

**Incomplete Gamma Kernels
Generalizing Locally Optimal Projection Operators**

Stotko, Patrick; Weinmann, Michael ; Klein, Reinhard

DOI

[10.1109/TPAMI.2024.3349967](https://doi.org/10.1109/TPAMI.2024.3349967)

Publication date

2024

Document Version

Final published version

Published in

IEEE Transactions on Pattern Analysis and Machine Intelligence

Citation (APA)

Stotko, P., Weinmann, M., & Klein, R. (2024). Incomplete Gamma Kernels: Generalizing Locally Optimal Projection Operators. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(6), 4075-4089. <https://doi.org/10.1109/TPAMI.2024.3349967>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Incomplete Gamma Kernels: Generalizing Locally Optimal Projection Operators

Patrick Stotko , Michael Weinmann , and Reinhard Klein 

Abstract—We present incomplete gamma kernels, a generalization of Locally Optimal Projection (LOP) operators. In particular, we reveal the relation of the classical localized L_1 estimator, used in the LOP operator for point cloud denoising, to the common Mean Shift framework via a novel kernel. Furthermore, we generalize this result to a whole family of kernels that are built upon the incomplete gamma function and each represents a localized L_p estimator. By deriving various properties of the kernel family concerning distributional, Mean Shift induced, and other aspects such as strict positive definiteness, we obtain a deeper understanding of the operator’s projection behavior. From these theoretical insights, we illustrate several applications ranging from an improved Weighted LOP (WLOP) density weighting scheme and a more accurate Continuous LOP (CLOP) kernel approximation to the definition of a novel set of robust loss functions. These incomplete gamma losses include the Gaussian and LOP loss as special cases and can be applied to various tasks including normal filtering. Furthermore, we show that the novel kernels can be included as priors into neural networks. We demonstrate the effects of each application in a range of quantitative and qualitative experiments that highlight the benefits induced by our modifications.

Index Terms—Kernels, Locally Optimal Projection, Mean Shift, point clouds, point cloud denoising, projection operators, robust loss functions, surface reconstruction, theory.

I. INTRODUCTION

DIGITAL 3D scene models have become a crucial prerequisite for numerous applications in entertainment, advertisement, design, architecture, autonomous systems, and cultural heritage. In this context, the accurate digitization of real-world objects and scenes is of great relevance and offers new opportunities regarding a variety of tasks including AR/VR-based inspection and collecting realistic training data for tasks in robotics, autonomous driving, aerial or satellite surveys. Aside from professional scanning campaigns with expensive laser scanning equipment, there has also been an increasing trend

towards more practical scene capture with consumer-grade hardware such as passive purely image-based scene scanning using Structure-from-Motion and Multi-view Stereo approaches, or with respective cheaper active time-of-flight depth sensors that have meanwhile even been integrated into numerous mobile devices. However, the use of passive scene scanning or active scanning based on cheap hardware with low sensor quality and low sensor resolution induces noise in the capture process and thereby results in noisy point clouds and a low number of points that might not preserve finer geometric details, which, in turn, may lead to artifacts in the registration and subsequent surface reconstruction procedures. Furthermore, the limited accessibility of capture conditions as well as occlusions induce holes and highly irregular samplings and distributions of the captured data. These challenges result in an increasing interest in robust filtering techniques that are capable of handling noise, outliers, registration artifacts, as well as irregularly-sampled and missing data and can provide a clean, denoised, and uniformly resampled point cloud suitable for high-fidelity surface reconstruction.

Among others, the Locally Optimal Projection (LOP) operator [1] has gained a lot of attention in recent years due to its benefit of not relying on a well-defined surface parametrization or a piecewise planar approximation and, meanwhile, there has been a whole series of further extensions of this approach [2], [3], [4], [5]. Furthermore, many learning-based approaches also aim at projecting the noisy data onto a (latent) denoised manifold [6], [7]. Therefore, investigations towards the unification of traditional approaches with their respective regularization techniques might be of great relevance for future learning-based approaches as well. Even further, traditional techniques, and particularly those approaches that incorporate probabilistic modeling of the data such as the density-based Mean Shift clustering method [8], [9], [10], become more and more relevant in modern deep learning methods. Besides their application to structure the latent space representation of the data within encoder-decoder approaches [11], there is even a direct relation between the Mean Shift approach and denoising autoencoders [12]. In particular, as the output of an optimal denoising autoencoder corresponds to the local mean of the true data density [13], the autoencoder loss can be interpreted as a Mean Shift vector [12]. However, to the best of our knowledge, this observation has not yet been explored in the context of point cloud denoising. Hence, relating traditional concepts to modern deep learning methods might not only lead to a more explainable behavior of the latter but also allow increasing the resulting performance. In turn, this relies

Manuscript received 4 May 2022; revised 2 November 2023; accepted 19 December 2023. Date of publication 4 January 2024; date of current version 7 May 2024. This work was supported in part by the DFG projects KL 1142/11-2 (DFG Research Unit FOR 2535 Anticipating Human Behavior) and NFDI4Culture (DFG Project) under Grant 441958017. Recommended for acceptance by C.S. Ong. (Corresponding author: Patrick Stotko.)

Patrick Stotko and Reinhard Klein are with the Institute of Computer Science II – Visual Computing, University of Bonn, 53115 Bonn, Germany (e-mail: stotko@cs.uni-bonn.de; rk@cs.uni-bonn.de).

Michael Weinmann is with the Department of Intelligent Systems, Delft University of Technology, 2628 XE Delft, The Netherlands (e-mail: m.weinmann@tudelft.nl).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TPAMI.2024.3349967>, provided by the authors.

Digital Object Identifier 10.1109/TPAMI.2024.3349967




Fig. 1. Relation between LOP and Mean Shift in the example of the 2D Fish model. Minimizing the *localized* L_1 attraction energy with the Gaussian kernel K_{Gaussian} (left) results in the same trajectory $q^{(t)}$ as applying Mean Shift on a *global* kernel density estimate with the kernel K_{LOP} (right).

on the better understanding of the relationship between previous (traditional) techniques.

In this paper, we investigate the theoretical relationship of projection-based point cloud denoising approaches with their respective properties and show that these are unified within the common probabilistic Mean Shift framework. In particular, the key contributions of our work are:

- We reveal the relation of the classical localized L_1 estimator used in LOP to the Mean Shift framework via a novel kernel K_{LOP} and introduce the family of *incomplete gamma kernels* K_{Γ} as a generalization of this result where each kernel represents a localized L_p estimator (see Section III).
- We derive various properties of the kernel family concerning distributional, Mean Shift induced, and other aspects such as strict positive definiteness to obtain a deeper understanding of the operator's projection behavior (see Section IV).
- We demonstrate that leveraging the derived theoretical insights enables several applications including an improved Weighted LOP (WLOP) density weighting scheme, a more accurate Continuous LOP (CLOP) kernel approximation, the derivation of *incomplete gamma losses*, a set of novel robust loss functions, as well as neural network priors (see Section V).

In our evaluation, we demonstrate the benefits induced by our modifications in a range of quantitative and qualitative experiments. Furthermore, the theoretical insights of our investigations with their proven effect may be of great relevance also for future learning-based approaches. The source code is available at <https://github.com/stotko/incomplete-gamma-kernels>.

II. RELATED WORK

In the following, we provide a review of geometric and learning-based denoising approaches. Furthermore, we also review seminal work regarding the theory and application of the Mean Shift framework due to its relationship to LOP approaches that we will demonstrate later.

A. Geometric Denoising Approaches

Following early approaches such as the local fitting of tangent planes [14] or using radial basis functions [15], respective

developments particularly focused on projection-based methods, sparsity-based methods and non-local methods.

Projection-Based Methods: These approaches rely on the assumption of an underlying smooth surface and the projection of noisy data points onto the estimated local surface. For this purpose, respective approaches apply moving least squares (MLS) based methods [16], [17], [18], [19], robust principal component analysis (RPCA) [20] and moving robust principal component analysis (MRPCA) [21], or locally optimal projection based operators where the LOP operator [1] has been extended in terms of Weighted LOP (WLOP) [2], Feature LOP (FLOP) [3], Continuous LOP (CLOP) [5], Edge-Aware Resampling (EAR) [4] and a Gaussian mixture model inspired projection operator [22]. The latter has been demonstrated to be capable of resampling point clouds while preserving features due to the additional guidance of filtered normals.

Sparsity-Based Methods: This class of approaches relies on the assumption that objects can be represented in terms of piecewise smooth surfaces with sparse features. Respective denoising techniques include L_0 -norm [23], [24] and L_1 -norm minimization [21], [25], [26], sparse dictionary learning [27] as well as patch-based or feature-based graph Laplacian regularization [28], [29], [30], graph-based point cloud denoising based on jointly leveraging geometry and color information [31], guided filtering based on normal information followed by a L_1 -medial skeleton extraction to get the sharp structure of the surface [32] as well as leveraging gravitational feature functions [33]. In the context of denoising dynamic point clouds, Hu et al. [34] explored the temporal coherence of spatio-temporal graphs with respect to the underlying surface, where a respective manifold-to-manifold distance has been introduced. Furthermore, data-driven exemplar priors have been used for surface reconstruction [35], where the sparsity of local shapes from a collection of 3D objects has been explored.

Non-Local Methods: In contrast to the previous classes, these approaches rely on the assumption that geometric statistics are (approximately) shared by certain surface patches of a 3D model, i.e. local surface denoising is conducted based on collected neighborhoods with similar geometry [36], [37], [38], [39]. However, the definition of a suitable metric as well as the regular representation of local surface structures remain

challenging. Furthermore, density-based point cloud denoising has been approached by first applying particle-swarm based optimization for kernel density estimation followed by a Mean Shift clustering-based outlier removal and a final bilateral mesh filtering [40].

B. Learning-Based Denoising Approaches

Recent works more and more leverage deep learning for point cloud denoising as well as surface reconstruction from point clouds. Examples include approaches for point cloud consolidation and resampling such as PointNet [41], PointNet++ [42], patch-based progressive point cloud upsampling [43] as well as the unification of the considerations of densifying, denoising and completing point clouds [44]. Other approaches followed the principles of initially projecting the points onto coarse-level local reference planes and applying a subsequent refinement [45] or the initial removal of outliers before conducting the denoising [46]. Further approaches include edge-aware point cloud consolidation [47], adversarial defense [48], graph-convolutional methods [49], unsupervised approaches such as Total Denoising [50], gradient field based denoising [51], [52], [53], differentiable approaches [54], [55], [56] as well as manifold learning based on encoder-decoder architectures [6], [7]. Non-local self-similarities have also been considered to define neural self-priors that capture geometric repetitions [57], capture semantically related non-local features [58], or apply self-correction by allowing the model to capture structural and contextual information from initially disorganized parts [59]. Furthermore, normalizing flows have been applied to learn the distribution of noisy points and disentangle noise from the latent space [60]. In addition, the feature-aware recurrent point cloud denoising network (RePCD-Net) [61] combines a recurrent network architecture for noise removal with multi-scale feature aggregation and propagation and a feature-aware Chamfer distance loss.

C. Mean Shift Approaches

The Mean Shift approach [8] is a well-studied local mode-seeking method with diverse applications including data clustering [9], [10], [62], [63], image filtering [10], segmentation [10], [64], denoising [12], [65], and object tracking [64]. Tremendous effort has been spent to study its convergence behavior [9], [10], [66], [67], [68], [69] which culminated in a rigorous set of properties proven by Yamasaki and Tanaka [70]. Recently, Mean Shift clustering has also been applied in the latent space of neural encoder-decoder approaches to achieve a better structured data representation [11]. Furthermore, the connection between the Mean Shift approach and denoising autoencoders [71] has been revealed by Bigdeli et al. [12], who leveraged the observation that the output of an optimal denoising autoencoder (DAE) is a local mean of the true data density [13] to show that the autoencoder loss is a Mean Shift vector and to use the respective magnitude to define a prior for image restoration.

