

Diffusion unit: Interpretable edge enhancement and suppression learning for 3D point cloud segmentation

Haoyi Xiu^{a,b}, Xin Liu^{a,*}, Weimin Wang^{a,1}, Kyoung-Sook Kim^a, Takayuki Shinohara^c, Qiong Chang^d, Masashi Matsuoka^b

^a Artificial Intelligence Research Center, AIST, Tokyo, Japan

^b Department of Architecture and Building Engineering, Tokyo Institute of Technology, Tokyo, Japan

^c Innovation Technology Office Research Center, PASCO Corporation, Tokyo, Japan

^d Department of Computer Science, Tokyo Institute of Technology, Tokyo, Japan

ARTICLE INFO

Communicated by Y. Liu

Dataset link: <https://github.com/martianxiu/DiffusionUnit>

Keywords:

3D point clouds
Diffusion
Edge awareness
Edge enhancement
Deep learning
Segmentation

ABSTRACT

3D point clouds are discrete samples of continuous surfaces which can be used for various applications. However, the lack of true connectivity information, i.e., edge information, makes point cloud recognition challenging. Recent edge-aware methods incorporate edge modeling into network designs to better describe local structures. Although these methods show that incorporating edge information is beneficial, *how* edge information helps remains unclear, making it difficult for users to analyze its usefulness. To shed light on this issue, in this study, we propose a new algorithm called Diffusion Unit (DU) that handles edge information in a principled and interpretable manner while providing decent improvement. First, we theoretically show that DU learns to perform task-beneficial edge enhancement and suppression. Second, we experimentally observe and verify the edge enhancement and suppression behavior. Third, we empirically demonstrate that this behavior contributes to performance improvement. Extensive experiments and analyses performed on challenging benchmarks verify the effectiveness of DU. Specifically, our method achieves state-of-the-art performance in object part segmentation using ShapeNet part and scene segmentation using S3DIS. Our source code is available at <https://github.com/martianxiu/DiffusionUnit>.

1. Introduction

A 3D point cloud is a basic and flexible shape representation where the surface is represented as a set of discrete points in the 3D space. Powered by the recent advancement of cost-effective sensor technology, numerous large-scale datasets are publicly available to research communities, facilitating deep learning-based point cloud understanding. The application field of such research can be autonomous driving [1–3], remote sensing [4,5], and so on.

Among diverse types of deep neural networks (DNNs), Convolutional Neural Network (CNN) is the most effective algorithm for computer vision. However, the unstructured and unordered nature of point clouds prevents researchers from directly applying CNN to point clouds. Early works often project point clouds to regular grids before applying CNN [6–9] or utilizing special data structures [10–12]. However, most of these methods may suffer from information loss caused by projections. This issue is resolved by PointNet [13], which consists of permutation-invariant operations and can handle point clouds in a

lossless manner. Further, PointNet++ extends PointNet to achieve hierarchical learning [14]. Based on the PointNet++ framework, varieties of convolution methods [15–19] and transformers [20–24] for point clouds are proposed.

On the other hand, point clouds naturally lack structures because they are discrete point samples of continuous surfaces. This means that true connectivity information, i.e., edge information, is not available. Hence, incorporating edge modeling into model designs can help to learn local structures of continuous surfaces, which in turn leads to better performance in varieties of high-level tasks [18,25,26], especially in those that benefit from optimized connectivity. Edge information, or the edge feature, is often represented by the output of neural network functions (e.g., multi-layer perceptions (MLPs)) that take features of a center point and its neighbor as input [25]. The resulting edge features are used, for instance, as extra descriptors describing the local structure [26,27] or as similarity measures reflecting spatial consistency [21,28]. Although these edge-aware methods empirically

* Corresponding author.

E-mail address: xin.liu@aist.go.jp (X. Liu).

¹ Also is affiliated with Dalian University of Technology.

show that the incorporation of edge features improves performance, how they work to improve performance remains unclear. This lack of interpretability can be problematic because (1) it makes it difficult for users to identify fruitful research directions when it works well and (2) it also makes it difficult for users to diagnose the problem when they fail.

In this study, we propose the *Diffusion Unit (DU)* that handles edges in a principled and interpretable manner while providing decent performance gain. DU extends the classic diffusion theory [29, 30] that is often used for edge-preserving smoothing. An important motivation for extending the diffusion theory is that it provides a mathematically transparent framework that makes the behavior of DU interpretable. Specifically, we theoretically show that DU learns to enhance task-beneficial edges (e.g., part boundaries of objects) and suppress unhelpful discontinuities. Furthermore, we experimentally observe and verify the edge enhancement and suppression behavior of DU through intuitive visualizations. Moreover, we empirically demonstrate that this behavior contributes to performance improvement.

To validate the interpretability and effect of DU, the task used must simultaneously meet two conditions: (1) interpretability can be validated. That is, the task has clear definitions of task and non-task edges so that we can empirically verify the edge enhancement and suppression behavior of DU; (2) Modeling the connectivity/edges between points is important. That is, the optimization of edges must be closely related to the task objective so that the effect of DU can be fairly and confidently evaluated. In this study, we adopt part and scene segmentation tasks for the validation of DU, as they clearly satisfy the above two conditions. Therefore, we construct *DU-Net*, a network specialized for point cloud segmentation, with DU as its core building block. Extensive experiments are performed to verify the edge enhancement/suppression behavior of DU as well as the practical effectiveness of DU-Net. In particular, we achieve state-of-the-art performance in object part segmentation using ShapeNet part [31] and scene segmentation using S3DIS [32]. The effect of DU is shown in Fig. 1.

Our main contributions are summarized as follows:

- We propose DU that performs edge enhancement and suppression learning in a principled manner.
- We theoretically analyze and experimentally demonstrate the edge enhancement and suppression behavior of DU.
- We design DU-Net for point cloud segmentation and perform extensive experiments to verify its practical effectiveness and design choices.

The remainder of this paper is organized as follows. Section 2 revisits the related works concerning deep learning methods for 3D point clouds. In particular, we provide a detailed review of edge-aware methods that use edge features to better model local structures. Since diffusion theory is often used for edge-preserving smoothing and our work is closely related to it, diffusion-related methods are subsequently reviewed to show how it is developed and applied to various tasks. Section 3 first elaborates on the motivation to introduce explicit edge modeling. Next, we introduce the diffusion equation which forms the basis of DU. Then, we explain the challenges of applying the diffusion equation directly to point clouds and propose DU in its continuous form. We subsequently conduct a theoretical analysis and show that DU is guaranteed to perform adaptive edge enhancement and suppression. To apply DU to discrete point clouds, we introduce discretization schemes. Finally, we construct DU-Net for point cloud segmentation with DU as its core component. Section 4 empirically verifies the theoretical results of Section 3 through visualizations. Moreover, experiment results on part and scene segmentation tasks are also reported to validate its practical effectiveness. Then, ablation studies are also made to validate our design choices. We conclude this study in Section 5.




Fig. 1. The effect of DU. The heat map indicates the local feature smoothness. A low response indicates a smooth region while a high response indicates a non-smooth region. Red and orange circles indicate enhanced and suppressed features after applying DU. As can be seen, task-beneficial edges (in this case, object boundaries) are enhanced/preserved whereas intra-region discontinuities are suppressed.

2. Related work

In this section, we first provide a brief review of deep learning methods for point cloud analysis. Then, we review approaches that improve point cloud understanding with the explicit awareness of edge information. As the diffusion equation is often used for edge-preserving smoothing and we also use it as DU's basis, we subsequently provide a brief review of methods that use diffusion equations.

2.1. Deep learning for 3D point clouds

Projection-based methods Owing to the unstructured and unordered nature of point clouds, early methods that deal with point clouds are projection-based methods. Those methods project point clouds to regular grids (e.g., 2D pixels [6,7,33] and 3D voxels [8–11,34,35]) to enable matured regular convolution to operate on point clouds. The problem with these methods is that, however, they lose fine-grained geometric details through projections.

MLP-based methods MLP-based methods operate directly on the raw point clouds, thereby preventing information loss due to projections. The representative method is PointNet [13] whose main operations consist of shared MLPs and symmetric functions (e.g., max-pooling). Both operations are permutation-invariant, and thus the irregularity of point clouds is well-resolved. To further enhance PointNet, PointNet++ [14] applies PointNets to local subsets of points grouped according to different scales or resolutions to better model the local geometric details. By repeating such a process multiple times with increasing grouping scales, PointNet++ achieves efficient CNN-like hierarchical learning. Afterward, various advancements are made based on PointNet and PointNet++ [36,37]. For instance, ShellNet [38] further divides a local neighborhood into multiple groups and applies MLPs and symmetric functions to all groups, aiming to leverage fine-grained details. PointNeXt [19], on the other hand, constructs a modern version of PointNet++ by carefully combining the advanced training and scaling strategies. RepSurf [39] uses an explicit representation of surfaces based on several geometric priors and successfully enhances PointNet++. PointMixer [40] which is inspired from MLP-Mixer [41] is proposed. It realizes intra-set, inter-set, and hierarchical-set mixing to adapt the original mixing operation to point cloud processing.

