be kept constant within each scenario group. If not, individual timings for each controller and scenario could be a confound in the experiment.
- Resolution advisories must provide a valid conflict-free solution. This means that resolution advisories cannot suggest resolution commands in a yellow or red no-go area in the SSD. Even though such resolutions could solve the designed conflict, they would cause a secondary conflict with a noise aircraft. This would be potentially detrimental for controller trust in the automation, and could provoke a "reject all" strategy.
- Resolution advisories must appear when both aircraft are in the sector. This is important, as controllers should be able to explore alternative solutions by inspecting the SSD of both involved aircraft. In addition, they should be able to solve the conflict by interacting with any of the involved aircraft.

One risk is that the timing for advisories can be perceived as disturbing or otherwise annoying by controllers and negatively affect the acceptance and agreement with advisories. This problem is generally known as poor automation "etiquette." Here, etiquette refers to the degree to which automation adheres to human social behavioral "rules" for communication and interaction. Moreover, it should be noted that the use of a generalized resolution advisory by default is nonconformal and may constitute a confound. Therefore, analyzes were carried out as to determine to what extent controllers' solutions for the same designed conflict varied depending on when it was solved (see Chapter 6).

Note that several controllers' problem-solving styles consisted of interacting with both aircraft to solve the conflict. The simulator, however, did not allow the presentation of a dual interaction as a resolution advisory. As such, only the first interaction of a controller's problem-solving style was considered when determining a conformal solution. Finally, in the First empirical study (Chapter 3), controllers' general problem-solving styles were not defined. In this study, the conformal and nonconformal resolution advisories were based on solutions specific to each scenario replicate. For the subsequent empirical studies (Chapters 4 and 5), however, the procedure with problem-solving style was implemented in order to ascertain that a conformal resolution advisory represented in fact represented a consistent problem-solving preference, and not a randomly selected solution.

B

Solution Space Diagram

B-1 Background

This appendix provides the theoretical foundation of the *Solution Space Diagram* (SSD), and the development of the baseline SSD as used in this thesis. The ATC SSD visualizes the safe "go" areas and contained "no-go" areas within a horizontal representation of an aircraft's control space in terms of heading and speed. The shape and size of these areas are defined by constraints affecting the aircraft's flight path, such as the aircraft's own flight-technical performance and environmental obstructions representing other aircraft, terrain, and weather.

The inspiration for the SSD, as a representation of an aircraft's safe fields of travel, was inspired by Gibson's work in the context of automobile driving.[291] These principles were further cultivated into guidelines for interface design driven by the Cognitive Systems Engineering and Ecological Interface Design (EID) paradigms.[292] A first prototype, incorporating the SSD with an enhanced navigation display, was used by Stijn van Dam to support pilots in self-separation activities in a futuristic free-flight environment (Figure B-1(a)).[232] A first version for ATC was used by Mercado Velasco to support controllers in CD&R, and potentially reduce their workload (Figure B-1(b)).[102]

B-2 Constructing the solution space

The SSD representation utilizes the (relative) velocity plane of aircraft to provide an easily observable overview of how the current position, track, and speed of aircraft interact. A basic understanding of the SSD, and the information visualized, is best conveyed by considering the diagram in its simplest form. Consider the situation in Figure B-2 that illustrates a conflict between two aircraft on straight paths.

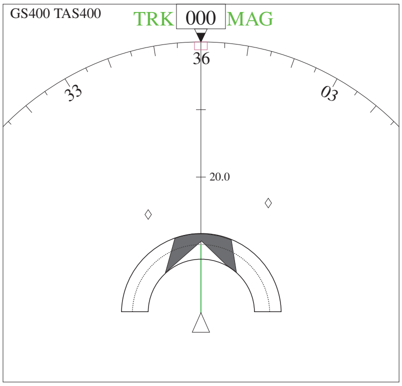

(a) SSD prototype with enhanced navigation display

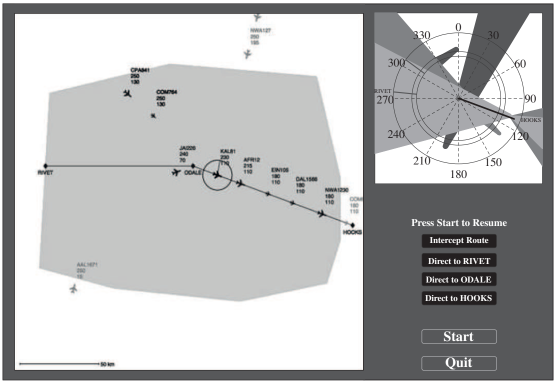

(b) SSD prototype with ATC radar display

FIGURE B-1: SSD a) integrated with a navigation display (flight deck),[293] and b) as an extension to a radar display (ATC work station).[102]

In Figure B-2(a) shows the conflict geometry in relative space. Aircraft *A* has been selected and is therefore designated the *controlled* aircraft. Aircraft *B* is considered an *intruder* in conflict with the controlled aircraft. The circle around aircraft *B* is the separation minimum, or *protected zone*, which represents the threshold distance required for safe separation (typically 5 nmi in en-route airspace). The grey area, or *conflict zone* illustrates how the SSD visualization is calculated by processing velocity plane information of both aircraft (i.e., V_1, V_2). The aircraft are in conflict. which is evident from the relative velocity vector ($V_{relative}$) of the controlled aircraft *A* being within the triangle formed by the tangent lines of the protected zone surrounding the intruder aircraft *B*. The *conflict zone* area comprises all relative velocity vectors $V_{relative}$ that result in a loss of separation.

In Figure B-2(b), the conflict geometry is considered in the absolute plane for the purpose of translating the aircraft's relative conflict representation (based on $V_{relative}$) to the aircraft's actual (i.e., absolute) velocity vector (V_2). To avoid conflict, the velocity vector of the controlled aircraft *A* V_2 must be outside of the triangular area. The three circles around the controlled aircraft *A* represents the triangle SSD. The dashed circle, intersecting with the velocity vector V_2, represents the current speed of aircraft *A*. The inner and outer circles represent the aircraft's maneuvering envelope in the horizontal plane (heading and speed). The visualized maneuver envelope provides an overview of all possible heading and speed options available to solve the conflict.

The SSD presentation is derived from visualizing the conflict zone within the maneuvering envelope of the controlled aircraft *A*. Figure B-2(c) shows the result-

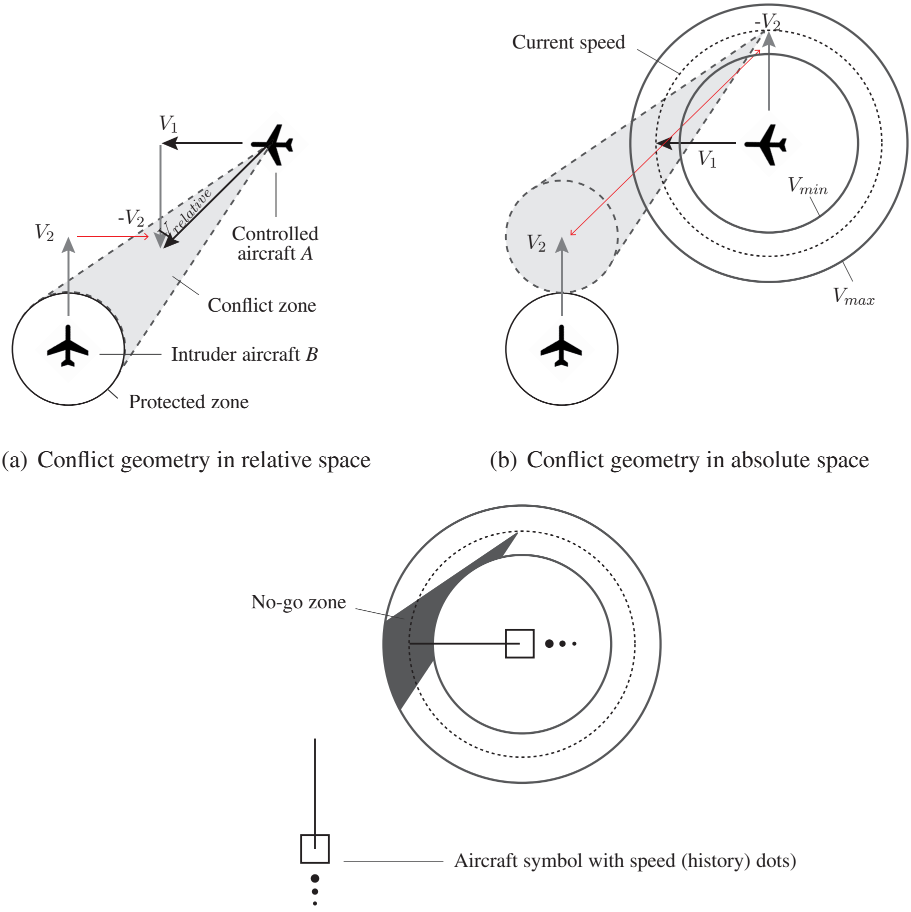

(a) Conflict geometry in relative space

(b) Conflict geometry in absolute space

(c) SSD for controlled aircraft

FIGURE B-2: The SSD representation in a conflict between two aircraft.

ing abstract SSD visualization, as an example of how it can be integrated with an ATC radar display. The partly shown conflict zone of aircraft B, referred to as the *no-go zone*, is visualized within the maximum V_{max} and minimum V_{min} speeds of aircraft *A's* maneuver envelope. The no-go zones provide the boundaries for safe travel. The size and position of the no-go zone reflect the relative position, velocity, and proximity of aircraft B. The angle between the legs of the triangle provides information on the relative proximity of aircraft B. The direction to which the triangle is convex (i.e., expands in width) represents the direction from which aircraft B is approaching. The color coding indicates time to separation loss, with the red color indicating less time than the orange color.

B-3 Using the SSD

For a given aircraft, the SSD provides the conflict areas, and conflict-free areas, imposed by surrounding traffic. Provided all intruder aircraft maintain their speed and heading, any velocity vector (of the controlled aircraft) placed outside the *no-go areas* implies safe passage. For example, in Figure B-3, all three velocity vectors point to *go areas* that ascertain conflict-free flight paths. In real life, however, controllers' are often very selective of the velocity vectors given. Several parameters are known to affect this decision, such as the deviation required from the intended flight path, rules of the air, and environmental factors (e.g., wind direction and strength).[29] Moreover, internal (subjective) parameters are known to influence conflict solutions, such as controllers' training, experience, strategies and preferences, and physical and psychological states.[25, 29, 80]

B-4 Towards higher degrees of automation

A number of broadly accepted taxonomies of human-automation system control have categorized levels of automation (LOA), from fully manual to fully automatic. Sheridan and Verplank's automation model proposed in 1978 is generally considered the first LOA taxonomy.[44] Parasuraman *et al.* proposed a LOA continuum, independently crossed with four stages of human information processing they categorized as: information acquisition: information analysis, decision selection, and

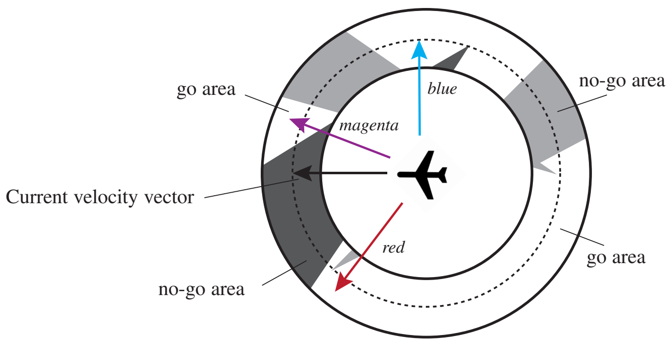

FIGURE B-3: The current position of the velocity vector indicate a conflict with another aircraft. Any velocity vector (magenta, blue, or red arrow) would ascertain safe separation, all pointing to conflict-free ares (white space).

action implementation.[43] A high level of integration automation (Stage 2) is qualitatively different from a high level of action automation (Stage 4). The levels and stages of automation are, however, intertwined in that a higher automation level at one stage may depend on a sufficiently high automation level at the preceding stage. For example, a conflict resolution system providing recommendations for solving a conflict (Stage 3) require a high level of both information acquisition (Stage 1) and information integration (Stage 2).

Figure B-4 provides and overview of the SSD's capabilities in a matrix of levels and stages of automation. While the levels of automation (LOA) axis refer to the degree to which a task is automated, the stages of automation axis considers these levels in association to human information processing functions.

The accuracy of the SSD, and hence its value, depend on a high level of information acquisition (Stage 1). Moreover, the SSD itself integrates critical data about aircraft, including their position, velocity vector information, and performance constraints (Stage 2). In turn, the integrated information can be used by an automated decision aid to provide resolution advisories (Stage 3), and even implement them as needed (Stage 4). The SSD also presents information in a diagram where triangular conflict zones provide the emergent perceptual features supporting human decision-making and conflict-solving.

The SSD can be considered representative of a high level Stage 2 automation, provided CD&R is the primary task for which it is used. The spatial and geometric representation of the CD&R problem provides a foundation against which the SSD can be incorporated with a decision aid offering explicit support, reflecting a high level of Stage 3 and Stage 4 automation capabilities. Resolution advisories can easily be depicted inside the SSD visualizations. For example, by providing recommendations on suitable (conflict-free) areas in which a solution can be found, offering several solution alternative, or suggesting a specific "best" solution (Stage 3). The decision aid's authority to execute an advisory can be varied, although this final stage of action implementation does not require the SSD visualization.

With *management by consent*, the automation selects a solution to the conflict and performs it only if the controller approves. At the next level, *management by exception*, the automation selects a solution and performs that solution, unless the controller specifically intervenes to cancel or veto it. At the highest level for Stage 4, in Figure B-4, the decision aid has full authority to both choose and perform a given solution, although it still informs the controller of its activities (as such, not the most extreme LOA).

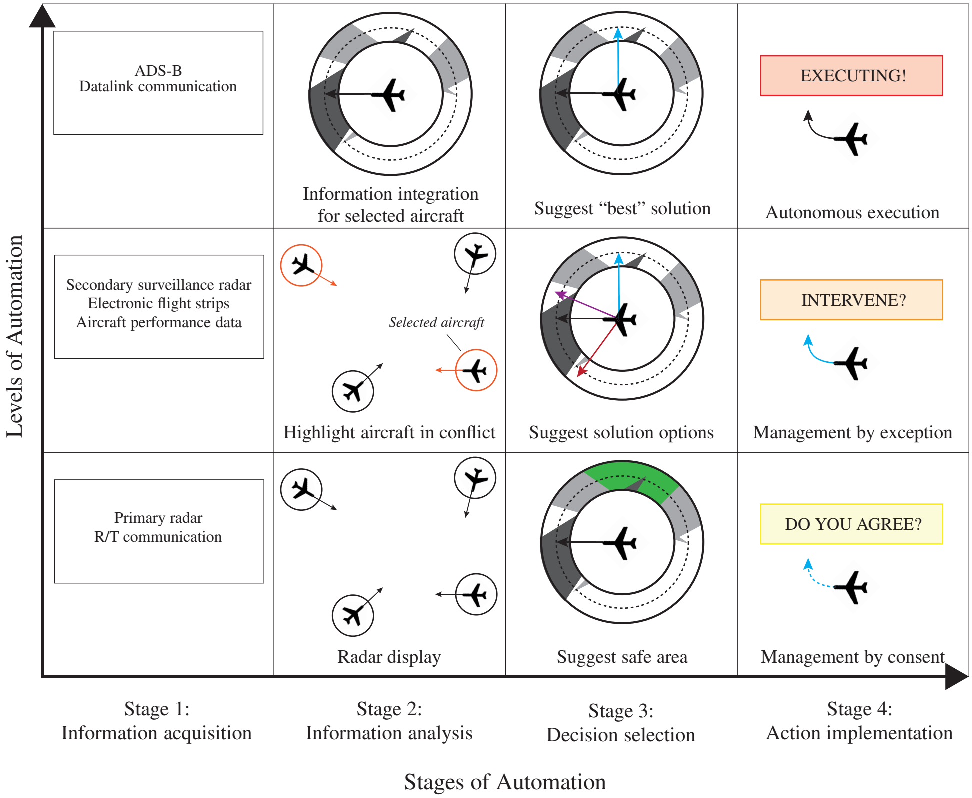

FIGURE B-4: The placement of the SSD in a matrix of levels and stages of automation.

B-5 Modified SSD

Mercado Velasco's ATC SSD prototype[102] was modified for purposes of using it in empirical simulations in this thesis. Modifications were based on feedback following visits to the Irish Aviation Authority control centers in Dublin and Shannon (Ireland), encompassing observations and interviews. A simplified SSD was developed in consultation with controllers working at these centers, better tailored to their specific workplace and needs. Among several comments, these were the most valuable:

1. Integrate the SSD with the radar plan view display as controllers do not like to switch between two separate displays.

2. Highlight all aircraft in conflict to support controllers in directing their attention.

3. Incorporate a conflict zone (i.e., no-go areas) look-ahead time of eight minutes as controllers tend to resolve conflicts en-route on a short time frame.

4. Simplify the visual representation of the SSD, in particular the triangle-shaped no-go areas, to reduce information complexity and overload.

5. Visualize how critical no-go areas are (i.e., time-to-conflict) through color coding.

The integration of the SSD and radar display support controllers in relating no-go areas for aircraft to their spatial positions (first comment). In addressing the second comment, simple state-based predictions were used to determine aircraft in conflict and highlight them on the radar display. This modification increased the SSD's level of information analysis (Stage 2). The third comment was addressed by introducing a time-horizon that affected the triangle-shaped visualization. By defining a finite time-horizon (set to five minutes), the tip of the triangle is rounded off. The smaller the look-ahead time, the more the triangle tip is cut off. Prior to this modification, when there was no limit to the time-horizon, pointing the velocity vector to the tip of the triangle would postpone a conflict indefinitely ($T=\infty$) as the controlled aircraft would end up flying parallel to the other aircraft.

The change is reflected in Figure B-5. In Figure B-5(b) the velocity vector is outside the constrained no-go area, indicating that separation is maintained for the time being. However, caution is needed since the current situation only ascertains that a conflict will not occur within the next eight minutes. As time progress, however, the current track will result in a conflict provided both aircraft leave their velocity vectors unchanged.

The fifth and final comment was addressed by limiting the maneuver envelope to the aircraft's current speed. Figure B-6 shows (a) the full SSD, which shows the

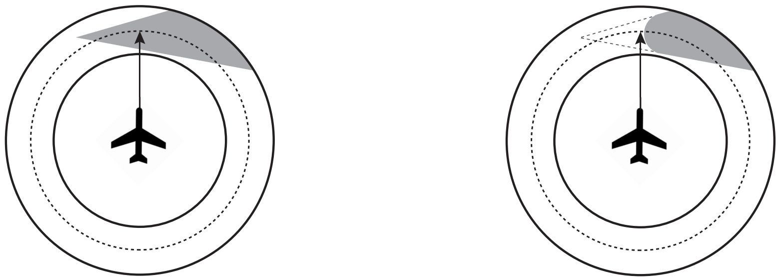

(a) SSD without time-horizon (b) SSD with time-horizon

FIGURE B-5: The effect of time look-ahead on SSD no-go area.