III. BACKGROUND

Before deriving our proposed kernel family as a generalization of LOP in the context of Mean Shift, we first provide a brief introduction into the concepts of both approaches.

A. Mean Shift

The basic objective of Mean Shift [8] is to find the modes of a probabilistic distribution f which has been observed by a (sparse) set of points $\mathcal{P} = \{\mathbf{p}_i \in \mathbb{R}^d\}$ in a d -dimensional vector space. Due to this abstract formulation, it has been applied to various computer vision problems by choosing an appropriate application-specific feature space, e.g. the $L^*u^*v^*$ color space for image filtering or segmentation [10]. In order to analyze the unknown distribution f at any point $\mathbf{q} \in \mathbb{R}^d$, it is modeled by a kernel density estimate:

$$\hat{f}_{\mathcal{P},K}(\mathbf{q}) = \frac{1}{|\mathcal{P}|h^d} \sum_i K\left(\frac{\mathbf{p}_i - \mathbf{q}}{h}\right) \quad (1)$$

Here, h denotes the kernel window size and K a kernel that is non-negative ($K(\mathbf{x}) \geq 0$), normalized ($\int_{\mathbb{R}^d} K(\mathbf{x}) d\mathbf{x} = 1$), and radially symmetric ($K(\mathbf{x}) = c_K k(\|\mathbf{x}\|^2)$). The 1-dimensional function k defined in the symmetry constraint, where c_K denotes a normalization constant, is called the kernel profile of K and plays an important role in the analysis of Mean Shift [70]. Furthermore, the gradient of the kernel density estimate

$$\nabla \hat{f}_{\mathcal{P},K}(\mathbf{q}) = \frac{2}{|\mathcal{P}|h^{d+2}} \frac{c_K}{c_G} \sum_i G\left(\frac{\mathbf{p}_i - \mathbf{q}}{h}\right) (\mathbf{p}_i - \mathbf{q}) \quad (2)$$

can be derived using the kernel $G(\mathbf{x}) = c_G g(\|\mathbf{x}\|^2)$ with the normalization constant c_G and its corresponding profile $g(x) = -\frac{d}{2x} k(x)$. Based on these two functions, the main component of the algorithm for finding the modes is the Mean Shift vector [10]

$$\mathbf{m}_{\mathcal{P},G}(\mathbf{q}) = \frac{h^2 c_G}{2 c_K} \frac{\nabla \hat{f}_{\mathcal{P},K}(\mathbf{q})}{\hat{f}_{\mathcal{P},K}(\mathbf{q})} = \frac{\sum_i G\left(\frac{\mathbf{p}_i - \mathbf{q}}{h}\right) (\mathbf{p}_i - \mathbf{q})}{\sum_i G\left(\frac{\mathbf{p}_i - \mathbf{q}}{h}\right)} \quad (3)$$

which describes the gradient vector normalized with respect to the kernel G . In particular, it directly determines the updated point at each time step t by the corresponding fixed-point iteration $\mathbf{q}^{(t+1)} = \mathbf{q}^{(t)} + \mathbf{m}_{\mathcal{P},G}(\mathbf{q}^{(t)})$ which performs gradient ascent on the kernel density estimate $\hat{f}_{\mathcal{P},K}$.

B. Locally Optimal Projection

Unlike approaches that are based on the probabilistic concept of Mean Shift to tackle various different computer vision tasks, the problem of denoising point clouds has also been of great relevance for the computer graphics community. Many advanced solutions have been independently developed there, especially methods leveraging the robust L_1 median. To efficiently compute the unique global solution in these median-based formulations, the iterative Weiszfeld algorithm [72] has become a popular and commonly used choice. Specifically, it also marks the foundation of LOP [1] as a robust localized 3D projection operator. Given a set of noisy 3D target points $\mathcal{P} = \{\mathbf{p}_i \in \mathbb{R}^3\}$ sampled from a smooth surface $\mathcal{S}$, the task here consists in

projecting and uniformly distributing an independent set of 3D points $\mathcal{Q} = \{\mathbf{q}_j \in \mathbb{R}^3\}$ onto the unknown surface $\mathcal{S}$ which is defined by the observations $\mathcal{P}$ only. This can be expressed in terms of an energy formulation

$$E(\mathcal{Q}) = \sum_j E_{\text{LOP}}(\mathbf{q}_j) + E_{\text{rep}}(\mathbf{q}_j) \quad (4)$$

based on an attraction and a repulsion term

$$E_{\text{LOP}}(\mathbf{q}_j) = \sum_i \theta(\|\mathbf{p}_i - \mathbf{q}_j^{(t)}\|) \|\mathbf{p}_i - \mathbf{q}_j\| \quad (5)$$

$$E_{\text{rep}}(\mathbf{q}_j) = \lambda_j \sum_{i, i \neq j} \theta(\|\mathbf{q}_i^{(t)} - \mathbf{q}_j^{(t)}\|) \eta(\|\mathbf{q}_i^{(t)} - \mathbf{q}_j\|) \quad (6)$$

where $\theta(x) = e^{-x^2/(h/4)^2}$ denotes a compact localization kernel and η a decreasing regularization function penalizing small distances between projection points to ensure a uniform distribution of the projected points. Common choices of η include the originally proposed function $\eta_{\text{LOP}}(x) = 1/(3x^3)$ [1] as well as the less rapidly decreasing function $\eta_{\text{WLOP}}(x) = -x$ [2]. Both energy terms are balanced by weights λ_j which are chosen such that they only depend on a single, global parameter $\mu \in [0, 1/2)$. Based on the Weiszfeld algorithm, the solution to this optimization problem can be obtained by the fixed-point iteration

$$\begin{aligned} \mathbf{q}_j^{(t+1)} = & \frac{\sum_i \alpha(\|\mathbf{p}_i - \mathbf{q}_j^{(t)}\|) \mathbf{p}_i}{\sum_i \alpha(\|\mathbf{p}_i - \mathbf{q}_j^{(t)}\|)} \\ & + \mu \frac{\sum_{i, i \neq j} \beta(\|\mathbf{q}_i^{(t)} - \mathbf{q}_j^{(t)}\|) (\mathbf{q}_j^{(t)} - \mathbf{q}_i^{(t)})}{\sum_{i, i \neq j} \beta(\|\mathbf{q}_i^{(t)} - \mathbf{q}_j^{(t)}\|)} \end{aligned} \quad (7)$$

with kernels $\alpha(x) = \theta(x)/x$ and $\beta(x) = \theta(x)/x \lfloor \frac{d}{dx} \eta(x) \rfloor$.

C. Generalization via Incomplete Gamma Kernels

Although Mean Shift and Locally Optimal Projection have been separately developed from different contexts and mathematical concepts with the aim of solving a distinct problem, we can link both approaches by rewriting the update step in (7):

$$\mathbf{q}_j^{(t+1)} = \mathbf{q}_j^{(t)} + \mathbf{m}_{\mathcal{P}, G_{\text{LOP}}}(\mathbf{q}_j^{(t)}) - \mu \mathbf{m}_{\mathcal{Q}_j^{(t)}, G_{\text{rep}}}(\mathbf{q}_j^{(t)}) \quad (8)$$

This reveals that the LOP operator is a combination of two Mean Shift steps in the 3D space: 1) a standard Mean Shift with respect to the target set $\mathcal{P}$ and G_{LOP} being the normalized kernel α ; and 2) a reverse applied Blurring Mean Shift [9] where the Mean Shift vector is instead subtracted and computed from the shifted and, in turn, blurred source set $\mathcal{Q}_j^{(t)} = \mathcal{Q}^{(t)} \setminus \{\mathbf{q}_j^{(t)}\}$ as well as from the kernel G_{rep} corresponding to the normalized kernel β . Therefore, we can interpret the localized L_1 attraction energy minimization with a Gaussian kernel in (5) as a maximization of a global kernel density estimate (1) with respect to a different kernel K_{LOP} . An example of this relation is shown in Fig. 1.

To derive K_{LOP} as part of a novel kernel family K_Γ for the general case of the $\mathbb{R}^d$ space, we consider the 1-dimensional profile of the involved kernel G_{LOP} , i.e. α in (7), which resembles a gamma distribution $f_\Gamma(x|a, b) \propto x^{a-1} e^{-x/b}$ with support $x \in (0, \infty)$ and parameters $a > 0, b > 0$. The profile of the actual kernel then resembles the distribution $\bar{F}_\Gamma(x|a, b) \propto \Gamma(a, x/b)$ which is the complementary CDF of f_Γ and based on the *upper incomplete gamma function* $\Gamma(a, x) = \int_x^\infty t^{a-1} e^{-t} dt$. Therefore, the d -dimensional kernel has the general form $K_\Gamma(\mathbf{x}|a, b) = c_{K_\Gamma} \Gamma(a, \|\mathbf{x}\|^2/b)$.

Since we also need to compute the respective normalization constant c_{K_Γ} , we switch the integration domain to spherical coordinates and substitute $s = r^2/b$:

$$\begin{aligned} \frac{1}{c_{K_\Gamma}} &= \int_{\mathbb{R}^d} \Gamma(a, \frac{\|\mathbf{x}\|^2}{b}) d\mathbf{x} = \int_\Omega \int_0^\infty \Gamma(a, \frac{r^2}{b}) r^{d-1} dr d\Omega \\ &= \frac{b^{\frac{d}{2}}}{2} \left[\int_\Omega d\Omega \right] \left[\int_0^\infty \Gamma(a, s) s^{\frac{d}{2}-1} ds \right] \end{aligned} \quad (9)$$

Due to radial symmetry, both integrals can be solved independently. The former one describes the surface area of the d -dimensional unit sphere Ω and has the closed form $\int_\Omega d\Omega = 2\pi^{d/2}/\Gamma(d/2)$. Using the relation $\int_0^\infty \Gamma(a, x) x^{b-1} dx = \Gamma(a+b)/b$ [73], we get an expression for the latter one in terms of the ordinary gamma function. We can also apply the recursive relation of the gamma function $\Gamma(a+1) = a\Gamma(a)$ and conclude that $1/c_{K_\Gamma} = (\pi b)^{d/2} \Gamma(d/2+a)/\Gamma(d/2+1)$.

Finally, we change the parametrization by setting $a = p/2, b = 2\sigma^2$ to obtain the final kernel:

$$K_\Gamma(\mathbf{x}|p, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{\frac{d}{2}}} \frac{\Gamma(\frac{d+2}{2})}{\Gamma(\frac{d+p}{2})} \Gamma(\frac{p}{2}, \frac{\|\mathbf{x}\|^2}{2\sigma^2}) \quad (10)$$

These *incomplete gamma kernels* span a family of Mean Shift kernels corresponding to L_p estimators of the attraction energy localized by a Gaussian kernel. An important special case of this family is the LOP kernel for which we choose $p = 1, \sigma^2 = 1/32$ and apply the identity $\Gamma(1/2, x) = \sqrt{\pi} \operatorname{erfc}(\sqrt{x})$ to get

$$K_{\text{LOP}}(\mathbf{x}) = \frac{4^d}{\pi^{\frac{d-1}{2}}} \frac{\Gamma(\frac{d+2}{2})}{\Gamma(\frac{d+1}{2})} \operatorname{erfc}(4\|\mathbf{x}\|) \quad (11)$$

where erfc denotes the *complementary error function*. Another special case is the corresponding Gaussian kernel K_{Gaussian} obtained by setting $p = 2$ which is a common choice in Mean Shift and has been extensively analyzed as the localized L_2 estimator of the geometric mean. Fig. 2 shows an interpolation between these kernels by varying the p -norm.

