Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag



CloudWalker: Random walks for 3D point cloud shape analysis

Adi Mesika^{a,*}, Yizhak Ben-Shabat^{a,b,1}, Ayellet Tal^{a,2} ^aTechnion - Israel Institute of Technology, Haifa, Israel ^bAustralian National University, Canberra, Austrlaia

ARTICLE INFO

Article history:

ABSTRACT

Point clouds are gaining prominence as a method for representing 3D shapes, but their irregular structure poses a challenge for deep learning methods. In this paper we propose CloudWalker, a novel method for learning 3D shapes using random walks. Previous works attempt to adapt Convolutional Neural Networks (CNNs) or impose a grid or mesh structure to 3D point clouds. This work presents a different approach for representing and learning the shape from a given point set. The key idea is to impose structure on the point set by multiple random walks through the cloud for exploring different regions of the 3D object. Then we learn a per-point and per-walk representation and aggregate multiple walk predictions at inference. Our approach achieves state-ofthe-art results for two 3D shape analysis tasks: classification and retrieval.

© 2023 Elsevier B.V. All rights reserved.

Keywords: point cloud, neural networks, 33d shape classification, 3d shape retrieval **1. Introduction**Three dimensional data acquire ometric information. Various for Three dimensional data acquired by sensors provides rich gedata, including meshes, volumetric grids, multi-view images, implicit functions and point clouds [1, 2, 3]. Our focus is on point cloud representation, which is a sampling of a continuous surface. It is often the rawest form of data obtained by modern 3D scanners.

Point clouds manage to preserve the original geometric information in 3D space without the need of discretization. However, there are challenges associated with this representation, particularly in the context of deep learning approaches. These include the lack of connectivity information, structure and order, the varied number of points, and corruptions such as noise, incompleteness, and density variations. An additional difficulty is the lack of large-scale datasets.

Several approaches have been recently proposed for point cloud deep learning. Point-wise *Multilayer perceptron (MLP)* methods explore global/local structures to enhance feature learning and overcome the permutation variance (order) using a symmetric function [4, 5, 6]. Others propose convolutional based methods, which can be divided into continuous [7, 8, 9, 10, 11] and discrete [12, 13]. These methods overcome some of the structural challenges using local connectivity information. Some works consider each point as a vertex of the graph, by adding edges that connect points, and apply graph networks [14], while others construct a hierarchical data structures (e.g., Octree and KD-tree) [15, 16]. While all of the above works handle the challenges of lack of structure and order, none has shown to be as effective as CNNs are for images. Therefore, the challenge of learning a representation for point clouds remains an active field of research.

Our work, illustrated in Figure 1, takes a different avenue. We propose to represent a point cloud by random walks through the cloud, which "explore" the shape. Intuitively, such a walk moves from point to point, through the cloud's valleys and ridges, and on its way discovers the geometry of the underlying surface. A walk is essentially a sequence of points. In a way, this sequence imposes some structure, which is essential for deep learning. A point cloud is represented by multiple walks, which partially cover the underlying surface. Due to

^{*}Corresponding author: E-mail: sadim@campus.technion.ac.il;

¹E-mail: sitzikbs@technion ac il

²E-mail: ayellet@ee.technion.ac.il



Fig. 1: CloudWalker's pipeline. A set of random walks is generated for a given point cloud. Each walk explores the object (walk points are color-coded according to the walk sequence, from blue to yellow). We use CloudWalker for two tasks: 3D object retrieval and classification and achieve SOTA results.

their randomness, some walks provide a more global view of the object, whereas others are local and focus on specific parts.

Recently, random walks were used for a variety of 3D model representations. MeshWalker [17] proposed to use random walks on meshes, however, random walks on point clouds is significantly different and poses several challenges, addressed in our work. First, point clouds lack the connectivity information that is inherently available in meshes, therefore there is no notion of local environment for each point. We solve this issue using a KD-tree data structure that uses spatial proximity to propose local points as candidates for walk generation. Moreover, since a mesh vertex differs from a point (due to connectivity information), our per-point representation is somewhat different from MeshWalker's vertex representation, by adding parameters, layers and additional possible neighbour connections. Lastly, the walks for each shape are aggregated differently than the walks for each mesh, as we are less certain regarding consecutive points than in the case of vertices. The fact that our local environment is not unequivocal has an impact on performance, which we demonstrate in the ablation study.

Xiagn *et al.* [18] proposed to use guided walks on point clouds. They generated curves based on a given set of rules and heuristics and then learn the shape representation, for that purpose, they get the entire point set as input. Our approach is significantly different and simpler since we only get part of the data (walk per shape) and utilize the power of randomness. In spite of the simplicity of our approach, we show that we can produce good results.

We show how to overcome major challenges in deep learning based point cloud processing: First, we address the generation of walks despite the lack of connectivity information and show how to do that efficiently. Furthermore, unlike some previous algorithms [4, 5, 8], our approach does not require normal vectors as inputs to improve performance. This is an important property since normal estimation is particularly inaccurate for the available datasets that contain synthetic CAD models or noisy real-world point clouds without normal [19, 20, 21]. Moreover, we learn a per walk and per point representation and propose an effective cross-walk aggregation module to produce a global shape representation.

Our approach is general in the sense that it can be used for various shape analysis tasks. We demonstrate its performance on two tasks: 3D object classification and 3D object retrieval. Our results are compared against the reported SOTA results for commonly-used datasets *3D-Future* [22], *ScanObjectNN* [23], and *ModelNet40* [24].

This paper makes two main contributions:

- We propose a novel representation for point clouds using random walks. The walks impose order and overcome the lack of connectivity and structure.
- We present an end-to-end learning framework that realizes this representation. We show that it works well even when the dataset has few unique objects and lacks normals. We achieves SOTA results for 3D shape classification and retrieval.