Point convolution-based methods The point convolution-based method defines a convolution operation on unstructured point clouds. For instance, some studies construct a pseudo grid for each convolution local neighborhood on which input points are projected [17,26,42]. Since the constructed grid is structured, a standard convolution operation becomes applicable to projected point features. On the other hand, some methods predict convolution filters on the fly [15,16,43]. For example, PointConv [16] generates convolution kernels for a local neighborhood by transforming relative positions (relative to local centers). On the other hand, RSCNN [43] uses more features including relative positions, absolute positions, and euclidean distances to generate kernels so that kernels are explicitly 3D geometry-aware. PAConv [18] additionally learns a kernel bank from which convolution kernels are generated. Though effective, these methods do not explicitly consider the local structures characterized by edges which can describe the underlying surface more accurately.

Transformer-based methods Inspired by the success of Transformers [44] in the Natural Language Processing field, many studies have tried to construct transformer-based methods for point cloud data [45,46]. Here, we focus on transformers that are designed for point cloud segmentation. Since self-attention is essentially a set operator, it can be straightforwardly applied to point clouds. For example, PAT [20] replaces the original multi-head self-attention with parameter-efficient group shuffle attention so that self-attention can be applied to massive point cloud data. PointASNL [47], on the other hand, uses self-attention for learning the global characteristics of point clouds. The obtained global features are then combined with local ones learned by point convolutions so that the resulting features are both aware of the global and local characteristics of point clouds. Point Transformer [21] applies vector attention instead of original self-attention to the local subset of points to increase the expressiveness while reducing computation costs. Inspired by the success of BERT [48], Point-BERT [49] pre-trains the transformer architecture by masked point modeling and demonstrates its good transferability of the learned representations. Stratified Transformer [22] achieves efficient modeling of long-range dependencies using the stratified strategy which samples nearby points densely and distant points sparsely. Recently, geodesic self-attention [50] is proposed to capture the long-range dependencies by introducing a metric on the Riemannian manifold. Point Transformer V2 [23] improves the efficiency of Point Transformer [21] by proposing efficient vector attention and lightweight partition-based pooling.

2.2. Edge-aware methods

To better model the local structure, edge-aware methods incorporate edge information into network design in various ways. As is described in Section 1, we regard methods that explicitly use the output of neural network functions that receive features of a center point and its neighbor as input as edge-aware methods. In general, there are three types of edge-aware methods. First, some methods perform convolution on edge features. The representative of such methods is EdgeConv [25] which performs convolution on edge features to better model the local geometric structure. This idea is generally applicable and thus is adopted in numerous following works (e.g., [26,27,51]). Despite that explicit modeling of edges brings performance improvements, the above methods are hardly interpretable, making users wonder how edges are modeled to improve performance. Second, some methods regard edge information as a measure of semantic distance. This idea uses edges as spatial weights that indicate the relationships between a center point and its neighbors. Moreover, it is often combined with the idea of self-attention [21,23,52,53] which converts edges into normalized weights. For example, Point Transformer [21] replaces the global dot product in the original self-attention with local edge features for modeling more fine-grained geometric relationships. This idea is proven particularly useful in point cloud analysis which leads to state-of-the-art

performance in segmentation tasks. However, forced normalization of edge information may undermine its usefulness because absolute values of edge features contain rich structural information like smoothness. On the other hand, edge information may be used as an extra source of supervision. For instance, [54] constructs a dedicated branch for learning edge features and pairs it with a loss function that encourages spatial consistency.

Compared to the above edge-aware methods, DU is theoretically guaranteed and empirically verified to perform edge enhancement and suppression. Furthermore, DU is designed to model edge features directly rather than converting them to normalized weights. In addition, DU realizes automatic edge modeling without additional supervision.

2.3. Diffusion methods

The diffusion methods model the smoothing of data (e.g., images) as diffusion processes. The pioneering work of such methods is anisotropic diffusion [29] which is designed to smooth out irrelevant discontinuities while preserving significant edges, achieving edge-preserving smoothing. Specifically, anisotropic diffusion performs smoothing in a locally adaptive manner by considering the edge strength of a center point. Points with high edge strength are preserved while those with low edge strength are smoothed. An important result of [29] is that such a behavior can be explained theoretically, making the method interpretable. Therefore, anisotropic diffusion and its theoretical/empirical behavior are extensively studied in image processing and numerous extensions are proposed [30,55–57]. Furthermore, the success of this idea is also adapted to other research fields (e.g., computer graphics [58–60]).

The idea of diffusion is also adopted in modern deep learning studies, in particular the development of various graph neural networks (GNNs). For example, DCNN [61] proposes a new architecture that learns diffusion-based representations for graph-structured data. GDC [62] similarly models the information propagation as generalized graph diffusion aiming at modeling long-range relationships. ADC [63] further improves the flexibility of GDC in choosing the neighborhood size by learning the optimal neighborhood size from data. [64] analyzes the existing architectures that can be interpreted as information diffusion on graphs using tools from partial differential equations (PDEs). Furthermore, the notion of diffusion is also used in probabilistic point cloud generation or point cloud completion. Inspired by the diffusion process in thermodynamics in which particles diffuse from the original distribution to a noise distribution, [65] models the point cloud generation as learning reverse diffusion process which transforms a noise distribution to a desired shape. [66] further integrates the denoising diffusion models with the point-voxel representation. PDR [67] provides a new paradigm to better capture point density distribution when performing completion.

Despite there exist many applications of the diffusion equation, extending the diffusion equation for adaptive edge enhancement and suppression for 3D point cloud segmentation remains under-explored.

3. Method

In this section, we first explain the motivation for introducing explicit edge modeling. Then, we provide a concise background about the diffusion equation on which we build our DU. Then, we give the continuous definition of DU and provide a theoretical analysis to show how DU achieves edge enhancement and suppression learning. We subsequently introduce the discretized DU which can be applied to discrete point clouds and easily implemented on modern machines. Lastly, we design DU-Net for point cloud segmentation.




Fig. 2. Examples of task and non-task edges. Red ellipses indicate task edges while blue ones indicate non-task edges.

3.1. Motivation

In this study, the term “edge” refers to one of the connections between two points. An edge indicates whether there is a change in attributes when moving from one point to another. For example, a large geometric change may indicate that the area is a corner. A change in color may also indicate the boundary of an object. Furthermore, these edges can be divided into two types, task edges and non-task edges, depending on the given task. Examples of such edges are shown in Fig. 2. Task edges in tasks such as segmentation and edge detection are generally region boundaries, while non-task edges in these tasks are other discontinuities that do not correspond to region boundaries (e.g. noise or insignificant edges).

Point clouds naturally lack connectivity information, or edges. Therefore, by explicitly modeling edges, the network can learn to optimize the connectivity between points and thus better describe the local geometry. Explicit edge modeling is useful for tasks where connectivity between points is important. For example, in semantic segmentation, the network must distinguish between task and non-task edges to make high-fidelity predictions (see Fig. 2), which is challenging because both low-level variations (e.g., geometric) and high-level semantics (e.g., object layout) must be considered simultaneously. To achieve this, feature representations of task edges should be enhanced/preserved, while non-task edges within region boundaries should be suppressed, since features of the same class should be similar to each other.

In this work, we accomplish explicit edge modeling in an *interpretable* way using DUs that learn to enhance edges that are beneficial to the task and suppress edges that are irrelevant (Section 3.4).

3.2. Preliminary

The diffusion equation describes the movement of diffusive substances from regions of higher concentration to lower concentration without creating or destroying mass [30]. For example, when sugar is put into a glass of water, it diffuses evenly and sweetens the water.

Let $u(p, t)$ denote the concentration at the position p and time t . The amount of substances that flow through per unit area per unit time (the flux s) is described by Fick’s law:

$$s = -g \cdot \nabla u, \quad (1)$$

where ∇ denotes the gradient operator and g denotes the diffusivity. This diffusivity denotes the rate of diffusion. The fact that diffusion processes do not create or destroy mass is expressed by the continuity equation:

$$\partial_t u = -\text{div}(s). \quad (2)$$

The continuity equation indicates that the change of concentration over time is caused only by the flux, which is measured by the divergence operator (div). Finally, the diffusion process is described by combining Eqs. (1) and (2):

$$\partial_t u = \text{div}(g \cdot \nabla u) \quad t \geq 0, \quad (3)$$

with the initial condition $u(p, 0) = u_0(p)$ and the boundary condition as appropriate.

3.3. Continuous DU

Here we first adapt the general concept of diffusion to point cloud analysis. Informally, let us consider a case where $u(p, t)$ denotes a feature value or vector (e.g., sensor intensity or colors) of the position p on a surface at time t . The strength of the edge is then described by ∇u . Therefore, Eq. (3) describes the change of the feature at position p in a short period of time (or in a discrete sense, the change of the feature after one step of diffusion). As can be seen, the behavior of the diffusion process is governed by the diffusivity g . g is often a hand-crafted function of u , i.e., $g = g(u)$, and thus the diffusion process is adaptive to evolving feature values. This adaptability enables classic diffusion methods to realize edge-preserving filtering (e.g., [29,55]). On the other hand, the behaviors of most existing choices of g are theoretically restricted [30], that is, there is a limit to the total amount of smoothing or enhancement that can be performed by these methods.