IV. KERNEL PROPERTIES

In the following, we derive several theoretical properties of the family of incomplete gamma kernels in $\mathbb{R}^d$ which are summarized in Table I and later leveraged in the applications (see Section V).




Fig. 2. Interpolation between 2D incomplete gamma kernels K_Γ with varying $p \in [1, 2]$ and fixed $\sigma^2 = 1/32$. Each kernel corresponds to a localized attraction energy minimization with the respective p -norm.

TABLE I
PROPERTIES OF INCOMPLETE GAMMA KERNELS $K_\Gamma(\mathbf{x} | p, \sigma^2)$ IN $\mathbb{R}^d$ FOR $p > 0$

Distribution	Mean	$\mathbf{0}$
	Covariance	$\frac{d+p}{d+2} \sigma^2 \mathbf{I}$
	Characteristic function	${}_1F_1\left(\frac{d+p}{2}, \frac{d+2}{2}, -\frac{\sigma^2 \ \omega\ ^2}{2}\right)$
Mean Shift	Moment-generating function	${}_1F_1\left(\frac{d+p}{2}, \frac{d+2}{2}, \frac{\sigma^2 \ \omega\ ^2}{2}\right)$
	Differentiable profile	$\checkmark$ except for $x = 0$ if $p \in (0, 2)$
	Strictly decreasing profile	$\checkmark$
Other	Convex profile	$\checkmark$ for $p \in (0, 2]$
	Analytic	$\checkmark$
	Bounded	$\checkmark$
Other	Completely monotonic profile	$\checkmark$ for $p \in (0, 2]$
	Strictly positive definite	$\checkmark$ for $p \in (0, 2]$

A. Characteristic Function and Fourier Transform

Feature preservation is a highly relevant aspect in the development of denoising approaches (see Fig. 4) and will later be taken into account in the definition of density weighting schemes (see Section V-A), loss functions (see Section V-C), and neural network priors (see Section V-D). In order to gain a deeper understanding of the proposed kernel family K_Γ in this context, we are interested in its characteristic function $\varphi_\Gamma(\omega)$ which can also be interpreted as the Fourier transform $\mathcal{F}$ of K_Γ at the angular frequency $\omega \in \mathbb{R}^d$.

First, we can apply the relation between the d -dimensional Fourier transform of a radially symmetric function $f(\mathbf{x})$ in terms of the Hankel transform of order $d/2 - 1$ of the function $\|\mathbf{x}\|^{d/2-1} f(\|\mathbf{x}\|)$ [74] to reduce the dimensionality of the integral to the radial component

$$\begin{aligned} \varphi_\Gamma(\omega) &= \mathcal{F}[K_\Gamma(\mathbf{x} | p, \sigma^2)](\omega) = \int_{\mathbb{R}^d} K_\Gamma(\mathbf{x} | p, \sigma^2) e^{i(\omega | \mathbf{x})} d\mathbf{x} \\ &= c_{K_\Gamma} \frac{(2\pi)^{\frac{d}{2}}}{\|\omega\|^{\frac{d}{2}-1}} \int_0^\infty \Gamma\left(\frac{p}{2}, \frac{r^2}{2\sigma^2}\right) J_{\frac{d}{2}-1}(\|\omega\| r) r^{\frac{d}{2}} dr \end{aligned} \quad (12)$$

where $J_q(x)$ denotes the *Bessel function of the first kind* of order q . This integral has the closed-form solution (see the Appendix, available online, for a more detailed derivation):

$$c_{K_\Gamma} \frac{(2\pi)^{\frac{d}{2}}}{\|\omega\|^{\frac{d}{2}-1}} \int_0^\infty \Gamma\left(\frac{p}{2}, \frac{r^2}{2\sigma^2}\right) J_{\frac{d}{2}-1}(\|\omega\| r) r^{\frac{d}{2}} dr$$




Fig. 3. Comparison of LOP and Gaussian kernels in spatial and frequency domain. Filtering with the LOP kernel K_{LOP} better preserves higher frequency information.

$$\begin{aligned} &= c_{K_\Gamma} (2\pi\sigma^2)^{\frac{d}{2}} \frac{\Gamma(\frac{d+p}{2})}{\Gamma(\frac{d+2}{2})} {}_1F_1\left(\frac{d+p}{2}, \frac{d+2}{2}, -\frac{\sigma^2 \|\omega\|^2}{2}\right) \\ &= {}_1F_1\left(\frac{d+p}{2}, \frac{d+2}{2}, -\frac{\sigma^2 \|\omega\|^2}{2}\right) \end{aligned} \quad (13)$$

Therefore, the characteristic function of the incomplete gamma kernel can be written in terms of the *confluent hypergeometric function of the first kind* ${}_1F_1$. Fig. 3 shows a comparison between the Gaussian kernel ($p = 2$) and the LOP kernel ($p = 1$) both in spatial and in frequency domain.

If we consider the special case ${}_1F_1(a, a, x) = e^x$, we can observe that this result is consistent with the Fourier transform of the Gaussian kernel. Furthermore as $d \rightarrow \infty$, the entire family of localized L_p kernel estimators converges to the L_2 estimator since distances become increasingly similar in higher dimensions due to the curse of dimensionality.

B. Moment-Generating Function

Another closely related and useful quantity to consider is the moment-generating function M_Γ of the kernel K_Γ which can be used to compute its mean vector $\mu_\Gamma \in \mathbb{R}^d$ and its covariance matrix $\Sigma_\Gamma \in \mathbb{R}^{d \times d}$. Especially the insights about the magnitude of Σ_Γ are crucial to define consistent density weights (see Section V-A) and an accurate CLOP approximation (see Section V-B). Although there is a direct connection to the characteristic function in terms of

$$M_\Gamma(\omega) = \varphi_\Gamma(-i\omega) = {}_1F_1\left(\frac{d+p}{2}, \frac{d+2}{2}, \frac{\sigma^2 \|\omega\|^2}{2}\right) \quad (14)$$

it does not necessarily exist in general, so we have to prove this property for all $\omega \in \mathbb{R}^d$. For this purpose, we repeatedly apply the comparison theorem of calculus to derive finite bounds of the

integral. We first observe that $M_\Gamma(\omega) > 0$ since the integrand consisting of K_Γ and an exponential term is positive. To obtain an upper bound of M_Γ , we switch to spherical coordinates and bound $\cos(\angle(\omega, x)) \leq 1$, where $\angle(\omega, x)$ denotes the angle between ω and x , to decouple the radial component from the angular one:

$$\begin{aligned} M_\Gamma(\omega) &= \int_{\mathbb{R}^d} K_\Gamma(x | p, \sigma^2) e^{\langle \omega, x \rangle} dx \\ &= c_{K_\Gamma} \int_{\Omega} \int_0^\infty \Gamma\left(\frac{p}{2}, \frac{r^2}{2\sigma^2}\right) r^{d-1} e^{\|\omega\| r \cos(\angle(\omega, x))} dr d\Omega \\ &\leq c_1 \int_0^\infty \Gamma\left(\frac{p}{2}, \frac{r^2}{2\sigma^2}\right) r^{d-1} e^{\|\omega\| r} dr \end{aligned} \quad (15)$$

For brevity, we put finite terms into constants c_i . Next, we combine the two individual upper bounds

$$\Gamma(a, x) \leq \begin{cases} a x^{a-1} e^{-x}, & a \in [1, \infty), x \in [a, \infty) \text{ [75]} \\ x^{a-1} e^{-x}, & a \in (0, 1], x \in (0, \infty) \text{ (partial int.)} \end{cases} \quad (16)$$

and apply them after splitting the integral at $r_0 = \max(1, p/2)$. Furthermore, we simplify the expression by completing the square in the exponential term:

$$\begin{aligned} &c_1 \int_0^\infty \Gamma\left(\frac{p}{2}, \frac{r^2}{2\sigma^2}\right) r^{d-1} e^{\|\omega\| r} dr \\ &\leq c_1 c_2 + c_1 \int_{r_0}^\infty r_0 \left(\frac{r^2}{2\sigma^2}\right)^{\frac{p}{2}-1} e^{-\frac{r^2}{2\sigma^2}} r^{d-1} e^{\|\omega\| r} dr \\ &= c_1 c_2 + c_1 c_3 \int_{r_0}^\infty r^{p+d-3} e^{-\frac{(r-\sigma^2\|\omega\|)^2}{2\sigma^2}} dr \end{aligned} \quad (17)$$

Since $r \in [1, \infty)$, the remaining polynomial can be bound by a higher order $k = \max(0, \lceil p + d - 3 \rceil) \in \mathbb{N}_0$. Therefore, this integral describes the (incomplete) k -th raw moment of a 1D normal distribution and is finite for any $k \in \mathbb{N}_0$ which, in turn, proves that $M_\Gamma(\omega) < \infty$ exists.

1) *Mean*: We can now use the moment-generating function M_Γ to directly compute all raw moments of the kernel K_Γ by evaluating the respective derivative at $\omega = \mathbf{0}$. For the mean vector $\mu_\Gamma \in \mathbb{R}^d$ of K_Γ , we consider the first-order derivative

$$\frac{\partial}{\partial \omega} M_\Gamma(\omega) = \frac{d+p}{d+2} \sigma^2 {}_1F_1\left(\frac{d+p+2}{2}, \frac{d+4}{2}, \frac{\sigma^2 \|\omega\|^2}{2}\right) \omega \quad (18)$$

and get $\mu_\Gamma = \frac{\partial}{\partial \omega} M_\Gamma(\mathbf{0}) = \mathbf{0}$ as the expected result for a radially symmetric kernel.

2) *Covariance*: Similarly, we compute the second-order derivative

$$\begin{aligned} \frac{\partial^2}{\partial \omega \partial \omega^\top} M_\Gamma(\omega) &= \frac{d+p}{d+2} \sigma^2 \left[{}_1F_1\left(\frac{d+p+2}{2}, \frac{d+4}{2}, \frac{\sigma^2 \|\omega\|^2}{2}\right) \mathbf{I} \right. \\ &\quad \left. + \frac{d+p+2}{d+4} \sigma^2 {}_1F_1\left(\frac{d+p+4}{2}, \frac{d+6}{2}, \frac{\sigma^2 \|\omega\|^2}{2}\right) \omega \omega^\top \right] \end{aligned} \quad (19)$$

and obtain the covariance matrix $\Sigma_\Gamma \in \mathbb{R}^{d \times d}$ of the kernel K_Γ using the first-order and second-order raw moments as $\Sigma_\Gamma = \frac{\partial^2}{\partial \omega \partial \omega^\top} M_\Gamma(\mathbf{0}) - \mu_\Gamma \mu_\Gamma^\top = [(d+p)/(d+2)] \sigma^2 \mathbf{I}$.

C. Mean Shift Properties

In addition to the distribution-specific properties above, we can get further insights into the kernel family by exploiting the comprehensive theory that has been developed for the Mean Shift algorithm [70]. This requires proving several additional properties including that the kernel K_Γ is bounded and analytic and that its profile k_Γ is differentiable, strictly decreasing, and convex.