2. Related work

A variety of representations of 3D point cloud have been proposed in the context of deep learning. The main challenge is how to re-organize the set description such that it could be processed within deep learning frameworks. Hereafter we briefly review the main representations; see [1] for a recent thorough survey.

Generally, 3D data for shape classification methods is represented using one of the following representations: (a) Multiview, (b) Volumetric grid, (c) Mesh, (d) Point clouds. Each representation requires a different approach for modifying the data to the form required by deep learning methods.

Multi-view approaches. Multi-view-based methods project the 3D data into 2D images and extract view-wise features, and then fuse these features for accurate shape classification. MVCNN [25] is a pioneering work, which essentially maxpools multi-view features into a global descriptor. Some information is lost in the projection process, but using multiple projections partially compensates. In addition, View-GCN [26] represented multiple views of a 3D shape by a view-graph. The view-graph representation enables to design Graph Convolutional Neural Network (GCN) to aggregate multi-view features by investigating relations of views.

Volumetric grid approaches. An alternative approach to transforming irregular point clouds to regular representations is 3D voxelization [27, 28, 24], followed by 3D Convolution Neural Network (CNN). If applied naively, this strategy can incur massive computation and memory costs due to the cubic growth of the number of voxels as a function of resolution. Since most

voxels are empty, the solution is to take advantage of sparsity. For example, OctNet [29] uses unbalanced octrees with hierarchical partitions. An approach based on sparse convolutions that only evaluate occupied voxels would significantly reduce computation and memory requirements.

Triangle mesh approaches. A mesh is represented as a set of vertices, edges and faces. To handle meshes directly, novel convolutions vertex neighborhoods have been defined [30, 31]. Other works parameterize the mesh in 2D [32, 33, 34]. In MeshCNN [35], a unique idea of using the edges of the mesh to perform pooling and convolution, is introduced. Our work is inspired by MeshWalker, which takes a different approach of representing a mesh by a set of random walks over its vertices, along the edges [17].

Point Cloud approaches. The point cloud representation is challenging because it is both unstructured and point-wise unordered. To overcome these challenges, PointNet [4, 5] utilizes permutation-invariant operators such as pointwise MLPs and pooling layers to aggregate features across a set. A number of approaches connect the point set into a graph and conduct message passing on this graph. DGCNN [36] performs graph convolutions on kNN graphs. KCNet [37] utilizes kernel correlation and graph pooling. DeepGCNs [38] explore the advantages of depth in graph convolutional networks for 3D scene understanding. GDANet [39] use an attention network and introduces a Geometry-Disentangle Module to dynamically disentangle point clouds into the contour and flat part of 3D objects. A voxel-based point cloud transformer, PVTNet [40], introduced a Sparse Window Attention module to alleviate the problem of expensive computations In addition, many other schemes have also been proposed. PCEDNet [41] focused on classification of edges in point clouds. They proposed a new parameterization for each point in a set containing millions of points. RCNet [42] utilizes standard RNN to construct a permutation-invariant network for 3D point cloud processing. Point2Sequences [43] is another RNN-based model that captures correlations between different areas in local regions of point clouds. Here we proposed the CloudWalker which is directly linked to the point cloud representation and can be used as input for a neural network.

Traditional random walks methods. Several methods utilized random walks for shape analysis in both the computer vision and graphics communities. Prior to the deep learning era they were frequently used for image clustering and segmentation [44, 45, 46, 47, 48, 49]. They were also used for generating super-pixels [50] and for visual tracking [51]. For a comprehensive survey of random walks applications see [52].

3. CloudWalker

Given a point cloud, our goal is to learn a representation that will capture both global and local geometric information about the underlying shape, for analysis tasks. Previous works showed that learning 3D object's representation is improved by combining both local and global features [5]. As discussed above, point clouds are challenging for deep learning approaches, due to the lack of spatial structure, which prohibits the use of spatial filters on a 3D lattice. Furthermore, point clouds come in various sizes.

We propose to address these challenges by using multiple random walks through the clouds. Each walk wanders through the object and "explores" it on the way. Jointly, multiple walks provide both local shape information, as well as global shape information; the very same region can be visited multiple times from different directions. This satisfies the above desired properties since the walk length is independent w.r.t. the number of points in the cloud (up to full coverage) and the points have a sequential order, that covers both local and global information.

Our proposed model consists of three modules, as illustrated in Figure 2: The *Random Walk Generator Module* generates multiple random walks per point set. The *CloudWalker Core* module extracts a single walk representation. It consists of two sub-modules: the *Point Representation Sub-module (PRM)* learns a shared, per-point representation, and the *Walk Representation sub-module (WRM)* combines the information along the walk and learns the relations between the points. Finally, the *Multi-walk Aggregation Module* produces the final shape representation, by aggregating multiple walk representations. This representation can then be used for shape analysis tasks. Hereafter we elaborate on each of these modules.

Random walk generator. The concept of randomness is very powerful [53, 54, 55]. This module generates multiple random walks for a given point set. In our setup, thanks to randomness some walks explore a single region of the point set and learn its fine details alongside some broader context, while other walks cover more regions and provide the model with additional information, not seen by other walks.

A walk is a sequence of points, where each point is associated with basic local geometric information. Should the underlying surface of the point cloud existed, it would be possible to utilize the adjacency information to determine, for each point, the next point on the walk. However, since connectivity information is unavailable, we propose a fast and simple approach to generate random walks.