Extending the diffusion equation using existing choices of g to our setting is nontrivial. First, g is often a function with hand-crafted parameters that require considerable domain knowledge to determine. This is difficult in our setting due to the large variability of input point clouds. Second, adopting existing forms of g may significantly limit expressiveness because of their theoretical restrictions. We want the method to be able to smooth and enhance features without restrictions if it is the right thing to do.

We now propose a new way to extend the diffusion equation to our setting which is not affected by the above issues. To this end, we propose *Diffusion Unit (DU)*, an algorithm that performs edge enhancement and suppression learning without restrictions. Formally, suppose a continuous spatial-temporal multi-channel point cloud $\mathbf{u} = \mathbf{u}(\mathbf{p}, t) = (u_1(\mathbf{p}, t), u_2(\mathbf{p}, t), \dots, u_d(\mathbf{p}, t))$, where d is the number of channels, t is time, \mathbf{p} denotes the position vector, and the initial condition is $\mathbf{u}(\mathbf{p}, 0) = \mathbf{h}$. We define (continuous) DU as:

$$\partial_t \mathbf{u} = \text{div}(\Phi(\nabla \mathbf{u})), \quad t \geq 0, \quad (4)$$

where $\nabla \mathbf{u} \in \mathbb{R}^{3 \times d}$ encodes channel-wise spatial (3D) gradient. We define $\Phi: \mathbb{R}^{3 \times d} \rightarrow \mathbb{R}^{3 \times d}$ as a trainable mapping of a form $\Phi(\nabla \mathbf{u}) = (\phi(\mathbf{u}_x), \phi(\mathbf{u}_y), \phi(\mathbf{u}_z))$ where $\phi: \mathbb{R}^d \rightarrow \mathbb{R}^d$. In other words, Φ applies ϕ to all d -dim components of the gradient $\nabla \mathbf{u}$. In practice, we use an MLP to implement ϕ .

As can be seen, our idea is to replace the diffusivity weighted gradient with a trainable mapping that takes the gradient as input. In doing so, the equation is optimized with respect to data and no longer depends on hand-crafted parameters. Furthermore, the above form of Φ naturally makes the equation permutation-invariant when it is discretized, a useful property for point cloud analysis (this point will be explained in more detail in Section 3.5). In addition, the use of MLP enables the equation to learn and exploit channel correlations which are proven to be extremely helpful for vision tasks [68]. Lastly, and most importantly, the design of Φ not only allows the equation to perform edge enhancement and suppression without restrictions but also enables us to understand the behavior of the equation in a mathematically transparent manner, making DU’s behavior interpretable. This point will be fully discussed in the following section.

3.4. Theoretical analysis of edge enhancement and suppression learning

In this part, we explain how DU performs edge enhancement and suppression learning. For simplicity, we consider a step edge (i.e., an abrupt change in feature values) convolved by a Gaussian. Ideally, such a structure can be detected by the spatial gradient $\nabla \mathbf{u}$. Without loss of generality, we assume that the edge is aligned with x axis ($\mathbf{u}_y = \mathbf{u}_z = \mathbf{0}$). The profile of a step edge and its derivatives are illustrated in Fig. 3. Now we focus on a single output channel i of Eq. (4) for brevity. In this setting, Eq. (4) can be simplified as:

$$(\partial_t \mathbf{u})_i = \frac{\partial}{\partial x} (\phi_i(\mathbf{u}_x)) = \nabla \phi_i \cdot \mathbf{u}_{xx} \quad (5)$$




Fig. 3. Profiles of the smoothed step edge (top) and its derivatives. At the inflection point, the second derivative is zero while the third derivative is smaller than 0.

$$= \sum_{j=1}^d (\phi'_i)_j \cdot (\mathbf{u}_{xx})_j, \quad i = 1, 2, \dots, d, \quad (6)$$

where j indexes the input channel. We are interested in how the edge, i.e., $(\mathbf{u}_x)_i$, changes during the diffusion process. In other words, we want to investigate the behavior of $(\partial_t \mathbf{u}_x)_i$. This can be realized using the above equations:

$$(\partial_t \mathbf{u}_x)_i = \left(\frac{\partial}{\partial t} \frac{\partial \mathbf{u}}{\partial x} \right)_i = \frac{\partial}{\partial x} \left(\frac{\partial \mathbf{u}}{\partial t} \right)_i \quad (7)$$

$$= \frac{\partial}{\partial x} \left(\sum_{j=1}^d (\phi'_i)_j \cdot (\mathbf{u}_{xx})_j \right) \quad (8)$$

$$= \sum_{j=1}^d (\phi''_i)_j \cdot (\mathbf{u}_{xx})_j^2 + (\phi'_i)_j \cdot (\mathbf{u}_{xxx})_j. \quad (9)$$

As shown in Fig. 3, at the inflection point $(\mathbf{u}_{xx})_j = 0$ and $(\mathbf{u}_{xxx})_j < 0$. Therefore, the behavior of $(\partial_t \mathbf{u}_x)_i$ is solely decided by the second term of Eq. (9). Since $(\mathbf{u}_{xxx})_j < 0$, the effect of a particular input channel j to the output channel i is determined by the sign of $(\phi'_i)_j$. Specifically, channel j has a positive impact if $(\phi'_i)_j < 0$, whereas it has a negative impact if $(\phi'_i)_j > 0$. Therefore, edge enhancement $((\partial_t \mathbf{u}_x)_i > 0)$ or suppression $((\partial_t \mathbf{u}_x)_i < 0)$ of a particular feature channel i can be achieved by ϕ through summing up contributions from all input channels. Note that only task-beneficial edges are enhanced/suppressed because ϕ is trained w.r.t data and tasks.

The empirical verification of edge enhancement and suppression behavior is presented in Section 4.1.

3.5. Discrete DU

The discretization of Eq. (4) is necessary because the 3D point clouds are discrete samples of continuous surfaces. Let s and $n \in \mathcal{N}_s$ denote the center point and its neighbor, respectively. Then, $\mathbf{u}_s, \mathbf{u}_n \in \mathbb{R}^d$

denote the features of the center point and its neighbor. Eq. (4) is continuous both in time and space, so both parts must be discretized. For the time discretization, we adopt the simple explicit scheme [30]:

$$\partial_t \mathbf{u}_s \approx \mathbf{u}_s^{\tau+1} - \mathbf{u}_s^{\tau}, \quad (10)$$

where τ denotes the discrete time index (iteration). For the space discretization, we use the following discretization scheme:

$$\text{div}(\Phi(\nabla \mathbf{u})) \approx \sum_{n \in \mathcal{N}_s} \phi(\mathbf{u}_n^{\tau} - \mathbf{u}_s^{\tau}), \quad (11)$$

where the gradient ∇u and divergence operator $\text{div}(\cdot)$ are approximated by the element-wise difference $\mathbf{d}_{sn} = \mathbf{u}_n^{\tau} - \mathbf{u}_s^{\tau}$ and summation $\sum_{n \in \mathcal{N}_s} \mathbf{d}_{sn}$, respectively, similar to [64]. In addition, this discretization scheme effectively makes DU an edge-aware method with the edge feature $\phi(\mathbf{u}_n^{\tau} - \mathbf{u}_s^{\tau})$.

Combining Eqs. (4), (10), (11), we have

$$\mathbf{u}_s^{\tau+1} = \mathbf{u}_s^{\tau} + \sum_{n \in \mathcal{N}_s} \phi(\mathbf{u}_n^{\tau} - \mathbf{u}_s^{\tau}). \quad (12)$$

Eq. (12) has two important characteristics for deep learning-based point cloud analysis. First, the resulting equation belongs to the residual learning framework [69], which makes optimization easier. Second, as we mentioned in Section 3.3, the equation is permutation-invariant because ϕ is a shared/pointwise MLP. This property is particularly useful in point cloud analysis, where the order of points is arbitrary and therefore the structured functions cannot be applied [13].

Furthermore, we add some improvements from the perspective of engineering. First, we average the influence of neighbors in the second term of Eq. (12) as this slightly increases performance. We speculate that the averaging enhances the outlier robustness. Second, we apply Batch Normalization [70] and ReLU [71] $\varphi = \text{ReLU} \circ \text{BatchNorm}$ to the second term to further facilitate optimization. As a result, the discretized DU is defined as:

$$\mathbf{u}_s^{\tau+1} = \mathbf{u}_s^{\tau} + \varphi \left(\frac{1}{|\mathcal{N}_s|} \sum_{n \in \mathcal{N}_s} \phi(\mathbf{u}_n^{\tau} - \mathbf{u}_s^{\tau}) \right). \quad (13)$$

The above definition of DU can be efficiently computed on modern computers and is used throughout this study. Its detailed computation flow is shown in Fig. 4 (middle).

3.6. Network architecture of DU-net

Using DU as a building block, we construct *DU-Net* to tackle point cloud segmentation. An overview of the network architecture is shown in Fig. 4 (top). In designing DU-Net, we follow the most common encoder-decoder style to achieve efficient hierarchical learning. Such an architecture style is adopted by numerous previous works (e.g., [14, 17, 27]).