1) *Differentiability, Monotonicity and Convexity*: In order to show that the profile is strictly decreasing, we consider its first-order derivative

$$\frac{d}{dx} k_\Gamma(x | p, \sigma^2) = \frac{d}{dx} \Gamma\left(\frac{p}{2}, \frac{x}{2\sigma^2}\right) = -\left(\frac{1}{2\sigma^2}\right)^{\frac{p}{2}} x^{\frac{p}{2}-1} e^{-\frac{x}{2\sigma^2}} \quad (20)$$

which is defined for all $x \in (0, \infty)$ as well as for $x = 0$ if $p \in [2, \infty)$. Since the involved polynomial and exponential terms are always positive, it follows that the derivative must be negative, that is $\frac{d}{dx} k_\Gamma(x | p, \sigma^2) < 0$, and the profile strictly decreasing.

Similarly, we see that the second-order derivative is given by

$$\frac{d^2}{dx^2} k_\Gamma(x | p, \sigma^2) = \frac{d}{dx} k_\Gamma(x | p, \sigma^2) \left[\frac{\frac{p}{2} - 1}{x} - \frac{1}{2\sigma^2} \right] \quad (21)$$

where, in order to ensure that $\frac{d^2}{dx^2} k_\Gamma(x | p, \sigma^2) > 0$, the latter term must be non-positive which is equivalent to the condition $x \geq (p-2)\sigma^2$. Since this should hold for all $x \in (0, \infty)$, convexity is only guaranteed for kernels with $p \in (0, 2]$ which, in particular, includes the Gaussian kernel ($p = 2$) as well as the LOP kernel ($p = 1$).

2) *Boundedness and Analyticity*: Instead of showing the properties of boundedness and analyticity for the kernel K_Γ itself, it is sufficient to show them for its profile k_Γ . Since k_Γ is non-negative and monotonically decreasing, we only have to consider the case $x = 0$. For this value, $\Gamma(a, x)$ reduces to the gamma function $\Gamma(a)$ which is finite for $a > 0$. Furthermore, analyticity directly follows from the fact that $\Gamma(a, x)$ is holomorphic in $x \in (0, \infty)$ for any fixed $a > 0$.

3) *Consequences for the LOP Operator*: The aforementioned properties have several direct implications [70] on the behavior of the LOP operator (with zero repulsion) as well as to Mean Shift applied with the incomplete gamma kernel K_Γ for $p \in (0, 2]$. With the exception of the finite set of target points $\mathcal{P}$ where singularities are introduced in the kernel G_Γ , the following properties hold:

Non-Zero Gradient: The gradient of the kernel density estimate $\nabla \hat{f}_{\mathcal{P}, K_\Gamma}$ in (2) is non-zero outside the convex hull of the target point set $\mathcal{P}$. This implies that all solutions must lie within the convex hull.

Plateau-Free Density: In addition to non-zero gradients, the kernel density estimate on the set $\mathbb{R}^d \setminus \mathcal{P}$ has no plateaus. Since the set of target points $\mathcal{P}$ is finite, we can extend this property to the full space $\mathbb{R}^d$.

Non-Decreasing Density Estimate: Another interesting subset to consider is the improvement ball $\mathcal{I}(q_j^{(t)})$ which denotes a d -dimensional sphere centered at the point $q_j^{(t)} + m_{\mathcal{P}, G_\Gamma}(q_j^{(t)})$ with radius $\|m_{\mathcal{P}, G_\Gamma}(q_j^{(t)})\|$. In case of the LOP operator, it

follows that all points $\mathbf{x}$ within the improvement ball have non-decreasing kernel density estimates $\hat{f}_{\mathcal{P}, K_\Gamma}(\mathbf{x}) \geq \hat{f}_{\mathcal{P}, K_\Gamma}(\mathbf{q}_j^{(t)})$.

Convergence of Density Estimate Sequence: As a consequence of the above property, the sequence of kernel density estimates $\{\hat{f}_{\mathcal{P}, K_\Gamma}(\mathbf{q}_j^{(t)})\}$ obtained via the fixed-point iteration $\mathbf{q}^{(t+1)} = \mathbf{q}^{(t)} + \mathbf{m}_{\mathcal{P}, G_\Gamma}(\mathbf{q}^{(t)})$ is non-decreasing. Furthermore, this sequence always converges.

Convergence of Mode Estimate Sequence: Finally, we can conclude that the mode estimate sequence $\{\mathbf{q}^{(t)}\}$ converges to a single point. Depending on the window size h and the distribution of the target points $\mathcal{P}$, this solution could be either a point $\mathbf{p} \in \mathcal{P}$ due to the singularity (for very small window sizes) or a different point $\mathbf{p} \in \mathbb{R}^d \setminus \mathcal{P}$ in the corresponding convex hull (for larger window sizes). While a proof for the convergence of a modified version of LOP in $\mathbb{R}^d$ has only been shown very recently [76], this provides an alternative way through the comprehensive Mean Shift theory which, though only considering the attraction term, extends and generalizes to all kernels K_Γ with $p \in (0, 2]$.

D. Further Properties

Besides the theoretical results concerning basic distribution-related aspects as well as insights in the projection behavior of LOP from the perspective of Mean Shift, we want to derive further properties and, in particular, strict positive definiteness. This opens up a broader set of applications beyond improved density weights (see Section V-A) and may also be relevant, e.g., in the field of Gaussian Process Regression, which, however is out of the scope of this work.

1) *Complete Monotonicity:* For this purpose, we show that the kernel profile k_Γ is completely monotonic, that is $(-1)^n \frac{d^n}{dx^n} k_\Gamma(x | p, \sigma^2) \geq 0$ for all $n \in \mathbb{N}_0$ and $x \in (0, \infty)$. From the derivation of the Mean Shift properties, we already know that $k_\Gamma(x | p, \sigma^2) > 0$ and $\frac{d}{dx} k_\Gamma(x | p, \sigma^2) < 0$ holds for all $x \in (0, \infty)$. Since x^a with $a \leq 0$ as well as e^{-x} are both completely monotonic and the product of two completely monotonic functions retains that property [77], we can conclude that this also holds for $k_\Gamma(x | p, \sigma^2)$ with $p \in (0, 2]$.

2) *Strict Positive Definiteness:* A direct consequence of the complete monotonicity of its profile is that the kernel K_Γ is a positive definite function. Furthermore, since the profile is also not a constant function, K_Γ must be even strictly positive definite [78]. Therefore, for any set of distinct points $\mathcal{P}$, the matrix

$$\mathbf{C} = (K_\Gamma(\mathbf{p}_i - \mathbf{p}_j | p, \sigma^2))_{ij} \in \mathbb{R}^{|\mathcal{P}| \times |\mathcal{P}|} \quad (22)$$

is symmetric and positive definite for $p \in (0, 2]$, so it can be interpreted as a Gram matrix and any linear system with respect to $\mathbf{C}$ has a unique solution which can be computed by, e.g., conjugate gradient solvers. This also directly extends to any truncated version of K_Γ where the matrix $\mathbf{C}$ becomes sparse and more efficient to solve as vanishing derivatives of the truncated profile do not affect complete monotonicity.

V. APPLICATIONS

Besides the application of other localized L_p estimators for point cloud denoising via the incomplete gamma kernels K_Γ , as shown in Fig. 4, we illustrate several further applications to demonstrate the benefits of the theoretical results derived for the kernel family.

A. WLOP Density Weights

In addition to a high robustness to noise, another crucial requirement in point cloud denoising is a uniform distribution of the denoised points. Although the introduction of an additional repulsion energy (6) mitigates the clustering effect of the attraction term (5), the projection is still highly dependent on the distribution of the target points $\mathcal{P}$. This has been addressed in WLOP [2] by computing weights v_i for each target point $\mathbf{p}_i$ and $v_j^{(t)}$ for each projection point $\mathbf{q}_j^{(t)}$ as an additional regularization based on the reciprocal and ordinary density value respectively with the (unnormalized) localization kernel θ . However, our derived theoretical properties reveal two major limitations of this particular choice: 1) Although the Gaussian localization kernel θ could be considered a reasonable approximation of the actual kernel K_{LOP} given in (11), it significantly differs in terms of the frequency spectrum (see Section IV-A and Fig. 3) as well as covariance (see Section IV-B) and leads to oversmoothing of high-frequency density information and, hence, a lacking preservation of fine-scale details; 2) taking the reciprocal to invert the density of $\mathbf{p}_i$ ignores the dependencies between the weights which corresponds to the assumption of constant density in a window of size h . In order to achieve a more accurate normalization, we propose two novel weighting schemes.

Simple Scheme: A simple extension to the WLOP weights keeps the assumption of the latter limitation and addresses only the former one by applying the actual kernel K_{LOP} , that is estimating the weights

$$v_i = \frac{1}{\hat{f}_{\mathcal{P}, K_{\text{LOP}}}(\mathbf{p}_i)}, \quad v_j^{(t)} = \hat{f}_{\mathcal{Q}^{(t)}, K_{\text{LOP}}}(\mathbf{q}_j^{(t)}) \quad (23)$$

via the kernel density estimate (1) of the point clouds $\mathcal{P}$ and $\mathcal{Q}^{(t)}$ respectively. This scheme can be easily integrated into existing applications of WLOP as it only involves a different kernel function in the overall weight computation.

Full Scheme: To address both limitations, we consider the kernel density estimate $\hat{f}_{\mathcal{P}, K_{\text{LOP}}}$ with weights v_i applied to each term. We want to enforce constant density at the points $\mathbf{p}_i$ which can be formulated as a linear optimization problem in matrix form:

$$\frac{1}{|\mathcal{P}| h^d} \left(K_{\text{LOP}}\left(\frac{\mathbf{p}_i - \mathbf{p}_j}{h}\right) \right)_{ij} (\dots, v_i, \dots)^T = \mathbf{1} \quad (24)$$

This corresponds to radial basis function (RBF) interpolation with a Gram matrix similar to (22) and we can obtain a unique solution since K_{LOP} is strictly positive definite (see Section IV-D). Furthermore, we truncate the kernel at $h/2$ to drastically reduce the memory requirements of the matrix and use a sparse conjugate gradient solver. In case of the projection points $\mathbf{q}_j^{(t)}$,




Fig. 4. Exemplary point cloud denoising of the *Elephant* model (302 458 points) with 30 WLOP [2] iterations ($h = 6$, $\mu = 0.4$) for different incomplete gamma kernels K_Γ using varying $p \in (0, 2]$ and fixed $\sigma^2 = 1/32$. The model has been corrupted with $\sigma_{\text{noise}} = 0.3$ (80 % points) and $\sigma_{\text{outlier}} = 1.5$ Gaussian noise (20 % points) respectively to account for both typical sensor noise and heavy outliers. Higher p -norms result in more regular but oversmoothed point distributions whereas lower values better preserve features. Unit of h , σ_{noise} , σ_{outlier} : [% BB diagonal].

we consider the inverse of the matrix which leads to the same weights as in the simple scheme.