For an input point set S_i we generate the walk W_{ij} of length l as follows. First, the walk's origin point $p_o \in S_i$ is randomly selected from the set. Then, points are iteratively added to the walk by selecting a random point from the set of k nearest neighbors of the last point in the sequence (excluding neighbors that were added at an earlier stage). In the rare case where all k nearest neighbors were already added to the walk, a new random un-visited point is chosen and the walk generation proceeds as before. Choosing the closest neighbor might seem like a straight-forward solution, however, this imposes a strong constraint on the generation process and reduces the randomness and the ability to visit sparser regions.

Since a point cloud may contain many points, an efficient nearest neighbor search is required to get a sense of the environment and propose point candidates to generate the walk. In contrast to meshes, where this information is inherently available, point clouds require this critical step to generate the walks. We propose to construct a KDtree [56] as a preprocessing step. Briefly, a KDtree is a hierarchical space partitioning data structure, in which every node represents an axis-aligned



Fig. 2: CloudWalker's pipline. For an input point set S_i , the *Random Walks Generator* generates several walks W_{ij} . Our model learns a representation for each walk by learning a shared, per-point representation, which are then aggregated along the walk by the *Walk Representation* sub-module. This sub-module consists of three RNN layers and produces a walk representation r_{ij} . The *Multi-Walk Aggregation* module is used at inference time to combine several walks into either the final class prediction for classification or an average representation for retrieval.

hyper-rectangle and contains the set of points in it. This is a very efficient data structure for nearest neighbor queries. Using the nearest neighbors to construct the walk makes it possible to move between underlying connected components along the walk, which works to our advantage since it allows our model to learn such crossings.

CloudWalker Core. Given an input walk, CloudWalker core is the main learned part of our pipeline. It is designed to learn an informative representation for a single walk. it consists of two sub-modules, *Point Representation* sub-module and *Walk Representation* sub-module. Details for each sub-module are provided below.

Point representation sub-module. Given an input walk W_{ij} (the *j*th walk for point cloud S_i), each point p_t along the walk is embedded onto a high-dimensional feature space $\phi_{PRM}(p_t)$. We convert each point to a 512-dimension vector using a three layered Multi-Layer Perceptron (MLP) that shares weights across points. This module is different from MeshWalker's vertex representation module by an additional MLP layer. The additional parameters give the network more capacity to compensate for the lack of connectivity information and learn a distinctive perpoint representation. Then, we perform instance normalization and ReLU nonlinear activation. We preserve the original point context by concatenating the input coordinate to the per-point representation and then feeding it into the next sub-module (WRM); see Figure 2. The module's output is a sequence of vectors per walk, $\Phi_{PRM}(W_{ij})$.

Walk representation sub-module. The input to this stage is a sequence of high-dimensional feature vectors, a single vector for each point in the walk. The output of this module is the walk representation $r_{ij} = \Phi_{WRM}(\Phi_{PRM}(W_{ij}))$ in the form of a single

feature vector for the walk.

Similarly to [17], we use three *Gated Recurrent Units* (*GRU*) [57]. Briefly, a GRU is able to "remember" and accumulate knowledge in a sequence. In our setup, each GRU layer receives all the embedded features from all points of the walk until the current point. It outputs the last hidden state, which captures an abstract representation of the input sequence. In a sense, this is the descriptor that describes the entire walk. This architecture allows walk information to persist. It is especially powerful since the local geometric detail as well as global geometric context are critical for describing the shape.

Finally, the walk representation r_{ij} is fed into a fully connected (FC) classifier to produce the prediction vector $\tilde{r_{ij}}$. The number of elements in $\tilde{r_{ij}}$ is equal to the number of classes.

Multi-walk aggregation module. This module gets as input multiple walk predictions \tilde{r}_{ij} , $j \in [1, m]$ of a single shape S_i . To compute the probability vector we apply a softmax: $\mathcal{P}_{ij} = softmax(\tilde{r}_{ij})$.

Then, we apply a symmetric aggregation function over all of the input walks. For the symmetric function we chose the majority vote: Thus, for a given S_i , the output prediction y_i is:

$$y_i = Majority(C_{i1}, \dots, C_{ij}, \dots, C_{im}), \tag{1}$$

where $C_{ij} = argmax(\mathcal{P}_{ij})$ is the class prediction.

We have experimented with alternative options for symmetric function, including *max* and *mean*, and found the majority vote to work best. This is attributed to its robustness to cases where the network is overconfident in a small subset of walks. A quantitative ablation study is presented in Section 5.

Implementation details. At inference, we use m = 48 walks per object S_i , the length of the walk is l = 800, and k = 20 near-

est neighbors. In training, we use Adam optimizer with cyclic learning rate; the initial and the maximum learning rates are set to 10^{-6} and $5 \cdot 10^{-4}$ respectively. The cycle size is 20*K* iterations and we train for a total of 100K iterations. We record the number of parameters and time using TensorFlow2.4 on a single NVIDIA GeForce RTX 3090. The total number of network parameters is 13.5M and training process takes 100 epochs that last 97 seconds each on average (on ModelNet40 [24]). The average inference time is 0.22 second per shape. The distances between predictions and ground truth labels were minimized through cross entropy loss. Note that sometimes (*e.g.* ModelNet40) the shape scale is available and may be informative. In this case we add the shape's bounding box diagonal length to the feature vector and feed it to the classifier for a minor improvement. Code will be released upon acceptance.

4. Applications

We evaluate the performance of our model on two main tasks: 3D shape classification and 3D object retrieval.

4.1. 3D shape Classification

Given a 3D point set, the goal is to classify it into one of predefined classes. Our model outputs a classification prediction probability vector for each walk. At inference, a majority prediction of the walks' prediction vectors is calculated to produce the final result.