Similar to CNNs, the encoder is responsible for hierarchical feature abstraction. In each encoder stage, input point features are transformed and subsequently downsampled. The downsampling ensures an efficient encoding of multi-resolution characteristics, leading to a more discriminative feature representation. After downsampling, point features are passed to the next stage. We construct a CNN-like encoder based on KPConv [17], which is a standard convolution operator adopted in many previous works for its excellent performance and ease of implementation [22, 72]. Specifically, we select the depthwise [73] version to reduce the overall complexity. Following prior works [14, 21], we adopt the farthest point sampling as the downsampling method.

The decoder recovers the original resolution of point clouds by performing successive upsampling followed by a DU. An upsampling layer receives the output of the previous layer and recovers the resolution of points in the adjacent upper stage of the encoder. We adopt U-Net [74]-style skip connection to assist feature reconstruction. A

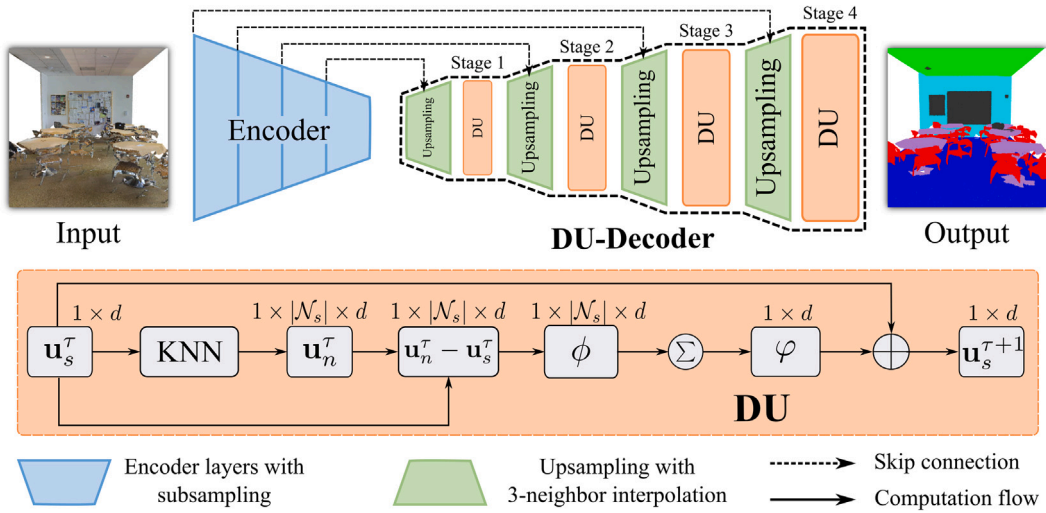


Fig. 4. The network architecture used in this study and the computation flow of the discretized DU (Eq. (13)). Top: the network architecture. Middle: computation flow of the DU. Here we assume the input is a point with a d -dimensional feature. Σ and \oplus denote neighborhood aggregation using summation and element-wise sum, respectively.

DU is subsequently applied to the upsampled features so that task-beneficial edges of each resolution are kept sharpened while irrelevant discontinuities are smoothed out.

We use DUs in the upsampling part because the upsampling part receives the features from the encoder part via skip connections (dotted arrows in Fig. 4). This means that the input features of the DU are a direct combination of the encoder and decoder features. Therefore, it is considered sufficient to place the DU only in the upsampling part. In addition, we confirm experimentally that including DUs in the encoder does not improve performance. We refer to our decoder as DU-decoder for brevity.

4. Experiments

In this section, we conduct experiments to answer the following questions:

1. Does DU really perform edge enhancement/suppression as described by the theoretical analysis in Section 3.4?
2. Can DU-Net compete with recent cutting-edge networks?
3. Is the design choices of DU and DU-Net reasonable?

To answer these questions, we use standard benchmarks of object part segmentation and scene segmentation tasks.

For object part segmentation, we use ShapeNet part dataset [31], which contains 16,881 3D models that are classified into 16 object categories. Each model is annotated with several parts (less than 6) from 50 object part classes. For instance, an airplane category usually consists of parts including airplane body, wings, tails, and engine classes.

For scene segmentation, we use Stanford large-scale 3D indoor spaces (S3DIS) [32]. The dataset contains 3D scans taken from six areas including 272 rooms. For instance, rooms such as lobby, office, and conference rooms are included. Each point is annotated with a class from 13 categories. A room, for example, in the dataset includes classes such as floor, wall, door, and table.

4.1. Verifying the behaviors of DU

We have theoretically analyzed that DU learns to enhance or suppress edges in Section 3.4. Here we empirically verify the edge enhancement and suppression behaviors of DU through visualizations. The features used for this experiment are taken from the last DU (stage 4 in Fig. 4) since it is the closest layer to the final classification layer.




Fig. 5. Examples of local feature smoothness before and after applying DU using the ShapeNet part test set. A low response indicates a smooth region while a high response indicates a non-smooth region. Red and orange rectangles indicate enhanced and suppressed edges after DU. The task-beneficial edges (in this case, part boundaries) are enhanced (become brighter) while within-part features are suppressed (become darker).

To verify the effect of DU, we examine *local feature smoothness* that summarizes how different the features are from their local neighbors. Formally, local feature smoothness is defined as $\|\sum_{n \in \mathcal{N}_s} \mathbf{f}_n - \mathbf{f}_s\|$, where




Fig. 6. Examples of local feature smoothness before and after applying DU using the S3DIS Area 5 test set. A low response indicates a smooth region while a high response indicates a non-smooth region. Red and orange rectangles indicate enhanced and suppressed features after DU. The task-beneficial edges (in this case, object boundaries) are enhanced (become brighter) while intra-region features are suppressed (become darker).

f_s and f_n denote the features of a center point and its neighbor from some layer. It can be observed that the smoothness value is closer to zero if the feature of the center point is similar to those of its neighbors while the value increases when they are dissimilar. Therefore, we can qualitatively grasp the edge enhancement/suppression behavior of DU by comparing smoothness values before and after applying DU.

Fig. 5 shows some examples of the change in local feature smoothness after applying DU from the ShapeNet part dataset. In this case, DU-Net is trained to assign part categories to each point. As can be seen, DU-Net is able to perform edge enhancement and suppression simultaneously. For instance, the third row shows a laptop whose boundary between its screen and keyboard is enhanced while the feature of its keyboards is smoothed. Therefore, DU successfully enhances the task-beneficial edges and suppressed the unhelpful discontinuities. Furthermore, we observe that DU can differentiate task-beneficial edges from other geometric edges (the top and last row). In the last row of Fig. 5, for example, the sharp edges on the boundary (but not on the part boundary) of the table are aggressively suppressed after applying DU. On the other hand, the points near the part boundary are enhanced by DU. Thus, DU can understand the semantics of the objects and perform edge enhancement/suppression selectively.

Fig. 6 demonstrates the effect of DU in complex scenes. The top row shows that DU successfully locates and enhances the boundary of a whiteboard. As a result, the boundary becomes more salient. In the second row, DU also manages to enhance the object boundary (e.g., the boundary between the window and wall) that is barely distinguishable before applying DU. Interestingly, only the boundaries between different categories are enhanced while other discontinuities such as the textures within the window and the window frame remain suppressed; therefore, the DU clearly understands and distinguishes between task-beneficial and other edges. The third row shows a case in which edge enhancement and suppression occur simultaneously. The discontinuous edges that appear on the floor are smoothed after the DU while the boundary between a box and wall is enhanced. In the last row, DU precisely enhances the point where the wall and the chair are attached, while nicely suppressing the edge of the table that is close to the chair but not attached to it.

Thus, the task-beneficial edge enhancement/suppression behavior of DU can be verified.

Table 1

Result of object part segmentation (%). The bold text indicates the best performance. Methods are arranged by year of publication.

Method	Year	ImIoU	CmIoU
PointNet [13]	CVPR'17	83.7	80.4
PointNet++ [14]	NIPS'17	85.1	81.9
PointCNN [15]	NIPS'18	86.1	84.6
PointConv [16]	CVPR'19	85.7	82.8
SSCN ^a [34]	CVPR'18	86.0	–
RSCNN ^a [43]	CVPR'19	86.2	84.0
KPCConv ^a [17]	ICCV'19	86.4	85.1
PAConv ^a [18]	CVPR'21	86.1	84.6
Point Transformer [21]	ICCV'21	86.6	83.7
CurveNet ^a [27]	ICCV'21	86.8	–
Stratified Transformer ^a [22]	CVPR'22	86.6	85.1
Ours (w/o DU)	–	86.2	83.6
Ours (w/ DU)	–	86.7	84.5
Ours (w/ DU) ^a	–	87.1	85.2

^a Denotes methods that adopt voting for post-processing.

4.2. Object part segmentation

For this experiment, we adopt the standard train-test split [14]. All available points are used for input. 3D coordinates along with surface normal information are used as the input features. We adopt the standard data augmentation strategy for point cloud analysis. Specifically, random anisotropic scaling in the range of [0.66, 1.5] and random translation in the range of [−0.2, 0.2] are used for data augmentation following prior works (e.g., [25]). We train the model using one NVIDIA Tesla V100 GPU. The model is trained for 150 epochs. SGD is used for optimization with an initial learning rate of 0.1. The learning rate is decayed by 0.1 when the epoch reaches 90 and 120. Following the common practice [17,18,43], we use voting for post-processing. We use the most widely adopted instance-wise mIoU (ImIoU) and category mIoU (CmIoU) defined in [13] for performance metrics. As for the loss function, we adopt a standard cross entropy loss.

