A Cognitive Conversational Agent for Training Child Helpline Volunteers

Al Owayyed, Mohammed; Despan, Alex; Tielman, Myrthe; Brinkman, Willem Paul

DOI
10.1145/3652988.3696197

Publication date
2024

Document Version
Final published version

Published in
Proceedings of the 24th ACM International Conference on Intelligent Virtual Agents, IVA 2024

Check for updates

A Cognitive Conversational Agent for Training Child Helpline Volunteers

Mohammed Al Owayyed*
Delft University of Technology
The Netherlands
King Saud University
Saudi Arabia
M.AlOwayyed@tudelft.nl

Alex Despan
Delft University of Technology
The Netherlands
a.despan@student.tudelft.nl

Myrthe Tielman
Delft University of Technology
The Netherlands
m.l.tielman@tudelft.nl

Willem-Paul Brinkman
Delft University of Technology
The Netherlands
W.P.Brinkman@tudelft.nl

Abstract

Child helplines offer a safe and private space for children to share their thoughts and feelings with volunteers. However, training these volunteers to help can be both expensive and time-consuming. In this demo, we present Lilobot, a conversational agent designed to train volunteers for child helplines. Lilobot's reasoning is based on the Belief-Desire-Intention (BDI) model, which simulates, for example, a bullied child who contacts the helpline through text. Users engage with Lilobot in a role-play format, taking on the volunteer's role. Through this system, volunteers can practice applying the Five Phase Model, a conversational strategy helplines use. The training tool includes a trainer interface for monitoring and modifying Lilobot's interactions. Trainers can also create new conversational scenarios through an authoring tool. An initial evaluation led to enhancements in Lilobot's knowledge base and intent recognition, addressing the main issues encountered by participants. The components used to implement the system were Java Spring for the BDI model and the authoring tool, Rasa for Natural Language Understanding, PostgreSQL for the database, and Vue.js for the front-end. This tool aims to provide volunteers with consistent, interactive training, enhancing their counselling skills in a controlled environment.

CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; • **Applied computing** → *Interactive learning environments*; • **Computing methodologies** → **Intelligent agents**.

Keywords

BDI, Training Simulation, Conversational Agents, Chatbot, Interactive Agents, Children Helpline

IVA '24, September 16–19, 2024, GLASGOW, United Kingdom

ACM ISBN 979-8-4007-0625-7/24/09
https://doi.org/10.1145/3652988.3696197

ACM Reference Format:
Mohammed Al Owayyed, Alex Despan, Myrthe Tielman, and Willem-Paul Brinkman. 2024. A Cognitive Conversational Agent for Training Child Helpline Volunteers. In *ACM International Conference on Intelligent Virtual Agents (IVA '24), September 16–19, 2024, GLASGOW, United Kingdom.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3652988.3696197

1 Introduction

Children worldwide face issues that require emotional and social support, which child helplines can offer in a safe and private environment [6]. Children can reach out to helplines by talking to a volunteer, typically through phone or a web-based chat interface. The volunteer's job is to counsel the child by applying specific communication strategies. A prominent strategy is the Five Phase Model [16]. This model supports the conversation's dynamics while ensuring that it remains child-centred. The first phase of the model is to build rapport with the child, then clarify their story, set the session's goal, work towards that goal, and finally conclude the conversation with the child.

Usually, volunteers are trained to apply the Five Phase Model through role-playing sessions, where a supervisor portrays a child. Although role-plays are useful and effective, they can be costly and time-consuming to set up and run, issues that an interactive agent can address [15]. Such an agent can be readily available and consistent, allowing repetitive practice opportunities in a safe, controlled environment [7, 13]. To this end, this demo shows a tool to train new volunteers to apply the Five Phase Model. This training tool mainly includes a chatbot simulation of a child called Lilobot. The volunteer's aim is to counsel Lilobot by applying the communication strategy correctly. Furthermore, the tool integrates a trainer interface that allows a supervisor to control Lilobot and craft new conversational agents to include in the training curriculum.

2 Training System Design

2.1 Lilobot's Design

Lilobot is a conversational agent that simulates a child who has been bullied at school [1]. Due to this experience, Lilobot contacts a children's helpline through a chat interface to discuss the bullying issue. We developed Lilobot's scenario and conversations in collaboration with the Dutch Child Helpline "De Kindertelefoon," where the scenario demonstrates behaviours like social anxiety, depression, and loneliness [3, 10]. Additionally, we chose the name "Lilobot" to avoid gender biases and stereotypes [5]. In this training scenario, the volunteer aims to assist the conversational agent in discussing the bullying with someone they trust, such as their parents or teachers.

