Robust Unsupervised Cleaning of Underwater Bathymetric Point Cloud Data

Cong Chen[1]
cong90@vt.edu
Abel Gawel[2]
abel.gawel@ethz-asl.ch
Stephen Krauss[1]
eagle138@vt.edu
Yuliang Zou[1]
ylzou@vt.edu
A. Lynn Abbott[1]
abbott@vt.edu
Daniel J.Stilwell[1]
stilwell@vt.edu

[1] The Bradley Department of Electrical and Computer Engineering
Virginia Tech, Blacksburg
Virginia, USA

[2] Autonomous Systems Lab
ETH, Zurich
Switzerland

Abstract

This paper presents a novel unified one-stage unsupervised learning framework for point cloud cleaning of noisy partial data from underwater side-scan sonars. By combining a swath-based point cloud tensor representation, an adaptive multi-scale feature encoder, and a generative Bayesian framework, the proposed method provides robust sonar point cloud denoising, completion, and outlier removal simultaneously. The condensed swath-based tensor representation preserves the point cloud associated with the underlying three-dimensional geometry by utilizing spatial and temporal correlation of sonar data. The adaptive multi-scale feature encoder identifies noisy partial tensor data without handcrafted feature labeling by utilizing CANDECOMP/PARAFAC tensor factorization. Each local embedded outlier feature under various scales is aggregated into a global context by a generative Bayesian framework. The model is automatically inferred by a variational Bayesian, without parameter tuning and model pre-training. Extensive experiments on large scale synthetic and real data demonstrate robustness against environmental perturbation. The proposed algorithm compares favourably with existing methods.

1 Introduction

Modern underwater applications, such as bathymetric mapping [15, 26] and autonomous navigation [18, 25], rely on accurate 3D reconstructions of the seafloor. While the community has made impressive progress in other domains such as automated driving [7, 21], reconstructing underwater surfaces remains challenging due to limited visibility and sensor noise. Side-scan sonar is a standard tool for seafloor imaging. A conventional side-scan

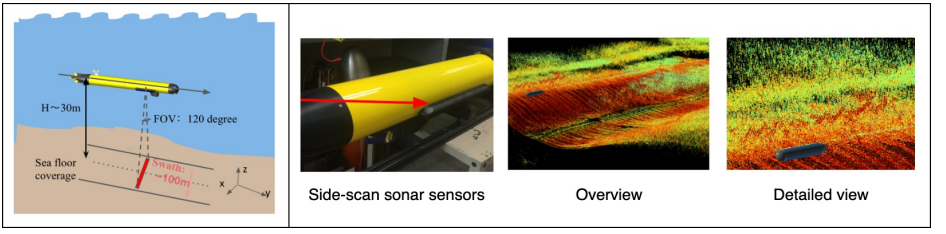

Figure 1: (left) Robotic system overview. (middle left) Side-scan sonar sensor. (middle right) Scanning data from side-scan sonar. (right) Detailed view of side-scan sonar data. The red color represents real seafloor data, including noise. The yellow color represents noise due to particles in the water and changes in water density.

sonar sensor generates a 2D image of sonar-returned intensity values without explicit depth information. In contrast, our work is inspired by the development of a novel underwater sonar system that generates a 3D point cloud of the seafloor. However, unlike point cloud measurements generated in terrestrial applications (e.g., car-mounted LiDAR), underwater point clouds generated from sonar are extremely noisy. Thus, new approaches to cleaning and inferring information from point clouds are required. Compared to electromagnetic signal based sensors, like cameras and LiDAR, sound waves can travel up to 200 meters in various water qualities, making sonar a robust sensor for large-scale underwater sensing. However, underwater point cloud data has challenges like sparsity, highly variable density, particles in the water, generally higher measurement noise, and noise due to changes in water density, salinity and temperature. So the sonar data can be noisy and sparse just like Fig. 1. However, underwater point cloud data has challenges like sparsity, highly variable density, particles in the water, generally higher measurement noise, and noise due to changes in water density, salinity and temperature. Thus, sonar point cloud data could be sparse and noisy (as shown in Fig. 1).

To handle the aforementioned challenges, common approaches for partially observed data completion capitalize on supervised learning. For example, volumetric convolutional networks are used to infer the complete 3D geometry from partial point cloud data or from several partial views from large synthetic datasets [8, 38]. However, there is limited labelled underwater point cloud data available, rendering these approaches less applicable to bathymetric sonar data. Furthermore, heavy underwater noises and outliers create biases that disrupt data completion inferencing. On the other hand, point cloud denoising methods [11, 14] fail to complete sparse data. Thus, there exists no method to address all challenges of underwater point clouds, i.e., completing missing data, reducing noise and removing outliers.