Evaluation measures. For each dataset we report both *in-stance accuracy (IA)* and *class accuracy (CA)*. Instance accuracy is defined as the percentage of the correctly-classified objects,

$$IA = \frac{\#correctly_classified}{\#test_instances}$$
(2)

Class accuracy is the mean of per-class instance accuracy,

$$CA = \frac{1}{C} \sum_{i \in C} IA(i), \tag{3}$$

where C is the number of classes.

Datasets. We present our results on three commonly-used datasets, where each contains different types of objects, different numbers of classes, a different number of objects per class, and different challenges. Figure 3 presents some examples from the datasets.

3D-FUTURE [22] is a recent dataset that contains 9, 992 industrial CAD models of furniture . It consists of 7 super-categories, with 1-12 sub-categories each, for a total of 34 categories. The train/test split is 6, 699/3, 293. This dataset is challenging both due to the objects it contains and its hierarchical structure, as objects in the same category but different sub-categories may resemble each other, requiring fine-grained information in the shape representation.

ScanObjectNN [23] is a real-world dataset that contains 2902 unique object scans from 15 classes. The objects are corrupted in various ways to create a set of $\sim 15k$ objects that are used for training and testing. We follow the official data split provided in [23]. This dataset is highly challenging due to the data corruptions found in scans which include: background points, object



(a) Examples from 3D-FUTURE dataset [22]



(b) Examples from ScanObjectNN dataset [23]



(c) Examples from ModelNet40 dataset [24]

Fig. 3: **Datasets snapshot.** Datasets vary in density, noise and available classes. Presented here are typical objects from a 3D-FUTURE which contains uniformly sampled industrial CAD models (chair and lamp classes), b ScanObjectNN which contains real-world objects, with their background in cyan (chair and table classes), and c ModelNet40 which contains CAD models (guitar and airplane classes).

partiality, and different transformation variants (translation, rotation, and scale). We performed our experiment on its most challenging variant, *PB T50 RS*, with and without background points.

ModelNet40 [24] contains CAD models. There are 12, 311 meshes divided into 40 classes, with 9, 843/2, 468 shapes for training/testing. Corresponding point clouds are generated by uniformly sampling, following the sampling protocol of [5].

Quantitative results. The results for 3D-FUTURE [22] are presented in Table 1. It shows that our method outperforms not only previous point-based, but also multi-view methods, reported for this dataset. We followed the same training protocol for point-based methods as specified by [22] and trained our model with 1K points per object. In particular, we uniformly sampled 1024 points for each shape. Our results outperform

Table 1: **Classification results on** *3D-FUTURE* **dataset.** Our results outperform those reported in [22] using both multi-view and point clouds as input. We note PointNet++ uses normals while we only use points. (* Our reproduction, the orginal paper did not report on this dataset)

Method	Input	IA (%)	CA (%)
MVCNN [25]	Multi-views	69.2	65.4
CurveNet* [18]	Point cloud	69.6	65.1
PointNet++ [5]	Point cloud	<u>69.9</u>	<u>66.0</u>
Ours	Point cloud	70.6	67.8



Fig. 4: ModelNet40 [24] difficult examples. Objects from different classes have similar geometric properties.

the others in both IA and CA evaluation measures.

Table 2 presents quantitative result for classification on 3D-FUTURE for each category. It shows that our method outperforms both, point-based and multi-view methods reported for this dataset. Our model leads on average and in 18/32 categories. It shows how challenging this dataset is because of its similar sub-categories. We note that while [5] uses normals as additional information, we do not. The mere vertices locations suffice for our method. This is especially important in cases where normals are unavailable, as normals are not outputs of most scanners and normal estimation is possible, but is not a trivial problem [19, 20, 21]. After all, normal estimation is "almost" surface reconstruction, and the reconstruction of a surface is exactly what we wish to avoid.

For ScanObjectNN, Table 3 presents the results with and without background points of real-world scans. It shows that the overall accuracy is better compared to state of the art methods [23]. When trained and tested without the background points, our average class accuracy 78.5% is higher than all reported results while maintaining competitiveness on the overall accuracy.

Table 4 reports the results on the ModelNet40 [24] dataset. For this dataset we achieved competitive results. Model-Net40 is known to be a difficult and saturated dataset, partially because of cross-labeled classes (desk/table, plant/flowerpot/vase, night-stand/dresser), as illustrated in Figure 4.

We report that when we choose the second best prediction for the mis-classified objects in Modelnet40 dataset, 97% of the test instances are correctly classified. Similarly, for the 3DFuture dataset we get 89% accuracy (~19% improvement). This findTable 2: Classification results for each category of *3D-FUTURE*. We note PointNet++ (PN++) uses normals and Multi-scale grouping (MSG) and we do not, while we use only the points. MVCNN uses 12 view.