The agent's decision-making process is based on the Belief-Desire-Intention (BDI) model [14], which has been used to simulate various interactive agents for training purposes (e.g., [2, 4]). Lilobot integrates BDI into its conversational model to represent the simulated child's state, based on which responses are chosen. Thus, this BDI model reflects the agent's 'mind,' which is a common concern for intelligent virtual agents [9, 11].

The agent holds 17 different beliefs about itself and the world (e.g., "Lilobot thinks the volunteer can be trusted"). Each belief has a numerical value ranging from 0 to 1, indicating its intensity. The belief values are influenced by a predefined mapping with the volunteer's inputs recognized through a Natural Language Understanding component. Lilobot's desires align with the Five Phase Model, representing a child's goals during the phases e.g., "Lilobot wants to talk about their problem", while the intentions reflect the current active desire in the conversation.

Lilobot's reasoning flow starts when a volunteer writes an input. Once this input is recognized, the system updates the corresponding belief values. Then, the system evaluates whether the current beliefs align with a (new) desire rule by checking if the belief values meet the desire's threshold values. If they do, the system adopts this new desire. If no new desire is triggered, the current desire remains active. Lilobot decides on an appropriate response based on the last volunteer's input and the current active desire. Using these two criteria, the agent retrieves a suitable response from the knowledge base and sends it to the volunteer.

The conversation can unfold in two ways: either the volunteer successfully navigates the Five Phase Model, or the agent perceives the interaction as unhelpful or intimidating and leaves. We modelled the latter to show the possible consequences of deviating from the Five Phase Model. At the end of a session, feedback is provided to the volunteer through a conversation transcript and a summary of belief changes, helping in understanding the child's thinking and the impact of applying the Five Phase Model. Figure 1 shows an example of the interaction with the agent, where a volunteer successfully navigates from phase three, i.e., setting up the conversation goal, to the start of phase four, i.e., working toward the goal.

We conducted a pilot evaluation of the initial Lilobot prototype [8]. Although participants highlighted the need for the agent to understand questions better and vary its responses, many were optimistic about using Lilobot as an additional training tool with some improvements. Based on this feedback, we expanded Lilobot's knowledge base and improved its intent recognition to cover more topics and possible user inputs.

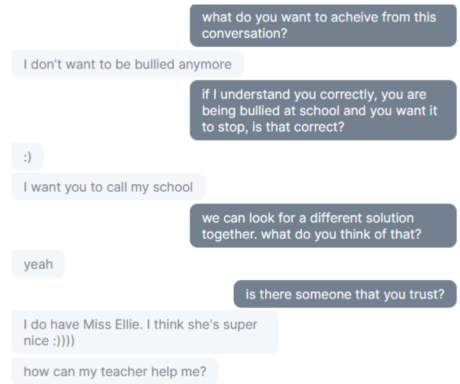

Figure 1: An example of a chat between a volunteer and Lilobot, showing the transition between phases 3 and 4.

2.2 Building a Training Environment

To enrich the training experience, we developed a trainer interface. Through it, trainers, e.g., supervisors, can monitor and control Lilobot's interactions, allowing them to view and dynamically modify the numerical values of beliefs during a volunteer's conversation with Lilobot or update Lilobot's utterances. Additionally, trainers can use an authoring tool to craft new scenarios. This tool enables trainers to add and edit simulated children scenarios directly. Therefore, even users unfamiliar with coding can create and update scenarios based on the volunteers' educational needs. To create a new scenario, trainers need to define three elements: 1) the simulated child's BDI values and rules, 2) possible volunteer inputs in the NLU model, and their mapping to changes in beliefs, and 3) the simulated child's replies, mapped to the desires and volunteer inputs, from which the agent can retrieve an appropriate response.

In addition to Lilobot's scenario, we created 11 other scenarios to support learning through different experiences. These scenarios are similar in their background scenario, i.e., being bullied, but differ in three aspects. First, they have slightly different BDI starting values, requiring volunteers to perform differently for the scenario to progress, e.g., show less or more empathy. Second, the scenario settings vary, such as being bullied at school, in the neighbourhood, or at a football match. Third, there are differences in the simulated child's misconceptions in the conversation, e.g., wanting the volunteer to call the school or seeking revenge on the bullies.

Implementation-wise, we used Java Spring to implement the BDI reasoning and the authoring tool. For the NLU components, we used Rasa to recognize volunteer inputs. PostgreSQL was used as the database to store the agent's BDI states and knowledge base, while Vue.js handled the program's front-end.