In this paper, we propose a unsupervised point cloud analysis method to accomplish denoising, outlier removal, and seafloor data completion in a single tensor Bayesian framework. To the best of our knowledge, we propose the first framework that simultaneously addresses the three aforementioned point cloud challenges. The proposed approach utilizes variational Bayesian learning for point cloud analysis directly from the observed data without any pretraining. Since sonar point cloud data is unordered and noisy, the proposed framework first tensorizes the sonar point cloud in a swath-based manner, and then exploits a multi-scale feature encoding based on the tensor factorization. By extracting a feature map from the partially observed sonar data, the framework maintains and updates the continuing global seafloor tensor, and detects sparse local outlier noises. The model selections are performed

automatically without any tuning. This framework enables generation of clean underwater maps solely from noisy sonar data. The proposed method could potentially be useful for general bathymetric sonar data shape completion, further seafloor feature registration, and other robotic perception problems.

In summary, we make the following key contributions:

- A robust and efficient unified one-stage unsupervised learning framework for simultaneous sonar point cloud completion, denoising, and outlier removal.
- A novel swath-based tensor representation designed for sonar point clouds.
- A multi-scale point cloud feature encoder by tensor factorization.
- The first algorithm to analyze underwater point clouds generated by a 3D side-scan sonar.
- Evaluation of the proposed method on synthetic and real large-scale underwater point cloud data, demonstrating the stability and efficiency of the proposed algorithm.

2 Related Work

3D Shape Completion. Existing techniques for 3D shape completion can be categorized into geometry-based, symmetry-driven and learning-based approaches. Geometry-based shape approaches use partial data to infer the missing data points. For instance, surface reconstruction methods [4, 9, 23, 36, 39] normally generate smooth interpolations in a 3D mesh to fill in missing data in locally incomplete scans. But these methods heavily depends on the prior information. Secondly, symmetry-driven methods [28, 37, 40, 41] were proposed, which add input to identify symmetry and partially repeating regular structure in objects. These approaches make assumptions of moderately complete inputs, where the missing shape regions can be copied from the observed regions. However, this assumption often does not hold in incomplete data of real-world scenarios. Finally, learning-based methods utilize end-to-end trained neural networks for mapping partial point clouds to complete 3D geometry [8, 29, 30, 38]. However, such supervised algorithms rely on labeled training data and cannot be easily transferred to new underwater scenarios. In contrast, our technique enables unsupervised generation of cleaned underwater data from noisy sensor data.

Point Cloud Denoising. 3D point cloud denoising methods aim to filter positional noise. They can be categorized into mesh-based, projection-based, and graph-based methods. Mesh-based denoising methods typically use partial differential equations (PDE) and bilateral filters (BF) [11, 14], but sometimes could lead to shrinkage and deformation [10]. The projection-based algorithm in [13] achieves denoising by estimating the reference planes of a point cloud and projecting the noisy point cloud to the reference planes iteratively until convergence. But these methods are very sensitive to parameter tuning. Graph-based denoising (GBD) methods utilize the graph Laplacian to approximate the surface shape's Laplace-Beltrami operator [3, 12, 44]. However, GBD is only practical for the manifold when consistency graph construction is available. In contract, we explore a unified framework for more general cases without any specific geometrical assumptions.

Outlier Removal. Outlier removal is essential for underwater point clouds. Outlier removal techniques for point clouds can be categorized into statistical approaches and deep learning data-driven approaches. Traditional statistical methods utilize robust local statistics with the

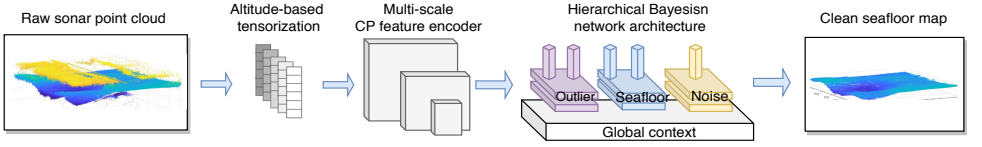

Figure 2: The overall structure of the proposed sonar data cleaning framework.

rigid structure assumption [5, 27, 34]. The drawback of these methods is that purely statistical methods do not always retain the desired geometric properties in a 3D shape. The second type of outlier removal methods leverages 3D point cloud databases [6, 16, 24]. These methods utilize knowledge about expected outliers, but require careful parameter tuning for new noise characteristics. In this paper, we propose a parameter-tuning-free framework for point cloud approximation that infers occluded regions directly from the observed data without any prior information.

3 Method

3.1 System Setup and Problem Formulation

Fig. 1 illustrates the 3D side-scan sonar and an autonomous underwater vehicle (AUV) on which it is mounted. As the vehicle moves forward, row-wise bathymetric point cloud data are collected by the sonar in a swath underneath the AUV that can be up to 200 meters wide. Sample output data is depicted in Fig. 1.

The proposed framework is illustrated in Fig. 2. First, an observation tensor is constructed from the observed point cloud. An unsupervised generative Bayesian network then provides a CANDECOMP/PARAFAC (CP) decomposition for automatic model selection. Inference of the models starts from random initialization. Three models are sparse outliers, a low-rank seafloor tensor, and noise. More detail of the network can be found in Fig. 4. Each local embedded feature under various scales is simultaneously aggregated into a global continuous seafloor context by a generative Bayesian framework. The inference converges when output of the network matches the observation.

