

POINT CLOUD DENOISING VIA MOMENTUM ASCENT IN GRADIENT FIELDS

Yaping Zhao^{1,3,*}, Haitian Zheng^{2,*}, Zhongrui Wang^{1,3}, Jiebo Luo², Edmund Y. Lam^{1,3,†}

¹ The University of Hong Kong, Pokfulam, Hong Kong SAR ² University of Rochester, USA

³ACCESS — AI Chip Center for Emerging Smart Systems, Hong Kong SAR




Fig. 1: Given a noisy point cloud as input, we utilize a neural network to estimate the implicit gradient fields, and iteratively update the position of each point by momentum gradient ascent towards the underlying surface. The color of each 3D point varies according to its z coordinate (perpendicular to the plane).

1. INTRODUCTION

ABSTRACT

To achieve point cloud denoising, traditional methods heavily rely on geometric priors, and most learning-based approaches suffer from outliers and loss of details. Recently, the gradient-based method was proposed to estimate the gradient fields from the noisy point clouds using neural networks, and refine the position of each point according to the estimated gradient. However, the predicted gradient could fluctuate, leading to perturbed and unstable solutions, as well as a long inference time. To address these issues, we develop the momentum gradient ascent method that leverages the information of previous iterations in determining the trajectories of the points, thus improving the stability of the solution and reducing the inference time. Experiments demonstrate that the proposed method outperforms state-of-the-art approaches with a variety of point clouds, noise types, and noise levels. Code is available at: <https://github.com/IndigoPurple/MAG>.

Index Terms— point cloud denoising, point cloud processing, 3D vision

*Equal contribution. †Corresponding author.

This work is supported in part by the Research Grants Council (GRF 17201620), by the Research Postgraduate Student Innovation Award (The University of Hong Kong), and by ACCESS — AI Chip Center for Emerging Smart Systems, Hong Kong SAR.

Point cloud denoising aims to restore clean point clouds from noise-corrupted ones. Due to the inherent limitations of acquisition devices or matching ambiguities in the 3D reconstruction, noise inevitably degrades the quality of scanned or reconstructed point clouds, for which point cloud denoising is favored. Moreover, the quality of point clouds affect the performance of downstream 3D vision tasks, *e.g.*, detection and segmentation. Therefore, point cloud denoising offers crucial preprocessing for relevant 3D vision applications.

Unlike image and video denoising [1], point cloud denoising is challenging because of the intrinsic unordered characteristic of point clouds. Point clouds consist of discrete 3D points irregularly sampled from continuous surfaces. Perturbed by noise, the 3D points can deviate from their original positions and yield the wrong coordinates. To tackle this issue, both traditional [2, 3, 4, 5, 6, 7] and deep learning [8, 9, 10, 11] methods have been explored but show limited performance. Recently, Luo and Hu propose score-based denoising (Score) [12] to iteratively update the point positions according to the estimated gradient fields. However, the predicted gradients can fluctuate, leading to perturbed and unstable solutions, as well as a large inference time.

To improve the performance and efficiency of the gradient-based method, as Figure 1 shows, we propose a novel iterative




Fig. 2: Pipeline of point cloud denoising via momentum ascent in gradient fields.

paradigm of point cloud denoising motivated by the classical momentum method [13] in optimization. Specifically, we employ Score [12] to estimate the gradient fields from the noisy point clouds, and compute the displacement of each point according to the estimated gradient. To avoid the fluctuated gradient of Score, we apply a momentum gradient ascent method that utilizes the previous iterations to seek promising directions to move forward, thus improving the solution stability and inference time. Experiments demonstrate that the proposed method outperforms state-of-the-art approaches.

Our main contributions are as follows:

1. We propose a **simple yet effective** iterative paradigm of point cloud denoising, which leverages past gradients to seek promising directions to move forward.
2. Our method effectively **prevents outliers** that is much more likely to occur for previous approaches.
3. Our method **accelerates optimization** and reduces the inference time of the previous gradient-based approach.
4. Experiments on synthetic and real-world datasets demonstrate the **superior performance** of our method with a variety of point clouds and noise types.