B. CLOP Kernel Approximation

While the major aspect of ensuring a uniform distribution of the projected point set can be addressed with an appropriate weighting as shown in Section V-A, obtaining these results efficiently becomes challenging especially on large datasets as all target points $\mathcal{P}$ will be taken into account for denoising. For this purpose, CLOP [5] first computes a more compact representation of the points $\mathcal{P}$ in terms of a smaller set of normal distributions $\mathcal{P}_N = \{(w_i, \mu_i, \Sigma_i)\}$ with weights $w_i \in \mathbb{R}$, mean vectors $\mu_i \in \mathbb{R}^3$, and local covariance matrices $\Sigma_i \in \mathbb{R}^{3 \times 3}$ and then extends the discrete attraction energy to the continuous space. However, since the integral in the respective update step cannot be directly solved, the kernel $\alpha(\|\mathbf{x}\|)$ is approximated by a radially symmetric Gaussian mixture model $\hat{\alpha}(\mathbf{x}) = (1/h) \sum_{k=1}^3 \hat{w}_k \hat{c}_k \mathcal{N}(\mathbf{x}/h | \mathbf{0}, \hat{\sigma}_k^2 \mathbf{I})$ consisting of three components with fitted parameters $\{(\hat{w}_k, \hat{\sigma}_k) \in \mathbb{R} \times \mathbb{R}\}$ and dimension-dependent constants $\hat{c}_k = |2\pi\hat{\sigma}_k^2 \mathbf{I}|^{1/2}$. In the context of Mean Shift, this implies that the kernel G_{LOP} is in fact approximated which directly allows us to derive

$$\hat{K}_{\text{LOP}}(\mathbf{x}) = \frac{\sum_{k=1}^3 \hat{\sigma}_k^2 \hat{w}_k \hat{c}_k \mathcal{N}(\mathbf{x} | \mathbf{0}, \hat{\sigma}_k^2 \mathbf{I})}{\sum_{k=1}^3 \hat{\sigma}_k^2 \hat{w}_k \hat{c}_k} \quad (25)$$

as an approximation of the kernel K_{LOP} in $\mathbb{R}^d$ with the same set of fitted parameters $\{(\hat{w}_k, \hat{\sigma}_k)\}$.

Kernel Fit: Finding the optimal parameter set is highly challenging due to the singularity of G_{LOP} at $x = 0$ and, thereby, the unbounded ratio between the smallest and largest sampling value in the half-open fitting interval $(0, 1]$ for $h = 1$. In contrast, we directly optimize on the kernel K_{LOP} which does not suffer from these limitations. We fix the parameter $w_3 = 1$ to constrain the remaining degree of freedom and obtain the solution from 10^7 uniformly sampled points in the interval $[0, 1]$ via the Levenberg-Marquardt algorithm (see Table II). Although the LOP operator is scale-invariant in terms of the kernel α , we nevertheless estimate a global scaling factor for the weights $\hat{w}_k$ via Levenberg-Marquardt optimization in the interval $(0.01, 1]$ for a better comparability with CLOP.

TABLE II
PARAMETER SETS OF KERNEL APPROXIMATION $\hat{K}_{\text{LOP}}$

k	CLOP [5]		Ours		Ours (Consistent)	
	$\hat{w}_k$	$\hat{\sigma}_k$	$\hat{w}_k$	$\hat{\sigma}_k$	$\hat{w}_k$	$\hat{\sigma}_k$
1	97.761	0.01010	61.509	0.02102	46.409	0.03118
2	29.886	0.03287	11.932	0.07289	9.635	0.10582
3	11.453	0.11772	5.069	0.15700	2.674	$\sqrt{1/32}$

Consistent Fit: In addition to the derivation of the kernel approximation $\hat{K}_{\text{LOP}}$, we can further apply the insights about the covariance of K_{LOP} (see Section IV-B). From (25), we can also see that the covariance matrix $\hat{\Sigma}_{\text{LOP}} \in \mathbb{R}^{d \times d}$ of the approximation consists of a convex combination of the parameters $\hat{\sigma}_k^2$:

$$\hat{\Sigma}_{\text{LOP}} = \frac{\sum_{k=1}^3 \hat{w}_k \hat{\sigma}_k^{d+4}}{\sum_{k=1}^3 \hat{w}_k \hat{\sigma}_k^{d+2}} \mathbf{I} \quad (26)$$

In the limit $d \rightarrow \infty$, this combination degenerates to $\hat{\Sigma}_{\text{LOP}} \rightarrow \max_k \hat{\sigma}_k^2 \mathbf{I}$ which is similar to the maximum norm L_∞ being the limit of the L_p norms. Therefore, we can enforce an additional consistency constraint in the parameter optimization process by fixing the parameter $\hat{\sigma}_3 = \sqrt{1/32}$ such that $\hat{\Sigma}_{\text{LOP}}$ matches the expected covariance $\Sigma_{\text{LOP}} = 1/32 [(d+1)/(d+2)] \mathbf{I} \rightarrow 1/32 \mathbf{I}$.

C. Robust Loss Functions

A closely related, yet slightly different task is mesh denoising where the quality of the noisy vertex set (similar to point sets in point cloud denoising) should be improved while preserving the mesh connectivity and avoiding self-intersections of the faces. Respective approaches tackle this problem in a two-stage approach where the face normals are initially denoised and subsequently used as a guidance in the second stage to consistently adjust the vertex positions. Obtaining a reliable estimate of a denoised normal $\mathbf{n} \in \mathbb{R}^3$ can be performed by leveraging gradient descent or M-estimation from the broad field of robust statistics [79] which is also related to the concept of anisotropic

diffusion [80]. Here, the objective function

$$L(\mathbf{n}) = \sum_i \rho(\|\mathbf{n}_i - \mathbf{n}\|) \quad (27)$$

defined with a robust loss function ρ is considered and a solution can be found based on the corresponding influence function $\Psi(x) = \frac{d}{dx}\rho(x)$ and anisotropic weight function $\tilde{g}(x) = \Psi(x)/x$ [81]:

$$\mathbf{n}^{(t+1)} = \frac{\sum_i \tilde{g}(\|\mathbf{n}_i - \mathbf{n}^{(t)}\|) \mathbf{n}^{(t)}}{\|\sum_i \tilde{g}(\|\mathbf{n}_i - \mathbf{n}^{(t)}\|) \mathbf{n}^{(t)}\|} \quad (28)$$

This result is closely related to the derivation of Mean Shift and shares many properties with it [10]. While the Gaussian loss is a well-known choice in this context, we can generalize it to the family of *incomplete gamma losses* along with the respective influence and anisotropic weight functions and corresponding frequency-related properties (see Section IV-A):

$$\rho_\Gamma(x | p, \sigma^2) = \frac{1}{\Gamma(\frac{p}{2})} \gamma(\frac{p}{2}, \frac{x^2}{2\sigma^2}) \quad (29)$$

$$\Psi_\Gamma(x | p, \sigma^2) = \frac{2}{(2\sigma^2)^{\frac{p}{2}}} |x|^{p-2} e^{-\frac{x^2}{2\sigma^2}} x \quad (30)$$

$$\tilde{g}_\Gamma(x | p, \sigma^2) = \frac{2}{(2\sigma^2)^{\frac{p}{2}}} |x|^{p-2} e^{-\frac{x^2}{2\sigma^2}} \quad (31)$$

Here, the losses ρ_Γ are built upon the *lower incomplete gamma function* $\gamma(a, x) = \int_0^x t^{a-1} e^{-t} dt$ which is connected to the upper incomplete gamma function via the relation $\gamma(a, x) + \Gamma(a, x) = \Gamma(a)$. By choosing $p = 1$ and applying the identity $\gamma(1/2, x) = \sqrt{\pi} \operatorname{erf}(\sqrt{x})$, we get the LOP loss

$$\rho_{\text{LOP}}(x | \sigma^2) = \operatorname{erf}\left(\frac{|x|}{\sqrt{2\sigma^2}}\right) \quad (32)$$

where erf denotes the *error function* and is related to its complementary counterpart via $\operatorname{erfc}(x) = 1 - \operatorname{erf}(x)$. Considering $\sigma^2 = 1/32$, the relation $\tilde{g}_{\text{LOP}}(x | 1/32) \propto g_{\text{LOP}}(x^2)$ further highlights the close similarity to Mean Shift. Due to the broad applicability of the concepts of robust statistics [79], i.e. the definition and optimization of loss functions, to traditional approaches, machine learning, and many other fields such as Mean Shift, our losses ρ_Γ may also be highly relevant for deep learning techniques and extend previous generalizations [82]. Fig. 5 shows a comparison between the Gaussian M-estimator ($p = 2$) and the LOP M-estimator ($p = 1$).

D. Neural Network Priors

In the context of learning-based point cloud denoising, many approaches rely on clean groundtruth data in order to learn how various types of synthetically generated noise such as Gaussian, uniform, or simulated sensor noise can be effectively removed from the given point set. However, they may not necessarily generalize to real-world data exhibiting different possibly systematic yet unknown noise patterns and, hence, lead to less reliable and accurate results. For this purpose, TotalDenoising [50] formulates the learning objective in an unsupervised manner, which only requires noisy input data without further knowledge about the noise itself or corresponding clean data. Given noisy




Fig. 5. Comparison of LOP and Gaussian M-estimators for $\sigma^2 = 1/2$. Due to the close relation to Mean Shift, these robust loss functions ρ do not only share the shape of the corresponding kernels K but also have similar properties and form localized versions of the common global L_2 and L_1 loss functions.

points $\mathbf{p} \in \mathbb{R}^3$, the parameters Θ of a neural network f_Θ are optimized such that the observations $\mathbf{p}$ are projected to the unknown surface $\mathcal{S}$ represented by the distribution of modes $p(\mathbf{x} | \mathcal{S})$. In particular, this can be formulated by the objective function

$$L(\Theta) = \mathbb{E}_{\mathbf{p} \sim p(\mathbf{x} | \mathcal{S})} \mathbb{E}_{\mathbf{q} \sim q(\mathbf{x} | \mathbf{p})} \rho(f_\Theta(\mathbf{p}, \mathbf{q})) \quad (33)$$

where ρ denotes a loss function, e.g. the commonly used L_2 loss or other choices such as robust losses (see Section V-C), and $\mathbf{q}$ denotes the generated samples during training. Since any point $\mathbf{x} \in \mathcal{S}$ on the surface is a valid target onto which the noisy input $\mathbf{p}$ can be projected, the unsupervised training procedure will not converge to a unique solution. Thus, the prior distribution

$$q(\mathbf{x} | \mathbf{p}) = p(\mathbf{x} | \mathcal{S}) K(\mathbf{W}(\mathbf{x} - \mathbf{p})) \quad (34)$$

regularizes this problem by convolving the distribution $p(\mathbf{x} | \mathcal{S})$ with a smoothing kernel K and a diagonal weight matrix $\mathbf{W} \in \mathbb{R}^{3 \times 3}$ to favor the closest mode as the unique solution for projection. This convolution-based regularization is similar to the formulation of the kernel density estimate (1) in the Mean Shift approach. Therefore, instead of using a Gaussian kernel as in the original TotalDenoising approach, we can apply the gained insights of our kernel family K_Γ regarding feature preservation (see Section IV-A and Fig. 3) and use the LOP kernel K_{LOP} given in (11) in the prior distribution.

VI. EXPERIMENTAL RESULTS

In the following, we demonstrate the effectiveness of our proposed extensions in Section V that are derived from the theoretical properties of the kernel family (see Section IV).