By identifying the three kinds of data from a noisy observation, the proposed framework enables low-rank tensor estimation, outlier removal, and denoising simultaneously. Learning a mapping from the noisy point cloud data represents an unsupervised recovery directly from the noisy sonar data. In order to find the unique model of a seafloor manifold, given a random initialization, the model selection is conducted automatically via optimizing the following energy minimization problems:

$$\theta^* = \arg\min_{\theta} \mathbb{E}_{\mathcal{Y}_\Omega \sim p(\mathcal{X})} \mathbb{E}_{f_\theta \sim q(f_\theta;\theta)} l\left(f_\theta(n), \mathcal{Y}_\Omega\right) \tag{1}$$

Here, f_θ represents a tunable function with given parameters θ, and we implement f_θ using a tensor Bayesian completion framework. $\mathbb{E}$ represents the mathematical expectation. Loss l is the Kullback–Leibler divergence between the model and observed noisy point cloud $\mathcal{Y}$, and Ω is the set of observed elements inside tensor. Random initialization of Gaussian priors and tensor rank for f_θ is represented by n. The distribution q is a Gaussian prior that is not related to noise model p. The noise model of the cleaned tensor $\mathcal{X}$ is $p(\mathcal{X})$. The learning converges to the maximized similarity between the predicted seafloor manifold from f_θ and the low-rank representation of the seafloor surface.

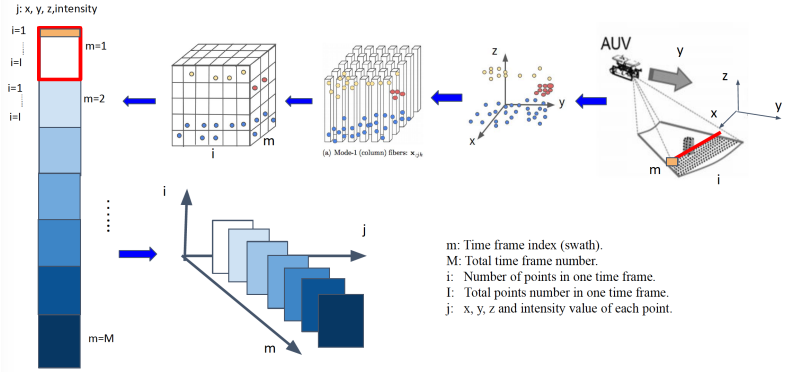

Figure 3: The tensor structure is created from bathymetric sonar data.

The observation model of the tensor sonar point cloud can be expressed as

$$\mathcal{Y}_\Omega = \mathcal{X} + \mathcal{S} + \tau \tag{2}$$

where $\mathcal{X}$ is the cleaned tensor with the global continuing seafloor information. Local outlier tensor $\mathcal{S}$ represents the local information with disconnected components, τ is the noise tensor representing the scan noise.

3.2 Point Cloud Adaptive Tensorization Representation.

The three main steps included in the adaptive altitude-based tensorization are voxel construction, closest element quantization, and adaptive altitude-based filtering.

Point cloud representation. A set of raw side-scan sonar point cloud data can be represented by a set of unordered points $\{\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_N\}$ with $\mathbf{p}_i \in \mathbb{R}^3$ representing the ith point. Each $\mathbf{p}_i$ has world-frame coordinates (x, y, z). N is the number of total points in a batch. The physical dimension of the underwater scene within the field of view (FOV) is a rectangular volume of size $L \times W \times H$, from the minimum $(x_{min}, y_{min}, z_{min})$ to the maximum $(x_{max}, y_{max}, z_{max})$. Each point's information is stored inside the corresponding tensor grid.

Swath-based tensorization of point cloud. Typically, a raw sonar point cloud is composed of millions of points. Directly mapping all the points into a voxel grid imposes computational burdens due to many empty entries [43]. Thus, we propose a compact swath-based tensorization for the point cloud representation (as illustrated in Fig. 3). First, the sonar points are grouped according to voxelization in space. Then the voxel grid is sliced and reconstructed into a swath-based tensor as in Fig. 3. When AUV drives in the y direction, the sensor scans the seafloor row by row. In each time frame, the sensor generates a row of points. The red strip in fig. 3 represents one row of point cloud data that are collected in one time frame. The value M represents the total number of frames, and m is the time frame index. Inside each frame, there are I points, each with corresponding index i. The value of i is normally aligned with x. Each point cloud entry in the tensor has a 5-valued coordinate: m, i, x, y, z.