The quantitative result is listed in Table 1. DU-Net achieves state-of-the-art performance both in terms of ImIoU and CmIoU. This reveals




Fig. 7. Qualitative results of object part segmentation. The red rectangles show the improvements brought by DU. As can be seen, DU encourages the model to produce smoother and more boundary-aware predictions.

that DU-Net not only performs well in the major classes but also achieves satisfactory accuracy for each class. Moreover, although the plain network (w/o DU in Table 1) fails to compete with the cutting-edge networks, DU succeeds in improving its performance, achieving the best performance.

Next, we compare the predictions of DU-Net and the plain network to provide some insights into how DU improves performance. The results are shown in Fig. 7. First, we find that DU-Net produces smoother predictions compared with the plain counterpart. For instance, the example of an airplane in the top row shows that the plain network has difficulty in distinguishing between the tail and body, resulting in ragged predictions. In contrast, DU-Net is able to produce smooth predictions even though the shape of the tail is similar to that of the body. We believe that the adaptability of DU in handling edges successfully suppresses the unhelpful edges, which in turn produces smoother predictions. Second, it is found that DU-Net tends to be more boundary-aware. In the second row, although both networks misclassify some of the points on the bag handle as bag body points, the result of DU-Net is much more accurate than the one of the plain counterpart, revealing that DU-Net has a better localization of the boundary between the bag body and handle. We conjecture that the ability to enhance task-beneficial edges makes features near boundaries more discriminative. As a result, DU-Net can detect part boundaries more precisely, thereby facilitating smooth predictions within boundaries and mitigating cross-boundary misclassifications (the third and last row).

4.3. Scene segmentation

Similar to [75], we advocate using Area five for testing and others for training. Following the common practice [17,21,36], we first grid-subsample each room with a grid size of 4 cm for training. For testing, all points are evaluated. We form an input point cloud by taking at most 24,000 points from a room. We use 3D coordinates and color information as input features. Random vertical rotation, random anisotropic scaling in the range of [0.66, 1.5], random jitter (Gaussian noise with

zero mean and 0.01 standard deviation), random color drop with a probability of 0.2, and random color auto-contrasting with a probability of 0.2 are used for data augmentation, following [21,22]. Similar to the previous work [21], we train the model for 76,500 iterations. For this dataset, four NVIDIA Tesla V100 GPUs are used for training. For optimization, AdamW [77] algorithm is used with an initial learning rate of 0.01. The learning rate is decayed based on the cosine annealing schedule. We use mIoU to quantitatively assess the performance. We use the cross entropy loss with label smoothing to train the model. The smoothing factor is set to 0.2.

The results are reported in Table 2. DU-Net achieves state-of-the-art performance in terms of mIoU, which means DU-Net performs well in terms of average performance. We find that DU-Net performs especially well in classifying classes such as door and bookcase (book in Table 2). Since these two classes often appear on the wall, one possible reason for the strong performance is that the edge enhancement capability of DU makes the boundary feature more discriminative, easing the detection of subtle but crucial boundaries. As can be seen, the plain network (w/o DU) fails to compete with the recent strong methods, while incorporating DUs achieves state-of-the-art performance. Therefore, the edge enhancement and suppression behavior of DU plays a key role in achieving good performance. As for the class-wise performance, we can verify that DUs successfully improve the performance by around 0.1–8.3, demonstrating the general usefulness of edge enhancement and suppression learning in scene segmentation.

To intuitively understand how DU improves performance, we qualitatively analyze the effect of DU. The results are shown in Fig. 8. We observe that DU-Net (w/ DU in Fig. 8) generally better discriminates between similar classes compared with the plain counterpart (w/o DU in Fig. 8). For instance, the example in the top row shows that the plain network is greatly confused with the window, wall, and column classes. DU-Net, on the other hand, correctly identifies those classes and produces much more accurate and smoother predictions. We believe that the adaptability of DU leads to enhanced sensitivity to object boundaries. Consequently, DU-Net can detect boundaries more accurately, which in turn facilitates spatial consistency within the detected boundaries. Further, the example in the second row shows that DU makes the model more shape-aware. The plain network shows significant confusion in differentiating between pipes and the ceiling as the pipes are attached to the ceiling. In contrast, DU-Net shows much better predictions. We conjecture that explicit consideration of edges enables DU-Net to obtain more shape-aware representations, resulting in better discrimination. Moreover, we observe that DU-Net can make effective use of semantic information when geometric information is inadequate. In the third row, the plain net completely fails to identify the door likely because the door is embedded in the wall, making it geometrically difficult to differentiate the door from the wall. On the contrary, DU-Net succeeds in producing far better predictions. Similarly, in the last row, although the plain net manages to detect the centers of drawings, it fails to identify precise boundaries. In contrast, DU-Net produces fine predictions with accurate boundaries, demonstrating its advanced capability to utilize semantic information.

4.4. Design analysis

We further study the design of DU and DU-Net. All experiments are conducted on the S3DIS dataset and the experiment settings are the same as described in Section 4.3.

Ablation study on DU The results are shown in Table 3. First, the effect of functions ϕ and φ are investigated. The performance drops significantly by 2.8 when ϕ is removed. This is reasonable since the adaptability of DU is mostly provided by ϕ . Therefore, the importance of ϕ in the design of DU is verified. Next, we observe that removing φ from DU results in a reduction of 1.4. Recall that $\varphi = \text{ReLU} \circ \text{BatchNorm}$ is mainly applied to ease optimization. Thus, we believe that removing




Fig. 8. Qualitative results of scene segmentation using S3DIS. The red rectangles show the improvements brought by DU. In general, DU-Net (w/ DU) tends to better discriminate between similar classes (the first row) and produce smoother and boundary-aware predictions (the second row). In some cases, DU-Net correctly identifies objects that the plain counterpart (w/o DU) almost completely failed to recognize (the third and fourth rows).

Table 2

Result of scene segmentation (%). The bold text indicates the best performance. Methods are arranged by year of publication.

Method	Year	mIoU	ceil.	floor	wall	beam	col.	wind.	door	chair	table	book.	sofa	board	clut.
PointNet [13]	CVPR'17	41.1	88.8	97.3	69.8	0.5	3.92	46.3	10.8	52.6	58.9	40.3	5.9	26.4	33.2
SegCloud [75]	3DV'17	48.9	90.1	96.1	69.9	0.0	18.4	38.4	23.1	75.9	70.4	58.4	40.9	13.0	41.6
PointCNN [15]	NIPS'18	57.3	92.3	98.2	79.4	0.0	17.6	22.8	62.1	80.6	74.4	66.7	31.7	62.1	56.7
PointWeb [52]	CVPR'19	60.3	92.0	98.5	79.4	0.0	21.1	59.7	34.8	88.3	76.3	69.3	46.9	64.9	52.5
MinkowskiNet [35]	CVPR'19	65.4	91.8	98.7	86.2	0.0	34.1	48.9	62.4	81.6	89.8	47.2	74.9	74.4	58.6
KPConv [17]	ICCV'19	67.1	92.8	97.3	82.4	0.0	23.9	58.0	69.0	91.0	81.5	75.3	75.4	66.7	58.9
JSENet [72]	ECCV'20	67.7	93.8	97.0	83.0	0.0	23.2	61.3	71.6	89.9	79.8	75.6	72.3	72.7	60.4
CBL [76]	CVPR'22	69.4	93.9	98.4	84.2	0.0	37.0	57.7	71.9	91.7	81.8	77.8	75.6	69.1	62.9
Point Trans. [21]	ICCV'21	70.4	94.0	98.5	86.3	0.0	38.0	63.4	74.3	89.1	82.4	74.3	80.2	76.0	59.3
Stratified Trans. [22]	CVPR'22	72.0	96.2	98.7	85.6	0.0	46.1	60.0	76.8	84.5	92.6	75.2	77.8	78.1	64.0
PointMixer [40]	ECCV'22	71.4	94.2	98.2	86.0	–	43.8	62.1	78.5	82.2	90.6	79.8	73.9	78.5	59.4
PointNeXt [19]	NeurIPS'22	70.8	–	–	–	–	–	–	–	–	–	–	–	–	–
Point Trans. V2 [23]	NeurIPS'22	71.6	–	–	–	–	–	–	–	–	–	–	–	–	–
Ours (w/o DU)	–	69.7	94.8	98.5	85.0	0.0	31.7	58.6	75.7	80.8	91.8	76.1	75.3	79.5	58.8
Ours (w/ DU)	–	72.2	95.6	98.6	85.2	0.0	40.0	60.7	82.7	83.1	90.8	83.5	78.5	75.9	64.1

Table 3

Ablation study on DU. Model 1 indicates the default design of DU. Numbers in parentheses denote the performance reduction compared with the default DU. All experiments are performed on the S3DIS dataset.