(a) Triangle SSD
with entire maneuver envelope

(b) Heading band SSD
that only shows current speed

FIGURE B-6: Simplification of the SSD for en-route ATC.

aircraft's entire maneuver envelope, and (b) the simplified heading band SSD, which only depicts heading options for the current speed. This representation was believed to better fit en-route CD&R. Heading commands are often preferred at high (flight) levels (as opposed to speed, and unless vertical options are available), because aircraft often operate close to their maximum performance envelope and therefore have a narrow control space for speed. To support controllers' understanding of time, no-go areas were color coded to better reflect the proximity of other aircraft and time available before conflict (fourth comment). Red areas were defined to reflect that separation loss would occur in less than one minute. Yellow areas (sometimes perceived as orange) reflected that separation loss would occur between one and four minutes ahead in time.

C

Simulation Briefing Packages

C-1 Prequel briefing

C-1-1 Introduction

Current economic and technological realities are driving a fundamental and global modernization of the ground systems used for ATC. These modernization programs are working toward a ground-up redesign of the tools, systems and methods used. Important changes include high-bandwidth digital datalinks, better data quality, and most importantly the introduction of advanced automated systems to support air traffic controllers in their separation tasks. In this experiment you will work with a new support tool to help you separate aircraft: the Solution Space Diagram (SSD). Details on the SSD are provided below. The simulation will be realistic in some respects, however, we have simplified several aspects of the ATC task. You can imagine yourself controlling a future airspace sector using new support tools that rely on digital datalinks. As such, radio communication with the aircraft is not necessary anymore. Another significant difference from current work practices is the omittance of current distance measuring equipment used for measuring separation distances between aircraft. Instead, distance measurement is incorporated in the SSD tool. Further, by using the SSD tool separation will always be assured.

C-1-2 Scenarios

In this simulation session, you will see a number of short air traffic scenarios (see Figure C-1), each lasting approximately 2 minutes. We have made some simplifying assumptions. First, all aircraft are in level flight at the same altitude. During the simulation you are not able to change the flight levels of aircraft to resolve conflicts. Second, you will interact with aircraft using the mouse and keyboard, rather than

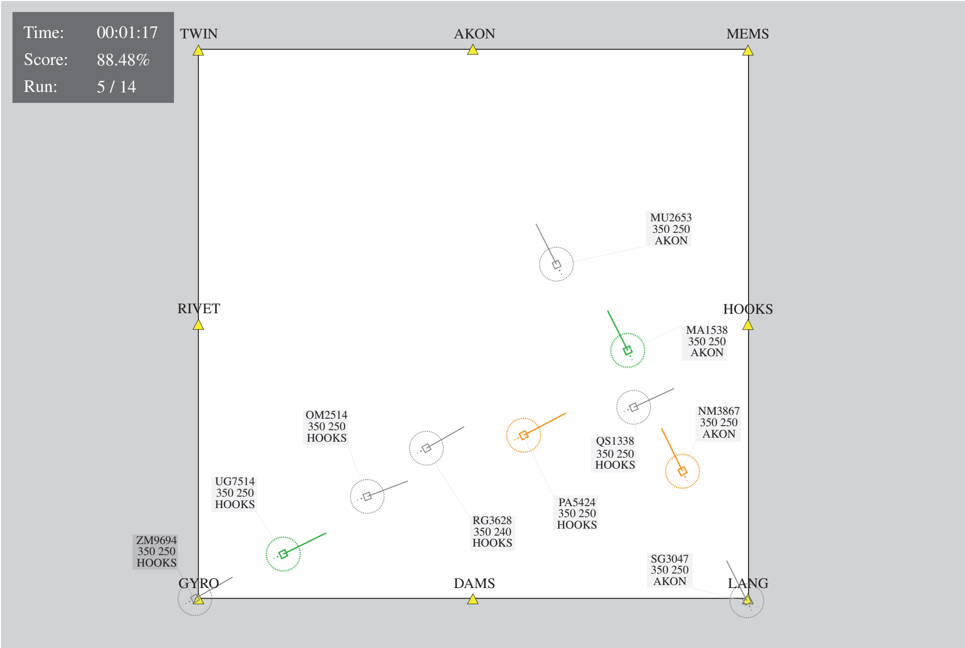

FIGURE C-1: MUFASA simulator screen.

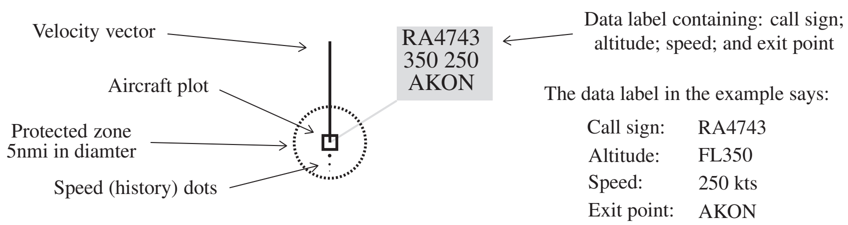

FIGURE C-2: Aircraft plot.

having to talk to each aircraft on the radio.

All scenarios display a hypothetical sector in a squared format, 80 x 80 nmi in size. The radar screen is updated in 1 second intervals. The simulator runs at 2x speed. This means that aircraft move 2 times faster than normal. This also means that the trend vector of the aircraft blip (i.e., the line in front of the aircraft blip) looks ahead 30 seconds instead of the normal 60 seconds. Aircraft are represented by squared symbols and an associated data label as described in Figure C-2.

C-1-3 Your tasks

Your task will be to play the role of an air traffic controller. As described below, your two main tasks will be: 1) to resolve conflicts between aircraft, and, 2) also to clear aircraft to their respective sector Exit Point. **Each task is considered to be equally important**. The Exit Points are shown as yellow triangles with four-letter names around the perimeter of the sector. The designated Exit Point of each aircraft is also shown the flight label of the aircraft.

Exit task. When you begin each session, you may notice that one or more aircraft are not heading toward their cleared Exit Point. The deviating aircraft are displayed in a *light grey color* (Figure C-1). The *green aircraft* are the ones that are already heading toward their designated Exit Point. When clicking near the plot of an aircraft, a ring appears around the aircraft (Figure C-3). This ring is a graphical representation of the SSD mentioned above. On this ring, a magenta line can be seen that indicates the direction toward the designated Exit Point.

If the current aircraft velocity vector is not aligned with the magenta line, the aircraft is not heading toward its intended Exit Point. In this case, you must vector the aircraft. Try to be as accurate as possible in this heading clearance. To activate the new vector, press the ENTER key on the keyboard. For a complete overview of the interaction with the SSD, see Table C-1.

C-1-4 Resolving conflicts

From time to time there may be conflicts between aircraft. A conflict occurs when aircraft are *predicted* to close within 5 nmi of one another. The protected zones of 5 nmi around aircraft are shown as a dotted circle (Figure C-2). When the system predicts a loss of separation, a warning will be provided. Given uncertainties in trajectory prediction, it could be that the system *seems* to alert“late.” A potential loss of separation is first visually signaled by both conflicting aircraft turning **amber**. This occurs roughly *60 seconds* before the loss of separation. Both aircraft turn **red** when a loss of separation is predicted within *30 seconds*, and is **accompanied by an aural alert**. A conflict involves at least two aircraft, but might sometimes involve more than two. Your task is to resolve such conflicts, by issuing a heading and / or speed clearances.

When (or before) an alert occurs, you may open (one at a time) the SSD for a given aircraft, to examine the nature of the conflict. In the simulator, the SSD ring shows three colors: grey, amber, and red (Figure C-4). **Grey zones** indicate a possible conflict **more than 60 seconds** ahead in time. **Amber zones** indicate the regions that will result in a loss of separation occurring between **30 to 60 seconds** ahead in time. **Red zones** indicate the regions that will result in a loss of separation

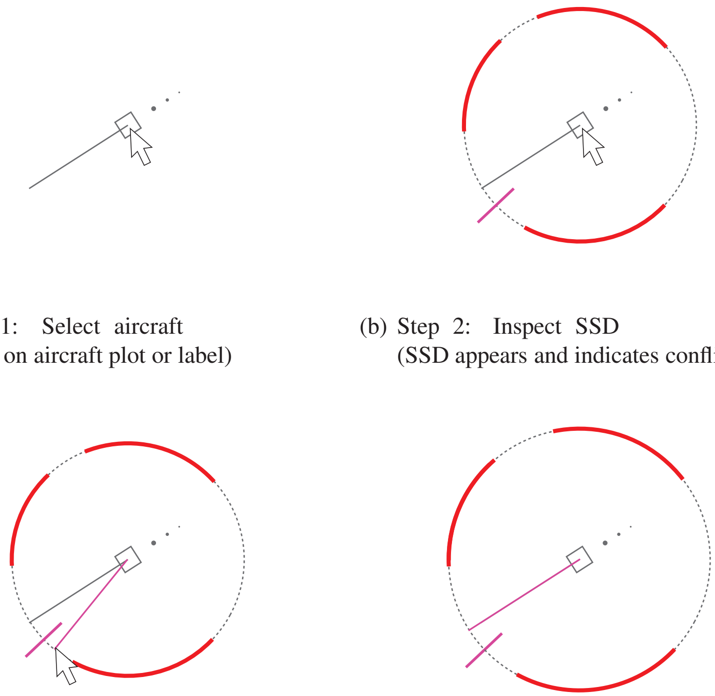

(a) Step 1: Select aircraft (click on aircraft plot or label)

(b) Step 2: Inspect SSD (SSD appears and indicates conflict areas)

(c) Step 3: Create new vector command (point and click within SSD)

(d) Step 4: Create new speed command (mouse scroll)

FIGURE C-3: Generating a new vector or velocity command with the SSD.

TABLE C-1: Overview of controller actions and their implementations in the simulator

Action	Implementation
Select aircraft	*Left Mouse Button (LMB) click near aircraft plot*
Generate new heading command	*LMB click and point on velocity vector*
Generate new velocity command	*Mouse scroll wheel*
Execute new velocity or vector command	*ENTER key*
Abort new velocity or vector command	*Click well outside SSD*

occurring between **now and 30 seconds** ahead in time. In essence, the grey, amber, and red zones indicate the "no go" areas for a resolution, whereby these zones are formed by the surrounding aircraft. So any "hole" in the ring represents a "go" area, or a potential solution, to resolve a conflict.

Conflicts can be resolved by either issuing a heading clearance, or by issuing a speed clearance, or a combined heading and speed clearance. Dragging the aircraft trajectory vector to a non-red or non-amber region will only change the heading of the aircraft and resolve the conflict. Notice also that by using the mouse scroll wheel the diameter of the ring will increase / decrease. The diameter of the ring indicates the speed setting of the aircraft. Notice that changing the speed also changes the location of the "no go" bands (because of the interaction between the required solution speed and heading). As mentioned before, you may solve a given conflict by either or both heading and speed, as you wish.

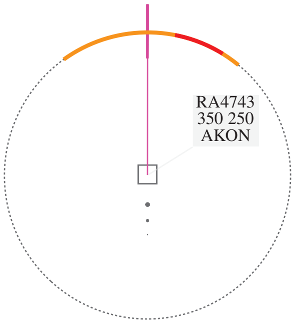

FIGURE C-4: The SSD ring showing that the aircraft is heading toward its Exit Point, but also that it is in a conflict and that a loss of separation will occur between 30 - 60 seconds ahead in time.

C-1-5 Performance score

In the upper left corner of the screen (Figure C-1) you will see your instantaneous performance score for the current session. This performance score indicates how accurately you are currently vectoring aircraft toward their respective exit points, and also how well you are doing at avoiding and resolving conflicts. The performance score is based on the average track angle errors of each aircraft (inside the sector) from their intended Exit Point. Further, each time a conflict occurs, a penalty (of 10%) is subtracted from the score for each involved aircraft (so 3 aircraft in conflict is worse (-30%) than just 2 aircraft (-20%) in conflict). Any speed change from the initial speed will affect the performance score. The larger deviation, the larger the penalty. The speed change penalty is permanent in that once a change is made it has a resulting effect on the performance score.

The theoretical maximum performance score is 100%, although this is essentially impossible. Nonetheless, you can and should try to keep your performance score as high as possible. At the end of each scenario, you will see displayed your average performance score for that scenario.

As a reference, a score of at least 80% should be possible to maintain for each scenario.

C-1-6 Workload/difficulty rating

At the end of each scenario, you are asked to give a difficulty rating for that scenario.

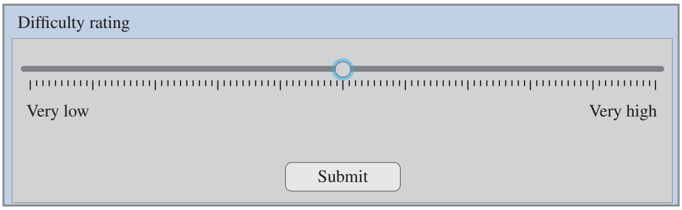

FIGURE C-5: Difficulty rating dialog.

C-2 Experiment briefing

C-2-1 Introduction

Welcome back to the second round of experiments in the MUFASA research project. Today you will continue to work with the same simulator that you encountered last week. In contrast to last week when system interaction was fully manual, you will in today's simulation work with an advanced conflict resolution advisory system. In general, the simulator is very similar to the previous one and you should therefore remember and recognize most of the environment and interactions.

C-2-2 Scenarios

Again, you will encounter a number of short air traffic scenarios (see Figure C-6), each lasting approximately 2 minutes.

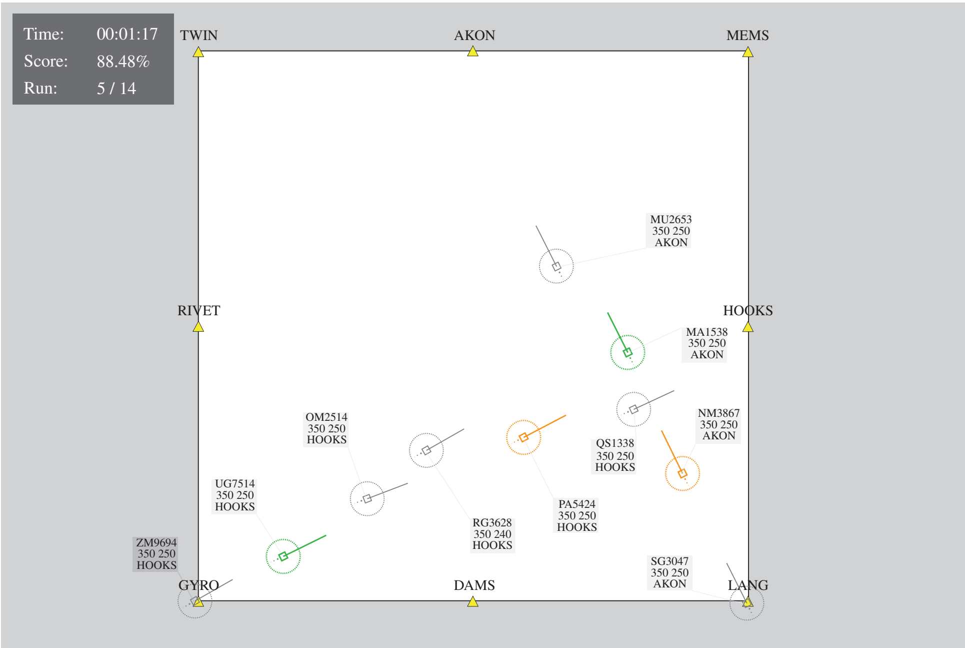

FIGURE C-6: MUFASA simulator screen.

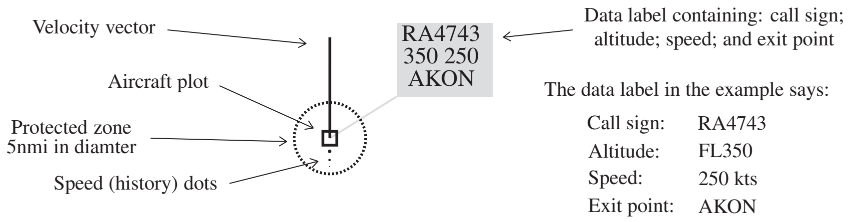

FIGURE C-7: Aircraft plot.

Rehearsal:.

- Sector size is 80 x 80 nmi;
- Screen updated at 1 second intervals;
- Simulation runs a 2x real speed;
- Vector line equals 30 seconds ahead in time;
- Amber zone means a loss of separation will occur between in less than 1 minute;
- Red zone means a loss of separation will occur in less than 30 seconds; and
- A go-zone (no color) is always safe.

C-2-3 Your tasks

Like in the previous simulation your two main tasks are: 1) to resolve conflicts between aircraft, and 2) to clear aircraft to their respective sector Exit Point. **Each task is considered to be equally important.**

C-2-4 Aural conflict warning alert

From the previous simulation you will remember the conflict warning alert. A conflict is defined as a loss of separation, and exists when aircraft are *predicted* to close within 5 nmi of one another. The conflict warning alert is triggered when the **loss of separation occurs within 30 seconds**. Simultaneously the plot and trajectory vector for each of the involved aircraft will **turn red** and you will hear an **aural alert** ("whoop").

C-2-5 Automation advisory

A new feature in this simulation is a midterm strategic aid - the automated resolution **advisory solution system (ASS)**. The system will help you to identify and proactively avoid conflicts by suggesting vector commands. The resolution advisory will always precede a conflict warning alert. However, the conflict warning alert is prioritized over the automation advisory. As such there may be instances where no advisory is provided before the conflict warning is triggered. This occurs, for example, in instances where the time to a loss of conflict is too short to provide an advisory.

When a conflict is detected by the automation advisory logic the following occurs (Figure C-8):

- Advisory sound consisting of 4 chimes: "beep, beep, beep, beep"
- Resolution advisory SSD display appears for aircraft to which the resolution advisory applies
- Advisory window appears in upper right corner prompting you to accept or reject the resolution advisory.

When a resolution advisory is shown all aircraft involved in the conflict will turn amber. You will not be able to interact with the aircraft involved in the advisory. You will, however, be able to interact with all other aircraft. First after the resolution advisory disappears (either by clicking 'accept' or 'reject') you will be able to interact with the aircraft involved in the advisory again.

Resolution advisory SSD display. The resolution advisories consist of a change in heading, speed, or a combination of both. When a resolution advisory is given, the solution diagram of the affected aircraft is automatically shown. The resolution advisory is indicated by an **orange vector line**, an **orange ring** (with a diameter greater or less than the green ring), or a **combination of both**. In the example provide in Figure C-9, the resolution advisory is to increase the speed of RG3628 and to turn him slightly left.

Advisory dialog. In the upper right corner of the display an advisory dialog appear with an Advisory agreement rating scale (1-100 scale from disagree to agree), two buttons with 'accept' and 'reject', and a countdown timer set to 30 seconds. If you do not accept or reject, the advisory dialog will automatically close after 30 seconds. The advisory dialog allows you to either **"Accept"** of **"Reject"** the resolution advisory.