A. Evaluation of WLOP Density Weights

In order to evaluate the performance of our density weighting schemes defined in (23) and (24) (see Section V-A), we measured the regularity of the point cloud $\mathcal{Q}$ after projection onto a highly irregular target $\mathcal{P}$ [2]. For this purpose, we sampled 74 000 target points from a 3D surface patch inversely proportional to the intensity of the mapped *Bird* image and took a random subset of 3700 points for projection. Then, we applied 100 iterations of the LOP operator as well as its weighted versions and computed




Fig. 6. Kernel density estimate $\hat{f}$ for $h = 3$ of a planar target surface patch $\mathcal{P}$ (74 000 points) that is sampled inversely proportional to the intensity of the *Bird* image. The regularity σ_Q of any projected point set $\mathcal{Q}$ onto this target directly depends on the uniformity of $\hat{f}$. Whereas WLOP [2] and our simple weighting scheme cannot fully remove high-frequency variations, our full weighting scheme leads to a significantly better normalization and more uniform density. Unit of h : [%BB diagonal].




Fig. 7. Regularity σ_Q of the input subset $\mathcal{Q}$ (3700 points) projected on a planar target surface patch $\mathcal{P}$ (74 000 points) that is sampled inversely proportional to the intensity of the *Bird* image. Due to the better density normalization, our weighting schemes further improve the regularity across various combinations of the window size h and the repulsion weight μ . Unit of h and σ_Q : [% BB diagonal].

the regularity

$$\sigma_Q = \left[\frac{1}{|\mathcal{Q}|} \sum_i \left(d(\mathbf{q}_i, \mathcal{Q} \setminus \{\mathbf{q}_i\}) - \overline{d(\mathbf{q}_i, \mathcal{Q} \setminus \{\mathbf{q}_i\})} \right)^2 \right]^{\frac{1}{2}} \quad (35)$$

which is defined as the standard deviation of the nearest neighbor distances $d(\mathbf{x}, \mathcal{Y}) = \min_j \|\mathbf{x} - \mathbf{y}_j\|$ within the point cloud $\mathcal{Q}$. Fig. 7 shows the quantitative results for 60×50 combinations of h and μ . Throughout 76.7% of all combinations, our simple weighting scheme (23) performs better than WLOP with a slightly lower value of σ_Q on average. Our full scheme (24) outperforms WLOP in 99.3% and the simple scheme in 98.7% of all combinations, especially in configurations with low repulsion weights $\mu \in [0, 0.2]$.

These improvements in point cloud regularity directly correspond to a more evenly distributed density along the surface. Fig. 6 depicts a comparison of 1000×1000 evenly sampled density values on the respective 3D surface patch. Both WLOP

and our simple scheme (23) normalize the lower frequency components of the density, but still retain high-frequency variations due to the independent computation of each weight. On the other hand, our full scheme (24) does not suffer from these artifacts and only leads to underestimated densities at the boundary and in sparsely sampled regions where the window size h is not sufficiently large to bridge these gaps.

Since Mean Shift and, thereby, LOP and its variants are scale-invariant with respect to a global normalization constant, we computed the mean density $\bar{f}$ for each weighting scheme and used this value to normalize each density distribution for a fair comparison. Whereas this value is close to one for both of our schemes due to the correct handling of the normalization constants, we can derive a theoretical estimate of this value for WLOP

$$\bar{f}_{\text{WLOP}} \approx \frac{1}{|\mathcal{P}| h^3} c_{\theta}^{(d=3)} \frac{c_{\text{LOP}}^{(d=3)} c_{\theta}^{(d=2)}}{c_{\text{LOP}}^{(d=2)} c_{\theta}^{(d=3)}} = 2.396 \quad (36)$$

which consists of three terms: 1) the normalization constant of the kernel density estimate in (1); 2) the missing normalization constant of the kernel θ ; and 3) a dimension-dependent correction factor. The last term models the different domains from which the density is accumulated as we consider a surface patch that corresponds to a 2D subspace embedded in the 3D space. Therefore, the integration domain of the density differs by one dimension which can be accounted for by the ratio of the normalization constants of both the actual density kernel K_{LOP} as well as the chosen kernel θ for density weight computation.

B. Evaluation of CLOP Kernel Approximation

We evaluated the approximation error of our fitted parameter set in Table II (see Section V-B) against the original one proposed by CLOP [5]. First, we quantified systematic errors of the kernel approximation $\hat{K}_{\text{LOP}}$ by analyzing its standard deviation $\|\hat{\Sigma}_{\text{LOP}}\|^{1/2}$ defined by the square root of the magnitude of its covariance matrix (26). Table III indicates that both CLOP and our approximation underestimate the actual value $\|\Sigma_{\text{LOP}}\|^{1/2}$ and that the bias increases in higher dimensions. Although these errors are significantly lower for our approximation throughout all dimensions, they may still be noticeable. Our consistent approximation always overestimates the actual standard deviation and has a slightly higher error than the unconstrained variant in

TABLE III
RATIO OF STANDARD DEVIATIONS $\|\hat{\Sigma}_{\text{LOP}}\|^{1/2}/\|\Sigma_{\text{LOP}}\|^{1/2}$ BETWEEN THE
KERNEL APPROXIMATIONS $\hat{K}_{\text{LOP}}$ AND THE ORIGINAL KERNEL K_{LOP} IN $\mathbb{R}^d$

	$d = 1$	$d = 2$	$d = 3$	$d = 4$	$d \rightarrow \infty$
CLOP [5]	0.7931	0.7632	0.7430	0.7291	0.6659
Ours	0.9917	0.9834	0.9737	0.9640	0.8881
Ours (Consistent)	1.0137	1.0252	1.0362	1.0440	1




Fig. 8. Analysis of bias in the width of the kernel approximations $\hat{K}_{\text{LOP}}$ to the original kernel K_{LOP} . A correction by scaling the parameters $\hat{\sigma}_k$ with a global factor $1/b_{\text{opt}}$ lowers the error for CLOP [5]. Nevertheless, our approximation still better follows K_{LOP} with a significantly lower error and is almost unbiased in the 1-dimensional case.

low dimensions up to $d = 5$. However, it becomes unbiased in the limit $d \rightarrow \infty$ and should be preferred in higher dimensions. We also considered minimizing the L_1 distance of the kernel approximation $\hat{K}_{\text{LOP}}$ (25) to the actual kernel K_{LOP} (11) by scaling the values $\hat{\sigma}_k$ with correction factors $1/b_{\text{opt}}$ to obtain an improved set of parameters which is shown in Fig. 8. Here, the error of our approximation is significantly lower than for CLOP both before and after optimal correction. Furthermore, the optimal scaling factors are similar to the ratios of the standard deviation for $d = 1$.

In addition to the theoretical analysis of the kernel approximations, we also measured the reconstruction error when replacing the actual kernel in the update step (7) of the projection with the respective approximation. For this purpose, we chose the *Block* model and uniformly sampled 50 000 target points $\mathcal{P}$ and 25 000 projection points $\mathcal{Q}$ respectively. We applied 100 iterations of WLOP as a smoothing operator with a large window size of $h = 25$ percent of the bounding box diagonal of $\mathcal{P}$ and a repulsion weight $\mu = 0.4$. Then, we measured the distance of each point to the (triangulated) surface of the reference point cloud as well as the mean point-surface distance

$$d_{\text{surface}}(\mathcal{X}, \mathcal{Y}) = \frac{1}{|\mathcal{X}|} \sum_i \min_j d(\mathbf{x}_i, t(\mathbf{y}_j)) \quad (37)$$

where $t(\mathbf{y}_j)$ denotes the j -th triangle of $\mathcal{Y}$. Fig. 9 shows the results of this point cloud smoothing operation. Whereas the




Fig. 9. Bias of the kernel approximations $\hat{K}_{\text{LOP}}$ when applying WLOP [2] to smooth the *Block* model (25 000 points) with the window size $h = 25$. Whereas the CLOP [5] approximation introduces systematic errors at the edges due to the bias in the width, our variant closely resembles the behavior of the original kernel K_{LOP} . Unit: [% BB diagonal].

CLOP approximation introduces higher errors at the edges of the sampled model due to the significantly underestimated standard deviation of the kernel, our approximation does not suffer from these artifacts.

C. Evaluation of Robust Loss Functions

We tested the LOP M-estimator with the corresponding loss ρ_{LOP} (32) against other popular choices, i.e. L_2 , L_1 , and Gaussian M-estimators, for normal filtering in the context of mesh denoising based on the iterative update step in (28) (see Section V-C). For this, we used the *Gargoyle* (86 311 vertices, 172 610 faces) and *Box* (70 134 vertices, 140 259 faces) models and corrupted the vertices in random directions by $0.25 \bar{l}_e$ uniform noise where $\bar{l}_e$ denotes the average face edge length. Then, we applied 50 iterations of normal filtering with $\sigma = 0.3$ for each face normal $\mathbf{n}$ within its geometric neighborhood of size $r = 1.5 \bar{l}_e$, that is all normals whose face centers are traversable along the surface within a ball of size r . To avoid the singularity of the L_1 and LOP losses at $x = 0$, we only considered the neighboring face normals in the initial iteration and used all normals subsequently. For the second stage of the mesh denoising framework, we used the vertex update by Zhang et al. [83] with their default parameters of 20 iterations and $w = 0.001$ which avoids the triangle flipping problem. We evaluated the reconstruction error by the mean angular distance

$$d_{\text{angle}}(\mathcal{X}, \mathcal{Y}) = \frac{1}{|\mathcal{X}|} \sum_i \left[\frac{360}{2\pi} \arccos(\langle \mathbf{n}(\mathbf{x}_i) | \mathbf{n}(\mathbf{y}_i) \rangle) \right] \quad (38)$$

to the face normals of the ground truth mesh. Figs. 10 and 11 show comparisons between the L_2 , L_1 , Gaussian, and LOP loss. Whereas the L_2 loss leads to a very smooth surface, its L_1 counterpart is less sensitive to large normal variations within the local neighborhood and better preserves features. However, sharp edges cannot be reconstructed since all collected normals are considered in a global fashion. The Gaussian and LOP loss functions can be viewed as localized versions of the former




Fig. 10. Mesh denoising of the *Gargoyle* model (86 311 vertices, 172 610 faces) corrupted with $0.25 \bar{l}_e$ uniform noise. Although the mean angular distance d_{angle} is slightly higher for the LOP loss ρ_{LOP} , features and finer details are better preserved. Unit: $[\circ]$.




Fig. 11. Mesh denoising of the *Box* model (70 134 vertices, 140 259 faces) corrupted with $0.25 \bar{l}_e$ uniform noise. Filtering with the LOP loss ρ_{LOP} results in the lowest mean angular distance d_{angle} and reconstructs fine details best. Unit: $[\circ]$.

losses and do not suffer from this limitation. Finer details being at a similar scale as the applied noise are hard to reconstruct and mostly smoothed out by all variants, but can be partially recovered by the LOP loss.

D. Evaluation of Neural Network Priors

We tested the application of different neural network priors for unsupervised point cloud denoising (see Section V-D). For this purpose, we used the reference implementation of TotalDenoising [50] and trained the method in an unsupervised manner on the real-world large-scale *Paris-rue-Madame* dataset [84], which was captured by a LiDAR sensor and has no groundtruth data available. In particular, we employed the same training procedure and default parameter set in all cases, that is the scaling parameter $\alpha = 0.5$ as well as the standard deviation $\sigma = 0.5657$ for the kernel in the prior (34), and only replaced the Gaussian kernel by the LOP kernel. Finally, we used the trained models to denoise the test split of the dataset. Fig. 12 shows a comparison between the Gaussian and LOP kernel priors. While the LiDAR noise is reliably removed from the test datasets for both kernels, filtering with the Gaussian kernel leads to significant oversmoothing of smaller features and, in addition, to a noticeable shrinkage of the point cloud. Furthermore, holes may be introduced at sharp features like edges since the Gaussian prior tends to move nearby points into clusters apart from these features. In contrast to this, the LOP kernel with equal standard deviation σ preserves higher frequency information and, in turn, sharp features without introducing holes. The shrinkage effect also is greatly mitigated which becomes apparent at thin structures, like the sash bars in the windows, where the width is

significantly better preserved and closely matches the expected width from the input data.