Model	$ \mathcal{N}_s $	Edge feat.	ϕ	φ	mIoU
1 (default)	16	$\mathbf{u}_n^r - \mathbf{u}_s^r$	✓	✓	72.2
2	16	$\mathbf{u}_n^r - \mathbf{u}_s^r$	–	✓	69.4 (–2.8)
3	16	$\mathbf{u}_n^r - \mathbf{u}_s^r$	✓	–	70.8 (–1.4)
4	16	$\mathbf{u}_n^r - \mathbf{u}_s^r$	–	–	69.8 (–2.4)
5	8	$\mathbf{u}_n^r - \mathbf{u}_s^r$	✓	✓	70.9 (–1.3)
6	24	$\mathbf{u}_n^r - \mathbf{u}_s^r$	✓	✓	71.2 (–1.0)
7	28	$\mathbf{u}_n^r - \mathbf{u}_s^r$	✓	✓	70.6 (–1.6)
8	16	$\mathbf{u}_n^r - \mathbf{u}_s^r$	✓	✓	69.4 (–2.8)

φ increases the optimization difficulty, which adversely affects the performance. On the other hand, the reduction of the performance is less severe compared with removing ϕ ; hence, adaptability is relatively more important than optimization in the design of DU. As expected, removing both functions results in a significant reduction in performance.

Interestingly, the reduction of performance caused by removing both functions is less severe than the one that removes ϕ . We conjecture that the direct application of φ without any transformation may lead to information loss since φ contains ReLU. Then, we vary $|\mathcal{N}_s|$ from 8 to 28 to analyze the effect of neighborhood size in the design of DU. As can be observed, the lower scores are obtained when $|\mathcal{N}_s|$ takes too small (8) or too large values (28). On the other hand, better performance is achieved when $|\mathcal{N}_s|$ takes intermediate values. Therefore, we set $|\mathcal{N}_s|$ to 16 throughout this study. Finally, we replace the neighbor difference in Eq. (13) ($\mathbf{u}_n^r - \mathbf{u}_s^r$) with the point feature (\mathbf{u}_n^r) to validate the edge feature we used. The point feature is often used in layers that perform neighbor aggregation (e.g., graph convolution layers). As a result, we observe a 2.8 decrease in mIoU when we replace the difference with the point feature. We speculate that the use of the point feature causes DU to perform neighborhood aggregation, thereby losing the ability to explicitly model edges.

Influence of DUs at different stages Since DUs are applied to various stages of the decoder, we investigate the individual influence of DU at each stage in terms of performance and computational complexity. The

Table 4

Influence of DUs at different stages. Stages indicate different resolutions (see Fig. 4). The numbers of parameters (M) and floating-point operations (FLOPs) (G) are also reported. The numbers in parentheses denote the performance reduction compared with the default DU-Net (DU at all stages).

Model	#param. (M)	FLOPs (G)	mIoU
DU at all stages	8.07	11.94	72.2
DU at stage 1	7.80	10.05	70.3 (−1.9)
DU at stage 2	7.18	10.06	70.6 (−1.6)
DU at stage 3	7.03	10.07	71.0 (−1.2)
DU at stage 4	6.99	10.07	70.3 (−1.9)
w/o DU	6.98	9.43	69.7 (−2.5)

Table 5

Comparison with other methods. Various methods and DU are compared by replacing DU with other methods in the decoder. All decoders are paired with the encoder of DU-Net. Type denotes the main operation in each decoder layer. The values in the parentheses indicate the performance reduction compared with DU-decoder.

Decoder layer	Type	mIoU
DU (Ours)	Edge-aware	72.2
Feature propagation [14]	MLP	69.7 (−2.5)
PointDeconv [16]	Deconvolution	70.0 (−2.2)
Vector attention [21]	Edge-aware	70.3 (−1.9)

results are listed in Table 4. From stages 1 to 4, the resolution increases because of upsamplings while the feature dimension decreases. First, it is found that a single DU can immediately improve performance. As can be seen, DU at each stage improves upon the plain network (w/o DU) by 0.6–1.3. Therefore, although the amount of improvement varies, each DU can make effective use of information at the corresponding stage. Second, we observe that the benefits brought by DU at each stage can be accumulated. Specifically, applying DUs at all stages achieves the best performance while the effect of a single DU is limited. We speculate that the adaptability of the DUs allows them to be tailored to each resolution, resulting in a cumulative improvement. Third, we find that because of the lightweight nature of the DUs, each DU produces improvement with only a slight increase in computation cost. For example, applying DU at stage 3 improves mIoU by 1.3, but only increases #param. and FLOPs by 0.7% and 6.7%, respectively. Thus, an acceptable cost-performance tradeoff can be achieved when computational resources are limited.

Comparison with other methods We compare DU with different methods to validate its effectiveness. Several methods of different types (second column in Table 5) are selected for comparison. Specifically, methods are compared with DU by replacing DU with those methods. The encoder of DU-Net is used as the default encoder. The results are listed in Table 5. We first compare DU with the feature propagation layer, which is a core layer of PointNet++ [14]. This layer is a standard layer that has been adopted by numerous previous works (e.g., [17,18,22,43]). As can be seen, DU-decoder outperforms the feature propagation layer by 2.5, demonstrating its effectiveness. The second method is PointDeconv, a deconvolution layer based on PointConv [16]. We can see that DU-decoder outperforms pointDeconv by 2.2 mIoU. The third method is vector attention [21], which is the core component of Point Transformer [21]. Vector attention uses edges as a similarity measure for generating attention weights. The result shows that using DU can provide a 1.9 mIoU gain compared with vector attention.

Effect of DU on voxel-based methods Voxel-based methods are another popular type of network for point cloud analysis. They use regular grids to approximate the connectivity between points; therefore, we expect that the explicit modeling of edges provided by DUs is also valid for those networks. To validate this point, we use two voxel-based methods, submanifold sparse convolutional networks (SSCN) and MinkowskiNet, as baselines. We reimplemented the two networks and conducted experiments according to our experimental setup. In addition, we also make sure that the reimplemented networks performed

Table 6

Effect of DU on voxel-based methods. We report ImIoU for SSCN and mIoU for MinkowskiNet. Numbers in parentheses indicate improvements brought by DU.

Method	Dataset	w/o DU	w/ DU
SSCN [34]	ShapeNet	86.1	86.4 (+0.3)
MinkowskiNet [35]	S3DIS	70.4	71.3 (+0.9)

as well as or better than the original networks. Experiments related to SSCN and MinkowskiNet are performed using ShapeNet and S3DIS, respectively. The results are reported in Table 6. As can be seen, DU improves SSCN by 0.3% ImIoU and improves MinkowskiNet by 0.9%. Consequently, we confirm that DU is applicable to voxel-based networks.

5. Conclusion

3D point clouds naturally lack structural information that describes the underlying continuous surfaces. To alleviate this issue, edge information has been used to describe the local structures of point clouds. While edges have proven useful, it is still unclear how they help to improve. In this study, we propose DU that allows us to handle edges in a principled and interpretable manner. First, we theoretically figure out that DU performs task-beneficial edge enhancement and suppression. Second, we verify the result of theoretical analysis by intuitive visualizations. Third, we validate that the incorporation of DUs into the network design improves performance. DU-Net is constructed to tackle point cloud segmentation. Specifically, DU-Net achieves state-of-the-art performance in object part segmentation and scene segmentation. We believe that DU can provide users with a clear idea of how to use this method, what results to expect, and how to diagnose problems when it fails, which is valuable since most of the methods based on DNNs operate in a black-box manner.

One possible limitation of DU is that it may fail when point clouds are significantly corrupted (e.g., by noise, density variations, etc.), since DU relies heavily on the spatial gradient of point features (Eq. (4)). In addition, we suspect that DU may not be appropriate for tasks that do not benefit from explicit edge modeling. For example, unlike segmentation, DU may have a limited impact on classification tasks because optimizing point connectivity or edges is not closely related to the task objective.

From an application point of view, it will be interesting to apply DU to other domains where the modeling of edges is crucial. An example is damage detection using point clouds [78]. Damage is often characterized by non-smooth distributions; thus, it would be interesting to see how enhanced adaptability in handling edges influences the damage detection result.

CRedit authorship contribution statement

Haoyi Xiu: Conceptualization, Methodology, Software, Writing – original draft, Validation, Writing – review & editing, Visualization. **Xin Liu:** Conceptualization, Writing – original draft, Writing – review & editing, Supervision, Funding acquisition. **Weimin Wang:** Conceptualization, Writing – original draft, Writing – review & editing. **Kyoung-Sook Kim:** Writing – review & editing, Supervision, Funding acquisition. **Takayuki Shinohara:** Writing – review & editing. **Qiong Chang:** Writing – review & editing. **Masashi Matsuoka:** Resources, Data curation, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The code is publicly available at <https://github.com/martianxiu/DiffusionUnit>.

Acknowledgments

This paper is based on results obtained from a project, JPNP20006, commissioned by the New Energy and Industrial Technology Development Organization (NEDO). We would also like to acknowledge the support from JSPS Grant-in-Aid for Scientific Research (21K12042).