2. RELATED WORK

To perform point cloud denoising, traditional methods [2, 3, 4, 5, 6, 7] heavily rely on geometric priors but show limited performance. Deep learning-based approaches break the performance limit of the point cloud denoising. Among them, some [8, 9, 10, 11, 14] denoise by estimating the deviation of noisy points from the clean surface, but often results in outliers due to the coarse one-step estimation. Others [15] predict the underlying manifold of a noisy point cloud for reconstruction, which loses details in the downsampling stage.

Recently, Score [12] is proposed to tackle the aforementioned issues by iteratively updating the point position in implicit gradient fields learned by neural networks. However, the predicted gradients suffer fluctuation, leading to perturbed and unstable solutions, as well as a large inference time. Moreover, once its estimated gradient deviates from the correct direction, the gradient estimation errors could be accumulated and result in serious outliers.

3. METHOD

Given a noisy point cloud $\mathbf{X} = \{\mathbf{x}\}_{i=1}^N$, we first implement a network that inputs noisy point clouds and outputs point-wise gradients, and then utilize the estimated gradients to denoise point clouds by momentum ascent, as Figure 2 shows.

3.1. Neural Network and Loss Function

The network aims at estimating the gradient in the neighborhood space around $\mathbf{x}_i$, denoted as $\hat{g}_i(\mathbf{x})$. We adopt the network and loss function reported in Score [12], which inputs point coordinates $\mathbf{x} \in \mathbb{R}^3$ surrounding $\mathbf{x}_i$, learns point-wise features f_i , and outputs the estimated gradient $\hat{g}_i(\mathbf{x})$:

$$\hat{g}_i(\mathbf{x}) = G(\mathbf{x} - \mathbf{x}_i, f_i), \quad (1)$$

where $G(\cdot)$ is the gradient estimation network implemented by a multi-layer perceptron. Then the loss function [12] is

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_i = \frac{1}{N} \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{x}_i)} [\|g(\mathbf{x}) - \hat{g}_i(\mathbf{x})\|_2^2], \quad (2)$$

where $\mathcal{N}(\mathbf{x}_i)$ is a distribution concentrated in the neighborhood of $\mathbf{x}_i$ in $\mathbb{R}^3$ space, $g(\mathbf{x})$ is a vector from $\mathbf{x}$ to the ground truth clean surface. Following Score [12], the ensemble gradient is calculated as

$$\mathbf{z}_i(\mathbf{x}) = \frac{1}{k} \sum_{\mathbf{x}_j \in H(\mathbf{x}_i; k)} \hat{g}_j(\mathbf{x}), \quad \mathbf{x}_j \in H(\mathbf{x}_i; k), \quad \mathbf{x} \in \mathbb{R}^3, \quad (3)$$

where $H(\mathbf{x}_i; k)$ is the k -nearest neighborhood of $\mathbf{x}_i$.

Though Score [12] uses the ensemble gradient for robustness, it still suffers in fluctuated gradients. Moreover, once its estimated gradient deviates from the correct direction, the errors could be accumulated and result in serious outliers.

3.2. Momentum Ascent Denoising Algorithm

To alleviate those problems faced by previous gradient-based methods, we propose to perform point cloud denoising by updating point coordinates with momentum gradient ascent. At the beginning of point cloud denoising, we initialize the coordinate for each point according to the input point cloud:

$$\mathbf{x}_i^{(0)} = \mathbf{x}_i, \quad \mathbf{x}_i \in \mathbf{X}. \quad (4)$$




Fig. 3: Comparison of classical gradient ascent and momentum gradient ascent.