In addition to these qualitative results, we also performed a quantitative comparison between recent learning-based approaches. For a fair comparison, we only considered unsupervised methods, that is TotalDenoising (TD) [50] with the original Gaussian kernel prior as well as with our LOP kernel prior, and Score-based denoising [51] which, although being supervised, also supports training in an unsupervised manner. We used the *ModelNet40* dataset [85] provided by TD [50] that consists of 15 collected object classes with 5 training meshes and 2 test meshes per class, from which point clouds were sampled using Poisson disk sampling and corrupted by Gaussian noise with standard deviations of $\sigma_{\text{noise}} = 0.5, 1, 1.5$ percent of the bounding box diagonal. After training all methods with default parameters on this dataset, we applied them to denoise the test data using 1, 4, 7 iterations with TD and 1, 1, 2 iterations with Score-based denoising respectively for the different noise levels. We measured the performance with the commonly used Chamfer distance

$$d_{\text{CD}}(\mathcal{X}, \mathcal{Y}) = \frac{1}{|\mathcal{X}|} \sum_i d(\mathbf{x}_i, \mathcal{Y})^2 + \frac{1}{|\mathcal{Y}|} \sum_i d(\mathbf{y}_i, \mathcal{X})^2 \quad (39)$$

where $d(\mathbf{x}_i, \mathcal{Y}) = \min_j \|\mathbf{x}_i - \mathbf{y}_j\|$ denotes the nearest neighbor distance of $\mathbf{x}_i$ to the point cloud $\mathcal{Y}$ as well as with the point-to-mesh distance

$$d_{\text{P2M}}(\mathcal{X}, \mathcal{Y}) = \frac{1}{|\mathcal{X}|} \sum_i \min_j d(\mathbf{x}_i, t(\mathbf{y}_j))^2 + \frac{1}{|\mathcal{Y}|} \sum_j \min_i d(\mathbf{x}_i, t(\mathbf{y}_j))^2 \quad (40)$$




Fig. 12. Point cloud denoising of the *Paris-rue-Madame* dataset with TotalDenoising [50] for the Gaussian and LOP kernels as neural network priors. Performing the unsupervised training with the LOP prior better preserves features and edges and leads to significantly less oversmoothing and shrinkage. Furthermore, using the LOP prior also better handles the systematic pattern induced by the scanning process.

TABLE IV
COMPARISON BETWEEN UNSUPERVISED LEARNING-BASED POINT CLOUD DENOISING METHODS

	$\sigma_{\text{noise}} = 0.5$		$\sigma_{\text{noise}} = 1$		$\sigma_{\text{noise}} = 1.5$	
	d_{CD}	d_{P2M}	d_{CD}	d_{P2M}	d_{CD}	d_{P2M}
Score (Unsup.) [51]	0.150	0.115	0.404	0.315	0.821	0.609
TD (K_{Gaussian}) [50]	0.173	0.150	0.343	0.303	0.466	0.428
TD (K_{LOP})	0.175	0.148	0.282	0.246	0.424	0.389

Unit: [10^{-4}].

where $d(x_i, t(y_j))$ is the distance of x_i to the triangle $t(y_j)$ as in (37). Table IV shows the results of all methods averaged over the test data. At the low noise level $\sigma_{\text{noise}} = 0.5$, Score-based denoising achieves slightly better results than TotalDenoising where the performance is very similar between the Gaussian and the LOP kernel prior. On the other hand, applying the LOP kernel prior significantly increases the robustness of TotalDenoising

at higher noise levels and also outperforms all other approaches with respect to both the Chamfer and the point-to-mesh distance.

VII. CONCLUSION

We presented incomplete gamma kernels, a novel family of kernels generalizing LOP operators. By revisiting the classical localized L_1 estimator used in LOP, we revealed its relation to the Mean Shift framework via a novel kernel K_{LOP} and generalized this result to arbitrary localized L_p estimators. We derived several theoretical properties of the kernel family K_Γ concerning distributional, Mean Shift induced, and other aspects such as strict positive definiteness to obtain a deeper understanding of the operator's projection behavior. Furthermore, we illustrated several applications including an improved WLOP density weighting scheme, a more accurate kernel approximation for CLOP, incomplete gamma losses ρ_Γ as a novel set of robust loss functions, as well as better neural network priors and confirmed their effectiveness in a variety of quantitative and qualitative experiments. We expect that building upon the insights provided by our work will be beneficial for future developments on point cloud denoising as well as in many other related fields, including applications beyond pure denoising.

ACKNOWLEDGMENT

The authors would like to thank the reviewers for their valuable comments which helped them to improve the quality of their manuscript.

REFERENCES

- [1] Y. Lipman, D. Cohen-Or, D. Levin, and H. Tal-Ezer, "Parameterization-free projection for geometry reconstruction," *ACM Trans. Graph.*, vol. 26, no. 3, pp. 22-es, 2007.
- [2] H. Huang, D. Li, H. Zhang, U. Ascher, and D. Cohen-Or, "Consolidation of unorganized point clouds for surface reconstruction," *ACM Trans. Graph.*, vol. 28, no. 5, pp. 1–7, 2009.
- [3] B. Liao, C. Xiao, L. Jin, and H. Fu, "Efficient feature-preserving local projection operator for geometry reconstruction," *Comput. Aided Des.*, vol. 45, no. 5, pp. 861–874, 2013.
- [4] H. Huang, S. Wu, M. Gong, D. Cohen-Or, U. Ascher, and H. Zhang, "Edge-aware point set resampling," *ACM Trans. Graph.*, vol. 32, no. 1, 2013, Art. no. 9.
- [5] R. Preiner, O. Mattausch, M. Arikian, R. Pajarola, and M. Wimmer, "Continuous projection for fast L_1 reconstruction," *ACM Trans. Graph.*, vol. 33, no. 4, 2014, Art. no. 47.
- [6] D. Zhang, X. Lu, H. Qin, and Y. He, "Pointfilter: Point cloud filtering via encoder-decoder modeling," *IEEE Trans. Vis. Comput. Graphics*, vol. 27, no. 3, pp. 2015–2027, Mar. 2021.
- [7] X. Xu, G. Geng, X. Cao, K. Li, and M. Zhou, "TDNet: Transformer-based network for point cloud denoising," *Appl. Opt.*, vol. 61, no. 6, pp. C80–C88, 2022.
- [8] K. Fukunaga and L. Hostetler, "The estimation of the gradient of a density function, with applications in pattern recognition," *IEEE Trans. Inf. Theory*, vol. 21, no. 1, pp. 32–40, Jan. 1975.
- [9] Y. Cheng, "Mean shift, mode seeking, and clustering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 17, no. 8, pp. 790–799, Aug. 1995.
- [10] D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 5, pp. 603–619, May 2002.
- [11] P. Madaan, A. Maiti, S. Anand, and S. Mittal, "Deep mean shift clustering," Bachelor's Thesis, Indraprastha Inst. Inf. Technol., 2019.
- [12] S. A. Bigdeli and M. Zwicker, "Image restoration using autoencoding priors," in *Proc. Int. Joint Conf. Comput. Vis. Imag. Comput. Graph. Theory Appl.*, 2018, pp. 33–44.