References

- [1] C.R. Qi, W. Liu, C. Wu, H. Su, L.J. Guibas, Frustum pointnets for 3d object detection from rgb-d data, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 918–927.
- [2] Y. Cui, R. Chen, W. Chu, L. Chen, D. Tian, Y. Li, D. Cao, Deep learning for image and point cloud fusion in autonomous driving: A review, *IEEE Trans. Intell. Transp. Syst.* (2021).
- [3] L.-H. Wen, K.-H. Jo, Deep learning-based perception systems for autonomous driving: A comprehensive survey, *Neurocomputing* (2022).
- [4] T. Shinohara, H. Xiu, M. Matsuoka, FWNNet: Semantic segmentation for full-waveform LiDAR data using deep learning, *Sensors* 20 (12) (2020) 3568.
- [5] H. Xiu, X. Liu, W. Wang, K.-S. Kim, T. Shinohara, Q. Chang, M. Matsuoka, DS-net: A dedicated approach for collapsed building detection from post-event airborne point clouds, *Int. J. Appl. Earth Obs. Geoinf.* 116 (2023) 103150.
- [6] A. Kanazaki, Y. Matsushita, Y. Nishida, Rotationnet: Joint object categorization and pose estimation using multiviews from unsupervised viewpoints, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 5010–5019.
- [7] H. Su, S. Maji, E. Kalogerakis, E. Learned-Miller, Multi-view convolutional neural networks for 3d shape recognition, in: Proceedings of the IEEE International Conference on Computer Vision, 2015, pp. 945–953.
- [8] D. Maturana, S. Scherer, Voxnet: A 3d convolutional neural network for real-time object recognition, in: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2015, pp. 922–928.
- [9] Y. Zhou, O. Tuzel, Voxelnet: End-to-end learning for point cloud based 3d object detection, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 4490–4499.
- [10] P.-S. Wang, Y. Liu, Y.-X. Guo, C.-Y. Sun, X. Tong, O-cnn: Octree-based convolutional neural networks for 3d shape analysis, *ACM Trans. Graph.* 36 (4) (2017) 1–11.
- [11] M. Tatarchenko, A. Dosovitskiy, T. Brox, Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs, in: Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 2088–2096.
- [12] R. Klovov, V. Lempitsky, Escape from cells: Deep kd-networks for the recognition of 3d point cloud models, in: Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 863–872.
- [13] C.R. Qi, H. Su, K. Mo, L.J. Guibas, Pointnet: Deep learning on point sets for 3d classification and segmentation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 652–660.
- [14] C.R. Qi, L. Yi, H. Su, L.J. Guibas, PointNet++: Deep hierarchical feature learning on point sets in a metric space, *Adv. Neural Inf. Process. Syst.* 30 (2017).
- [15] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, B. Chen, Pointcnn: Convolution on χ -transformed points, in: Proceedings of the 32nd International Conference on Neural Information Processing Systems, 2018, pp. 828–838.
- [16] W. Wu, Z. Qi, L. Fuxin, Pointconv: Deep convolutional networks on 3d point clouds, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 9621–9630.
- [17] H. Thomas, C.R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, L.J. Guibas, Kpconv: Flexible and deformable convolution for point clouds, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 6411–6420.
- [18] M. Xu, R. Ding, H. Zhao, X. Qi, Paconv: Position adaptive convolution with dynamic kernel assembling on point clouds, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 3173–3182.
- [19] G. Qian, Y. Li, H. Peng, J. Mai, H. Hammoud, M. Elhoseiny, B. Ghanem, PointNeXt: Revisiting PointNet++ with improved training and scaling strategies, in: Advances in Neural Information Processing Systems (NeurIPS), 2022.
- [20] J. Yang, Q. Zhang, B. Ni, L. Li, J. Liu, M. Zhou, Q. Tian, Modeling point clouds with self-attention and gumbel subset sampling, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 3323–3332.
- [21] H. Zhao, L. Jiang, J. Jia, P.H. Torr, V. Koltun, Point transformer, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 16259–16268.
- [22] X. Lai, J. Liu, L. Jiang, L. Wang, H. Zhao, S. Liu, X. Qi, J. Jia, Stratified transformer for 3D point cloud segmentation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 8500–8509.
- [23] X. Wu, Y. Lao, L. Jiang, X. Liu, H. Zhao, Point transformer V2: Grouped vector attention and partition-based pooling, in: Advances in Neural Information Processing Systems (NeurIPS), 2022.
- [24] S. Guo, J. Li, Z. Lai, S. Han, Ctpoint: A novel local and global features extractor for point cloud, *Neurocomputing* 511 (2022) 273–289.
- [25] Y. Wang, Y. Sun, Z. Liu, S.E. Sarma, M.M. Bronstein, J.M. Solomon, Dynamic graph cnn for learning on point clouds, *ACM Trans. Graph.* 38 (5) (2019) 1–12.
- [26] Z. Liu, H. Hu, Y. Cao, Z. Zhang, X. Tong, A closer look at local aggregation operators in point cloud analysis, in: European Conference on Computer Vision, Springer, 2020, pp. 326–342.
- [27] T. Xiang, C. Zhang, Y. Song, J. Yu, W. Cai, Walk in the cloud: Learning curves for point clouds shape analysis, in: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2021.
- [28] L. Wang, Y. Huang, Y. Hou, S. Zhang, J. Shan, Graph attention convolution for point cloud semantic segmentation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 10296–10305.
- [29] P. Perona, J. Malik, Scale-space and edge detection using anisotropic diffusion, *IEEE Trans. Pattern Anal. Mach. Intell.* 12 (7) (1990) 629–639.
- [30] J. Weickert, *Anisotropic Diffusion in Image Processing*, Vol. 1, Teubner Stuttgart, 1998.
- [31] L. Yi, V.G. Kim, D. Ceylan, I.-C. Shen, M. Yan, H. Su, C. Lu, Q. Huang, A. Sheffer, L. Guibas, A scalable active framework for region annotation in 3d shape collections, *ACM Trans. Graph.* 35 (6) (2016) 1–12.
- [32] I. Armeni, O. Sener, A.R. Zamir, H. Jiang, I. Brilakis, M. Fischer, S. Savarese, 3D semantic parsing of large-scale indoor spaces, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 1534–1543.
- [33] Y. Feng, Z. Zhang, X. Zhao, R. Ji, Y. Gao, GVCNN: Group-view convolutional neural networks for 3D shape recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 264–272.
- [34] B. Graham, M. Engelcke, L. Van Der Maaten, 3D semantic segmentation with sub-manifold sparse convolutional networks, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 9224–9232.
- [35] C. Choy, J. Gwak, S. Savarese, 4D spatio-temporal convnets: Minkowski convolutional neural networks, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 3075–3084.
- [36] S. Lan, R. Yu, G. Yu, L.S. Davis, Modeling local geometric structure of 3d point clouds using geo-cnn, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 998–1008.
- [37] S. Qiu, S. Anwar, N. Barnes, Geometric back-projection network for point cloud classification, *IEEE Trans. Multimed.* (2021).
- [38] Z. Zhang, B.-S. Hua, S.-K. Yeung, Shellnet: Efficient point cloud convolutional neural networks using concentric shells statistics, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 1607–1616.
- [39] H. Ran, J. Liu, C. Wang, Surface representation for point clouds, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 18942–18952.
- [40] J. Choe, C. Park, F. Rameau, J. Park, I.S. Kweon, PointMixer: MLP-mixer for point cloud understanding, in: Proceedings of the European Conference on Computer Vision (ECCV), 2022.
- [41] I.O. Tolstikhin, N. Houlsby, A. Kolesnikov, L. Beyer, X. Zhai, T. Unterthiner, J. Yung, A. Steiner, D. Keysers, J. Uszkoreit, et al., Mlp-mixer: An all-mlp architecture for vision, *Adv. Neural Inf. Process. Syst.* 34 (2021) 24261–24272.
- [42] J. Mao, X. Wang, H. Li, Interpolated convolutional networks for 3d point cloud understanding, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 1578–1587.
- [43] Y. Liu, B. Fan, S. Xiang, C. Pan, Relation-shape convolutional neural network for point cloud analysis, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 8895–8904.
- [44] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, *Adv. Neural Inf. Process. Syst.* 30 (2017).
- [45] Z. Liu, Z. Zhang, Y. Cao, H. Hu, X. Tong, Group-free 3d object detection via transformers, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 2949–2958.
- [46] Z. Qin, H. Yu, C. Wang, Y. Guo, Y. Peng, K. Xu, Geometric transformer for fast and robust point cloud registration, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 11143–11152.
- [47] X. Yan, C. Zheng, Z. Li, S. Wang, S. Cui, Pointasnl: Robust point clouds processing using nonlocal neural networks with adaptive sampling, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 5589–5598.
- [48] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional transformers for language understanding, in: J. Burstein, C. Doran, T. Solorio (Eds.), Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2–7, 2019, Volume 1 (Long and Short Papers), Association for Computational Linguistics, 2019, pp. 4171–4186, <http://dx.doi.org/10.18653/v1/n19-1423>.