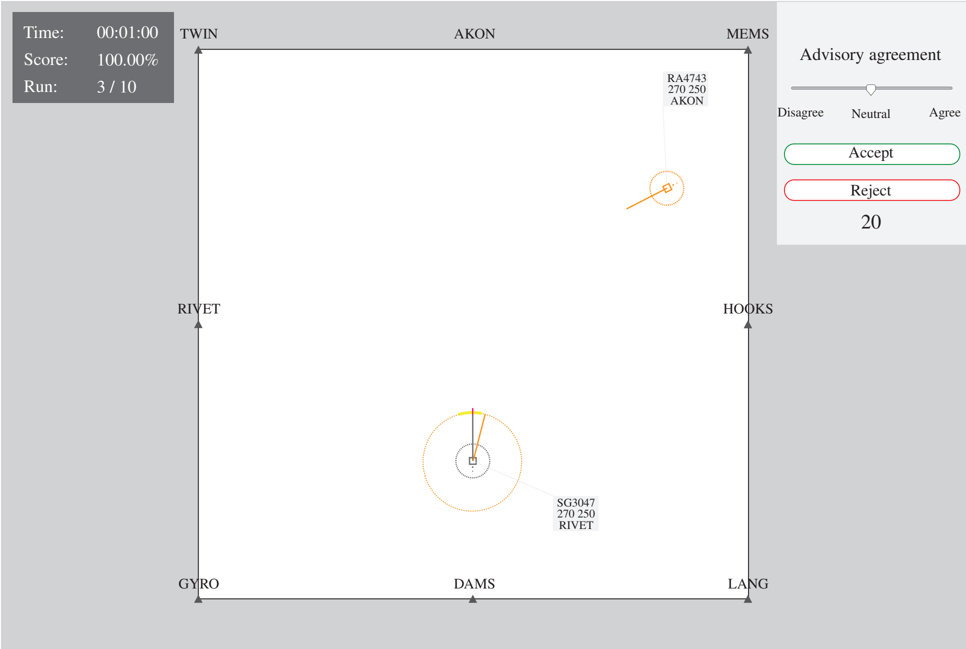

FIGURE C-8: MUFASA simulator screen with automation advisory.

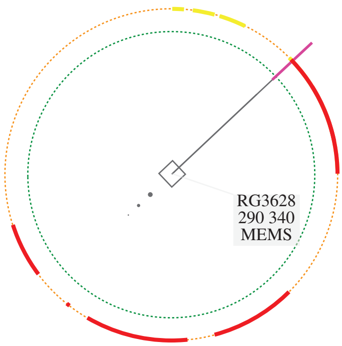

FIGURE C-9: Resolution advisory SSD.

Before accepting or rejecting the advisory we want you to consider how you feel about the advisory. You indicate to what extent you agree with the advisory, ranging from disagree to agree, by clicking somewhere on the associated scale. In order to be more precise when giving a rating you should click on the round selector and drag it to the desired value (if trying to click on a higher or lower value the round selector will jump in intervals of 10). Note that you have to rate your agreement with the advisory before being able to accept or reject the advisory.

When the automated advisory appears you should **carefully inspect the advisory** before choosing to accept or reject it. The automation does not necessarily advise the most optimal resolution. You may want to solve the conflict otherwise.

FIGURE C-10: Advisory dialog.

C-2-6 Workload/difficulty rating

Like last week, you are asked to give a difficulty rating for that scenario. In order to be more precise when giving a rating you should click on the round selector and drag it to the desired value (if trying to click on a higher or lower value the round selector will jump in intervals of 10).

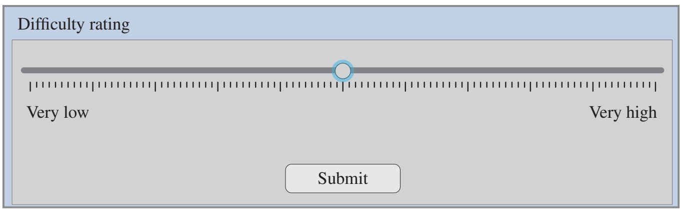

FIGURE C-11: Difficulty rating dialog.

C-2-7 Source bias study

The effect of two different advisory sources on the acceptance of a decision-aiding system is investigated. The two sources are human and automation. The order of the source condition is varied between participants. Group 1 start with a human adviser, while group 2 starts with automation. Both the human and automation source condition measurement sessions consist of 10 training scenarios and five measurement scenarios. During each session, all resolution advisories are either human or automation generated. The source of resolution advisories are not mixed within the measurement sessions. Before each session, participants are provided with instructions about the source behind the resolution advisories. The human is presented as an air traffic controller. The following instructions are given:

Air traffic controller. *"All resolution advisories suggested in this session are made by an air traffic controller."*

Automation. *"All resolution advisories suggested in this session are generated by automation."*

C-2-8 Automation transparency study

Participants are divided into separate groups. One group starts using the heading band representation. The other group starts using the triangle representation. Both groups interact with both representations.

Note that the no-go zones are shown for the current state of the aircraft. It is a presentation updated in real time. That means, that if any aircraft is in transition, e.g., increasing or decreasing speed, or turning left or right, the no-go zone will be updated for every second the simulator is updated. As such, if you execute a command, you have to consider the time it takes to complete the command and how it affects the no-go zone.

Heading band representation (baseline). The heading band representation is the standard representation. As you can see, this is the interface that you have used so far. The ring indicates conflict zones and the solution spaces available for the aircraft at the current speed of the aircraft. You can investigate how the conflict zones and solution spaces change by scrolling through various speeds using the mouse scroll wheel.

Triangle representation. In the triangle representation, the entire speed envelope of the aircraft is visible. The inner diameter represents the lower boundary of the speed envelope. The outer diameter represents the upper boundary of the speed

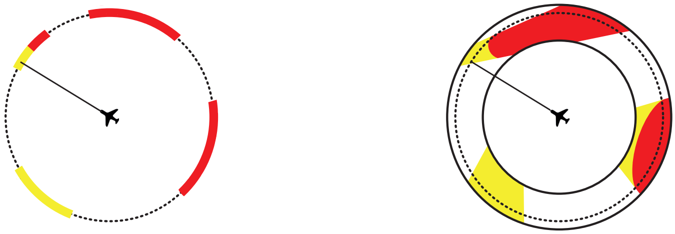

(a) Heading band representation (b) Full solution space representationD

FIGURE C-12: The two SSD representations.

envelope. In this view, the conflict areas are shown for the entire speed envelope of the aircraft.

D

Literature review of human-automation trust

Table D-1: Antecedents of human-automation trust

Antecedent of trust	Finding
Trust before use	People often indicate higher *liking* in human sources than in automated sources.[49,177]
	People often indicate higher trust, higher reliability, greater expertise, higher authority, and higher performance expectations of automated sources as opposed to human sources (e.g., automation bias and perfect automation schema).[63,165,166]
	Much trust in external advice may be attributed to a heuristic of relying in order to reduce mental effort.[63,173]
	Trust in humans and automated sources varies with pedigree: Operators may trust novice automation more than novice human, while trust in expert humans is greater than trust in expert automation.[167]
Trust during use	Determinants for intentions to use often do not translate well to determinants of actual use.[167,174]
	Trust in automated aids is primarily influenced by its performance.[171]
	Trust in human aids continues to be influenced by the dispositional attitudes such as knowledge, effort and expertise.[171]
	People's confidence in a source does not determine agreement with advisories from that source.[171,185]
	Automated sources are judged more negatively if they fail than are human sources,[49,50,167,174] especially if the task appears easy to the operator.[294]
	People are more forgiving of errors made by a human source than errors made by an automated source.[49,165]
	Reliance on an automated aid is not affected by advisory timing (i.e., before or after the operator has made a decision).[50]
	Experts are more likely to critically consider arguments, while novices are more likely to rely on external clues when given advice.[173]
Influence of perceptions	The expertise of an automated aid is mainly associated with knowledge, while expertise of a human aid more broadly is associated with knowledge, experience, and education.[171]
	Advice from an automated source perceived more rational and objective than the same advice from a human source.[172]
	Trust in other humans may be biased by a worry of participants to admit distrust. This does not apply to automaton.[185]
	Human sources are generally perceived as imperfect and have less credibility, which negatively affects trust.[55,161]
	Motivation to perform well is critical to appropriate trust.[173]
Framing effects	Anthropomorphic features (e.g., automation's appearance and behavior) can increase trust and reliance.[48,114,118,176]
	Reliance is driven more by how credible a source is framed, than by consideration of the actual problem at hand.[173,175]
	Attitudes of trust, confidence, and liking are sensitive to framing.[48,110,171]

E

Solution distributions

E-1 Study 1 solution distributions

E-1-1 Solution parameter hierarchy, Study 1

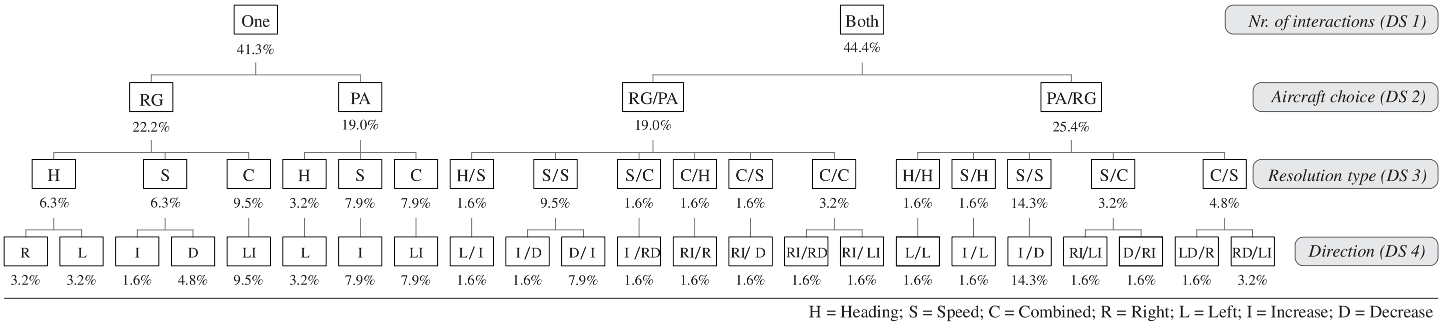

FIGURE E-1: Solution distribution for the Solution parameter hierarchy in proportion (%) of total solutions, Scenario 2.

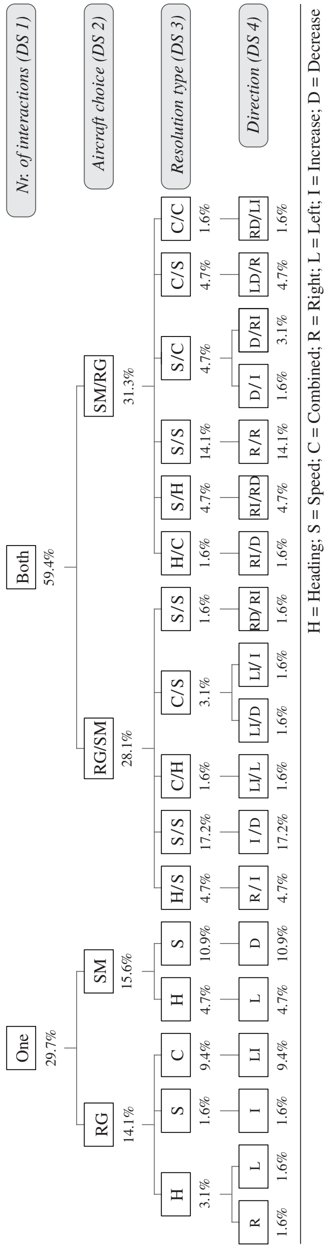

FIGURE E-2: Solution distribution for the Solution parameter hierarchy in proportion (%) of total solutions, Scenario 3.

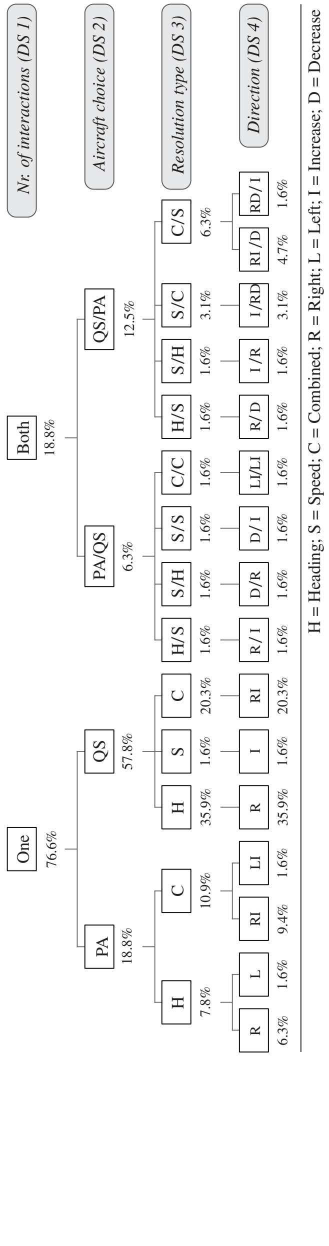

FIGURE E-3: Solution distribution for the Solution parameter hierarchy in proportion (%) of total solutions, Scenario 4.

E-1-2 Control problem classification, Study 1

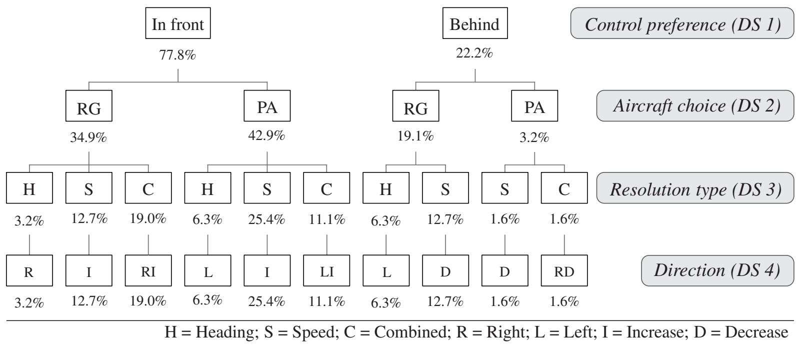

FIGURE E-4: Solution distribution for the Control problem classification in proportion (%) of total solutions, Scenario 2.

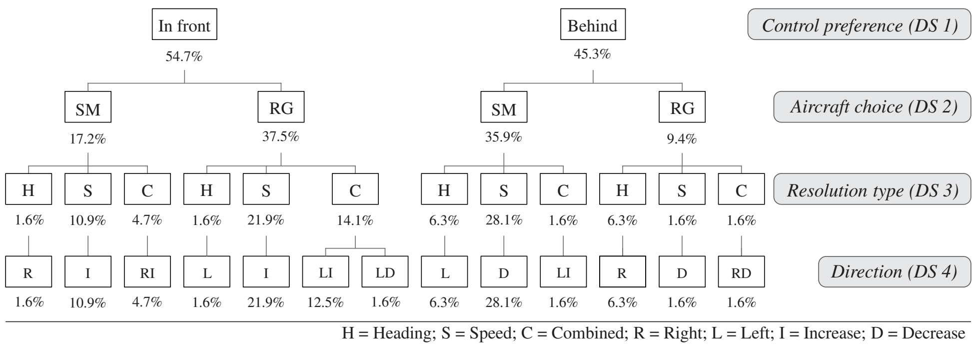

FIGURE E-5: Solution distribution for the Control problem classification in proportion (%) of total solutions, Scenario 3.

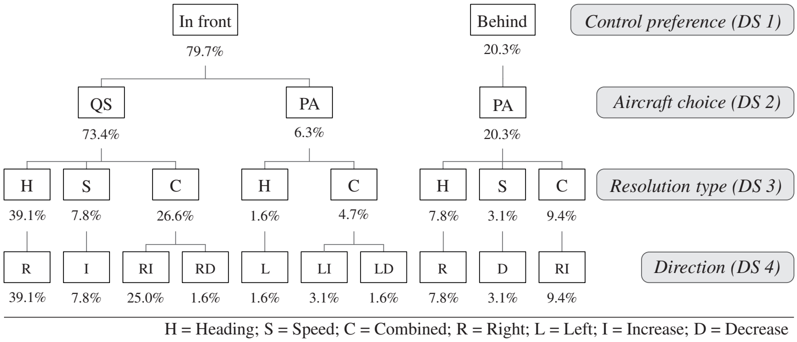

FIGURE E-6: Solution distribution for the Control problem classification in proportion (%) of total solutions, Scenario 4.

E-1-3 Solution geometry, Study 1

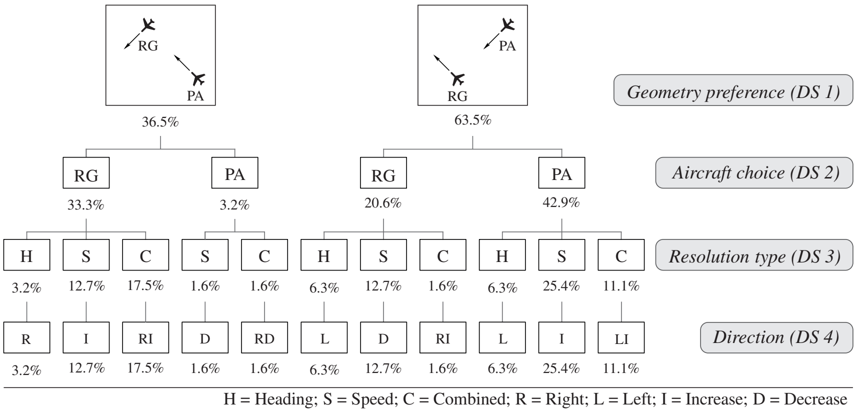

FIGURE E-7: Solution distribution for the Solution geometry classification in proportion (%) of total solutions, Scenario 2.

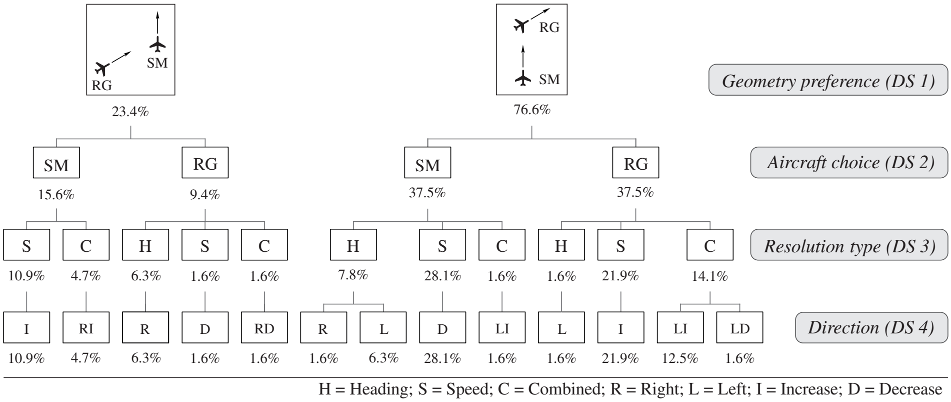

FIGURE E-8: Solution distribution for the Solution geometry classification in proportion (%) of total solutions, Scenario 3.