# Points		10K (Sparse)						50K (Dense)					
Type	Method	1%		2%		3%		1%		2%		3%	
		CD	P2M	CD	P2M	CD	P2M	CD	P2M	CD	P2M	CD	P2M
Gaussian	Bilateral [16]	3.646	1.342	5.007	2.018	6.998	3.557	0.877	0.234	2.376	1.389	6.304	4.730
	GLR [17]	2.959	1.052	3.773	1.306	4.909	2.114	0.696	0.161	1.587	0.830	3.839	2.707
	PCNet [11]	3.515	1.148	7.467	3.965	13.067	8.737	1.049	0.346	1.447	0.608	2.289	1.285
	DMR [15]	4.482	1.722	4.982	2.115	5.892	2.846	1.162	0.469	1.566	0.800	2.432	1.528
	Score [12]	2.521	0.463	3.686	1.074	4.708	1.942	0.716	0.150	1.288	0.566	1.928	1.041
	Ours	2.498	0.459	3.629	1.054	4.686	1.923	0.706	0.146	1.287	0.557	1.931	1.045
Laplace	GLR [17]	3.223	1.121	4.751	2.090	7.977	4.773	0.962	0.374	3.269	2.325	8.675	7.162
	PCNet [11]	4.616	1.940	11.082	7.218	20.981	15.922	1.190	0.458	2.854	1.868	7.555	6.020
	DMR [15]	4.600	1.811	5.441	2.469	6.918	3.714	1.243	0.537	1.881	1.077	3.609	2.634
	Score [12]	2.915	0.674	4.601	1.799	6.332	3.271	0.823	0.231	1.658	0.869	2.728	1.681
	Ours	2.887	0.669	4.521	1.754	6.237	3.257	0.816	0.212	1.648	0.857	2.717	1.679
Uniform	GLR [17]	1.850	1.015	2.948	1.052	3.400	1.109	0.485	0.071	0.656	0.132	0.903	0.293
	PCNet [11]	1.205	0.337	3.378	1.018	5.044	1.995	0.806	0.228	1.064	0.358	1.218	0.451
	DMR [15]	4.307	1.640	4.445	1.693	4.685	1.857	1.064	0.391	1.159	0.464	1.287	0.572
	Score [12]	1.277	0.248	2.467	0.418	3.079	0.654	0.506	0.047	0.690	0.129	0.917	0.282
	Ours	1.243	0.239	2.404	0.397	3.052	0.638	0.487	0.045	0.679	0.122	0.905	0.278

Table 1: Point cloud denoising comparisons on the PU-Net dataset [18]. CD and P2M are multiplied by 10^4 .

To leverage past gradients, we introduce an auxiliary vector, which is initialized as zero and updated with a leaky average over past gradients:

$$\begin{aligned} \mathbf{v}_i^{(0)} &= \mathbf{0} \in \mathbb{R}^3, \\ \mathbf{v}_i^{(t)} &= \alpha \mathbf{z}_i(\mathbf{x}_i^{(t-1)}) + (1 - \alpha) \mathbf{v}_i^{(t-1)}, \quad t = 1, \dots, T, \end{aligned} \quad (5)$$

where t is the iterative step, α is the momentum weight, $\mathbf{v}_i$ serves to relieve gradient variance and obtain more stable directions of ascent. Finally, the point cloud is denoised by updating point coordinates with momentum gradient ascent:

$$\mathbf{x}_i^{(t)} = \mathbf{x}_i^{(t-1)} + \beta \gamma^t \mathbf{v}_i^{(t)}, \quad t = 1, \dots, T, \quad (6)$$

where t is the iterative step, T is the total number of iteration steps, β is the step size, γ^t is the decay coefficient decreasing towards 0 to ensure convergence. While Score requires a relatively large number of total steps $T = 30$ to conduct their experiments, our method applies momentum gradient ascent, thus reducing the step number to $T = 15$ while achieving better performance.

As Figure 3 shows, compared to classical gradient ascent, momentum gradient ascent leverages the information of previous iterations in determining the trajectories of points, thereby benefiting solution stability and inference time.




Fig. 4: Comparisons on the dense point cloud duck of the PU-Net dataset [18]. The color of each 3D point varies according to its z coordinate (perpendicular to the plane).

4. EXPERIMENT

Experiment Setting. The commonly used PU-Net dataset [18] is adopted for training and testing following previous works [11, 12, 15]. Moreover, we use the real-world dataset Paris-rue-Madame [19] for qualitative evaluation. We utilized a single model trained on the PU-Net training set to conduct both synthetic and real-world experiments. Two traditional methods, bilateral filtering [16], GLR [17], and three deep-learning-based approaches including PCNet [11], DMR [15] and Score [12] are used for comparison. Two commonly used metrics are adopted for quantitative evaluation, the Chamfer distance (CD) [20] and point-to-mesh distance (P2M) [21].