- [49] X. Yu, L. Tang, Y. Rao, T. Huang, J. Zhou, J. Lu, Point-bert: Pre-training 3d point cloud transformers with masked point modeling, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 19313–19322.
- [50] Z. Li, X. Tang, Z. Xu, X. Wang, H. Yu, M. Chen, X. Wei, Geodesic self-attention for 3D point clouds, in: *Advances in Neural Information Processing Systems*, 2022.
- [51] G. Li, M. Muller, A. Thabet, B. Ghanem, Deepgcns: Can gcns go as deep as cnns? in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9267–9276.
- [52] H. Zhao, L. Jiang, C.-W. Fu, J. Jia, Pointweb: Enhancing local neighborhood features for point cloud processing, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5565–5573.
- [53] H. Xiu, X. Liu, W. Wang, K. Kim, T. Shinohara, Q. Chang, M. Matsuoka, Enhancing local feature learning for 3D point cloud processing using unary-pairwise attention, in: *32nd British Machine Vision Conference 2021, BMVC 2021*, Online, November 22–25, 2021, BMVA Press, 2021, p. 49, URL: <https://www.bmvc2021-virtualconference.com/assets/papers/0678.pdf>.
- [54] L. Jiang, H. Zhao, S. Liu, X. Shen, C.-W. Fu, J. Jia, Hierarchical point-edge interaction network for point cloud semantic segmentation, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 10433–10441.
- [55] M.J. Black, G. Sapiro, D.H. Marimont, D. Heeger, Robust anisotropic diffusion, *IEEE Trans. Image Process.* 7 (3) (1998) 421–432.
- [56] J. Weickert, Coherence-enhancing diffusion filtering, *Int. J. Comput. Vis.* 31 (2) (1999) 111–127.
- [57] T. Brox, J. Weickert, B. Burgeth, P. Mrazek, Nonlinear structure tensors, *Image Vis. Comput.* 24 (1) (2006) 41–55.
- [58] M. Desbrun, M. Meyer, P. Schröder, A.H. Barr, Implicit fairing of irregular meshes using diffusion and curvature flow, in: *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, 1999, pp. 317–324.
- [59] U. Clarenz, U. Diewald, M. Rumpf, Anisotropic geometric diffusion in surface processing, in: *Visualization Conference, IEEE, IEEE Computer Society, Los Alamitos, CA, USA*, 2000, p. 70, <http://dx.doi.org/10.1109/VISUAL.2000.885721>, URL: <https://doi.ieeecomputersociety.org/10.1109/VISUAL.2000.885721>.
- [60] C.L. Bajaj, G. Xu, Anisotropic diffusion of surfaces and functions on surfaces, *ACM Trans. Graph.* 22 (1) (2003) 4–32.
- [61] J. Atwood, D. Towsley, Diffusion-convolutional neural networks, *Adv. Neural Inf. Process. Syst.* 29 (2016).
- [62] J. Gasteiger, S. Weissenberger, S. Günnemann, Diffusion improves graph learning, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, 32, Curran Associates, Inc., 2019, URL: <https://proceedings.neurips.cc/paper/2019/file/23c894276a2c5a16470e6a31f4618d73-Paper.pdf>.
- [63] J. Zhao, Y. Dong, M. Ding, E. Kharlamov, J. Tang, Adaptive diffusion in graph neural networks, *Adv. Neural Inf. Process. Syst.* 34 (2021).
- [64] B. Chamberlain, J. Rowbottom, M.I. Gorinova, M. Bronstein, S. Webb, E. Rossi, Grand: Graph neural diffusion, in: *International Conference on Machine Learning*, PMLR, 2021, pp. 1407–1418.
- [65] S. Luo, W. Hu, Diffusion probabilistic models for 3d point cloud generation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2837–2845.
- [66] L. Zhou, Y. Du, J. Wu, 3D shape generation and completion through point-voxel diffusion, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5826–5835.
- [67] Z. Lyu, Z. Kong, X. XU, L. Pan, D. Lin, A conditional point diffusion-refinement paradigm for 3D point cloud completion, in: *International Conference on Learning Representations*, 2022, URL: <https://openreview.net/forum?id=wqD6TfYkrm>.
- [68] J. Hu, L. Shen, G. Sun, Squeeze-and-excitation networks, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7132–7141.
- [69] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [70] S. Ioffe, C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: *International Conference on Machine Learning*, PMLR, 2015, pp. 448–456.
- [71] X. Glorot, A. Bordes, Y. Bengio, Deep sparse rectifier neural networks, in: *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, JMLR Workshop and Conference Proceedings*, 2011, pp. 315–323.
- [72] Z. Hu, M. Zhen, X. Bai, H. Fu, C.-I. Tai, Jsenet: Joint semantic segmentation and edge detection network for 3d point clouds, in: *European Conference on Computer Vision*, Springer, 2020, pp. 222–239.
- [73] L. Sifre, S. Mallat, Rigid-motion scattering for texture classification, 2014, arXiv preprint [arXiv:1403.1687](https://arxiv.org/abs/1403.1687).
- [74] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2015, pp. 234–241.
- [75] L. Tchapmi, C. Choy, I. Armeni, J. Gwak, S. Savarese, Segcloud: Semantic segmentation of 3d point clouds, in: *2017 International Conference on 3D Vision (3DV)*, IEEE, 2017, pp. 537–547.
- [76] L. Tang, Y. Zhan, Z. Chen, B. Yu, D. Tao, Contrastive boundary learning for point cloud segmentation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 8489–8499.
- [77] I. Loshchilov, F. Hutter, Decoupled weight decay regularization, in: *International Conference on Learning Representations*, 2019, URL: <https://openreview.net/forum?id=Bkg6RiCqY7>.
- [78] H. Xiu, T. Shinohara, M. Matsuoka, M. Inoguchi, K. Kawabe, K. Horie, Collapsed building detection using 3D point clouds and deep learning, *Remote Sens.* 12 (24) (2020) 4057.

Haoyi Xiu received M.M.G from Keio University and a Ph.D. degree from Tokyo Institute of Technology, respectively. He is currently working as a postdoctoral researcher at the Artificial Intelligence Research Center (AIRC) of the National Institute of Advanced Industrial Science and Technology (AIST). His research interests include 3D point cloud processing, deep learning, machine learning, computer vision, remote sensing, and building damage detection.

Xin Liu is a senior researcher at the Artificial Intelligence Research Center (AIRC) of the National Institute of Advanced Industrial Science and Technology (AIST) in Japan. He received Ph.D. and M.Sc in Computer Science from Tokyo Institute of Technology and Wuhan University, respectively. His research interests are in the fields of machine learning and data mining, especially graph learning, web recommendation, anomaly detection, and 3D point clouds processing.

Weimin Wang received the B.S. degree from Shanghai Jiao Tong University in 2009, the M.S. degree from Osaka University in 2012, and the Ph.D. degree from Nagoya University in 2017. He is currently a guest researcher at the Artificial Intelligence Research Center (AIRC) of the National Institute of Advanced Industrial Science and Technology (AIST). His research interests include computer vision, multi-modal data fusion, and scene understanding.

Kyoung-Sook Kim is a senior researcher at the Artificial Intelligence Research Center and a research planning officer at the Department of Artificial Intelligence of Information Technology and Human Factors, the National Institute of Advanced Industrial Science and Technology (AIST) in Japan. She worked at National Institute of Information and Communications Technology (NICT) in Japan from Nov. 2007 to Mar. 2014. She received my B.S., M.S., and Ph.D. degrees in Computer Science from Pusan National University in Korea in 1998, 2001, and 2007, respectively. Her research interests are in big data analysis, spatiotemporal data platforms, location-based services, etc. She currently contributes to international standardization related to big data and AI technologies, by serving as a co-chair of Moving Features SWG and GeoAI DWG of the Open Geospatial Consortium (OGC) and an expert of ISO/IEC/JTC1 SC42 (AI) WG2 (Data), IEC SyC Smart City WG3 (Reference Architecture), and ISO/TC204 (ITS) WG3(ITS database technology).

Takayuki Shinohara graduated from the University of Tsukuba, majoring in Earth Sciences (2013), and the Tokyo Institute of Technology, majoring in Architecture (2015), earning a Master of Engineering degree in Yokohama, Japan. In 2015, he joined Pasco Corporation as a software engineer in the R&D department. In 2019, he enrolled in the doctoral program at the Tokyo Institute of Technology, specializing in Architecture (Yokohama, Japan), and obtained a Doctor of Engineering degree in 2022. His research interests include airborne LIDAR data processing, image segmentation, and classification.

Qiong Chang is currently an assistant professor with the School of Computing, Tokyo Institute of Technology. He received a Ph.D. in intelligent interaction technologies in 2019 from the University of Tsukuba, Japan. His main research interests are GPU computing, both fundamental primitives and applications, especially in the fields of large-scale data processing and computer vision, and real-time implementation on embedded processors. He is a member of IEEE, ACM and IPSJ, and the Associate Editor of *Journal of Real-Time Image Processing*.

Masashi Matsuoka is a Professor of Department of Architecture and Building Engineering at Tokyo Institute of Technology (Tokyo Tech), Japan. He received his BS degree in architectural engineering from Muroran Institute of Technology in 1990. In 1992 and 1996, respectively, he received his MS and Ph.D. degrees, both from the Tokyo Tech. He was with Tokyo Tech, Remote Sensing Technology Center of Japan (RESTEC), Earthquake Disaster Mitigation Research Center, RIKEN and National Research Institute for Earth Science and Disaster Resilience (NIED) in Kobe, and National Institute of Advanced Industrial Science and Technology (AIST) in Tsukuba prior to returning to Tokyo Tech. He is a member of Earthquake Engineering Research Institute (EERI) and Architectural Institute of Japan (AIJ). His current research interests include geoinformatics such as remote sensing and GIS for disaster management and monitoring to the built environment.