3 Conclusions and Future Work

The tool provides volunteers with interactive training through a simulated child, enhancing their application of the Five Phase Model. Based on an initial evaluation, we updated Lilobot to address issues related to its interpretation and response capabilities. The educational tool allows trainers to tailor Lilobot's responses for individual volunteers and author new conversational agent scenarios to augment training. In the future, we plan to extend the system to provide personalized feedback and guidance, and to integrate an emotional model into the BDI [12]. Additionally, we will explore how Large Language Models can improve the BDI model and the NLU components. We also plan to conduct a comprehensive evaluation to assess the training outcomes of the tool.

Acknowledgments

Lilobot was based on Sharon Grundmann's MSc thesis [8]. We want to thank De Kindertelefoon for their invaluable help in the design, as well as Merijn Bruijnes, Sharon Grundmann, Fedde de Haas, Veena Madhu, and Tristan Schmidt for their valuable contributions.

Video material & source code

A video showing a sample interaction with Lilobot and the authoring tool can be found here:

https://surfdrive.surf.nl/files/index.php/s/DsUx0URwRC3sP1A

The source code for Lilobot is available here:

https://github.com/alowayyedm/LilobotTraining

References

[1] Mohammed Al Owayyed, Sharon Grundmann, Merijn Bruijnes, and Willem-Paul Brinkman. 2023. Training child helpline counsellors with a BDI-based conversational agent. In *BNAIC/BeNeLearn 2023: Joint International Scientific Conferences on AI and Machine Learning*.

[2] Keith Anderson, Elisabeth André, Tobias Baur, Sara Bernardini, Mathieu Chollet, Evi Chryssafidou, Ionut Damian, Cathy Ennis, Arjan Egges, Patrick Gebhard, et al. 2013. The TARDIS framework: intelligent virtual agents for social coaching in job interviews. In *International Conference on Advances in Computer Entertainment Technology*. Springer, 476–491.

[3] Michel Boivin and Shelley Hymel. 1997. Peer experiences and social self-perceptions: a sequential model. *Developmental psychology* 33, 1 (1997), 135.

[4] Merijn Bruijnes. 2016. *Believable suspect agents: response and interpersonal style selection for an artificial suspect.* Ph. D. Dissertation. University of Twente.

[5] Miruna-Valeria Craiut and Ioana Raluca Iancu. 2022. Is technology gender neutral? A systematic literature review on gender stereotypes attached to artificial intelligence. *Human Technology* 18, 3 (2022), 297–315.

[6] Ruben Georges Fukkink, Suzan Bruns, and Rudy Ligtvoet. 2016. Voices of children from around the globe; An international analysis of children's issues at child helplines. *Children & Society* 30, 6 (2016), 510–519.

[7] Jonathan Gratch, David DeVault, and Gale Lucas. 2016. The benefits of virtual humans for teaching negotiation. In *International Conference on Intelligent Virtual Agents*. Springer, 283–294.

[8] Sharon Grundmann. 2022. A bdi-based virtual agent for training child helpline counsellors.

[9] Patrick Kenny, Arno Hartholt, Jonathan Gratch, William Swartout, David Traum, Stacy Marsella, and Diane Piepol. 2007. Building interactive virtual humans for training environments. In *Proceedings of i/itsec*, Vol. 174. 911–916.

[10] Anat Brunstein Klomek, Frank Marrocco, Marjorie Kleinman, Irvin S Schonfeld, and Madelyn S Gould. 2007. Bullying, depression, and suicidality in adolescents. *Journal of the American Academy of Child & Adolescent Psychiatry* 46, 1 (2007), 40–49.

[11] Michael Kriegel, Ruth Aylett, Pedro Cuba, Marco Vala, and Ana Paiva. 2011. Robots meet ivas: A mind-body interface for migrating artificial intelligent agents. In *Intelligent Virtual Agents: 10th International Conference, IVA 2011, Reykjavik, Iceland, September 15-17, 2011. Proceedings 11*. Springer, 282–295.

[12] Dongxu Lu. 2023. Emotion Model for Child Helpline Training Tool. (2023).

[13] Samuel Mascarenhas, Joao Dias, Rui Prada, and Ana Paiva. 2010. A dimensional model for cultural behavior in virtual agents. *Applied Artificial Intelligence* 24, 6 (2010), 552–574.

[14] Anand S Rao, Michael P Georgeff, et al. 1995. BDI agents: From theory to practice.. In *Icmas*, Vol. 95. 312–319.

[15] Luke Rogers. 2011. Developing simulations in multi-user virtual environments to enhance healthcare education. *British Journal of Educational Technology* 42, 4 (2011), 608–615.

[16] Trine Natasja Sindahl. 2011. *Chat Counselling for Children and Youth - A Handbook.*