Quantitative and Qualitative Results. As Table 1 shows,




Fig. 5: Point cloud denoising comparisons on the dense point cloud chair of the PU-Net dataset [18]. The color of each 3D point varies according to its z coordinate (perpendicular to the plane).



Fig. 6: Point cloud denoising comparisons on the real-world dataset Paris-rue-Madame [18]. The color of each 3D point varies according to its z coordinate (perpendicular to the plane).

# Points	10K (Sparse)	50K (Dense)
Score [12]	0.66	6.27
Ours	0.40	4.76

Table 2: Average inference time per point cloud in minutes on PU-Net [18] dataset using identical environments.

our model significantly outperforms previous methods bilateral filtering [16], PCNet [11], DMR [15], and surpasses GLR [17], Score [12] in the majority of cases. To further illustrate the effectiveness of our method, we compare the qualitative performance with the competitive baseline Score [12]. As Figure 4 shows, Score takes 30 steps to achieve the denoising results and suffers in outliers, while ours takes only 15 steps and outputs results with fewer outliers and smoother outlines. In Figure 5, compared to Score [12], our denoised point clouds feature fewer outliers, better structures and smoother outlines in all the iterative steps. In Figure 6, Score results in cracks, while ours effectively reduces the outliers, maintains cleaner and smoother outlines.

Inference Time. As Table 2 shows, Score requires a longer inference time. In contrast, our method utilizes previous gradients to seek promising directions to move forward, thus reducing inference time by approximately 25% ~ 40%.

Ablation Study. We investigate hyper-parameters of the proposed algorithm formulated in Equation 5, 6. Other implementation details, including learning rates, decay coefficient γ , etc., are the same as the classical gradient-based method Score [12]. According to Table 3, we recommend the setting $T = 15$, $\alpha = 0.9$, and $\beta = 0.2$, which is used in this paper.

Gaussian Noise			1%		2%		3%	
T	α	β	CD	P2M	CD	P2M	CD	P2M
15	0.9	0.2	2.498	0.459	3.629	1.054	4.686	1.923
1	0.9	0.2	2.858	0.695	5.665	2.576	7.905	4.467
5	0.9	0.2	2.489	0.453	3.998	1.283	4.906	2.038
10	0.9	0.2	2.488	0.453	3.784	1.137	4.725	1.938
20	0.9	0.2	2.508	0.463	3.702	1.091	4.696	1.940
30	0.9	0.2	2.522	0.468	3.683	1.084	4.705	1.946
15	0.5	0.2	2.500	0.463	3.732	1.109	4.724	1.959
15	0.8	0.2	2.499	0.460	3.728	1.104	4.702	1.932
15	1.0	0.2	2.521	0.463	3.686	1.074	4.708	1.942
15	0.9	0.1	2.482	0.446	3.948	1.244	4.838	1.984
15	0.9	0.3	2.526	0.470	3.680	1.083	4.704	1.946
15	0.9	0.5	2.594	0.499	3.700	1.116	4.864	2.087

Table 3: Ablation study *w.r.t.* hyper-params T, α, β on sparse point clouds from PU-Net [18]. CD & P2M multiplied 10^4 .

5. CONCLUSION

In this paper, we propose point cloud denoising via momentum ascent in gradient fields. To improve the previous gradient-based method, we propose a simple yet effective iterative paradigm of point cloud denoising, which leverages past gradients to seek promising directions to move forward. Our method effectively prevents outliers that are much more likely to occur for previous methods. Moreover, our method accelerates optimization and reduces the inference time of the previous gradient-based method. On both synthetic and real-world datasets, extensive experiments demonstrate that our method outperforms state-of-the-art methods with a variety of point clouds, noise types and noise levels.