Suppose the data acquisition is ideal and assume that the AUV travels in the direction of the y axis in the world-coordinate frame. In that case, all points from the same frame will have the same y value, and these points are collected simultaneously. Thus m becomes the first mode of the tensor. We quantize the points with index i inside the swath based

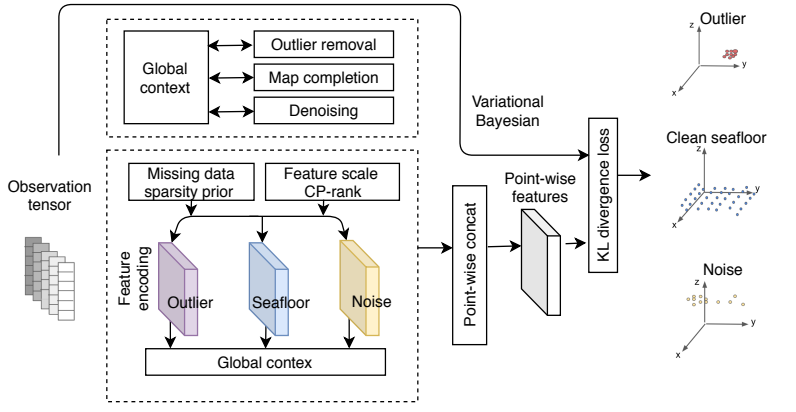

Figure 4: Detailed sonar point cloud cleaning framework, consisting of four parts: a) adaptive multi-scale feature extraction; b) generative Bayesian tensor network; c) variational Bayesian inference processing.

on the voxel they reside in. Thus i becomes the second mode of the tensor. The point location attributes (x, y, z) become the third tensor mode j. Thus observed sonar data can be constructed as a tensor $\mathcal{Y} \in \mathbb{R}^{M \times I \times 3}$.

Altitude-oriented dynamic quantization. In practice, gaps are present and point density is highly variable due to occlusions and sensor noise. These factors can bias the inference step. Thus, we utilize closest element quantization methods to create an clean tensor. Assuming that the point cloud is discretized into voxels. For each m, there will be sub-voxels of size $1 \times L \times H$, which contain L voxels of size $1 \times 1 \times H$. Variables L and H refer to the original length and height of the overall dimension of point cloud $\mathcal{Y}$ in metric units. In the presence of noise, each $1 \times 1 \times H$ voxel may contain multiple points. We therefore quantize the voxels by choosing the x and y values closest to the center of a voxel, and the z value from the point with minimum altitude.

Through this tensorization procedure, the irregular distribution of elements of $\mathcal{Y}$ are placed in a compact structure while retaining rich geometry information. Our method therefore results in an effective point cloud representation for large-scale, open-space scene learning problems. It is computationally more efficient than many voxel-based point cloud representation by avoiding many empty entries in space [43].

3.3 Adequate Multi-Scale Feature Encoding by Tensor Factorization

An efficient feature encoding method for point clouds relies on an optimal feature map size, which balances rich geometry and fast computation time. While conventional deep learning methods use fixed feature map dimensions [21, 46], this paper proposes a novel multi-scale point cloud feature encoding method by CP factorization [2, 32, 45], under the assumption that a seafloor can be characterized as a spatially-correlated low-rank tensor.

The CP-factorization lowers the tensor dimension by decomposing it into sets of latent factor matrices. All latent factors share the same sparsity due to the assumption of underlying spatial correlation in continuing seafloor data [31]. And finally, the CP-rank, which represented the model complexity, is optimized by the continuing global seafloor context and local outlier information automatically. A tensor can be CP-factorized [20] as follows:

$$\mathcal{X} = \sum_{r=1}^{R} a_r^{(1)} \circ a_r^{(2)} \circ a_r^{(3)} = [[A^{(1)}, A^{(2)}, A^{(3)}]] \tag{3}$$

where $\circ$ denotes the vector outer product, a_r denotes the factor vectors of $\mathcal{Y}$, and $\left\{\mathbf{A}^{(n)} | n = 1, \ldots, N\right\}$ are the latent factor matrices of $\mathcal{Y}$. The mode-n factor matrix can be denoted as row-wise vectors $A^{(n)} = \left[a_1^{(n)}, \ldots, a_{I_n}^{(n)}\right]^T$. Here we define tensor mode N as three, and $[[\cdots]]$ is the notation for CP factorization. CP-rank is represented as R. The tensor's CP model decides the rank-one tensors.

3.4 Generative Bayesian Tensor Network

In order to consider the global seafloor context as a low-CP-rank tensor along with outlier and sparsity simultaneously, the generative hierarchical Bayesian tensor framework is introduced.

Each local embedded feature under various scales is aggregated into a global context. For the feature scale selection, the corresponding CP-rank is determined by minimizing the latent space dimension. Corresponding column-wise sparsity is shared by all latent factor matrices using a covariance matrix Λ. Thus the latent space retains the seafloor's global information,

$$p\left(A^{(n)} | \lambda\right) = \prod_{i_n=1}^{I_n} \mathcal{N}\left(a_{i_n}^{(n)} | 0, \Lambda^{-1}\right) \tag{4}$$

where $p(\lambda) = \prod_{r=1}^{R} \mathrm{Ga}\left(\lambda_r | c_0, d_0\right)$. The independent identically distributed (i.i.d.) gamma distribution is denoted as Ga, and (c_0, d_0) represent the sparsity prior of i.i.d. gamma distributions. As for the local embedding information, the sparse outlier term $\mathcal{S}$ is designed to represent various outliers. It is also modeled by the sparsity inducing prior, but individual hyperparameters of each element are independent.