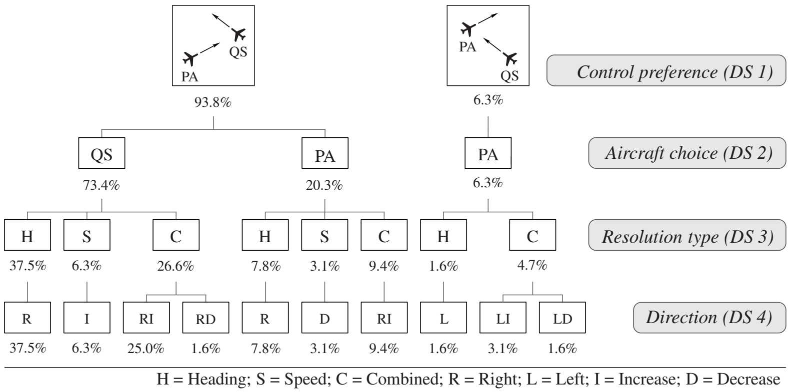

FIGURE E-9: Solution distribution for the Solution geometry classification in proportion (%) of total solutions, Scenario 4.

References

[1] Majumdar, A. and Ochieng, W. Y., "Factors affecting air traffic controller workload: Multivariate analysis based on simulation modeling of controller workload," *Transportation Research Record: Journal of the Transportation Research Board*, Vol. 1788, 2002, pp. 58–69.

[2] Majumdar, A., Ochieng, W. Y., Bentham, J., and Richards, M., "En-route sector capacity estimation methodologies: An international survey," *Journal of Air Transport Management*, Vol. 11, No. 6, 2005, pp. 375–387.

[3] SESAR, "European ATM master plan. The roadmap for sustainable air traffic management. Edition 2," Rep., EUROCONTROL, Brussels, Belgium, Oct. 2012.

[4] JPDO, "Concept of operations for the Next Generation Air Transportation System. Version 3.2," Rep., Joint Planning and Development Office, Washington, DC, 2011.

[5] Hilburn, B., "Cognitive complexity in air traffic control: A literature review," Eec note no. 04/04, EEC, Brétigny-sur-Orge, France, Mar. 2004.

[6] JPDO, "NextGen avionics roadmap Version 2.0," Rep., JPDO, Washington, DC, Sep. 2011.

[7] SESAR, "The ATM target concept: Deliverable D3," Rep. dlm-0612-001-02-00, EUROCONTROL, Brussels, Belgium, Sep. 2007.

[8] Bekier, M., Molesworth, B. R. C., and Williamson, A. M., "Tipping point: The narrow path between automation acceptance and rejection in air traffic management," *Safety Sci.*, Vol. 50, No. 2, 2012, pp. 259–265.

[9] Bolic, T. S., *Automation adoption and adaptation in air traffic control*, Ph.d. thesis, Dept. Civ. Envir. Eng., Univ. California, Berkeley, CA, 2006.

[10] Ehrmanntraut, R., *Full automation of air traffic management in high complexity airspace*, Ph.d. thesis, Dept. Transp. Traf. Sci. "Friedrich List", Univ. Dresden, Dresden, Germany, 2010.

[11] Kuchar, J. K. and Yang, L. C., "Survey of conflict detection and resolution modelling methods," *Proc. AIAA GNF Conf.*, New Orleans, LA, Aug. 11-13 1997, pp. 1388–1397.

[12] Kauppinen, S., Brain, C., and Moore, M., "European medium-term conflict detection field trials," *Proc. 21st DASC*, Vol. 1, IEEE, Irvine, CA, Oct. 27-31 2002, pp. 2C1–1–2C1–12.

[13] Parasuraman, R. and Riley, V., "Humans and automation: Use, misuse, disuse, abuse," *Hum. Factors*, Vol. 39, No. 2, 1997, pp. 230–253.

[14] Bekier, M., *Automation acceptance in air traffic management*, Ph.d. thesis, Dept. Aviat., Univ. New South Wales, Sydney, Australia, 2013.

[15] Lee, J. D. and See, K. A., "Trust in automation: Designing for appropriate reliance," *Int. J. Hum-Comput. St.*, Vol. 46, No. 1, 2004, pp. 50–80.

[16] Riley, V., "Operator reliance on automation: Theory and data," *Automation and human performance: Theory and applications*, edited by R. Parasuraman and M. Mouloua, Erlbaum, Hillsdale, NJ, 1996, pp. 19–35.

[17] Prevot, T., Homola, J. R., Martin, L. H., Mercer, J. S., and Cabrall, C. D., "Toward automated air traffic control: Investigating a fundamental paradigm shift in human/systems interaction," *Int. J. Hum-Comput. Int.*, Vol. 28, No. 2, 2012, pp. 77–98.

[18] Hilburn, B. and Flynn, M., "Air traffic controller and management attitudes toward automation: An empirical investigation," *4th USA/Europe ATM R&D Seminar*, Santa Fe, NM, Dec. 3-7. 2001.

[19] Kuchar, J. K. and Yang, L. C., "A review of conflict detection and resolution modeling methods," *IEEE Trans. Intell. Transp. Sys.*, Vol. 1, No. 4, 2000, pp. 179–189.

[20] Gigerenzer, G., Hertwig, R., and Pachur, T., *Heuristics: The foundations of adaptive behavior*, Oxford University Press, New York, NY, 2011.

[21] Simon, H. A., "Rational choice and the structure of the environment," *Psychol. Rev.*, Vol. 63, No. 2, 1956, pp. 129–138.

[22] Tversky, A. and Kahneman, D., "Judgement under uncertainty: Heuristics and biases," *Science*, Vol. 185, No. 4157, sep 1974, pp. 1124–1131.

[23] Averty, P., "Conflict perception by ATCS admits doubt but not inconsistency," *6th USA/Europe ATM R&D Seminar*, Baltimore, MD, Jun. 27-30 2005.

[24] Fothergill, S. and Neal, A., "Conflict-resolution heuristics for en route air traffic management," *Proc. HFES 57th Annual Meeting*, Vol. 57, No.1, San Diego, CA, Sep. 30-Oct. 4 2013, pp. 71–75.

[25] Histon, J., Li, L., and Hansman, R. J., "Airspace structure, future ATC systems, and controller complexity reduction," *Proc. 29th DASC*, IEEE, Salt Lake Cuty, UT, Oct. 3-7. 2010, pp. 4.A.4–1–4.A.4–14.

[26] Nunes, A. and Mogford, R. H., "Identifying controller strategies that support the 'Picture'," *Proc. HFES 47th Annual Meeting*, Vol. 47 of *1*, Denver, CO, Oct. 13-17 2003, pp. 71–75.

[27] Sridhar, B., Sheth, K. S., and Grabbe, S., "Airspace complexity and its application in air traffic management," *2nd USA/Europe ATM R&D Seminar*, Orlando, FL, Dec. 1-4 1998.

[28] Xing, J., "Measures of information complexity and the implications for automation design," Rep. dot/faa/am-04/17, FAA Civ. Aerospace Med. Inst., Oklahoma City, OK, Oct. 2004.

[29] Kirwan, B. and Flynn, M., "Towards a cognitive tool for conflict resolution assistance: A literature review," Rep. asa.01.cora.2.del04-a.lit, EUROCONTROL, Brussels, Belgium, Mar. 2002.

[30] Averty, P., Guittet, K., and Lezaud, P., "An ordered logit model of air traffic controller's conflict risk judgment," *ATC Quart.*, Vol. 16, No. 2, 2008, pp. 101–125.

[31] Flicker, R. and Fricke, M., "Improvement on the acceptance of a conflict resolution system by air traffic controllers," *6th USA/Europe ATM R&D Seminar*, Baltimore, MD, Jun. 27-30 2005.

[32] Lehman, S., Bolland, S., Remington, R., Humphreys, M. S., and Neal, A., "Using A* graph traversal to model conflict resolution in air traffic control," *Proc. 10th ICCM*, edited by D. D. Salvucci and G. Gunzelmann, Philadelphia, PA, Aug. 5-8 2010, pp. 139–144.

[33] Rantanen, E. M. and Nunes, A., "Hierarchical conflict detection in air traffic control," *Int. J. Aviat. Psychol.*, Vol. 15, No. 4, 2005, pp. 339–362.

[34] Kauppinen, S. and Flynn, M., "From business needds to field trials - a development of conflict resolution automation involving the controller," *Proc. 23rd DASC*, Vol. 1, IEEE, Salt Lake City, UT, Oct. 24-28 2004, pp. 5B1–1–5B1–12.

[35] Kirwan, B. and Flynn, M., "Investigating air traffic controller conflict resolution strategies," Rep. asa.01.cora.2.del04-b.rs, EUROCONTROL, Brussels, Belgium, Mar. 2002.

[36] Funder, D. C., "Personality," *Annu. Rev. Psychol.*, Vol. 52, No. 1, 2001, pp. 197–221.

[37] Kozhevnikov, M., "Cognitive styles in the context of modern psychology: Toward an integrated framework of cognitive style," *Psychol. Bull.*, Vol. 133, No. 3, 2007, pp. 464–481.

[38] Turing, A. M., "Computing machinery and intelligence," *Mind*, Vol. 59, No. 236, 1950, pp. 433–460.

[39] Fuld, R. B., Liu, Y., and Wickens, C. D., "The impact of automation on error detection: Some results from a visual discrimination task," *Proc. HFES 31st Annual Meeting*, Vol. 31, No. 2, New York, NY, Oct. 19-22 1987, pp. 156–160.

[40] Redding, R. E., Ryder, J. M., Seamster, T. L., Purcell, J. A., and Cannon, J. R., "Cognitive task analysis of en route air traffic control: Model extension and validation," Rep. ed340848, Hum. Technol., Inc., McLean, VA, Jan. 1992.

[41] Kahneman, D., *Thinking, fast and slow*, Farrar, Straus and Giroux, New York, NY, 2011.

[42] Endsley, M. R. and Kaber, D. B., "Level of automation effects on performance, situation awareness and workload in a dynamic control task," *Ergonomics*, Vol. 42, No. 3, 1999, pp. 462–492.

[43] Parasuraman, R., Sheridan, T. B., and Wickens, C. D., "A model of types and levels of human interaction with automation," *IEEE Trans. Syst., Man, Cybern. A, Syst. ,Humans*, Vol. 30, No. 3, 2000, pp. 286–297.

[44] Sheridan, T. B. and Verplank, W. L., "Human and computer control of undersea tele-operators," Tech. rep., MIT Man-Mach. Syst. Lab., Cambridge, MA, Jul. 1978.

[45] Wickens, C. D., Li, H., Santamaria, A., Sebok, A., and Sarter, N. B., "Stages and levels of automation: An integrated meta-analysis," *Human Factors*, Vol. 4, No. 4, 2010, pp. 389–393.

[46] Schaefer, K. E., Billings, D. R., Szalma, J. L., K, A. J., Sanders, T. L., Chen, J. Y. C., and Hancock, P. A., "A meta-analysis of factors influencing the development of trust in automation: Implications for human-robot interactions," Tech. Rep. ARL-TR-6984, Army Res. Lab., Aberdeen Proving Ground, MD, Jul. 2014.

[47] Schaefer, K. E., Chen, J. Y. C., Szalma, J. L., and Hancock, P. A., "A meta-analysis of factors influencing the development of trust in automation: Implications for understanding autonomy in future systems," *Hum. Factors*, Vol. 58, No. 3, 2016, pp. 377–400.

[48] de Visser, E. J., Monfort, S. S., McKendrick, R., Smith, M. A., McKnight, P. E., Krueger, F., and Parasuraman, R., "Almost human: Anthropomorphism increases trust resilience in cognitive agents," *J. Exp. Psychol.-Appl.*, Vol. 22, No. 3, 2016, pp. 331–49.

[49] Merritt, S. M., Sinha, R., Curran, P. G., and Ilgen, D. R., "Attitudinal predictors of relative reliance on human vs. automated advisors," *Int. J. Hum. Factors Ergon.*, Vol. 3, No. 3-4, 2015, pp. 327–345.

[50] van Dongen, K. and van Maanen, P.-P., "A framework for explaining reliance on decision aids," *Int. J. Hum-Comput. St., Volume = 71*, , No. 4, 2013, pp. 410–424.

[51] Billings, C. E., *Aviation automation: The search for a human-centered approach*, Erlbaum, Hillsdale, NJ, 1997.

[52] Dunnebeil, S., Sunyaev, A., Blohm, I., Leimeister, J. M., and Krcmar, H., "Determinants of physicians' technology acceptance for e-health in ambulatory care," *Int. J. Med. Inform.*, Vol. 81, No. 11, 2012, pp. 746–760.

[53] Poon, E. G., Blumenthal, D., Jaggi, T., Honour, M. M., Bates, D. W., and Kaushal, R., "Overcoming barriers to adopting and implementing computerized physician order entry systems in U.S. hospitals," *Health Affair.*, Vol. 23, No. 4, 2004, pp. 184–190.

[54] Yarbrough, A. K. and Smith, T. B., "Technology acceptance among physicians: A new take on TAM," *Med. Care. Res. Rev.*, Vol. 64, No. 6, 2007, pp. 650–672.

[55] Dzindolet, M. T., Beck, H. P., and Pierce, L. G., "Adaptive automation: Building flexibility into human-machine systems," *Understanding adaptability: A prerequisite for effective performance within complex environments*, edited by S. Burke, L. G. Pierce, and E. Salas, Vol. 6, Emerald, 2006, pp. 213–245.

[56] Dekker, S. W. A. and Woods, D. D., "To intervene or not to intervene: The dilemma of management by exception," *Cogn. Technol. Work.*, Vol. 1, No. 2, 1999, pp. 746–760.

[57] Kaber, D. B. and Endsley, M. R., "The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task," *Theor. Issues Ergon. Sci.*, Vol. 5, No. 2, 2004, pp. 113–153.

[58] Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., and Feltovich, P. J., "Ten challenges for making automation a "team player" in joint human-agent activity," *IEEE Intell. Syst.*, Vol. 19, No. 6, 2004, pp. 91–95.

[59] Rogers, E. M., *Diffusion of innovations*, Free Press, New York, NY, 3rd ed., 1983.

[60] Fischer, J. E., Steiner, F., Zucol, F., Berger, C., Martignon, L., Bossart, W., Altwegg, M., and Nadal, D., "Use of simple heuristics to target macrolide prescription in children with community-acquired pneumonia," *Arch. Pediatr. Adolesc. Med.*, Vol. 156, No. 10, 2002, pp. 1005–1008.

[61] Loft, S., Bolland, S., Humphreys, M. S., and Neal, A., "A theory and model of conflict detection in air traffic control: Incorporating environmental constraints," *J. Exp. Psychol.-Appl.*, Vol. 15, No. 2, 2009, pp. 106–124.

[62] BenMessaoud, C., Kharrazi, H., and MacDorman, K. F., "Facilitators and barriers to adopting robotic-assisted surgery: Contextualizing the unified theory of acceptance and use of technology," *PLoS ONE*, Vol. 6, No. 1, 2011, pp. e16395.

[63] Mosier, K. L. and Skitka, L. J., "Human decision makers and automated decision aids: Made for each other?" *Automation and human performance: Theory and applications*, edited by M. Mouloua and R. Parasuraman, Elrbaum, Mahwah, NJ, 1996, pp. 201–220.

[64] Ghazizadeh, M., Lee, J. D., and Boyle, L. N., "Extending the technology acceptance model to assess automation," *Cogn. Technol. Work*, Vol. 14, No. 1, 2012, pp. 39–49.

[65] Bigley, G. A. and Roberts, K. H., "The incident command system: High-reliability organizing for complex and volatile task environments," *Acad. Manage. J.*, Vol. 44, No. 6, 2001, pp. 1281–1300.

[66] Roberts, K. H., "Characteristics of one type of high reliability organization," *Organ. Sci.*, Vol. 1, No. 2, 1990, pp. 160–176.

[67] Dzindolet, M. T., Beck, H. P., Pierce, L. G., and Dawe, L. A., "A framework of automation use," Tech. rep. arl-tr-2412, Army Res. Lab., Aberdeen Proving Ground, MD, Mar. 2001.

[68] Lee, J. and Moray, N., "Trust, control strategies and allocation of function in human-machine systems," *Ergonomics*, Vol. 35, No. 10, 1992, pp. 1243–1270.

[69] Lee, K. K., Kerns, K., Bone, R., and Nickelson, M., "Development and validation of the controller acceptance rating scale (CARS): Results of empirical research," *4th USA/Europe ATM R&D Seminar*, Santa Fe, NM, Dec. 3-7. 2001.

[70] Muir, B. M., "Trust between humans and machines, and the design of decision aids," *Int. J. Man-Mach. Stud.*, Vol. 27, No. 5-6, 1987, pp. 527–539.

[71] Lee, J. and Moray, N., "Trust, self-confidence, and operators' adaptation to automation," *Ergonomics*, Vol. 40, No. 1, 1994, pp. 153–184.

[72] Madhavan, P. and Wiegmann, D. A., "Similarities and differences between human-human and human-automation trust: An integrative review," *Theor. Issues Ergon. Sci.*, Vol. 8, No. 4, 2007, pp. 277–301.

[73] Wiegmann, D. A., "Agreeing with automated diagnostic aids: A study of users' concurrence strategies," *Hum. Factors*, Vol. 44, No. 1, 2002, pp. 44–50.

[74] Goldmuntz, L., Kefaliotis, J. T., Kleiman, L. A., Rucker, R. A., Schuchman, L., and Weathers, D., "The AERA concept," Report no. faa-em-81-3, DOT/FAA Office Syst. Eng. Manage., Springfield, VA, Mar. 1981.

[75] Niedringhaus, W. P., "A mathematical formulation for planning automated aircraft separations for AERA 3," Rep. dot/faa/ds-89/20, MITRE, National Technical Information Service, Springfield, VA, Apr. 1989.

[76] Nicolaon, J.-P., Dean, G., and Miller, W., "An overview of ARC 2000 version 3 from the operational point of view," Rep. 286, EEC, Brétigny-sur-Orge, France, Dec. 1995.

[77] Jorna, P. G. A. M., Pavet, D., van Blanken, M., and Pichancourt, I., "PHARE Ground Human Machine Interface (GHMI) project: Summary report," Rep. DOC 99-70-02 99-70-02, EUROCONTROL, Brussels, Belgium, Aug. 1999.

[78] Langan-Fox, J., Canty, J. M., and Sankey, M. J., "Human-automation teams and adaptable control for future air traffic management," *Int. J. Ind. Erg.*, Vol. 39, No. 5, Sept. 2009, pp. 894–903.

[79] Loft, S., Sanderson, P., Neal, A., and Mooij, M., "Modeling and predicting mental workload in en route air traffic control: Critical review and broader implications," *Hum. Factors*, Vol. 49, No. 3, 2007, pp. 376–399.