Category	Ours (%)	PN++ (%)	MVCNN (%)
Children Cabinet	16.0	32.1	72.0
Nightstand	83.5	71.8	75.0
Bookcase	58.7	52.3	66.7
Wardrobe	85.0	82.0	56.7
Coffee Table	83.7	82.6	67.9
Corner/Side Table	73.0	74.7	64.5
Side Cabinet	69.5	<u>65.2</u>	47.9
Wine Cabinet	44.3	67.1	62.9
TV Stand	80.5	73.6	73.5
Drawer Chest	44.8	55.2	67.5
Shelf	48.4	48.4	51.9
Round End Table	93.7	<u>75.0</u>	52.2
King-size Bed	89.6	91.2	78.6
Bunk Bed	88.8	77.8	57.1
Bed Frame	95.4	<u>93.8</u>	<u>93.8</u>
Single bed	68.9	68.9	<u>67.7</u>
Kids Bed	00.0	14.3	<u>12.5</u>
Dining Chair	76.8	<u>63.9</u>	50.5
Lounge/Office Chair	54.0	60.5	<u>60.3</u>
Classic Chinese Chair	62.5	62.5	<u>57.1</u>
Barstool	88.8	<u>66.7</u>	32.0
Dressing Table	77.3	68.2	<u>73.7</u>
Dining Table	76.0	61.1	84.8
Desk	<u>30.9</u>	20.4	54.0
Three-seat Sofa	83.1	<u>82.6</u>	71.7
armchair	74.9	68.0	<u>72.5</u>
Loveseat Sofa	45.5	64.5	<u>62.9</u>
L-shaped Sofa	92.9	<u>85.9</u>	83.3
Lazy Sofa	42.8	<u>50.0</u>	66.7
Stool	70.9	<u>75.8</u>	91.9
Pendant Lamp	92.8	<u>90.9</u>	89.8
Ceiling Lamp	72.0	<u>70.7</u>	63.0
Average Class Accuracy	67.8	66.0	65.2
Instance Accuracy	70.6	69.9	69.2

ing further shows that some categories have fine-grained differences that are difficult to discriminate.

4.2. Retrieval

Given a query object, the goal is to retrieve objects from the dataset, ordered by their relevancy (similarity) to the query. There are various ways to evaluate shape retrieval. The most common evaluation is the *mean Average Precision (mAP)* over test queries:

$$mAP = \frac{1}{Q} \sum_{i=1}^{Q} AP(S_i), \qquad (4)$$

where *Q* is the number of queries in the set. *AP* is defined as $AP = \frac{1}{GTP} \sum_{k}^{N} P@k \times rel@k$, where *GTP* is the number of ground truth positives, *N* is the size of the ordered set, *P@k* refers to the precision at *k* and *rel@k* is an indicator function which equals 1 if the object at rank k is from the same class as

Table 3: **Classification results on** *ScanObjectNN* [23]. We use the difficult variant of the dataset, PB T50 RS, with (w/) and without (w/o) background (BG) points. We achieve SOTA results on the variant without the background; (*CA* results were not reported for objects without background). we achieve CA SOTA results on the variant with the background and competitive IA results. The second best is underlined. (*concurrent work)

Methods	w/o B	G (%)	w/ BG (%)	
Wiethous	IA	CA	IA	CA
PointNet [4]	74.4	-	68.2	63.4
PointNet++ [5]	80.2	-	77.9	75.4
3DmFV [58]	69.8	-	63.0	58.1
SpiderCNN [8]	76.9	-	73.7	69.8
SimpleView [59]	-	-	79.5	-
DGCNN [36]	<u>81.5</u>	-	78.1	73.6
PointCNN [12]	80.8	-	78.5	75.1
DRNet [60]	-	-	80.3	78.0
GBNet [61]	-	-	80.5	77.8
DeltaConv* [62]	-	-	84.7	-
BGA-PN++ [23]	-	-	80.2	77.5
BGA-DGCNN [23]	-	-	79.7	75.7
Ours	82.2	79.5	80.3	78.5

Table 4: Classification results on *ModelNet40* [24] For this dataset we show competitive results.

Method	IA (%)	CA (%)
PointNet [4]	89.2	86
PointNet++ [5]	90.7	-
PointCNN [12]	92.2	88.1
SpiderCNN [8]	92.4	-
KPConv [7]	92.9	-
SimpleView [59]	93.6	90.5
PAConv [63]	93.6	-
PointTrans. [64]	93.7	<u>90.6</u>
CurveNet [18]	93.8	-
GBNet [61]	93.8	91
PVT [40]	94.0	-
RPNet [65]	94.1	-
Ours	93.1	90.1

query S_i and 0 otherwise. Note that significantly fewer methods report their performance on this task than for classification.

We present our results on commonly used datasets for shape retrieval task, ModelNet40 and ModelNet10 [24]. We use the most common train/test splits: 9,843/2,468 for ModelNet40 and 3,991/908 for ModelNet10. We used the output feature vector from our classification network by averaging all walks probability vectors for each object and get a global feature vector \overline{P}_i ; see Figure 2.

Table 5 shows that our method achieves SOTA results on both ModelNet40 and ModelNet10 [24]. It outperforms approaches that use a variety of 3D representations, including meshes, multi-views, and point clouds.

Figure 5 presents qualitative results for object retrieval on ModelNet40 [24] dataset. It shows the top five retrieved objects by our model. As can be seen, the retrieved objects are indeed similar and thus belong to the same class. In the bottom row

Table 5: **3D Shape Retrieval results.** This table shows the *mAP* on Model-Net40 & ModelNet10 [24], sorted by the input type (Point Cloud, Multi-view, and mesh). Our CloudWalker outperforms other methods.

Method	input	ModelNet40	ModelNet10
GWCNN [33]	Mesh	59.0	74.0
MeshNet [30]	Mesh	81.9	-
MVCNN [25]	MV	79.5	-
SeqViews [66]	MV	89.1	-
PointNet [4]	PC	70.5	-
PointCNN [12]	PC	83.8	-
DGCNN [36]	PC	85.3	-
DensePoint [11]	PC	88.5	93.2
Ours	PC	92.9	93.7



Fig. 5: **Qualitative results for Object Retrieval.** The query test objects are on the left, where the top five retrieved objects are on the right. The falsely retrieved object is highlighted.

we show a particularly interesting case where the query object belongs to the *plants* category, but the highlighted model belongs to a different category and thus considered to be an error. However, this is in fact the exact same object with different set of points. This example demonstrates some of the challenges in the dataset.