6. REFERENCES

- [1] Yaping Zhao, Haitian Zheng, Zhongrui Wang, Jiebo Luo, and Edmund Y Lam, “MANet: improving video denoising with a multi-alignment network,” in *2022 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2022, pp. 2036–2040.
- [2] Marc Alexa, Johannes Behr, Daniel Cohen-Or, Shachar Fleishman, David Levin, and Claudio T Silva, “Point set surfaces,” in *Proceedings Visualization, 2001. VIS’01.*, 2001, pp. 21–29.
- [3] Haim Avron, Andrei Sharf, Chen Greif, and Daniel Cohen-Or, “ l_1 -sparse reconstruction of sharp point set surfaces,” *ACM Transactions on Graphics*, vol. 29, no. 5, pp. 1–12, 2010.
- [4] Frédéric Cazals and Marc Pouget, “Estimating differential quantities using polynomial fitting of osculating jets,” *Computer Aided Geometric Design*, vol. 22, no. 2, pp. 121–146, 2005.
- [5] Hui Huang, Shihao Wu, Minglun Gong, Daniel Cohen-Or, Uri Ascher, and Hao Zhang, “Edge-aware point set resampling,” *ACM Transactions on Graphics*, vol. 32, no. 1, pp. 1–12, 2013.
- [6] Enrico Mattei and Alexey Castrodad, “Point cloud denoising via moving rpca,” in *Computer Graphics Forum*, 2017, vol. 36, pp. 123–137.
- [7] Yujing Sun, Scott Schaefer, and Wenping Wang, “Denoising point sets via l_0 minimization,” *Computer Aided Geometric Design*, vol. 35, pp. 2–15, 2015.
- [8] Chaojing Duan, Siheng Chen, and Jelena Kovacevic, “3d point cloud denoising via deep neural network based local surface estimation,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 8553–8557.
- [9] Pedro Hermosilla, Tobias Ritschel, and Timo Ropinski, “Total denoising: Unsupervised learning of 3d point cloud cleaning,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 52–60.
- [10] Francesca Pistilli, Giulia Fracastoro, Diego Valsesia, and Enrico Magli, “Learning graph-convolutional representations for point cloud denoising,” in *European Conference on Computer Vision*. Springer, 2020, pp. 103–118.
- [11] Marie-Julie Rakotosaona, Vittorio La Barbera, Paul Guerrero, Niloy J Mitra, and Maks Ovsjanikov, “Pointcleannet: Learning to denoise and remove outliers from dense point clouds,” in *Computer Graphics Forum*, 2020, vol. 39, pp. 185–203.
- [12] Shitong Luo and Wei Hu, “Score-based point cloud denoising,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 4583–4592.
- [13] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep learning*, 2016.
- [14] Ying Li and Huankun Sheng, “A single-stage point cloud cleaning network for outlier removal and denoising,” *Pattern Recognition*, vol. 138, pp. 109366, 2023.
- [15] Shitong Luo and Wei Hu, “Differentiable manifold reconstruction for point cloud denoising,” in *Proceedings of the ACM International Conference on Multimedia*, 2020, pp. 1330–1338.
- [16] Julie Digne and Carlo De Franchis, “The bilateral filter for point clouds,” *Image Processing On Line*, vol. 7, pp. 278–287, 2017.
- [17] Jin Zeng, Gene Cheung, Michael Ng, Jiahao Pang, and Cheng Yang, “3d point cloud denoising using graph laplacian regularization of a low dimensional manifold model,” *IEEE Transactions on Image Processing*, vol. 29, pp. 3474–3489, 2019.
- [18] Lequan Yu, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng, “Pu-net: Point cloud upsampling network,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2790–2799.
- [19] Andrés Serna, Beatriz Marcotegui, François Goulette, and Jean-Emmanuel Deschaud, “Paris-rue-madame database: a 3D mobile laser scanner dataset for benchmarking urban detection, segmentation and classification methods,” in *International Conference on Pattern Recognition, Applications and Methods ICPRAM 2014*, 2014.
- [20] Haoqiang Fan, Hao Su, and Leonidas J Guibas, “A point set generation network for 3d object reconstruction from a single image,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 605–613.
- [21] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari, “Accelerating 3d deep learning with pytorch3d,” *arXiv preprint arXiv:2007.08501*, 2020.