$$p\left(\mathcal{S}_\Omega | \gamma\right) = \prod_{i_1, \ldots, i_N} \mathcal{N}\left(\mathcal{S}_{i_1,i_2,i_3} | 0, \gamma_{i_1,i_2,i_3}^{-1}\right) \tag{5}$$

where $p(\gamma) = \prod_{i_1,\ldots,i_N} \mathrm{Ga}\left(\gamma_{i_1,i_2,i_3} | a_0^\gamma, b_0^\gamma\right)$. Here $\gamma_{i_1 \ldots i_N}$ represent the individual parameters and (a_0^γ, b_0^γ) represent the sparsity prior of the i.i.d. gamma distribution. The noise term τ is also an i.i.d. gamma distribution with (a_0^τ, b_0^τ) as priors: $p(\tau) = \mathrm{Ga}\left(\tau | a_0^\tau, b_0^\tau\right)$. The observation model can be expressed as a joint distribution:

$$p\left(\mathcal{Y}_\Omega | \mathcal{X}, \mathcal{S}_\Omega, \tau\right) = \prod_{n=1}^{3} p\left(A^{(n)} | \lambda\right) p\left(\mathcal{S}_\Omega | \gamma\right) p(\lambda) p(\gamma) p(\tau) \tag{6}$$

Given the Bayesian model, we can use the varational Bayesian approach [42] to perform the inference, aiming to minimize the error in equation 1.

4 Experimental Results

4.1 Validation on Synthetic Bathymetric Data

We first quantitatively evaluate the proposed framework on synthetic data that is generated in a underwater simulation environment using Unreal Engine. In addition to the seafloor, the simulation models outlier objects and noise. An AUV trajectory of 100m length is simulated,

using bathymetric data collected with a simulated sonar. The result is a dataset containing a low-rank tensor $\mathcal{X}$ of size $1000 \times 500 \times 3$, i.e., $500,000$ points. In total, 30 sets of different terrain data are generated, with 20 object permutations. Noise is drawn from a Gaussian noise distribution with $\mathcal{N}(0,20)$ on the depth sensing as yellow noise in Fig. 6. The resulting point cloud is illustrated in Fig. 5. The proposed method takes average five minutes computation for each data using I3 CPU, and it does not require any pre-training.

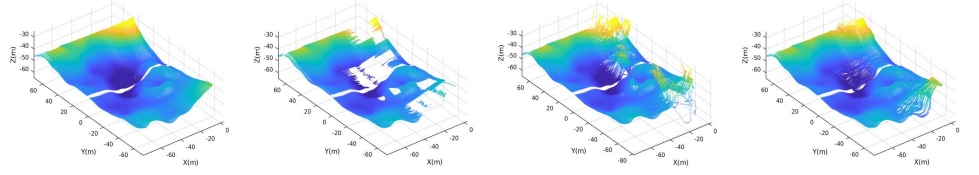

(a) Ground truth (b) Missing data (c) SB-LME (d) Ours

Figure 5: Exemplary completion result on 22% missing data. Left to right: ground truth, data input, SB-LME, our method.

Table 1: RMSE results with different missing data between our method and SB-LME [1].

	10%	13%	16%	19%	22%
SB-LME [1]	0.0179	0.0253	0.0225	0.0338	0.0268
Ours	**0.0068**	**0.0083**	**0.0122**	**0.0088**	**0.011**

As Fig. 5 shows, we randomly delete a portion of data to create artificial sparse block outliers and missing data for 20 permutations per data on the simulated bathymetric map for the observed $\mathcal{Y}$. To evaluate the map cleaning performance, we calculate the Root Mean Squared Error (RMSE) between the data completed with our approach and the ground truth. We compare the proposed technique with one classic unsupervised baseline, SB-LME [1]. Table 1 shows point-cloud completion results for the observed bathymetric data subject to increasing amounts of incomplete data. With an increase in missing data from 10% to 22%, which is the common missing data rate on real sonar data, the benchmarked method's RMSE varies between 1.79% and 2.68%, while our method consistently outperforms the benchmarks with RMSE between 0.68% and 1.1%. We also show a visual comparison in Fig. 5.

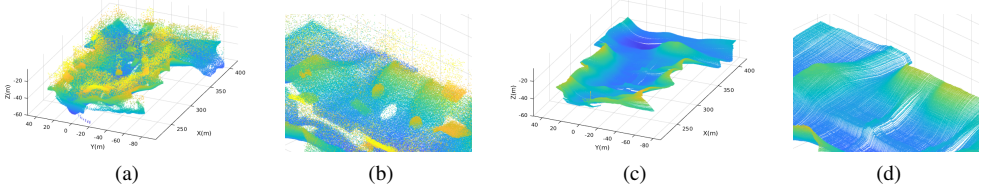

(a) (b) (c) (d)

Figure 6: Synthetic data with completion and outlier removal: (a) original data including outlier objects and noise, (b) detail view of original data, (c) Output using the proposed method, (d) detail view of output data

The simulation further generates bathymetric sonar data containing outlier objects and noise for evaluating the outlier removal and denoising performance of our approach. Fig. 6 shows the simulated seafloor bathymetric sonar data used for this experiment, which contains artificial static and moving outlier objects above the seafloor, noise and missing data due to object occlusions. Here, moving objects mimic moving fishes. The proposed method is

Table 2: RMSE results with different outliers and noise types.