5. Ablation study

Walk length. In this experiment we explore the effects of walk length, *i.e.* the number of points in each walk, on the performance of classification. Table 6 shows that a longer walk improves performance up to a certain length. This can be explained by the fact that longer walks are likely to have better coverage of the object, and thus better distinguish between different categories. However, when the walk is too long, the GRUs' ability to "remember" points in the sequences is reduced, which negatively impacts performance. For example, on ScanObjectNN [23], we get improvement when using up to 40% of the total number of vertices per walk. This quantity is dataset dependant, however we found 40% to perform well across multiple datasets.

Aggregation method. In our model, we aggregate predictions from multiple walks on the shape. Table 7 compares three symmetric aggregation functions. While the results are similar, there is a slight advantage to majority voting.

Table 6: Walk length ablation. Performance improves as length increases, up to a certain length. Tested on ScanObjectNN [23].

Walk length	IA (%)	CA (%)
0.1V	75.3	72.1
0.2V	78.0	75.2
0.4V	82.2	79.5
0.5V	79.3	76.4

Table 7: **Aggregation method ablation.** Majority voting is slightly better than other symmetric functions. Tested on ModelNet40 [24].

Aggregation method	IA (%)	CA (%)
Average	92.8	89.9
Max	92.4	89.2
Majority	93.1	90.1

Number of walks. We study how the number of walks (used at inference) influences classification accuracy. Table 8 shows that as the number of walks increases, the accuracy improves up to 48 walks where it saturates. Note that above 48 walks most points in the point cloud are visited at least once (~full coverage). Note that, even very few walks result in very good accuracy. This means that the number of walks may be a tune-able parameter that balances the trade-off between computational power and accuracy. In our method, we chose to use 48 walks.

Random walk generation method. Generating random walks can be performed in various ways. In the following we explore several options, in all of which the first point is randomly chosen and the next points are sequentially added from the last point's neighbors. The options are: (1) Random-randomly choosing among the point's unvisited neighbors with a uniform distribution. (2) High variance-calculating the change in variance for each neighbor of the most recently added point and selecting the one that increases the variance the most. (3)Combined—combining both random and high variance methods by choosing a neighbor who increases the variance 30% of the times and randomly otherwise. Note that the heuristic based on the walk variance is effective, but it is not sufficient since it reduces randomness. The results are presented in Table 9. We found that the random strategy is the best, reinforcing our claim of the power of randomness.

Stability. Cloudwalker relies on the power of randomness. To evaluate the consistency across different random seeds (stability), we evaluate our model 100 times, each time, we generate a

Table 8: **Number of walks ablation.** The accuracy improves with the number of walks per shape. Tested on 3D-FUTURE [22].

Number of walks	Coverage (%)	IA (%)	CA (%)
4	76	68.6	65.5
16	82	69.7	67.1
32	92	69.8	67.5
48	99	70.6	<u>67.8</u>
64	99	<u>70.5</u>	67.5
96	100	70.3	67.9

Table 9: Walk generation ablation. Random walk generation performs better than hand-crafted heuristics. Tested on ModelNet40 [24].

Generation	n method	IA (%)	CA (%)
Rand	Random High variance		90.1
High va			86.5
Comb	ined	91.7	87.9
	1013	0PA	
VA	E.		
	CALL.	Ĩ,	
	A F	a se	Ats State
	1 second	0	
2	Z		
	Ļ	1	
l l	LIS		
	X Anton	*	

Fig. 6: Walk comparison. CloudWalker and CurveNet walks on three different shapes. Top row: A single CloudWalker walk of length 500 points (color-coded according to the walk sequence, from blue to yellow), bottom row: A single CurveNet walk, composed of 100 curves of length 5 points (different color for each curve).

different set of random walks. We found that the average overall accuracy for ModelNet40 dataset is 92.8%, with a standard deviation of 0.1% and a maximum value of 93.1%.

Comparison of random and guided walks. Figure 6 compares our random walks to the guided walks generated by CurveNet [18] on the same object. It is evident that each method explores different parts of the model. Interestingly, although the walks are distinctly different, both methods manage to successfully classify the shape. Note that we restrict our walks from revisiting vertices, which is different from CurveNet, where there is no such hard restriction and curves may overlap in accordance with a crossover suppression strategy.

Limitations. Figure 7 shows failure cases in the retrieval task. In (a), the 5^{th} retrieved object is a vase, rather than a bottle. This vase is geometrically similar to a bottle, however the top differs. Most of the walks provide the geometry of a vase. There is a lower probability of capturing this particular geometry which distinguishes it from a bottle due to the few points in that region. In (b) we present five of the top ten objects retrieved for the leftmost nightstand. A dresser is placed at the 10th spot. These objects belong to different classes but have similar geometry.

Furthermore, our random walk generator process might be negatively impacted by strong anisotropic sampling, as occurs in 360 lidar acquisition. We have not explored this challenge and it will be an interesting direction for future work.

6. Conclusions

This paper introduces a novel approach for representing point clouds suited for deep learning architectures. The key idea is to represent the point cloud using multiple random walks on the shape. The randomness of the walks can be viewed as a form



(b) Geometrically similar objects are in different categories.

Fig. 7: Limitation. When objects are geometrically similar, yet belong to different classes, our method might err.

of data augmentation that does not require explicit manipulation of the point cloud. We utilize this representation and introduce CloudWalker, an end-to-end shape representation learning pipeline. The approach is general, yet simple, and relies on the power of randomness. We have shown our approach to be very effective for the tasks of classification and retrieval. An interesting avenue for future work may provide confidence measures for model's predictions based on cross-walk correlations. Additionally, random walks in scenes would be another interesting extension of our paper. Furthermore, alternative random walk generation processes and their use in other geometrical tasks is another path to explore.