[80] Sperandio, J.-C., "The regulation of working methods as a function of work-load among air traffic controllers," *Ergonomics*, Vol. 21, No. 3, 1978, pp. 195–202.

[81] Shah, A. K. and Oppenheimer, D. M., "Heuristics made easy: An effort–reduction framework," *Psychol. Bull.*, Vol. 134, No. 2, 2008, pp. 207–222.

[82] Gigerenzer, G., "Why heuristics work," *Perspect. Psychol. Sci.*, Vol. 3, No. 1, 2008, pp. 20–29.

[83] Goodrich, M. A., Stirling, W. C., and Boer, E. R., "Satisficing Revisited," *Mind. Mach.*, Vol. 10, No. 1, 2000, pp. 79–109.

[84] Goodrich, M. A. and Boer, E. R., "Model-based human-centered task automation: A case study in ACC system design," *IEEE Trans. Syst., Man, Cybern. A, Syst. ,Humans*, Vol. 33, No. 3, 2003, pp. 325–336.

[85] Rasmussen, J., "Skill, rules and knowledge: Signals, signs, and symbols, and other distinctions in human performance models," *IEEE Trans. Syst., Man, Cybern.*, Vol. 13, No. 3, 1983, pp. 257–266.

[86] Nijhuis, H. B., "Role of the human in the evolution of ATM systems (RHEA): Final report," Rep no. rhea/nl/wpr/8/04, European Commission/NLR, Amsterdam, Netherlands, Feb. 1998.

[87] Mogford, R. H., Murphy, E. D., Roske-Hofstrand, R. J., Yastrop, G., and Guttman, J. A., "Application of research techniques for documenting cognitive processes in air traffic control: Sector complexity and decision making," Tech. note dot/faa/ct-tn94/3, CTA Incorp., Atlantic City, NJ, Jun. 1994.

[88] Eyferth, K., Niessen, C., and Spaeth, O., “A model of air traffic controllers’ conflict detection and conflict resolution,” *Aerosp. Sci. Technol.*, Vol. 7, No. 6, 2003, pp. 409–416.

[89] Kallus, K. W., Barbarino, M., and VanDamme, D., “Integrated task and job analysis of air traffic controllers - Phase 2: Task analysis of en-route controllers,” Rep. hum.et.st01.1000-rep-04, EUROCONTROL, Brussels, Belgium, Oct. 1999.

[90] Rozen, N. E., Love, W. D., Kirk, D. B., and Heagy, W. S., “Problem analysis, resolution and ranking (PARR) initial functional performance assessment,” *AIAA GNC Conf.*, American Institute of Aeronautics and Astronautics, Monterey, CA, Aug. 5-6 2002, pp. 1926–1943.

[91] Özgür, M. and Cavcar, A., “A knowledge-based conflict resolution tool for en-route air traffic controllers,” *Aircr. Eng. Aerosp. Tec.*, Vol. 80, No. 6, 2008, pp. 649–656.

[92] Alavi, M. and Joachimsthaler, E. A., “Revisiting DSS implementation research: A meta-analysis of the literature and suggestions for researchers,” *MIS Quart.*, Vol. 16, No. 1, 1992, pp. 95–116.

[93] Cegarra, J. and Hoc, J.-M., “Cognitive styles as an explanation of experts’ individual differences: A case study in computer-assisted troubleshooting diagnosis,” *Int. J. Hum-Comput. St.*, Vol. 64, No. 2, 2006, pp. 123–136.

[94] Chakraborty, I., Hu, P. J.-H., and Cui, D., “Examining the effects of cognitive style in individuals’ technology use decision making,” *Decis. Support Syst.*, Vol. 45, No. 2, 2008, pp. 228–241.

[95] Lu, H.-P., Yu, H.-J., and Lu, S. S. K., “The effects of cognitive style and model type on DSS acceptance: An empirical stud,” *Eur. J. Oper. Res.*, Vol. 131, No. 3, 2001, pp. 649–663.

[96] Liu, Y., Lee, Y., and Chen, A. N. K., “Evaluating the effects of task-individual-technology fit in multi-DSS models context: A two-phase view,” *Decis. Support Syst.*, Vol. 51, No. 3, 2011, pp. 688–700.

[97] Langan-Fox, J., Sankey, M. J., and Canty, J. M., “Human factors measurement for future air traffic control systems,” *Hum. Factors*, Vol. 51, No. 5, 2009, pp. 595–637.

[98] Stankovic, S., Raufaste, E., and Averty, P., “Determinants of conflict detection: A model of risk judgments in air traffic control,” *Hum. Factors*, Vol. 50, No. 1, 2008, pp. 121–134.

[99] Willems, B. F. and Koros, A., “Advanced concept of the national airspace system 2015: Human factors considerations for air traffic control,” Dot/faa/tc-tn-07/21, William J. Hughes Tech. Cent., Atlantic City, NJ, Jun. 2007.

[100] Hilburn, B. and Jorna, P. G. A. M., “Workload in air traffic control,” *Stress, workload, and fatigue: Theory, research, and practice*, edited by P. A. Hancock and P. A. Desmond, Erlbaum, Hillsdale, NJ, 2001, pp. 384–394.

[101] Hilburn, B., Westin, C., and Borst, C., “Will controllers accept a machine that thinks like they think? The role of strategic conformance in decision aiding automation,” *ATC Quart.*, Vol. 22, No. 2, 2014, pp. 115–136.

[102] Mercado Velasco, G. A., Borst, C., Mulder, M., and van Paassen, M. M., “Solution space decision support for reducing controller workload in a route merging task,” *AIAA J. Aircraft.*, submitted for publication.

[103] Dillon, A. and Morris, M. G., “User acceptance of new information technology: Theories and models,” *Annual review of information science and technology*, edited by M. Williams, Vol. 31, Information Today, Medford NJ, 1996, pp. 3–32.

[104] Bagozzi, R. P., “The legacy of the technology acceptance model and a proposal for a paradigm shift,” *J. Assoc. Inform. Syst.*, Vol. 8, No. 4, 2007, pp. 244–254.

[105] King, W. R. and He, J., “A meta-analysis of the technology acceptance model,” *Inform. Manage.*, Vol. 43, No. 6, 2006, pp. 740–755.

[106] Venkatesh, V., Morris, M. G., Davis, G. B., and Davis, F. D., “User acceptance of information technology: Toward a unified view,” *MIS Quart.*, Vol. 27, No. 3, 2003, pp. 425–478.

[107] Davis, F. D., *A technology acceptance model for empirically testing new end-user information systems: Theory and results*, Ph.d. thesis, Dept. Sloan School Manage., MIT, Cambridge, MA, 1985.

[108] Davis, F. D., Bagozzi, R. P., and Warshaw, P. R., “User acceptance of computer technology: A comparison of two theoretical models,” *Manage. Sci.*, Vol. 35, No. 8, 1989, pp. 982–1003.

[109] Karahanna, E., Agarwal, R., and Angst, C. M., “Reconceptualizing compatibility beliefs,” *MIS Quart.*, Vol. 30, No. 4, 2006, pp. 781–804.

[110] Merritt, S. M. and Ilgen, D. R., “Not all trust is created equal: Dispositional and history-based trust in human-automation interactions,” *Hum. Factors*, Vol. 50, No. 2, 2008, pp. 194–210.

[111] Cramer, H., Evers, V., Ramlal, S., Someren, M., Rutledge, L., Stash, N., Aroyo, L., and Wielinga, B., “The effects of transparency on trust in and acceptance of a content-based art recommender,” *User Model. User-Adap. Inter.*, Vol. 18, No. 5, 2008, pp. 455–496.

[112] Dillon, A., “Human acceptance of information technology,” *International encyclopedia of human factors and ergonomics*, edited by W. Karwowski, Taylor and Francis, London, UK, 2001.

[113] Tornatzky, L. G. and Klein, K. J., “Innovation characteristics and innovation adoption implementation: A meta-analysis of findings,” *IEEE Trans. Eng. Manag.*, Vol. EM-29, No. 1, 1982, pp. 28–43.

[114] Parasuraman, R. and Miller, C. A., “Trust and etiquette in high-criticality automated systems,” *Commun. ACM*, Vol. 47, No. 4, 2004, pp. 5155.

[115] Kiesler, S., Powers, A., Fussell, S. R., and Torrey, C., “Anthropomorphic interactions with a robot and robot-like agent,” *Soc. Cognition*, Vol. 26, No. 2, 2008, pp. 169–181.

[116] Waytz, A., Cacioppo, J., and Epley, N., “Who sees human?: The stability and importance of individual differences in anthropomorphism,” *Perspect. Psychol. Sci.*, Vol. 5, No. 3, 2010, pp. 219–232.

[117] Karr-Wisniewski, P. and Prietula, M., “CASA, WASA, and the dimensions of us,” *Comput. Hum. Beh.*, Vol. 26, No. 6, 2010, pp. 1761–1771.

[118] Nass, C. and Moon, Y., “Machines and mindlessness: Social responses to computers,” *J. Soc. Issues*, Vol. 56, No. 1, 2000, pp. 81–103.

[119] Nass, C., Moon, Y., and Carney, P., “Are people polite to computers? Responses to computer-based interviewing systems,” *J. Appl. Soc. Psychol.*, Vol. 29, No. 5, 1999, pp. 1093–1109.

[120] Hayes, C. C. and Miller, C. A., *Human-computer etiquette: Cultural expectations and the design implications they place on computers and technology*, Auerbach Publications, Boca Raton, FL, 2010.

[121] Borst, C., Flach, J. M., and Ellerbroek, J., “Beyond Ecological Interface Design: Lessons from concerns and misconceptions,” *IEEE Trans. Human–Mach. Syst.*, Vol. 45, No. 2, 2015, pp. 164–175.

[122] Vicente, K. J. and Rasmussen, J., “Ecological interface design: Theoretical foundations,” *IEEE Trans. Syst., Man, Cybern.*, Vol. 22, No. 4, 1992, pp. 589–606.

[123] ICAO, *Rules of the air*, ICAO, 10th ed., Jul. 2005.

[124] Nass, C., Fogg, B. J., and Moon, Y., “Can computers be teammates?” *Int. J. Hum-Comput. St.*, Vol. 45, No. 6, 1996, pp. 669–678.

[125] Eyben, F., Wöllmer, M., Poitschke, T., Schuller, B., Blaschke, C., Färber, B., and Nguyen-Thien, N., “Emotion on the road: Necessity, acceptance, and feasibility of affective computing in the car,” *Adv. Hum-Comput. Inter.*, Vol. 2010, 2010, pp. 1–17.

[126] Adomavicius, G. and Tuzhilin, A., “Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions,” *IEEE Trans. Knowl. Data Eng.*, Vol. 17, No. 6, 2005, pp. 734–749.

[127] Sheridan, T. B., *Humans and automation: System design and research issues*, Wiley, New York, NY, 2002.

[128] Hilburn, B., “Evaluating human interaction with advanced air traffic management automation,” *The Role of Humans in Intelligent and Automated Systems, RTO-MP-088*, Warsaw, Poland, Oct. 7-9 2002, pp. 17–1–17–12.

[129] Hilburn, B., “Strategic decision aiding: A comprehensive bio-behavioral assessment of human interaction with real-time decision aids. Phase I final report,” Nato crg 94.104., NATO, Brussels, Belgium, 1996.

[130] Maehigashi, A., Miwa, K., Terai, H., Kojima, K., Morita, J., and Hayashi, Y., “Experimental investigation of misuse and disuse in using automation system,” *Proc. 14th HCI*, edited by J. A. Jacko, Jul. 9-14 2011, pp. 384–393.

[131] Bonzano, A., Cunningham, P., and Meckiff, C., “ISAC: A CBR system for decision support in air traffic control,” *Proc. 3rd EWCBR*, Vol. 1168, Lausanne, Switzerland, Nov. 14-16 1996, pp. 44–57.

[132] Niessen, C., Eyferth, K., and Bierwagen, T., “Modelling cognitive processes of experienced air traffic controllers,” *Ergonomics*, Vol. 42, No. 11, 1999, pp. 1507–1520.

[133] Hilburn, B., Parasuraman, R., Jha, P. D., and McGarry, K., “Emerging human factors issues in future air traffic management,” Tech. rep., The Hague, Netherlands, 2006.

[134] Westin, C., Hilburn, B., and Borst, C., “Report on initial 3D framework,” E.02.08-mufasa-d1, EUROCONTROL, Brussels, Belgium, Sep. 2011.

[135] Bahner, J. E., Hüper, A.-D., and Manzey, D., “Misuse of automated decision aids: Complacency, automation bias and the impact of training experiences,” *Int. J. Hum-Comput. St.*, Vol. 66, No. 9, 2008, pp. 688–699.

[136] Djokic, J., Lorenz, B., and Fricke, H., “Air traffic complexity as a workload driver,” *Transp. Res. Part C Emerg. Technol.*, Vol. 18, No. 6, 2010, pp. 930–936.

[137] Mogford, R. H., Guttman, J. A., Morrow, S. L., and Kopardekar, P., “The complexity construct in air traffic control: A review and synthesis of the literature,” Tech. rep., Springfield, VA, 1995.

[138] Kirwan, B., Rodgers, M., and Schäfer, D., “Human factors impacts in air traffic management,” 2005.

[139] Neal, A., Flach, J., Mooij, M., Lehman, S., Stankovic, S., and Hasenbosch, S., “Envisaging the future air traffic management system,” *Int. J. Aviat. Psychol.*, Vol. 21, No. 1, 2011, pp. 16–34.

[140] Oppermann, R. and Simm, H., “Adaptability: User-initiated individualization,” *Adaptive user support: Ergonomic design of manually and automatically adaptable software*, edited by R. Oppermann, Erlbaum, Hillsdale, NJ, 1994, pp. 14–66.

[141] Metzger, U. and Parasuraman, R., “The role of the air traffic controller in future air traffic management: An empirical study of active control versus passive monitoring,” *Hum. Factors*, Vol. 43, No. 4, 2001, pp. 519–528.

[142] Good, I. J., “Speculations concerning the first ultraintelligent machine,” *Advances in computers*, edited by F. L. Alt and M. Rubinoff, Vol. 6, Academic Press, New York, NY, 1966, pp. 31–88.

[143] Hawking, S., Russell, S., Tegmark, M., and Wilczek, F., “Stephen Hawking: ’Transcendence looks at the implications of artificial intelligence – but are we taking AI seriously enough?’,” *The Independent*, May 1 2014.

144 Alberdi, E., Ayton, P., Povyakalo, A. A., and Strigini, L., "Automation bias and system design: A case study in a medical application," *Proc. IEE MOD HFI DTC*, London, UK, Nov. 16-17 2005, pp. 53–60.

145 Goddard, K., Roudsari, A., and Wyatt, J. C., "Automation bias: A systematic review of frequency, effect mediators, and mitigators," *J. Amer. Med. Inform. Assoc.*, Vol. 19, No. 1, 2012, pp. 121–127.

146 Rieh, S. Y. and Danielson, D. R., "Credibility: A multidisciplinary framework," *Annual Rev. Info. Sci. & Technol.*, Vol. 41, No. 1, 2007, pp. 307–364.

147 Kelly, C., Boardman, M., Goillau, P., and Jeannot, E., "Guidelines for trust in future ATM systems: A literature review," Rep. hrs/hsp-005-gui-01, EUROCONTROL, Brussels, Belgium, May. 2003.

148 Wang, L., Jamieson, G. A., and Hollands, J. G., "Trust and reliance on an automated combat identification system," *Hum. Factors*, Vol. 51, No. 3, 2009, pp. 281–291.

149 Oleson, K. E., Hancock, P. A., Billings, D. R., and Schesser, C. D., "Trust in unmanned aerial systems: A synthetic, distributed trust model," *Proc. 16th ISAP*, Dayton, OH, May. 2-5 2011.

150 BFU, "Investigation report," Tech. Rep. AX001-1-2/02, German Federal Bureau of Aircraft Accidents Investigation, Braunschweig, Germany, May 2004.

151 NTSB, "Aircraft accident report: Descent below visual glidepath and impact with seawall," Tech. Rep. NTSB/AAR-14/01, National Transportation Safety Board, Washington, DC, Jun. 2014.

152 Lee, J. D., "Review of a pivotal human factors article: "Humans and automation: Use, misuse, disuse, abuse"," *Hum. Factors*, Vol. 50, No. 3, June 2008, pp. 404–410.

153 Wiener, E. L., "Cockpit automation," *Human factors in aviation*, edited by E. L. Wiener and D. C. Nagel, Harcourt Press, San Diego, CA, 1988, pp. 433–461.

154 Szalma, J. L., "Individual differences in human-technology interaction: Incorporating variation in human characteristics into human factors and ergonomics research and design," *Theor. Issues Ergon. Sci.*, Vol. 10, No. 5, 2009, pp. 381–397.

155 Westin, C., Borst, C., and Hilburn, B., "Strategic conformance: Overcoming acceptance issues of decision aiding automation?" *IEEE Trans. Human–Mach. Syst.*, Vol. 46, No. 1, 2016, pp. 41–52.

156 Muir, B. M., "Trust in automation: Part I. Theoretical issues in the study of trust and human intervention in automated systems," *Ergonomics*, Vol. 37, No. 11, 1994, pp. 1905–1922.

157 Tseng, S. and Fogg, B. J., "Credibility and computing technology," *Commun. ACM*, Vol. 42, No. 5, 1999, pp. 39–44.

158 Earle, T. C., "Trust in risk management: A model-based review of empirical research," *Risk anal.*, Vol. 30, No. 4, 2010, pp. 541–574.

159 Kim, J. and Moon, J. Y., "Designing towards emotional usability in customer interfaces – trustworthiness of cyber-banking system interfaces," *Interact. comput.*, Vol. 10, No. 1, 1998, pp. 1–29.

160 Lewandowsky, S., Mundy, M., and Tan, G. P., "The dynamics of trust: Comparing humans to automation," *J. Exp. Psychol.-Appl.*, Vol. 6, No. 2, 2000, pp. 104–123.

161 Dzindolet, M. T., Beck, H. P., Pierce, L. G., and Dawe, L. A., "A framework of automation use," Tech. rep. arl-tr-2412, Army Res. Lab., Aberdeen Proving Ground, MD, Mar. 2001.

162 Gao, J. and Lee, J. D., "Extending the decision field theory to model operators' reliance on automation in supervisory control situations," *IEEE Trans. Syst., Man, Cybern. A, Syst. ,Humans*, Vol. 36, No. 5, 2006, pp. 943–959.

163 Moray, N., Inagaki, T., and Itoh, M., "Adaptive automation, trust, and self-confidence in fault management of time-critical tasks," *J. Exp. Psychol.-Appl.*, Vol. 6, No. 1, 2000, pp. 44–58.

[164] Wrn, Y. and Ramberg, R., "People's perception of human and computer advice," *Comput. Hum. Beh.*, Vol. 12, No. 1, 1996, pp. 17–27.

[165] Dzindolet, M. T., Pierce, L. G., Beck, H. P., Dawe, L. A., and Anderson, B. W., "The perceived utility of human and automated aids in a visual detection task," *Hum. Factors*, Vol. 44, No. 1, 2002, pp. 79–94.