Methods	Static outliers		Moving outliers		Noise type			All
	5	10	5	10	Above	Inside	Under	
Object [35]	0.11	0.16	0.15	0.24	0.05	0.00	0.05	0.40
Guided [17]	0.27	0.59	0.43	0.35	0.40	0.47	0.69	0.57
SB-LME [1]	0.36	0.57	0.50	0.67	0.46	0.36	0.67	0.46
Bilateral [11]	1.00	1.01	1.01	1.07	0.98	0.98	0.98	1.07
Ours	**0.05**	**0.09**	**0.08**	**0.13**	**0.01**	**0.00**	**0.01**	**0.25**

able to identify all simulated outliers and additionally recovers the occluded regions on the seafloor surface, as further depicted in Fig. 6.

Table 2 illustrates the proposed method's outlier removal and denoising performance against state-of-the-art methods: ObjectMaps [35], GuidedFilter [17], SB-LME [1] and Bilateralfilter [11]. Static outliers include sunken objects, moving outliers include fish, and are generated through the Unreal Engine simulator. The RMSE results show that when noise and outliers are imposed individually or simultaneously, the proposed method always outperforms the benchmarks at the three functions. Furthermore, the benchmarked methods only perform sole functions and do not perform well in the underwater scenario.

4.2 Shape Recovery Results on Real Bathymetric Sonar Data

Field tests are conducted in a large lake. The data consists of multiple straight segments collected using a UAV equipped with a side-scan sonar sensor. Each segment covered a distance of about 200m in length, resulting in one point cloud covering 200m $\times$ 100m seafloor data and 5M points. 30 sets of data were collected.

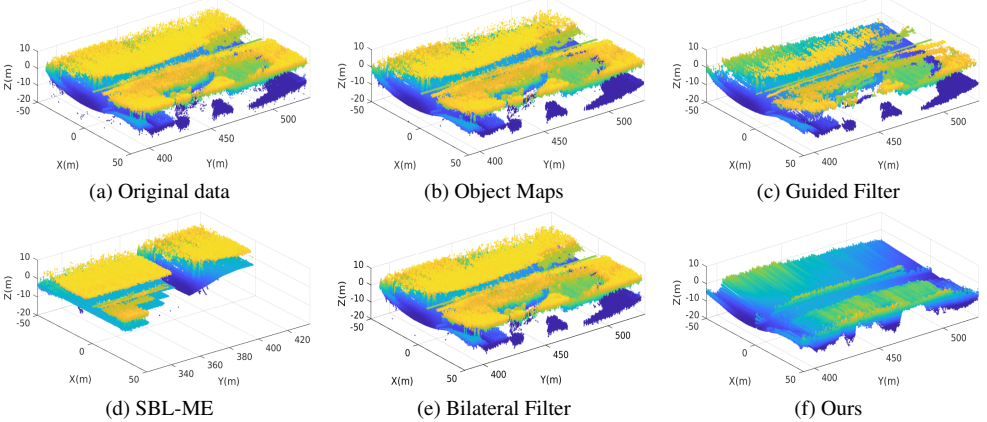

(a) Original data (b) Object Maps (c) Guided Filter

(d) SBL-ME (e) Bilateral Filter (f) Ours

Figure 7: Real bathymetric sonar data completion and denoising. The green points are the seafloor, the yellow points are the noise above the seafloor, the dark blue points are the noise under the seafloor, and the large hole is the missing part of the seafloor data. The qualitative results show that the proposed algorithm completed the missing data and remove all outlier noise, whereas other benchmarked methods fail in all tasks.

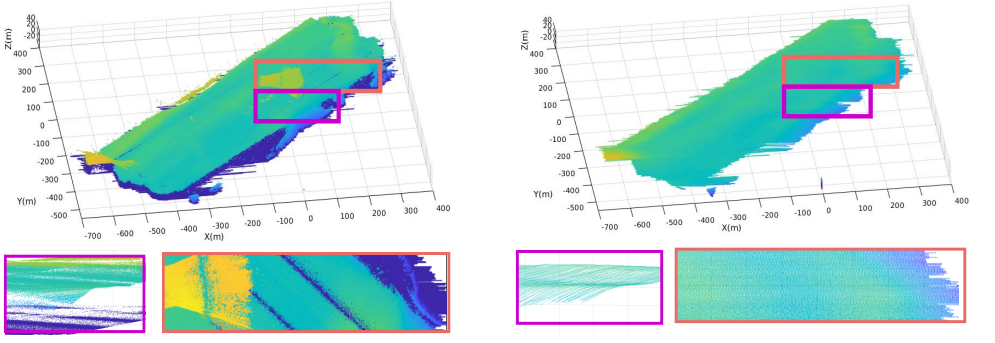

(a) (b)

Figure 8: Bathymetric map cleaning from a lawnmower-path field test. (a) Before processing, (b) After processing. The green points are the seafloor, yellow points are the noise above seafloor and dark blue color points are noise under seafloor.