Acknowledgments. This work was partially supported by the Israel Science Foundation (ISF) 1083/18, ADRI 117632, and the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curiegrant agreement No 893465.

References

- Guo, Y, Wang, H, Hu, Q, Liu, H, Liu, L, Bennamoun, M. Deep learning for 3d point clouds: A survey. IEEE TPAMI 2020;.
- [2] Forsyth, D, Ponce, J. Computer vision: A modern approach. Prentice hall; 2011.
- [3] Marschner, S, Shirley, P. Fundamentals of computer graphics. CRC Press; 2018.
- [4] Qi, CR, Su, H, Mo, K, Guibas, LJ. Pointnet: Deep learning on point sets for 3d classification and segmentation. In: CVPR. 2017, p. 652–660.
- [5] Qi, CR, Yi, L, Su, H, Guibas, LJ. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: Advances in Neural Information Processing Systems (NeurIPS); vol. 30. 2017,.
- [6] Zaheer, M, Kottur, S, Ravanbakhsh, S, Poczos, B, Salakhutdinov, RR, Smola, AJ. Deep sets. In: Advances in Neural Information Processing Systems (NeurIPS). 2017..
- [7] Thomas, H, Qi, CR, Deschaud, JE, Marcotegui, B, Goulette, F, Guibas, LJ. Kpconv: Flexible and deformable convolution for point clouds. In: ICCV. 2019, p. 6411–6420.
- [8] Xu, Y, Fan, T, Xu, M, Zeng, L, Qiao, Y. Spidercnn: Deep learning on point sets with parameterized convolutional filters. In: ECCV. 2018, p. 87–102.
- [9] Liu, Y, Fan, B, Xiang, S, Pan, C. Relation-shape convolutional neural network for point cloud analysis. In: CVPR. 2019, p. 8895–8904.
- [10] Wu, W, Qi, Z, Fuxin, L. Pointconv: Deep convolutional networks on 3d point clouds. In: CVPR. 2019, p. 9621–9630.
- [11] Liu, Y, Fan, B, Meng, G, Lu, J, Xiang, S, Pan, C. Densepoint: Learning densely contextual representation for efficient point cloud processing. In: ICCV. 2019, p. 5239–5248.

- [12] Li, Y, Bu, R, Sun, M, Wu, W, Di, X, Chen, B. PointCNN: Convolution on x-transformed points. Advances in Neural Information Processing Systems (NeurIPS) 2018;31:820–830.
- [13] Rao, Y, Lu, J, Zhou, J. Spherical fractal convolutional neural networks for point cloud recognition. In: CVPR. 2019, p. 452–460.
- [14] Simonovsky, M, Komodakis, N. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In: CVPR. 2017, p. 3693–3702.
- [15] Lei, H, Akhtar, N, Mian, A. Octree guided cnn with spherical kernels for 3d point clouds. In: CVPR. 2019, p. 9631–9640.
- [16] Klokov, R, Lempitsky, V. Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In: ICCV. 2017, p. 863–872.
- [17] Lahav, A, Tal, A. MeshWalker: Deep mesh understanding by random walks. ACM TOG 2020;39:1–13.
- [18] Xiang, T, Zhang, C, Song, Y, Yu, J, Cai, W. Walk in the cloud: Learning curves for point clouds shape analysis. In: ICCV. 2021,.
- [19] Guerrero, P, Kleiman, Y, Ovsjanikov, M, Mitra, NJ. Pcpnet learning local shape properties from raw point clouds. In: Computer Graphics Forum; vol. 37. Wiley Online Library; 2018, p. 75–85.
- [20] Ben-Shabat, Y, Gould, S. Deepfit: 3d surface fitting via neural network weighted least squares. In: ECCV. Springer; 2020, p. 20–34.
- [21] Lenssen, JE, Osendorfer, C, Masci, J. Deep iterative surface normal estimation. In: CVPR. 2020, p. 11247–11256.
- [22] Fu, H, Jia, R, Gao, L, Gong, M, Zhao, B, Maybank, S, et al. 3d-future: 3d furniture shape with texture. International Journal of Computer Vision 2021;129(12):3313–3337.
- [23] Uy, MA, Pham, QH, Hua, BS, Nguyen, DT, Yeung, SK. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In: ICCV. 2019,.
- [24] Wu, Z, Song, S, Khosla, A, Yu, F, Zhang, L, Tang, X, et al. 3d shapenets: A deep representation for volumetric shapes. In: CVPR. 2015, p. 1912–1920.
- [25] Su, H, Maji, S, Kalogerakis, E, Learned-Miller, E. Multi-view convolutional neural networks for 3d shape recognition. In: ICCV. 2015, p. 945–953.
- [26] Wei, X, Yu, R, Sun, J. View-gcn: View-based graph convolutional network for 3d shape analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020, p. 1850–1859.
- [27] Maturana, D, Scherer, S. Voxnet: A 3d convolutional neural network for real-time object recognition. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE; 2015, p. 922–928.
- [28] Qi, CR, Su, H, Nießner, M, Dai, A, Yan, M, Guibas, LJ. Volumetric and multi-view cnns for object classification on 3d data. In: CVPR. 2016, p. 5648–5656.
- [29] Riegler, G, Osman Ulusoy, A, Geiger, A. Octnet: Learning deep 3d representations at high resolutions. In: CVPR. 2017, p. 3577–3586.
- [30] Feng, Y, Feng, Y, You, H, Zhao, X, Gao, Y. Meshnet: Mesh neural network for 3d shape representation. In: AAAI; vol. 33. 2019, p. 8279– 8286.
- [31] Verma, N, Boyer, E, Verbeek, J. Feastnet: Feature-steered graph convolutions for 3d shape analysis. In: CVPR. 2018, p. 2598–2606.
- [32] Boscaini, D, Masci, J, Rodolà, E, Bronstein, M. Learning shape correspondence with anisotropic convolutional neural networks. Advances in Neural Information Processing Systems (NeurIPS) 2016;29.
- [33] Ezuz, D, Solomon, J, Kim, VG, Ben-Chen, M. GWCNN: A metric alignment layer for deep shape analysis. In: Computer Graphics Forum; vol. 36. Wiley Online Library; 2017, p. 49–57.
- [34] Sinha, A, Bai, J, Ramani, K. Deep learning 3d shape surfaces using geometry images. In: ECCV. Springer; 2016, p. 223–240.
- [35] Hanocka, R, Hertz, A, Fish, N, Giryes, R, Fleishman, S, Cohen-Or, D. MeshCNN: a network with an edge. ACM TOG 2019;38:1–12.
- [36] Wang, Y, Sun, Y, Liu, Z, Sarma, SE, Bronstein, MM, Solomon, JM. Dynamic graph cnn for learning on point clouds. Acm Transactions On Graphics (tog) 2019;38:1–12.
- [37] Shen, Y, Feng, C, Yang, Y, Tian, D. Mining point cloud local structures by kernel correlation and graph pooling. In: CVPR. 2018, p. 4548–4557.
- [38] Li, G, Muller, M, Thabet, A, Ghanem, B. Deepgcns: Can gcns go as deep as cnns? In: ICCV. 2019, p. 9267–9276.
- [39] Xu, M, Zhang, J, Zhou, Z, Xu, M, Qi, X, Qiao, Y. Learning geometrydisentangled representation for complementary understanding of 3d object point cloud. arXiv preprint arXiv:201210921 2021;2.
- [40] Zhang, C, Wan, H, Liu, S, Shen, X, Wu, Z. Pvt: Point-voxel transformer for 3d deep learning. arXiv preprint arXiv:210806076 2021;.