[166] Ma, R. and Kaber, D. B., "Effects of in-vehicle navigation assistance and performance on driver trust and vehicle control," *Int. J. Ind. Erg.*, Vol. 37, No. 8, 2007, pp. 665–673.

[167] Madhavan, P. and Wiegmann, D. A., "Effects of information source, pedigree, and reliability on operator interaction with decision support systems," *Hum. Factors*, Vol. 49, No. 5, 2007, pp. 773–785.

[168] Parasuraman, R. and Manzey, D. H., "Complacency and bias in human use of automation: An attentional integration," *Hum. Factors*, Vol. 52, No. 3, 2010, pp. 381–410.

[169] Wickens, C. D., Clegg, B. A., Vieane, A. Z., and Sebok, A. L., "Complacency and automation bias in the use of imperfect automation," *Hum. Factors*, Vol. 57, No. 5, 2015, pp. 728–739.

[170] Jian, J.-Y., Bisantz, A. M., and Drury, C. G., "Foundations for an empirically determined scale of trust in automated systems," *Int. J. Cog. Erg.*, Vol. 4, No. 1, 2000, pp. 53–71.

[171] Lerch, F. J., Prietula, M., and Kulik, C. T., "The Turing effect: The nature of trust in expert systems advice," *Expertise in context*, edited by J. F. Paul, M. F. Kenneth, and R. H. Robert, MIT Press, Cambridge, MA, 1997, pp. 417–448.

[172] Dijkstra, J. J., Liebrand, W. B. G., and Timminga, E., "Persuasiveness of expert systems," *Behav. Inform. Technol.*, Vol. 17, No. 3, 1998, pp. 155–163.

[173] Dijkstra, J. J., "User agreement with incorrect expert system advice," *Behav. Inform. Technol.*, Vol. 18, No. 6, 1999, pp. 399–411.

[174] Dzindolet, M. T., Peterson, S. A., Pomranky, R. A., Pierce, L. G., and Beck, H. P., "The role of trust in automation reliance," *Int. J. Hum-Comput. St.*, Vol. 58, No. 6, 2003, pp. 697–718.

[175] Lacson, F. C., Wiegmann, D. A., and Madhavan, P., "Effects of attribute and goal framing on automation reliance and compliance," *Proc. HFES 49th Annual Meeting*, Vol. 49, No. 3, Orlando, FL, Sep. 26-30 2005, pp. 482–486.

[176] Waytz, A., Heafner, J., and Epley, N., "The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle," *J. Exp. Soc. Psychol.*, Vol. 52, 2014, pp. 113–117.

[177] Merritt, S. M., "Affective processes in human-automation interactions," *Hum. Factors*, Vol. 53, No. 4, 2011, pp. 356–370.

[178] de Vries, P., Midden, C., and Bouwhuis, D., "The effects of errors on system trust, self-confidence, and the allocation of control in route planning," *Int. J. Hum.-Comput. St.*, Vol. 58, No. 6, 2003, pp. 719–735.

[179] Bradshaw, J. M., Beautement, P., Breedy, M. R., Bunch, L., Drakunov, S. V., Feltovich, P., Hoffman, R. R., Jeffers, R., Johnson, M., Kulkarni, S., Lott, J., Raj, A., Suri, N., and Uszok, A., "Making agents acceptable to people," *Intelligent Technologies for Information Analysis: Advances in Agents, Data Mining, and Statistical Learning*, edited by N. Zhong and J. Liu, Springer Berlin Heidelberg, 2004, pp. 361–406.

[180] Stedmon, A. W., Sharples, S., Littlewood, R., Cox, G., Patel, H., and Wilson, J. R., "Datalink in air traffic management: Human factors issues in communications," *Appl. Ergon.*, Vol. 38, No. 4, 2007, pp. 473–480.

[181] Hughes, S. M. and Harrison, M. A., "I like my voice better: Self-enhancement bias in perceptions of voice attractiveness," *Perception*, Vol. 42, No. 9, 2013, pp. 941–949.

[182] Culley, K. E. and Madhavan, P., "A note of caution regarding anthropomorphism in HCI agents," *Comput. Hum. Behav.*, Vol. 29, No. 3, 2013, pp. 577–579.

[183] Groom, V. and Nass, C., "Can robots be teammates?: Benchmarks and predictors of failure in human-robot teams," *Interact. Stud.*, Vol. 8, No. 3, 2008, pp. 219–226.

[184] Jha, P. D., Bisantz, A. M., Parasuraman, R., and Drury, C. G., "Air traffic controllers' performance in advance air traffic management system: Part I-performance results," *Int. J. Aviat. Psychol.*, Vol. 21, No. 3, 2011, pp. 283–305.

[185] Lyons, J. B. and Stokes, C. K., "Human-human reliance in the context of automation," *Hum. Factors*, Vol. 54, No. 1, 2012, pp. 112–121.

[186] Sarter, N. B., Woods, D. D., and Billings, C., "Automation surprises," *Handbook of human factors & ergonomics*, edited by G. Salvendy, Wiley, New York, NY, 2nd ed., 1997, pp. 19–35.

[187] Christoffersen, K. and Woods, D. D., "How to make automated systems team players," *Advances in human performance and cognitive engineering research*, Vol. 2, 2002, pp. 1–12.

[188] Le Goff, K., Rey, A., and Berberian, B., "Toward a model for effective human-automation interaction: The mediated agency," *6th DHM*, edited by V. G. Duffy, Applications in health, safety, ergonomics and risk management: Ergonomics and health, Springer, 2015, pp. 274–283.

[189] D'Arcy, J. F. and Della Rocco, P. S., "Air traffic control specialist decision making and strategic planning – A field survey," Dot/faa/ct-tn01/05, William J. Hughes Tech. Cent., Atlantic City, NJ, Mar. 2001.

[190] Dudley, R., Dorneich, M. C., Letsu-Dake, E., Rogers, W., Whitlow, S. D., Dillard, M., and Nelson, E., "Characterization of information automation on the Flight Deck," *Hum. Factors*, Vol. 1, No. 58, 2014, pp. 295–299.

[191] Helldin, T., *Transparency for future semi-automated systems: Effects of transparency on operator performance, workload and trust*, Ph.d. dissertation, Örebro univ., Örebro, Sweden, 2014.

[192] Höök, K., "Steps to take before intelligent user interfaces become real," *Interact. Comput.*, Vol. 12, No. 4, 2000, pp. 409–426.

[193] Jameson, A. and Gajos, K. Z., "System that adapt to their users," *Human computer interaction handbook: Fundamentals, evolving technologies, and emerging applications*, edited by J. A. Jacko, CRC Press, Boca Raton, FL, 3rd ed., 2012, pp. 431–456.

[194] Jameson, A., "Adaptive interfaces and agents," *Human-computer interaction handbook: Fundamentals, evolving technologies and emerging applications*, edited by A. Sears and J. J. A, CRC Press, Boca Raton, FL, 2nd ed., 2008, pp. 433–458.

[195] Sadler, G., Battiste, H., Ho, N., Hoffmann, L., Lyons, J., Johnson, W., Shively, R., and Smith, D., "Effects of transparency on pilot trust and agreement in the autonomous constrained flight planner," *35th IEEE/AIAA DASC*, Sacramento, CA, Sep. 25-29 2016.

[196] Lim, Brian Y amd Dey, A. K. and Avrahami, D., "Why and why not explanations improve the intelligibility of context-aware intelligent systems," *Proc. SIGCHI*, Boston, MA, Apr. 4-9 2009, pp. 2119–2128.

[197] Chen, J. Y. C. and Barnes, M. J., "Human–agent teaming for multirobot control: A review of human factors issues," *IEEE Trans. Human–Mach. Syst.*, Vol. 44, No. 1, 2014, pp. 13–29.

[198] Parasuraman, R. and Jiang, Y., "Individual differences in cognition, affect, and performance: Behavioral, neuroimaging, and molecular genetic approaches," *Neuroimage*, Vol. 59, No. 1, 2012, pp. 7082.

[199] Woods, D. D., Patterson, E. S., and Roth, E. M., "Can we ever escape from data overload? A cognitive systems diagnosis," *Cogn. Technol. Work*, Vol. 4, No. 1, 2002, pp. 22–36.

[200] Chen, J. Y. C., Procci, K., Boyce, M. W., Wright, J., Garcia, A., and Barnes, M. J., "Situation awareness-based agent transparency," Arl-tr-6905, Army Res. Lab., Aberdeen Proving Ground, MD, Apr. 2014.

[201] Ricci, F., Rokach, L., and Shapira, B., *Recommender systems handbook*, Springer, New York, NY, 2nd ed., 2015.

[202] Nielsen, J., *Usability engineering*, Academic, Boston, MA, 1993.

[203] Letsu-Dake, E., Rogers, W., Whitlow, S. D., Nelson, E., Dillard, M., Dorneich, M. C., and Dudley, R., "Flight deck information automation: A human-in-the-loop in trail procedure simulation study," *34th DASC*, Prague, Czech Republic, 2015, pp. 3D1–1–3D1–12.

[204] Brown, J. S., "Toward a new epistemology for learning," *Intelligent tutoring systems: At the crossroads of artificial intelligence and education*, edited by J. Frasson and J. Gauthiar, Ablex, Norwood, NJ, 1990, pp. 266–282.

[205] Endsley, M. R., "Toward a theory of situation awareness in dynamic systems," *Hum. Factors*, Vol. 37, No. 1, 1995, pp. 32–64.

[206] Rao, A. S. and Georgeff, M. P., "BDI agents: From theory to practice," *Proc. ICMAS*, San Fransisco, CA, Jun. 12-14 1995, pp. 312–319.

[207] Beck, H. P., Dzindolet, M. T., and Pierce, L. G., "Automation usage decisions: Controlling intent and appraisal errors in a target detection task," *Hum. Factors*, Vol. 49, No. 3, 2007, pp. 429–437.

[208] Thong, N. T. and Son, L. H., "HIFCF: An effective hybrid model between picture fuzzy clustering and intuitionistic fuzzy recommender systems for medical diagnosis," *Expert Syst. Appl.*, Vol. 42, No. 7, 2015, pp. 3682–3701.

[209] Wang, W. and Benbasat, I., "Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs," *J. Manage. Inform. Syst.*, Vol. 23, No. 4, 2007, pp. 217–246.

[210] Herlocker, J. L., Konstan, J. A., Terveen, L. G., and Riedl, J. T., "Evaluating collaborative filtering recommender systems," *ACM Trans. Inform. Syst.*, Vol. 22, No. 1, 2004, pp. 5–453.

[211] Tintarev, N. and Masthoff, J., "Explaining recommendations: Design and evaluation," *Recommender systems handbook*, edited by F. Ricci, L. Rokach, and B. Shapira, Springer, New York, NY, 2nd ed., 2015, pp. 353–382.

[212] Ricci, F., Rokach, L., and Shapira, B., "Recommender systems: Introduction and challenges," *Recommender systems handbook*, edited by F. Ricci, L. Rokach, and B. Shapira, Springer, New York, NY, 2nd ed., 2015, pp. 1–34.

[213] Herlocker, J. L., Konstan, J. A., and Riedl, J. T., "Explaining collaborative filtering recommendations," *Proc. CSCW*, ACM, Philadelphia, PA, Dec. 2-6 2000, pp. 241–250.

[214] Schaffer, J., Giridhar, P., Jones, D., Höllerer, T., Abdelzaher, T., and O'Donovan, J., "Getting the message?: A study of explanation interfaces for microblog data analysis," *20th ACM IUI*, ACM, Atlanta, GA, Mar. 29 - Apr. 1 2015, pp. 345–356.

[215] Boyce, M. W., Chen, J. Y. C., Selkowitz, A. R., and Lakhmani, S. G., "Agent transparency for an autonomous squad member," Arl-tr-7298, Army Res. Lab., Aberdeen Proving Ground, MD, May 2015.

[216] Selkowitz, A. R., Lakhmani, S. G., Larios, C. N., and Chen, J. Y. C., "Agent transparency and the autonomous squad member," *Proc. HFES 60th Annual Meeting*, Vol. 60, No. 1, Washington, DC, Sep. 19-23 2016, pp. 1319–1323.

[217] Göritzlehner, R., Borst, C., Ellerbroek, J., Westin, C., van Paassen, M. M., and Mulder, M., "Effects of transparency on the acceptance of automated resolution advisories," *IEEE SMC*, San Diego, CA, Oct. 5-8 2014, pp. 2965–2970.

[218] Bass, E. J. and Pritchett, A. R., "Human-automated judge learning: A methodology for examining human interaction with information analysis automation," *IEEE Trans. Syst., Man, Cybern. A, Syst. ,Humans*, Vol. 38, No. 4, 2008, pp. 759–776.

[219] Bass, E. J., Baumgart, L. A., and Shepley, K. K., "The effect of information analysis automation display content on human judgment performance in noisy environments," *J. Cogn. Eng. Decis. Making*, Vol. 7, No. 1, 2013, pp. 49–65.

[220] Bisantz, A. M., Stone, R. T., Pfautz, J., Fouse, A., Farry, M., Roth, E. M., Nagy, A. L., and Thomas, G., "Visual representations of meta-information," *J. Cogn. Eng. Decis. Making*, Vol. 3, No. 1, 2009, pp. 67–91.

[221] Pfautz, J., Fouse, A., Fichtl, T., Roth, E., Bisantz, A., and Madden, S., "The impact of meta-information on decision-making in intelligence operations," *Proc. HFES 49th Annual Meeting*, Vol. 49, No. 3, Orlando, FL, Sep. 26-30 2005, pp. 214–218.

[222] Duggan, G. B., Banbury, S., Howes, A., Patrick, J., and Waldron, S. M., "Too much, too little or just right: Designing data fusion for situation awareness," *Proc. HFES 48th Annual Meeting*, Vol. 48, No. 3, New Orleans, LA, Sep. 20-24 2004, pp. 528–532.

[223] Marois, R. and Ivanoff, J., "Capacity limits of information processing in the brain," *Trends Cogn. Sci.*, Vol. 9, No. 6, 2005, pp. 296–305.

[224] Moacdieh, N. and Sarter, N. B., "Display clutter: A review of definitions and measurement techniques," *Hum. Factors*, Vol. 1, No. 57, 2015, pp. 61–100.

[225] Hall, M. J., Hall, S. A., and Tate, T., "Removing the HCI bottleneck: How the human-computer interface affects the performance of data fusion systems," *Handbook of multisensor data fusion*, edited by D. L. Hall and J. Llinas, CRC Press, Boca Ration, FL, 2001, pp. 89–104.

[226] Flach, J. M., Tanabe, F., Monta, K., Vicente, K. J., and Rasmussen, J., "An ecological approach to interface design," *Proc. HFES 42th Annual Meeting*, Vol. 42, No. 3, Chicago, IL, Oct. 5-9 1998, pp. 295–299.

[227] Kilgore, R. and Voshell, M., "Increasing the transparency of unmanned systems: Applications of ecological interface design," *Proc. 6th VAMR*, edited by R. Shumaker and S. Lackey, Springer, Heraklion, Greece, Jun. 22-27 2014, pp. 378–389.

[228] Rasmussen, J., "The role of hierarchical knowledge representation in decision-making and system management," *IEEE Trans. Syst., Man, Cyberne.*, Vol. SMC-15, No. 2, 1985, pp. 234–243.

[229] Wong, B. L. W., Rozzi, S., Amaldi, P., Woodward, P., and Fields, B., "A Framework for considering spatial-temporal representation design in air traffic control," *Proc. HFES 50th Annual Meeting*, Vol. 50, No. 3, San Francisco, CA, Oct. 16-20 2006, pp. 250–254.

[230] Brázdilová, S. L., Cásek, P., and Kubalčík, J., "Air traffic complexity for a distributed air traffic management system," *J. Aerosp. Eng.*, Vol. 225, No. 6, 2011, pp. 665–674.

[231] van Doorn, B., Bakker, B., and Meckiff, C., "Evaluation of advanced conflict modelling in the highly interactive problem solver," *4th USA/Europe ATM R&D Seminar*, Santa Fe, NM, Dec. 3-7 2001.

[232] van Dam, S. B. J., Mulder, M., and van Paassen, M. M., "Ecological interface design of a tactical airborne separation assistance tool," *IEEE Trans. Syst., Man, Cybern. A, Syst. ,Humans*, Vol. 38, No. 6, 2008, pp. 1221–1233.

[233] Klomp, R., Borst, C., van Paassen, R., and Mulder, M., "Expertise level, control strategies, and robustness in future air traffic control decision aiding," *IEEE Trans. Human–Mach. Syst.*, Vol. 46, No. 2, 2016, pp. 255–266.

[234] Kirk, D. B., Bowen, K. C., Heagy, W. S., Rozen, N. E., and Viets, K. J., "Problem analysis, resolution and ranking (PARR) development and assessment," *4th USA/Europe ATM R&D Seminar*, Santa Fe, NM, Dec. 3-7. 2001.

[235] Hopkin, V. D., *Human factors is air traffic control*, Taylor & Francis, London, UK, 1995.

[236] Shanteau, J., "Why task domains (still) matter for understanding expertise," *J. Appl. Res. Mem. Cogn.*, Vol. 4, No. 3, 2015, pp. 169–175.

[237] Hancock, P. A., Hancock, G. M., and Warm, J. S., "Individuation: The N = 1 revolution," *Theor. Issues Ergon. Sci.*, Vol. 10, No. 5, 2009, pp. 481–488.

[238] Dewberry, C., Juanchich, M., and Narendran, S., "Decision-making competence in everyday life: The roles of general cognitive styles, decision-making styles and personality," *Pers. Indiv. Differ.*, Vol. 55, No. 7, 2013, pp. 783–788.

[239] Riding, R. and Cheema, I., "Cognitive styles–an overview and integration," *Educational Psychology*, Vol. 11, No. 3-4, 1991, pp. 193–215.

[240] Soane, E., Schubert, I., Lunn, R., and Pollard, S., "The relationship between information processing style and information seeking, and its moderation by affect and perceived usefulness: Analysis vs. procrastination," *Pers. Indiv. Differ.*, Vol. 72, 2015, pp. 72–78.

[241] Davidsson, S. and Alm, H., "Context adaptable driver information – Or, what do whom need and want when?" *Appl. Ergon.*, Vol. 45, No. 4, 2014, pp. 994–1002.

[242] Freyer, J., Deml, B., Maurer, M., and Färber, B., "ACC with enhanced situation awareness to reduce behavior adaptations in lane change situations," *IEEE IV*, IEEE, Istanbul, Turkey, Jun. 13-15 2007, pp. 999–1004.

[243] Inagaki, T., "Smart collaboration between human and machines based on mutual understanding," *Ann. Rev. Contr.*, Vol. 32, No. 2, 2008, pp. 253–261.

[244] Shanteau, J., "What does it mean when experts disagree?" *Linking expertise and naturalistic decision–making*, edited by E. Salas and G. Klein, Erlbaum, Mahwah, NJ, 2001, pp. 229–244.

[245] Klein, G., "A naturalistic decision making perspective on studying intuitive decision making," *J. Appl. Res. Mem. Cognit.*, Vol. 4, No. 3, 2015, pp. 164–168.