Furthermore, we performed a large scan data collection following a lawnmower path combining 6 segments, covering a total area of ca. 1000m $\times$ 1000m, and 50M points. In these data, common effects can be observed, such as underwater noise, outliers and missing data due to occlusions. Evaluation is qualitative assessment due to missing ground truth data.

Fig. 7 compares the result of the proposed method's sonar bathymetric data cleaning benchmarked against the baseline methods. The proposed method results in significantly cleaner bathymetric data of the seafloor, removing all the yellow noisy points and completing the seafloor. The baseline methods exhibit larger error in a qualitative comparison than the proposed methods. As explanation for the large errors, we believe that the real underwater sonar data is subject to larger and various uncertainties in noise, outliers, and missing data making it harder to assess than the synthetic benchmarks. We further tested several supervised deep-learned methods on the data, including [19], [33], and [22]. However, their models pre-trained on a small synthetic data set did not yield satisfactory results on our data.

Finally, Fig. 8 illustrates our results on the large lawnmower-path experiment. Our results demonstrate the efficacy of the approach in recovering a complete and organized seafloor map. Qualitative results show that our approach successfully recovers the seafloor. Ground truth of seafloor data shows the same average depth with the output as 20m, with the 40m noise above and 0m noise under seafloor being removed successfully.

5 Conclusion

This paper presents a robust technique for underwater bathymetric sonar data cleaning without prior training. The method achieves underwater bathymetric sonar data denoising, outlier removal, and data completion simultaneously in a single generative Bayesian tensor framework. The proposed framework has been evaluated with both simulated and real sonar point cloud data. Both quantitative and qualitative results on large-scale data demonstrate the efficacy of the proposed framework over previous methods. While we designed our method for the challenging task of seafloor point cloud data cleaning, we are interested in extending this work to point cloud cleaning in other applications in the future.

References

[1] S Derin Babacan, Martin Luessi, Rafael Molina, and Aggelos K Katsaggelos. Sparse bayesian methods for low-rank matrix estimation. *IEEE Transactions on Signal Processing*, 60(8):3964–3977, 2012.

[2] Juan Andrés Bazerque, Gonzalo Mateos, and Georgios B Giannakis. Rank regularization and bayesian inference for tensor completion and extrapolation. *IEEE transactions on signal processing*, 61(22):5689–5703, 2013.

[3] Mikhail Belkin and Partha Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. *Journal of Computer and System Sciences*, 74(8):1289–1308, 2008.

[4] Matthew Berger, Andrea Tagliasacchi, Lee Seversky, Pierre Alliez, Joshua Levine, Andrei Sharf, and Claudio Silva. State of the art in surface reconstruction from point clouds. In *Eurographics 2014 STAR Reports*, 2014.

[5] Antoni Buades, Bartomeu Coll, and J-M Morel. A non-local algorithm for image denoising. In *Conference on Computer Vision and Pattern Recognition*, 2005.

[6] Frédéric Chazal, David Cohen-Steiner, and Quentin Mérigot. Geometric inference for probability measures. *Foundations of Computational Mathematics*, 11(6):733–751, 2011.

[7] Siheng Chen, Baoan Liu, Chen Feng, Carlos Vallespi-Gonzalez, and Carl Wellington. 3D point cloud processing and learning for autonomous driving. *arXiv preprint arXiv:2003.00601*, 2020.

[8] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3D-encoder-predictor CNNs and shape synthesis. In *Conference on Computer Vision and Pattern Recognition*, 2017.

[9] James Davis, Stephen R Marschner, Matt Garr, and Marc Levoy. Filling holes in complex surfaces using volumetric diffusion. In *Proceedings. First IEEE International Symposium on 3D Data Processing Visualization and Transmission*, pages 428–441. IEEE, 2002.

[10] Jean-Emmanuel Deschaud and François Goulette. Point cloud non local denoising using local surface descriptor similarity. 2010.

[11] Julie Digne and Carlo De Franchis. The bilateral filter for point clouds. *Image Processing On Line*, 7:278–287, 2017.

[12] Chinthaka Dinesh, Gene Cheung, Ivan V Bajic, and Cheng Yang. Fast 3D point cloud denoising via bipartite graph approximation & total variation. *arXiv preprint arXiv:1804.10831*, 2018.

[13] Chaojing Duan, Siheng Chen, and Jelena Kovacevic. Weighted multi-projection: 3D point cloud denoising with tangent planes. In *IEEE Global Conference on Signal and Information Processing*, 2018.

[14] Abderrahim Elmoataz, Olivier Lezoray, and Sébastien Bougleux. Nonlocal discrete regularization on weighted graphs: a framework for image and manifold processing. *IEEE Transactions on Image Processing*, 17(7):1047–1060, 2008.

[15] Jay Gao. Bathymetric mapping by means of remote sensing: methods, accuracy and limitations. *Progress in Physical Geography*, 33(1):103–116, 2009.