- [13] G. Alain and Y. Bengio, "What regularized auto-encoders learn from the data-generating distribution," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 3563–3593, 2014.
- [14] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Surface reconstruction from unorganized points," in *Proc. Annu. Conf. Comput. Graph. Interactive Techn.*, 1992, pp. 71–78.
- [15] J. C. Carr et al., "Reconstruction and representation of 3D objects with radial basis functions," in *Proc. Annu. Conf. Comput. Graph. Interactive Techn.*, 2001, pp. 67–76.
- [16] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C. T. Silva, "Computing and rendering point set surfaces," *IEEE Trans. Vis. Comput. Graphics*, vol. 9, no. 1, pp. 3–15, First Quarter, 2003.
- [17] N. Amenta and Y. J. Kil, "Defining point-set surfaces," *ACM Trans. Graph.*, vol. 23, no. 3, pp. 264–270, 2004.
- [18] S. Fleishman, D. Cohen-Or, and C. T. Silva, "Robust moving least-squares fitting with sharp features," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 544–552, 2005.
- [19] A. C. Öztireli, G. Guennebaud, and M. Gross, "Feature preserving point set surfaces based on non-linear kernel regression," *Comput. Graph. Forum*, vol. 28, no. 2, pp. 493–501, 2009.
- [20] E. A. L. Narváez and N. E. L. Narváez, "Point cloud denoising using robust principal component analysis," in *Proc. Int. Conf. Comput. Graph. Theory Appl.*, 2006, pp. 51–58.
- [21] E. Mattei and A. Castrodad, "Point cloud denoising via moving RPCA," *Comput. Graph. Forum*, vol. 36, no. 8, pp. 123–137, 2017.
- [22] X. Lu, S. Wu, H. Chen, S.-K. Yeung, W. Chen, and M. Zwicker, "GPF: GMM-inspired feature-preserving point set filtering," *IEEE Trans. Vis. Comput. Graphics*, vol. 24, no. 8, pp. 2315–2326, Aug. 2018.
- [23] Y. Sun, S. Schaefer, and W. Wang, "Denoising point sets via L_0 minimization," *Comput. Aided Geometric Des.*, vol. 35/36, pp. 2–15, 2015.
- [24] X. Cheng, M. Zeng, J. Lin, Z. Wu, and X. Liu, "Efficient L_0 resampling of point sets," *Comput. Aided Geometric Des.*, vol. 75, 2019, Art. no. 101790.
- [25] H. Avron, A. Sharf, C. Greif, and D. Cohen-Or, " ℓ_1 -Sparse reconstruction of sharp point set surfaces," *ACM Trans. Graph.*, vol. 29, no. 5, 2010, Art. no. 135.
- [26] E. Leal, G. Sanchez-Torres, and J. W. Branch, "Sparse regularization-based approach for point cloud denoising and sharp features enhancement," *Sensors*, vol. 20, no. 11, 2020, Art. no. 3206.
- [27] J. Digne, S. Valette, and R. Chaine, "Sparse geometric representation through local shape probing," *IEEE Trans. Vis. Comput. Graphics*, vol. 24, no. 7, pp. 2238–2250, Jul. 2018.
- [28] J. Zeng, G. Cheung, M. Ng, J. Pang, and C. Yang, "3D point cloud denoising using graph Laplacian regularization of a low dimensional manifold model," *IEEE Trans. Image Process.*, vol. 29, pp. 3474–3489, 2020.
- [29] W. Hu, X. Gao, G. Cheung, and Z. Guo, "Feature graph learning for 3D point cloud denoising," *IEEE Trans. Image Process.*, vol. 68, pp. 2841–2856, 2020.
- [30] C. Dinesh, G. Cheung, and I. V. Bajić, "Point cloud denoising via feature graph Laplacian regularization," *IEEE Trans. Image Process.*, vol. 29, pp. 4143–4158, 2020.
- [31] M. A. Irfan and E. Magli, "Joint geometry and color point cloud denoising based on graph wavelets," *IEEE Access*, vol. 9, pp. 21149–21166, 2021.
- [32] Y. Zheng, G. Li, S. Wu, Y. Liu, and Y. Gao, "Guided point cloud denoising via sharp feature skeletons," *Vis. Comput.*, vol. 33, no. 6, pp. 857–867, 2017.
- [33] C. Shi et al., "Three-dimensional point cloud denoising via a gravitational feature function," *Appl. Opt.*, vol. 61, no. 6, pp. 1331–1343, 2022.
- [34] W. Hu, Q. Hu, Z. Wang, and X. Gao, "Dynamic point cloud denoising via manifold-to-manifold distance," *IEEE Trans. Image Process.*, vol. 30, pp. 6168–6183, 2021.
- [35] O. Remil, Q. Xie, X. Xie, K. Xu, and J. Wang, "Surface reconstruction with data-driven exemplar priors," *Comput. Aided Des.*, vol. 88, pp. 31–41, 2017.
- [36] G. Rosman, A. Dubrovina, and R. Kimmel, "Patch-collaborative spectral point-cloud denoising," *Comput. Graph. Forum*, vol. 32, no. 8, pp. 1–12, 2013.
- [37] X. Lu, S. Schaefer, J. Luo, L. Ma, and Y. He, "Low rank matrix approximation for 3D geometry filtering," *IEEE Trans. Vis. Comput. Graphics*, vol. 28, no. 4, pp. 1835–1847, Apr. 2022.
- [38] H. Chen, M. Wei, Y. Sun, X. Xie, and J. Wang, "Multi-patch collaborative point cloud denoising via low-rank recovery with graph constraint," *IEEE Trans. Vis. Comput. Graphics*, vol. 26, no. 11, pp. 3255–3270, Nov. 2020.
- [39] D. Zhu et al., "Non-local low-rank point cloud denoising for 3D measurement surfaces," *IEEE Trans. Instrum. Meas.*, vol. 71, 2022, Art. no. 5002214.
- [40] F. Zaman, Y. P. Wong, and B. Y. Ng, "Density-based denoising of point cloud," in *Proc. Int. Conf. Robot. Vis. Signal Process. Power Appl.*, 2017, pp. 287–295.
- [41] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 77–85.
- [42] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep hierarchical feature learning on point sets in a metric space," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 5105–5114.
- [43] W. Yifan, S. Wu, H. Huang, D. Cohen-Or, and O. Sorkine-Hornung, "Patch-based progressive 3D point set upsampling," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 5951–5960.
- [44] J. Choe, B. Joung, F. Rameau, J. Park, and I. S. Kweon, "Deep point cloud reconstruction," in *Proc. Int. Conf. Learn. Representation*, 2022.
- [45] C. Duan, S. Chen, and J. Kovacevic, "3D point cloud denoising via deep neural network based local surface estimation," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2019, pp. 8553–8557.
- [46] M.-J. Rakotosaona, V. La Barbera, P. Guerrero, N. J. Mitra, and M. Ovsjanikov, "PointCleanNet: Learning to denoise and remove outliers from dense point clouds," *Comput. Graph. Forum*, vol. 39, no. 1, pp. 185–203, 2020.
- [47] L. Yu, X. Li, C.-W. Fu, D. Cohen-Or, and P.-A. Heng, "EC-Net: An edge-aware point set consolidation network," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 398–414.
- [48] H. Zhou, K. Chen, W. Zhang, H. Fang, W. Zhou, and N. Yu, "DUP-Net: Denoiser and upsampler network for 3D adversarial point clouds defense," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 1961–1970.
- [49] F. Pistilli, G. Fracastoro, D. Valsesia, and E. Magli, "Learning graph-convolutional representations for point cloud denoising," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 103–118.
- [50] P. Hermosilla, T. Ritschel, and T. Ropinski, "Total denoising: Unsupervised learning of 3D point cloud cleaning," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 52–60.
- [51] S. Luo and W. Hu, "Score-based point cloud denoising," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 4563–4572.
- [52] H. Chen, B. Du, S. Luo, and W. Hu, "Deep point set resampling via gradient fields," 2021, *arXiv:2111.02045*.
- [53] Y. Zhao, H. Zheng, Z. Wang, J. Luo, and E. Y. Lam, "Point cloud denoising via momentum ascent in gradient fields," 2022, *arXiv:2202.10094*.
- [54] R. Roveri, A. C. Öztireli, I. Pandele, and M. Gross, "PointProNets: Consolidation of point clouds with convolutional neural networks," *Comput. Graph. Forum*, vol. 37, no. 2, pp. 87–99, 2018.
- [55] W. Yifan, F. Serena, S. Wu, C. Öztireli, and O. Sorkine-Hornung, "Differentiable surface splatting for point-based geometry processing," *ACM Trans. Graph.*, vol. 38, no. 6, 2019, Art. no. 230.
- [56] S. Luo and W. Hu, "Differentiable manifold reconstruction for point cloud denoising," in *Proc. ACM Int. Conf. Multimedia*, 2020, pp. 1330–1338.
- [57] R. Hanocka, G. Metzer, R. Gryes, and D. Cohen-Or, "Point2Mesh: A self-prior for deformable meshes," *ACM Trans. Graph.*, vol. 39, no. 4, 2020, Art. no. 126.
- [58] C. Huang, R. Li, X. Li, and C.-W. Fu, "Non-local part-aware point cloud denoising," 2020, *arXiv:2003.06631*.
- [59] Y. Chen et al., "Shape self-correction for unsupervised point cloud understanding," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 8362–8371.
- [60] A. Mao, Z. Du, Y.-H. Wen, J. Xuan, and Y.-J. Liu, "PD-Flow: A point cloud denoising framework with normalizing flows," 2022, *arXiv:2203.05940*.
- [61] H. Chen, Z. Wei, X. Li, Y. Xu, M. Wei, and J. Wang, "RePCD-Net: Feature-aware recurrent point cloud denoising network," *Int. J. Comput. Vis.*, vol. 130, pp. 615–629, 2022.
- [62] C. Grillenzoni, "Design of blurring mean-shift algorithms for data classification," *J. Classification*, vol. 33, no. 2, pp. 262–281, 2016.
- [63] G. Beck, T. Duong, M. Lebbah, H. Azzag, and C. Cérin, "A distributed and approximated nearest neighbors algorithm for an efficient large scale mean shift clustering," *J. Parallel Distrib. Comput.*, vol. 134, pp. 128–139, 2019.
- [64] J. Jang and H. Jiang, "MeanShift++: Extremely fast mode-seeking with applications to segmentation and object tracking," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 4100–4111.
- [65] S. A. Bigdeli, M. Zwicker, P. Favaro, and M. Jin, "Deep mean-shift priors for image restoration," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 763–772.

- [66] X. Li, Z. Hu, and F. Wu, "A note on the convergence of the mean shift," *Pattern Recognit.*, vol. 40, no. 6, pp. 1756–1762, 2007.
- [67] T.-L. Chen, "On the convergence and consistency of the blurring mean-shift process," *Ann. Inst. Statist. Math.*, vol. 67, no. 1, pp. 157–176, 2015.
- [68] Y. A. Ghassabeh, "A sufficient condition for the convergence of the mean shift algorithm with gaussian kernel," *J. Multivariate Anal.*, vol. 135, pp. 1–10, 2015.
- [69] K. Huang, X. Fu, and N. Sidiropoulos, "On convergence of epanechnikov mean shift," in *Proc. AAAI Conf. Artif. Intell.*, 2018, pp. 3263–3270.
- [70] R. Yamasaki and T. Tanaka, "Properties of mean shift," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 9, pp. 2273–2286, Sep. 2020.
- [71] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, "Extracting and composing robust features with denoising autoencoders," in *Proc. Int. Conf. Mach. Learn.*, 2008, pp. 1096–1103.
- [72] E. Weiszfeld, "Sur le point pour lequel la somme des distances de n points donnés est minimum," *Tohoku Math. J.*, vol. 43, pp. 355–386, 1937.
- [73] H. Bateman, *Higher Transcendental Functions*, vol. 2. New York, NY, USA: McGraw-Hill, 1953.
- [74] E. M. Stein and G. Weiss, *Introduction to Fourier Analysis on Euclidean Spaces*. Princeton, NJ, USA: Princeton Univ. Press, 1971.
- [75] P. Natalini and B. Palumbo, "Inequalities for the incomplete gamma function," *Math. Inequalities Appl.*, vol. 3, no. 1, pp. 69–77, 2000.
- [76] S. Faigenbaum-Golovin and D. Levin, "Manifold reconstruction and denoising from scattered data in high dimension," *J. Comput. Appl. Math.*, vol. 421, 2023, Art. no. 114818.
- [77] R. L. Schilling, R. Song, and Z. Vondracek, *Bernstein Functions: Theory and Applications*. Berlin, Germany: De Gruyter, 2012.
- [78] I. J. Schoenberg, "Metric spaces and completely monotone functions," *Ann. Math.*, vol. 39, no. 4, pp. 811–841, 1938.
- [79] F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel, *Robust Statistics: The Approach Based on Influence Functions*. Hoboken, NJ, USA: Wiley, 1986.
- [80] M. J. Black, G. Sapiro, D. H. Marimont, and D. Heeger, "Robust anisotropic diffusion," *IEEE Trans. Image Process.*, vol. 7, no. 3, pp. 421–432, Mar. 1998.
- [81] S. K. Yadav, M. Skrodzki, E. Zimmermann, and K. Polthier, "Surface denoising based on normal filtering in a robust statistics framework," in *Proc. Forum Math-for- Ind.*, 2021, pp. 103–132.
- [82] J. T. Barron, "A general and adaptive robust loss function," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4326–4334.
- [83] J. Zhang, B. Deng, Y. Hong, Y. Peng, W. Qin, and L. Liu, "Static/Dynamic filtering for mesh geometry," *IEEE Trans. Vis. Comput. Graphics*, vol. 25, no. 4, pp. 1774–1787, Apr. 2019.
- [84] A. Serna, B. Marcotegui, F. Goulette, and J.-E. Deschaud, "Paris-rue-madame database: A 3D mobile laser scanner dataset for benchmarking urban detection, segmentation and classification methods," in *Proc. Int. Conf. Pattern Recognit. Appl. Methods*, 2014, pp. 819–824.
- [85] Z. Wu et al., "3D ShapeNets: A deep representation for volumetric shapes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1912–1920.



Patrick Stotko received the BSc and MSc degrees in computer science from the University of Bonn, Germany, in 2014 and 2016, respectively. He is currently working toward the PhD degree with the Visual Computing Department, University of Bonn, Germany. His research interests include interactive geometry and appearance reconstruction techniques for AR and VR applications.



Michael Weinmann received the Dipl-Ing degree in electrical engineering and information technology from the University of Karlsruhe, Germany, in 2009, and the PhD degree in computer science from the University of Bonn, Germany, in 2016. He is an assistant professor with the Intelligent Systems Department, Delft University of Technology, The Netherlands. His current research interests include object and scene digitization, digital material appearance, scene representations, scene interpretation, data visualization, live telepresence, machine learning, computational

arts, and cultural heritage.



Reinhard Klein received the Dipl-Math degree in mathematics and the PhD degree in computer science from the University of Tübingen, Germany, in 1989 and 1995, respectively. He is a professor with the University of Bonn, Germany, and the director of the Visual Computing Department. In 1999, he became associate professor with the University of Darmstadt, Germany, and head of the Research Group Animation and Image Communication, Fraunhofer Institute for Computer Graphics. Since 2000, he is a professor with the University of Bonn, Germany, and director of the Institute of Computer Science II – Visual Computing. His research interests include computer graphics, material appearance, rendering, geometry processing, and architectural geometry.