[246] Malakis, S. and Kontogiannis, T., "A sensemaking perspective on framing the mental picture of air traffic controllers," *Appl. Ergon.*, Vol. 44, No. 2, 2013, pp. 327–339.

[247] Inoue, S., Furuta, K., Nakata, K., Kanno, T., Aoyama, H., and Brown, M., "Cognitive process modelling of controllers in en route air traffic control," *Ergonomics*, Vol. 55, No. 4, 2012, pp. 450–464.

[248] Weiss, D. J. and Shanteau, J., "Empirical assessment of expertise," *Hum. Factors*, Vol. 45, No. 1, 2003, pp. 104–116.

[249] Weiss, D. J. and Shanteau, J., "Who's the best? A relativistic view of expertise," *Appl. Cogn. Psychol.*, Vol. 28, No. 4, 2014, pp. 447–457.

[250] Thomas, R. P., Willem, B., Shanteau, J., Raacke, J., and Friel, B. M., "Measuring performance of expert controllers: An application to air traffic control," *Proc. HFES 45th Annual Meeting*, Vol. 45, No. 4, Minneapolis, MN, Oct. 8-12 2001, pp. 286–290.

[251] Saaty, T. L., "How to make a decision: The analytic hierarchy process," *Eur. J. Oper. Res.*, Vol. 48, No. 1, 1990, pp. 9–26.

[252] Westin, C., *Investigating controller reliability in air traffic complexity judgments*, M.s. thesis, Dept. Mech. Mater. Manuf. Eng. Univ. Nottingham, Nottingham, UK, 2012.

[253] Galin, D. and Ornstein, R., "Individual differences in cognitive style-I. Reflective eye movements," *Neuropsychologia*, Vol. 12, No. 3, 1974, pp. 367–376.

[254] Witkin, H. A., Dyk, R. B., Faterson, H. F., Goodenough, D. R., and Karp, S. A., *Psychological differentiation*, Wiley, New York, NY, 1962.

[255] Maliko-Abraham, H., "A study of field independence and United States military air traffic controllers," *Proc. HFES 48th Annual Meeting*, Vol. 48, No. 11, New Orleans, LA, Sep. 20-24 2004, pp. 1213–1217.

[256] van Eck, R. N., Fu, H., and Drechsel, P. V. J., "Can simulator immersion change cognitive style? Results from a cross-sectional study of field-dependence-independence in air traffic control students," *J. Comput. High. Educ.*, Vol. 27, No. 3, 2015, pp. 196–214.

[257] Chen, J. Y. C. and Barnes, M. J., "Supervisory control of multiple robots: Effects of imperfect automation and individual differences," *Hum. Factors*, Vol. 54, No. 2, 2012, pp. 157–174.

[258] Bradshaw, J. M., Hoffman, R. R., Johnson, M., and Woods, D. D., "The seven deadly myths "autonomous systems"," *IEEE Intell. Syst.*, Vol. 28, No. 3, 2013, pp. 54–61.

[259] DSB, "Task Force Report: The role of autonomy in DoD systems," Rep., Defense Science Board (DSB), Washington, DC, Jul 2012.

[260] Ruff, H. A., Narayanan, S., and Draper, M. H., "Human interaction with levels of automation and decision-aid fidelity in the supervisory control of multiple simulated unmanned air vehicles," *Presence-Teleop. Virt.*, Vol. 11, No. 4, 2002, pp. 335–351.

[261] Regtuit, R. M., *Strategic conformal automation for air traffic control. A machine learning approach*, M.s. thesis, Fac. Aerospace Eng. Delft Univ. Techn., Delft, The Netherlands, 2016.

[262] Klein, G., Calderwood, R., and Clinton-Cirocco, A., "Rapid decision making on the fire ground," *Proc. HFES 30th Annual Meeting*, Vol. 30, No. 6, Dayton, OH, Sep. 29 - Oct. 3 1986, pp. 576–580.

[263] Klein, G., "A recognition-primed decision (RPD) model of rapid decision making," *Decision making in action: Models and methods*, edited by G. Klein, J. Orasanu, R. Calderwood, and C. E. Zsambok, Ablex Publishing, Norwood, NJ, 1993, pp. 138–147.

[264] Klein, G., "Naturalistic decision making," *Hum. Factors*, Vol. 50, No. 3, 2008, pp. 456–460.

[265] Petty, R. E. and Cacioppo, J. T., "The elaboration likelihood model of persuasion," *Adv. Exp. Soc. Psychol.*, Vol. 19, 1986, pp. 123–205.

[266] Kahneman, D. and Klein, G., "Conditions for intuitive expertise: A failure to disagree," *Am. Psychol.*, Vol. 64, No. 6, 2009, pp. 515–526.

[267] Brunswik, E., *Perception and the representative design of psychological experiments*, Univ. of California Press, London, UK, 1956.

[268] Gigerenzer, G. and Goldstein, D. G., "Reasoning the fast and frugal way: models of bounded rationality," *Psychol. Rev.*, Vol. 103, No. 4, 1996, pp. 650–669.

[269] Newell, K. M. and Hancock, P. A., "Forgotten moments: A note on skewness and kurtosis as influential factors in inferences extrapolated from response distributions," *J. Mot. Behav.*, Vol. 16, No. 3, 1984, pp. 320–335.

[270] Abe, G. and Richardson, J., "The influence of alarm timing on braking response and driver trust in low speed driving," *Safety Sci.*, Vol. 43, No. 9, 2005, pp. 639–654.

[271] Abe, G. and Richardson, J., "The influence of alarm timing on driver response to collision warning systems following system failure," *Behav. Inf. Technol.*, Vol. 25, No. 5, 2006, pp. 443–452.

[272] Navarro, J., Deniel, J., Yousfi, E., Jallais, C., Bueno, M., and Fort, A., "Influence of lane departure warnings onset and reliability on car drivers' behaviors," *Appl. Ergon.*, Vol. 59, Part A, 2017, pp. 123–131.

[273] Wan, J., Wu, C., and Zhang, Y., "Effects of lead time of verbal collision warning messages on driving behavior in connected vehicle settings," *J. Safety Res.*, Vol. 58, 2016, pp. 89–98.

[274] Rasmussen, J., "Risk management in a dynamic society: A modelling problem," *Safety Sci.*, Vol. 27, No. 2, 1997, pp. 183–213.

[275] Hollnagel, E., *Safety-I and Safety-II: The past and future of safety management*, CRC Press, Boca Raton, FL, 2014.

[276] Nunes, A. and Kirlik, A., "An empirical study of calibration in air traffic control expert judgment," *Proc. HFES 49th Annual Meeting*, Vol. 49, No. 3, Orlando, FL, Sep. 26-30 2005, pp. 422–426.

[277] Armstrong, J. S., "Judgmental bootstrapping: Inferring experts' rules for forecasting," *Principles of forecasting: A handbook for researchers and practitioners*, edited by J. S. Armstrong, chap. 6, Kluwer, Norwell, MA, 2001, pp. 169–192.

[278] Brochard, M., Weber, R., Lygeros, J., Vaccaro, C., Garcia, J., and Hollnagel, E., "ERASMUS final report: En route air traffic soft management ultimate system," Erasmus-approved-wp4-d4.6-v1.1, EUROCONTROL, Brussels, May 2009.

[279] Inagaki, T., "Adaptive automation: Sharing and trading of control," *Handbook of cognitive task design*, edited by E. Hollnagel, Erlbaum, Mahwah, NJ, 2003, pp. 145–169.

[280] Rouse, W. B. and Rouse, S. H., "A framework for research on adaptive decision aids," Tech. rep. afamrl-tr-83-082, Search Technol. Inc., OH, Oct. 1983.

[281] Scerbo, M. W., "Theoretical perspectives on adaptive automation," *Automation and human performance: Theory and applications*, edited by R. Parasuraman and M. Mouloua, Erlbaum, Hillsdale, NJ, 1996, pp. 37–63.

[282] Tricot, N., Sonnerat, D., and Popieul, J. C., "Driving styles and traffic density diagnosis in simulated driving conditions," *IEEE IV*, Versailles, France, Jun. 17-21 2002, pp. 298–303.

[283] Fitts, P. M., "The information capacity of the human motor system in controlling the amplitude of movement," *J. Exp. Psychol.*, Vol. 47, No. 6, 1954, pp. 381–391.

[284] BEA, "Final report on the accident on 1st June 2009 to the Airbus A330-203 registered F-GZCP operated by Air France flight AF 447 Rio de Janeiro–Paris," Tech. Rep. BEA f-cp090601, French Civil Aviation Safety Investigation Authority, Paris, France, Jul. 2012.

[285] DSB, "Crashed during approach Boeing 737-800 near Amsterdam Schiphol Airport 25 February 2009," Tech. Rep. M2009LV0225-01, The Dutch Safety Board, The Hague, The Netherlands, May 2010.

[286] Sarter, N. B. and Woods, D. D., "How in the world did we ever get into that mode? Mode error and awareness in supervisory control," *Hum. Factors*, Vol. 37, No. 1, 1995, pp. 5–19.

[287] Verberne, F. M. F., *Trusting a virtual driver: Similarity as a trust cue*, Ph.d. thesis, Dept. Ind. Eng. Innov. Sci., Eindhoven Univ. Technol., Eindhoven, The Netherlands, 2015.

[288] Steichen, B., Carenini, G., and Conati, C., "User-adaptive information visualization: Using eye gaze data to infer visualization tasks and user cognitive abilities," *Proc. 18th IUI*, ACM, Santa Monica, CA, 2013, pp. 317–328.

[289] Spahn, A., "And lead us (not) into persuasion?? Persuasive technology and the ethics of communication," *Sci. and Eng. Ethics*, Vol. 18, No. 4, 2012, pp. 633–650.

[290] Brynjolfsson, E. and McAfee, A., *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*, W. W. Norton & Company Inc., New York, NY, 1997.

[291] Gibson, J. J., *The ecological approach to visual perception*, Houghton Mifflin, Boston, MA, 1979.

[292] Vicente, K. J. and Rasmussen, J., "The ecology of human-machine systems II: Mediating 'direct perception' in complex work domains," *Ecol. Psychol.*, Vol. 2, No. 3, 1990, pp. 207–249.

[293] Ellerbroek, J., Brantegem, K. C. R., van Paassen, M. M., and Mulder, M., "Design of a coplanar airborne separation display," *IEEE Trans. Human–Mach. Syst.*, Vol. 43, No. 3, 2013, pp. 277–289.

[294] Madhavan, P., Wiegmann, D. A., and Lacson, F. C., "Automation failures on tasks easily performed by operators undermine trust in automated aids," *Hum. Factors*, Vol. 48, No. 2, 2006, pp. 241–256.

Abbreviations

AAM	Automation Acceptance Model
ACC	Area Control Center
ADS-B	Automatic Dependent Surveillance - Broadcast
AERA	Automated En-Route Air Traffic Control
AHP	Analytic Hierarchy Process
AI	Artificial Intelligence
ASM	Autonomous Squad Member
ATC	Air Traffic Control
ATCC	Air Traffic Control Center
ATM	Air Traffic Management
AUD	Automation Usage Decision
Aut	Automation (source)
BDT	Beliefs, Desires, Intentions
CASA	Computers Are Social Actors
CD&R	Conflict Detection and Resolution
CE	Cognitive Engineering
cfY	Conformal
cfN	Nonconformal
COCOS	Controllers' Strategies Integrated into a Conflict Resolution System
CORA	Controller Resolution Assistant
CPA	Closest Point of Approach
CWS	Cochran-Weiss-Shanteau
CX	Complexity
DS	Decision stage

EID	Ecological Interface Design
ELM	Elaboration Likelihood Model
HB	Heading Band
HCI	Human-Computer Interaction
HIPS	Highly Interactive Problem Solver
HMI	Human Machine Interface
Hum	Human (source)
HZ	Hertz
ICAO	International Civil Aviation Organization
IDT	Innovation-Diffusion Theory
ISAC	Intelligent System for Aircraft Conflict Resolution
LFV	Luftfartsverket (Swedish Civil Aviation Administration)
LOA	Level of Automation
MbC	Management by Consent
MbE	Management by Exception
MUFASA	Multidimensional Framework for Advanced SESAR Automation
N	Number of data samples
NextGen	Next Generation Air Transportation System
nmi	Nautical miles
PHARE	The Programme for the Harmonised ATM Research in Eurocontrol
PVD	Plan View Display
R/T	Radiotelephony
SAT	Situation Awareness-based Agent Transparency
SBQ	Source Bias Questionnaire
SBVAS	Source Bias Visual Analogue Scale
SESAR	Single European Sky ATM Research
SRK	Skills, Rules, Knowledge
SSD	Solution Space Diagram
TAM	Technology Acceptance Model
TCC	Terminal Control Center
TFT	Thin-Film Transistor
TRI	Triangle
TSL	Time to Separation Loss
URET	User Request Evaluation Tool
VAS	Visual Analogue Scale

Samenvatting

Strategic Conformance: Exploring Acceptance of Individual-Sensitive Automation for Air Traffic Control
Carl Albert Lennart Westin

NET als vele complexe en tijd-kritische werkdomeinen wordt de luchtverkeersleiding (ATC) geconfronteerd met een fundamentele modernisering die gekarakteriseerd wordt door het gebruik van geavanceerde automatisering (zoals voorzien in SESAR in Europa en NextGen in de Verenigde Staten). De bedoeling is om de relatie tussen luchtverkeersleider en automatisering, wat nu gebaseerd is op een rigide functieverdeling, te laten evolueren naar een meer vloeiende en onderling gecoördineerde teamrelatie. Hierin wordt verwacht dat de luchtverkeersleider vooral een toezichthoudende rol vervult, terwijl hij of zij veel van de tactische "*hands-on*" taken aan de automatisering overlaat. Van automatisering wordt verwacht dat het zal groeien in intelligentie en cognitieve vaardigheden en daarmee meer een teamlid zal worden die beslissingsondersteuning biedt en autonoom kan optreden. Deze veranderingen brengen echter één van de meest dringende en uitdagende "*human factors*" problemen met zich mee, namelijk: hoe kunnen we automatisering ontwerpen die geaccepteerd en vertrouwd wordt door de luchtverkeersleider?

Geautomatiseerde systemen die cognitief rijper worden, zullen waarschijnlijk op vele manieren gaan lijken op een individu met eigen gedrag en persoonlijkheid. Mensen zullen deze systemen waarschijnlijk dan meer gaan beschouwen als een mens en minder als een technologisch instrument. Deze trend blijkt uit huidige intelligente persoonlijke assistenten, zoals *Siri* van Apple, *Cortana* van Microsoft en de Google *Assistant*. In fictie hebben we ook andere mogelijke toekomstvisies gezien, zoals de geestige en sarcastische *TARS*-robot in de film *Interstellar*, het nieuwsgierige en verleidelijke *Samantha*-besturingssysteem in de film *Her*, en de rustige en geruststellende *HAL 9000* in de film *2001: A Space Odyssey*, die helaas ook lijdt aan paranoia.

Mensen kunnen moeite hebben met het begrijpen van dit soort systemen, niet alleen omdat hun gedachtegang ondoorzichtig is, maar ook omdat zij wellicht anders zullen redeneren. Dientengevolge kunnen gebruikers zowel het systeem als haar adviezen niet vertrouwen en het dus als geheel afwijzen. Dit soort wantrouwen kan het waargenomen acceptatieprobleem van adviessystemen in de luchtverkeersleiding gedeeltelijk verklaren. Om het acceptatieprobleem aan te pakken en het ATC automatiseringsontwerp te ondersteunen, is het visionaire MUFASA (Multidimensional Framework for Advanced SESAR Automation) project in het leven geroepen met als doel een raamwerk te ontwikkelen voor toekomstige automatiseringsniveaus. Het project heeft verondersteld dat beslissingsondersteuning voor conflictdetectie en resolutie (CD&R) *conform* moet zijn aan de voorkeurs strategie van een luchtverkeersleider, teneinde de acceptatie te verhogen en de samenwerking tussen mens en automatisering te verbeteren. *Strategic conformance* (strategische conformiteit) werd geïntroduceerd als een compatibiliteitsconcept dat specifiek de mate beschrijft waarin de *schijnbare* probleemoplossende stijl van een adviessysteem overeenkomt met die van de menselijke verkeersleider.

Dit proefschrift is een extensie van het succesvolle MUFASA-project richting empirisch onderzoek over individueel gecentreerde automatisering. In het algemeen richt dit proefschrift zich op geautomatiseerde beslissingshulpmiddelen die expliciete, geïndividualiseerde oplossingen (of adviezen) aanbieden voor stuurtaken in zeer dynamische tijd en veiligheid kritische werkdomeinen. Het ambitieuze doel was om een fundamenteel begrip te krijgen in hoeverre de acceptatie van adviessystemen werd beïnvloed door hoe goed de aangedragen adviezen overeenstemde met die van een individuele luchtverkeersleider.

Om *strategic conformance* empirisch te bestuderen is een nieuwe methode ontwikkeld, wat gebaseerd is op het opnemen en afspelen van luchtverkeersleiders' eigen oplossingen en deze te vermommen als zijnde geautomatiseerde adviezen. Er werd dus geen algoritme gebruikt om conflictoplossingen te berekenen. Om de unieke oplossing stijl van een luchtverkeersleider vast te stellen, werd de luchtverkeersleider onbewust vier keer blootgesteld aan hetzelfde verkeersscenario en conflict in een "*prequel*" simulatie. Oplossingen werden vervolgens geanalyseerd, verpakt en gepresenteerd als conflict resolutie advies. In de daaropvolgende simulaties werd de *strategic conformance* gevarieerd door dezelfde luchtverkeersleider bloot te stellen aan haar/zijn eigen oplossing voor hetzelfde conflict (*conformal*) of een tegenstrijdige oplossing van een collega voor hetzelfde conflict (*non-conformal*).

Dit proefschrift onderzocht specifiek de *strategic conformance* effecten in het kader van selectie en implementatie van beslissingen en adviezen. Daarin werden luchtverkeersleiders ondersteund door het Solution Space Diagram (SSD), een prototype display dat een hoge mate van informatieverzameling en -integratie aanbiedt wat nodig was voor het faciliteren van hogere niveaus van besluitvorming. Terwijl het SSD zelf geen specifieke conflictoplossingen adviseert, vergemakkelijkt het de implementatie en presentatie van adviezen ten behoeve van de *strategic conformance* manipulatie. Het SSD is namelijk een ecologisch hulpmiddel die verschillende kritische parameters van het CD&R-probleem op een visuele manier integreert. Het werd gebruikt in alle simulaties als onderdeel van dit proefschrift, en verscheen wanneer een vliegtuig werd geselecteerd. Als zodanig vertegenwoordigde de simulaties een futuristische omgeving waarin, in tegenstelling tot de huidige ATC-operaties, meer verkeer aanwezig was, digitale datalink-communicatie beschikbaar was tussen de luchtverkeersleider en piloten en het SSD als hulpmiddel in CD&R gebruikt werd. Verder werden alle simulaties en conflictoplossingen beperkt tot het horizontale vlak.