[16] Leonidas Guibas, Dmitriy Morozov, and Quentin Mérigot. Witnessed k-distance. *Discrete & Computational Geometry*, 49(1):22–45, 2013.

[17] Kaiming He, Jian Sun, and Xiaoou Tang. Guided image filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(6):1397–1409, 2012.

[18] Tata Herbert, Nzelibe Ifechukwu Ogochukwu, and Ayodeji John Faneye. Bathymetric mapping for safe navigation: A case study of part of Lagos Lagoon, 2019.

[19] Pedro Hermosilla, Tobias Ritschel, and Timo Ropinski. Total denoising: Unsupervised learning of 3D point cloud cleaning. In *International Conference on Computer Vision*, 2019.

[20] Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.

[21] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Computer Vision and Pattern Recognition*, 2019.

[22] Minghua Liu, Lu Sheng, Sheng Yang, Jing Shao, and Shi-Min Hu. Morphing and sampling network for dense point cloud completion. *arXiv preprint arXiv:1912.00280*, 2019.

[23] Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. Laplacian mesh optimization. In *Proceedings of the 4th International Conference on Computer Graphics and Interactive Techniques in Australasia and Southeast Asia*, pages 381–389. ACM, 2006.

[24] Xiaojuan Ning, Fan Li, Ge Tian, and Yinghui Wang. An efficient outlier removal method for scattered point cloud data. *PLOS ONE*, 13(8):e0201280, 2018.

[25] Dongdong Peng, Jiaqi Gao, and Tian Zhou. Underwater terrain matching navigation based on Gaussian process regression with a multi-beam bathymetric sonar. *The Journal of the Acoustical Society of America*, 146(4):3089–3089, 2019.

[26] William D Philpot. Bathymetric mapping with passive multispectral imagery. *Applied Optics*, 28(8):1569–1578, 1989.

[27] R Pincus, V Barnett, and T Lewis. Outliers in statistical data. j. wiley & sons 1994, xvii. 582 pp.,£ 49.95. *Biometrical Journal*, 37(2):256–256, 1995.

[28] Joshua Podolak, Philip Shilane, Aleksey Golovinskiy, Szymon Rusinkiewicz, and Thomas Funkhouser. A planar-reflective symmetry transform for 3D shapes. In *ACM Transactions on Graphics (TOG)*, volume 25, pages 549–559. ACM, 2006.

[29] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *Computer Vision and Pattern Recognition*, 2017.

[30] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *Conference on Neural Information Processing Systems*, 2017.

[31] Stephan Rabanser, Oleksandr Shchur, and Stephan Günnemann. Introduction to tensor decompositions and their applications in machine learning. *arXiv preprint arXiv:1711.10781*, 2017.

[32] Piyush Rai, Yingjian Wang, Shengbo Guo, Gary Chen, David Dunson, and Lawrence Carin. Scalable bayesian low-rank decomposition of incomplete multiway tensors. In *ICML*, 2014.

[33] 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*. Wiley Online Library, 2019.

[34] Peter J Rousseeuw and Mia Hubert. Robust statistics for outlier detection. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 1(1):73–79, 2011.

[35] Radu Bogdan Rusu, Zoltan Csaba Marton, Nico Blodow, Mihai Dolha, and Michael Beetz. Towards 3D point cloud based object maps for household environments. *Robotics and Autonomous Systems*, 56(11):927–941, 2008.

[36] Kripasindhu Sarkar, Kiran Varanasi, and Didier Stricker. Learning quadrangulated patches for 3D shape parameterization and completion. In *3D Vision*, 2017.

[37] Ivan Sipiran, Robert Gregor, and Tobias Schreck. Approximate symmetry detection in partial 3D meshes. In *Computer Graphics Forum*, volume 33, pages 131–140. Wiley Online Library, 2014.

[38] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Conference on Computer Vision and Pattern Recognition*, 2017.

[39] Olga Sorkine and Daniel Cohen-Or. Least-squares meshes. In *Proceedings Shape Modeling Applications*, 2004.

[40] Minhyuk Sung, Vladimir G Kim, Roland Angst, and Leonidas Guibas. Data-driven structural priors for shape completion. *ACM Transactions on Graphics (TOG)*, 34(6): 175, 2015.

[41] Sebastian Thrun and Ben Wegbreit. Shape from symmetry. In *International Conference on Computer Vision*, 2005.

[42] Dimitris G Tzikas, Aristidis C Likas, and Nikolaos P Galatsanos. The variational approximation for Bayesian inference. *IEEE Signal Processing Magazine*, 25(6):131–146, 2008.

[43] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D ShapeNets: A deep representation for volumetric shapes. In *Computer Vision and Pattern Recognition*, 2015.

[44] 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. *arXiv preprint arXiv:1803.07252*, 2018.

[45] Qibin Zhao, Liqing Zhang, and Andrzej Cichocki. Bayesian cp factorization of incomplete tensors with automatic rank determination. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1751–1763, 2015.

[46] Yin Zhou and Oncel Tuzel. VoxelNet: End-to-end learning for point cloud based 3D object detection. In *Conference on Computer Vision and Pattern Recognition*, 2018.