Dit proefschrift beschrijft drie *human-in-the-loop* (mens-in-de-lus) studies. Daarnaast werden simulatiedata post-hoc geanalyseerd in een vierde studie. De eerste empirische studie culmineerde in een grootschalige real-time simulatie met zestien ervaren luchtverkeersleiders. De studie varieerde *strategic conformance* (d.w.z. *conformal* of *non-conformal* adviezen) naast de autoriteit van het systeem (d.w.z. besturing op basis van toestemming of uitzondering) en de taakcomplexiteit (hoog of laag). Resultaten gaven aan dat luchtverkeersleiders *conformal* adviezen vaker accepteerden (d.w.z. adviezen op basis van hun eigen unieke conflictoplossing stijl), in hogere mate eens waren met de adviezen en sneller reageerden op het advies dan bij *non-conformal* adviezen op basis van een collega's contrasterende, maar nog steeds werkbare en veilige oplossing. Verassend genoeg waren verkeersleiders het in 25% van de gevallen oneens met hun eigen conformal adviezen. De overige twee studies, en de post-hoc analyse, werden uitgevoerd om verdere aannemelijke oorzaken voor deze waargenomen meningsverschillen te onderzoeken, samen met overige vragen uit de eerste studie. De twee resterende human-in-the-loop studies herhaalden de experimentele aanpak uit de eerste studie, weliswaar met kleine verfijningen en aanpassingen.

De *Source bias* studie onderzocht in hoeverre de acceptatie van en vertrouwen in automatisering als hulpmiddel voor conflictoplossing beïnvloedt werd door vermoedelijke bron (mens of computer) van een gegeven advies. Vijf ervaren luchtverkeersleiders hebben deelgenomen aan een simulatie waarin de *strategic conformance* van een advies varieerde tezamen met de vermoedelijke bron van het advies. Terwijl responsies van subjectieve vragenlijsten een lichte voorkeur voor de menselijke adviseur aangaven, was dat niet zichtbaar in de objectieve simulatieresultaten.

De *Automation transparency* studie onderzocht, naast strategic conformance, effecten van het geboden inzicht in de adviezen (middels het interface) op luchtverkeersleiders' aanvaarding en begrip van de adviezen. Negen luchtverkeersleiders in opleiding hadden deelgenomen aan een real-time simulatie, waarin twee niveaus van transparantie werden gebruikt: *heading band* SSD (lage transparantie) en *triangle* SSD (hoge transparantie), die meer zinvolle informatie verstrekt over de relatie tussen conflicterende vliegtuigen. Resultaten toonden aan dat de transparantere *triangle* SSD beter begrepen werd. Hoewel er geen interactie tussen *conformance* en transparantie werd gevonden, werden *conformal* adviezen iets vaker geaccepteerd dan *non-conformal* adviezen, wat ook de resultaten van de eerste studie ondersteunt. Bovendien, bij het gebruik van de *triangle* SSD, werden conflicten vaker opgelost door koerswijzigingen en combinaties van koers en snelheid. Dit geeft aan dat de oplossingen van luchtverkeersleiders afhankelijk zijn van hoe een conflict op het interface wordt gepresenteerd. Aangezien *conformal* (en *non-conformal*) adviezen gebaseerd waren op oplossingen bij het gebruik van de *heading band* SSD, zijn deze adviezen mogelijk niet representatief genoeg geweest alziende *conformal* zodra de *triangle* SSD werd gebruikt.

De vierde en laatste *Consistency* studie analyseerde alle oplossingen die door de luchtverkeersleiders werden gegeven in de bovenstaande drie empirische studies. Het doel was om te bepalen of afwijzing van adviezen mogelijk verklaard kon worden door inconsistenties in de oplossingen van luchtverkeersleiders. De studie onderzocht de mate waarin luchtverkeersleiders de herhaalde conflicten (vier herhalingen) in de loop der tijd (*intra-rater* variabiliteit) consequent hadden opgelost, en in hoeverre zij onderling over oplossingen (*inter-rater* variabiliteit) overeenstemming hebben bereikt. Op basis van een herziening van ATC-conflictoplossing strategieën werd een classificatie model ontwikkeld waarmee de oplossingen van luchtverkeersleiders objectief en kwalitatief konden worden beschreven. Uit de resultaten bleek dat luchtverkeersleiders consistent waren, maar het onderling niet eens waren over hoe de conflicten opgelost konden worden. Echter, consistentie was beperkt tot besluitvorming op een hoger niveau, namelijk om een vliegtuig voor- of achterlangs een ander vliegtuig te sturen of met beide of slechts één vliegtuig interactie te hebben. Luchtverkeersleiders waren inconsistent ten opzichte van meer gedetailleerde oplossingsparameters, zoals de richting van een oplossing (bijvoorbeeld naar links of rechts sturen) en de exacte waarde (bijvoorbeeld een koerswijziging van exact 35 graden). Consistentie en overeenstemming waren ook niet veel beter voor conflicten die een bepaald soort oplossing bevoordeelden. Een verschil werd echter wel gevonden met betrekking tot de algemene oplossingsstrategie. Met conflicten die de oplossing in een bepaalde richting duwden, heeft de meerderheid van de luchtverkeersleiders overeenstemming bereikt over een gedeelde oplossingsgeometrie, terwijl bij het onbevooroordeelde conflict de oplossingen conform het classificatie model waren. Ervaren luchtverkeersleiders waren echter iets meer consistent dan de trainees in termen van het classificatie model.

Samengevat heeft dit proefschrift niet alleen bijgedragen aan kennis en empirisch inzicht over wat de acceptatie drijft van automatiseringsadviezen in de luchtverkeersleiding, maar ook over de samenwerking tussen mens en machine in het algemeen. Empirische resultaten toonden namelijk aan dat *conformal* ATC-automatisering de acceptatie van systeemadviezen kan verbeteren, en de responsietijd kan verlagen. Deze voordelen werden waargenomen op verschillende niveaus van expertise, met name ten opzichte van ervaren verkeersleiders. *Strategic conformance* kan daarmee dus het meest voordelig zijn bij de invoering van nieuwe geautomatiseerde beslissingshulpmiddelen, om zodoende initiële acceptatie te krijgen.

De ontwikkeling van conformal automatisering en andere gepersonaliseerde beslissingsondersteuning vereist dat de operator enigszins consistent is in zijn/haar oplossingen. Het ontwerpen van gepersonaliseerde systemen vereist echter ook een ethische overweging, omdat deze systemen de macht hebben om de acceptatie en het vertrouwen te benvloeden, onafhankelijk van de werkelijke prestaties en betrouwbaarheid van het systeem. Terwijl technologische vooruitgang het mogelijk maakt om het systeem zichzelf automatisch te laten aanpassen aan persoonlijke voorkeuren, behoeften en vaardigheden van een individu, zijn er verschillende technische uitdagingen die moeten worden overwonnen alvorens echt *conformal* automatisering kan worden ontwikkeld. De belangrijkste uitdaging is hoe men de unieke voorkeuren en oplossingsstijl van individuen kan identificeren en modelleren. Verder onderzoek is nodig om de juiste parameters vast te stellen die de probleemoplossende stijl van een persoon zouden karakteriseren.

Automatiseringsontwerpers in een specifiek werkdomein moeten goed overwegen welk doel een *conformal* systeem zou dienen voordat men dit introduceert. Veel werkdomeinen, zoals een vliegtuigcockpit, zijn meer gebaad bij homogeniteit, waarin individuele verschillen juist reduceert worden. Aanvullend aan *conformal* automatisering die specifieke adviezen verstrekt, kan ook het interface de variabiliteit in probleemoplossende stijlen ondersteunen. *Ecological Interface Design* kan bijvoorbeeld worden gebruikt om gepersonaliseerde probleemoplossingen mogelijk te maken door de ‘objectieve waarheid’ te visualiseren. Door alle mogelijkheden en beperkingen te laten zien die een situatie beïnvloeden, kan de operator een probleem zelf op haar of zijn gewenste manier oplossen.

Terugkerend naar de artificial intelligence (AI) aangedreven persoonlijke assistenten. Hoewel deze nog in de kinderschoenen staan (bijvoorbeeld Siri, Cortana, Assistant), zullen ze naar verwachting de interactie tussen mens en automatisering aanzienlijk veranderen richting een vloeiende samenwerking die vergelijkbaar is met de futuristische beelden van geautomatiseerde personages (bijvoorbeeld TARS, Samantha en HAL 9000). Hoewel de acceptatie van deze systemen afhankelijk is van veel aspecten, kan de compatibiliteit met de mens, en niet in de laatste plaats hun *conformance*, een significant effect hebben op hoe bereidwillig men is om hiermee te willen werken en het zullen accepteren en vertrouwen. Hoewel de mate waarin acceptatie en vertrouwen afhangt van de *strategic conformance* van de intelligente assistent nog steeds wordt bestudeerd, geeft dit proefschrift de mogelijke voordelen daarvan aan, in ieder geval in het kader van luchtverkeersleiding. De verleiding om automatisering te personaliseren, gewoonweg omdat het kan, moet echter vermeden worden. Met name in veiligheid kritische domeinen moeten zowel de voordelen als de nadelen van dergelijke systemen goed overwogen en geëvalueerd worden alvorens men overgaat tot implementatie.

Publications

The work of this thesis has resulted in the following publications. Papers are listed in reversed chronological order.

Journal publications

Westin, C., Borst, C., & Hilburn, B., "Strategic conformance: Overcoming acceptance issues of decision aiding automation?" *IEEE Transactions on Human-Machine Systems*, Vol. 46, No. 1, 2016, pp.41-52.

Hilburn, B., **Westin, C.**, & Borst, C., "Will controllers accept a machine that thinks like they think? The role of strategic conformance in decision aiding automation," *Air Traffic Control Quarterly*, Vol. 22, No. 2, 2014, pp. 115-136.

Conference publications

Westin, C., Borst, C., & Hilburn, B., "Automation transparency and personalized decision support: Air traffic controller interaction with a resolution advisory system," *13th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Human-Machine Systems (HMS)*, Kyoto, Japan, Aug. 30-Sep. 2, 2016.

Westin, C., Hilburn, B., & Borst, C., "Air traffic controller decision-making consistency and consensus in conflict solution performance," 5th Challenges in European Aerospace (CEAS) Air & Space Conference, Delft, The Netherlands, Sep. 7-11, 2015.

Westin, C., Borst, C., & Hilburn, B., "An empirical investigation into three underlying factors affecting automation acceptance," *5th SESAR Innovation Days (SID)*, Bologna, Italy, Dec. 1-3, 2015.

Hilburn, B., **Westin, C.**, & Borst, C., "Strategic conformance: An important concept in future automation design?" *6th International Conference on Research in Air Transportation (ICRAT)*, Istanbul, Turkey, May 26-30, 2014.

Göritzlehner, R., Borst, C., Ellerbroek, J., **Westin, C.**, van Paassen, M. M., & Mulder, M., "Effects of transparency on the acceptance of automated resolution advisories," *IEEE International Conference on Systems, Man, and Cybernetics (SME)*, San Diego, CA, Oct. 5-8, 2014.

Westin, C., Hilburn, B., & Borst, C., "The effect of strategic conformance on acceptance of automated advice: Concluding the MUFASA project," *3rd SESAR Innovation Days (SID)*, Stockholm, Sweden, Nov. 26-28, 2013.

Westin, C., Borst, C., & Hilburn, B., "Mismatches between automation and human strategies: An investigation into future air traffic management decision aiding," *17th International Symposium on Aviation Psychology (ISAP)*, Dayton, OH, May 6-9, 2013.

Borst, C., **Westin, C.**, & Hilburn, B., "An investigation into conflict detection and resolution strategies in air traffic management," *2nd SESAR Innovation Days (SID)*, Braunschweig, Germany, Nov. 27-29, 2012.

Westin, C., Hilburn, B., & Borst, C., "Mismatches between automation and human strategies: An investigation into future air traffic management (ATM) decision aiding," *1st SESAR Innovation Days (SID)*, Toulouse, France, Nov. 29-Dec. 1, 2011.

Acknowledgments

While my name on the cover of this thesis hints that I am the only author of this work, this is certainly not the case. This thesis would not have been possible without the inspiration, encouragements, and support of supervisors, colleagues, friends, and family.

My first roller-coaster ride provides an appropriate metaphor capturing my PhD journey. Standing at a distance the roller-coaster rises above the other entertainments as the crown of what the amusements park (i.e., University) has to offer. When embarking the ride, you have all these expectations of excitement and fear, but you really have a poor clue of what the ride entails (i.e., the PhD experience and the thesis). The ride starts with a long and steep dreaded climb (which goes on forever) that eventually brings you to an edge, after which you catapult downhill in a basic free-fall. The built-up momentum is needed to ascertain that you make it through the following loops and spins before the ride ends. You step out of the car, not entirely sure how you are feeling. A part of you would like to continue, while another part swear to never do it again.

At times, in facing challenges and setbacks, I have doubted myself and felt that the PhD never will be completed. I realized too late that a “perfect” thesis does not exist, and that a PhD is like a work of art: it is never finished, merely abandoned. In contrast, in light of praise and accolades when work has processed easy, I have been absurdly confident that I will be able to wrap up my PhD and this thesis in no time. Not so surprising, perhaps, that keeping deadlines (thankfully most of them internally) has been one of the most profound challenges of my PhD work, and something that I am still struggling with. To all who have have worked with me, thank you equally for you patience and pressure.

To Max Mulder, my supervisor: I'm forever indebted to you for taking a chance and entrusting me with a PhD. I remember the walk and discussion we had in a London park after the MUFASA kickoff in April 2011 as the very start of my doctoral journey with Delft. Thank you for giving me the opportunity to pursue my PhD, and for providing the intellectual and material resources and colleagues needed to achieve it. In the words of the famous chef: "hurdue durdu bork bork bork"

To Clark Borst, my co-supervisor: While this marks the end of an incredibly rewarding, fun, and successful collaboration between us, the three stooges (you, Brian, myself), it opens up new possibilities for future endeavors. I am extremely, almost inappropriately, thankful for the invaluable guidance and enlightenment you have afforded me over the years. It will be arduous to journey further without our regular discussions, your constructive critique and feedback, your motivational praise, and your self-sacrificing support. Although I hope we never again have to endure an all-nighter to make a journal deadline!

To Brian Hilburn, my mentor and eternal cheerleader: My guess would be that you, somehow, foresaw this many years ago, maybe even at around the time we first got to work together. I owe a great part of this accomplishment to the motivation, enthusiasm, and encouragement that you have invested in my academic and research career. I remember clearly how enthusiastic I was when we first met in Ljungbyhed over ten years ago. I was provided an opportunity to work with this really cool and recognized human factors researcher. Little did I know that you would become a lifelong mentor that not only inspired me to start pursuing a PhD, but also saw me through it. On that same note, thank you Jessica and Cassidy for the encouraging support and for welcoming me to the family with open arms.

Another dear friend who bears great responsibility for inspiring my academic endeavors is Roel van Winsen. Thank you Roel for engaging in endless scientific and philosophical debates (from utterly nonsense to some promising ideas) and teaching me how to think critically and ask good questions.

A huge thanks to all air traffic controllers and coordinators at Shannon (IAA), Ireland, and Norrköping and Arlanda (LFV), Sweden for facilitating and participating in numerous time-consuming experiments and simulations. This thesis would not have been possible without you contribution. Thank you Dirk Schaefer for your sincere interest and support of the MUFASA project.

To my dear love Johanna Andersson and my family: Thank you for all the heartwarming support over the years. Especially you, Johanna, can attest to the roller-coaster spectrum of deep lows and high highs that this PhD journey has been. Thank you for keeping me grounded and dragging me back up to the surface when I have gotten myself submerged in PhD work. I apologize for my periodically unfair manners and seemingly unexpected and bizarre behaviors over the years. I love you.

Also, my sincere gratitude to Shannon & Clark, Eline & Michiel, and Emma & David who have provided me shelter and comfort during my visits to Delft. Finally, I would like to thank all my friends over the world and my fellow pilots and other colleagues at Norwegian Air Shuttle who have shown interest and supported my studies over the years.

Curriculum Vitae

Carl Westin was born on January 26th, 1982, in Huddinge, Sweden. His mother is Dutch and his father is Swedish. From 1998 to 2002 he attended secondary school at Sjödalsgymnasiet in Huddinge. During this period he was part of starting and running a business, Proglobe UF, during the last academic year as part of the Junior Achievement (JA) educational program in Sweden. Proglobe UF won the Swedish JA championship and came second in the European JA championship in 2002.

Aspiring to fulfil his childhood dream, he enrolled 2005 as a student the School of Aviation at Lund University (Ljungbyhed, Sweden). In November 2006 he graduated a commercial air traffic pilot with a University Diploma. Directly following graduation he joined Sterling Airlines (Copenhagen, Denmark) 2007 as a First Officer on the Boeing 737NG. In 2009, following Sterling Airlines bankruptcy in 2008, he joined Maximus Air Cargo (Abu Dhabi, United Arab Emirates) as a First Officer on the Airbus 300. In 2010 he left the Middle East and moved closer to home when joining Norwegian Air Shuttle (Oslo, Norway) as a First Officer, again on the Boeing 737NG. While initially based in Helsinki, Finland, and then Trondheim, Norway, he has been working from home in Stockholm since 2013.

Parallel to his flying, Carl has pursued an academic career that started in 2006. His studies at the School of Aviation required an university thesis, which he wrote on the subject of human factors challenges in unmanned aerial vehicles. One of the supervisors, PhD Brian Hilburn, subsequently hired Carl as an aviation/ATM research consultant to work for the Center for Human Performance Research (CHPR). During the following years Carl participated in several EUROCONTROL-led research projects.

With a growing interest into research he soon started exploring academic options. In 2008 he started his part time M.Sc. studies in Applied Ergonomics and Human Factors at the school of Mechanical, Materials and Manufacturing Engineering, Nottingham University (United Kingdom). For his thesis he investigated air traffic controllers' reliability in complexity judgments. He obtained his MSc degree (with distinction) in July 2012.

In 2011 he (as a CHPR researcher) was part of an international consortium awarded research funds as part of the EUROCONTROL-led SESAR WP-E long term and innovative research program. The MUFASA project explored how ATM automation acceptance and usage was affected by three factors: traffic complexity, level of automation, and strategic conformance. Being highly successful, the MUFASA project went on to be extended for a second round, ultimately ended in 2015. The project brought him close to Delft University of Technology (The Netherlands), one of the consortium members, and in 2013 he enrolled as a PhD student at the Faculty of Aerospace Engineering. It provided an excellent opportunity to combine the MUFASA project and extend it with a PhD.

Today he is still working as a pilot for Norwegian Air Shuttle, with over 5,000 jet hours. Furthermore, he is still employed with CHPR. Since 2017 he is also employed as a research engineer with Linköpings University, the division of Media and Information Technology at the Department of Science and Technology. He is currently working with self-explanatory automation and eye-tracking technologies for the purpose of facilitating new and effective training concepts in highly automated safety-critical domains.