- [41] Himeur, CE, Lejemble, T, Pellegrini, T, Paulin, M, Barthe, L, Mellado, N. Pcednet: A lightweight neural network for fast and interactive edge detection in 3d point clouds. ACM Transactions on Graphics (TOG) 2021;41(1):1–21.
- [42] Wu, P, Chen, C, Yi, J, Metaxas, D. Point cloud processing via recurrent set encoding. In: Proceedings of the AAAI Conference on Artificial Intelligence; vol. 33. 2019, p. 5441–5449.
- [43] Liu, X, Han, Z, Liu, YS, Zwicker, M. Point2sequence: Learning the shape representation of 3d point clouds with an attention-based sequence to sequence network. In: Proceedings of the AAAI Conference on Artificial Intelligence; vol. 33. 2019, p. 8778–8785.
- [44] Meilă, M, Shi, J. A random walks view of spectral segmentation. In: International Workshop on Artificial Intelligence and Statistics. PMLR; 2001, p. 203–208.
- [45] Gorelick, L, Galun, M, Sharon, E, Basri, R, Brandt, A. Shape representation and classification using the poisson equation. IEEE Transactions on Pattern Analysis and Machine Intelligence 2006;28(12):1991–2005.
- [46] Grady, L. Random walks for image segmentation. IEEE transactions on pattern analysis and machine intelligence 2006;28(11):1768–1783.
- [47] Grady, L, Funka-Lea, G. Multi-label image segmentation for medical applications based on graph-theoretic electrical potentials. In: Computer Vision and Mathematical Methods in Medical and Biomedical Image Analysis. Springer; 2004, p. 230–245.
- [48] Grady, L. Multilabel random walker image segmentation using prior models. In: 2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05); vol. 1. IEEE; 2005, p. 763–770.
- [49] Lai, YK, Hu, SM, Martin, RR, Rosin, PL. Fast mesh segmentation using random walks. In: Proceedings of the 2008 ACM symposium on Solid and physical modeling. 2008, p. 183–191.
- [50] Shen, J, Du, Y, Wang, W, Li, X. Lazy random walks for superpixel segmentation. IEEE Transactions on Image Processing 2014;23(4):1451– 1462.
- [51] Li, X, Han, Z, Wang, L, Lu, H. Visual tracking via random walks on graph model. IEEE transactions on Cybernetics 2015;46(9):2144–2155.
- [52] Xia, F, Liu, J, Nie, H, Fu, Y, Wan, L, Kong, X. Random walks: A review of algorithms and applications. IEEE Transactions on Emerging Topics in Computational Intelligence 2019;4(2):95–107.
- [53] Motwani, R, Raghavan, P. Randomized algorithms. Cambridge university press; 1995.
- [54] Chazelle, B. The discrepancy method: randomness and complexity. Cambridge University Press; 2001.
- [55] Sawhney, R, Crane, K. Monte carlo geometry processing: A grid-free approach to pde-based methods on volumetric domains. ACM Transactions on Graphics 2020;39(4).
- [56] De Berg, M, Van Kreveld, M, Overmars, M, Schwarzkopf, O. Computational geometry. In: Computational geometry. Springer; 1997, p. 1–17.
- [57] Cho, K, Van Merriënboer, B, Gulcehre, C, Bahdanau, D, Bougares, F, Schwenk, H, et al. Learning phrase representations using RNN encoderdecoder for statistical machine translation. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP) 2014;.
- [58] Ben-Shabat, Y, Lindenbaum, M, Fischer, A. 3DMFV: Threedimensional point cloud classification in real-time using convolutional neural networks 2018;3:3145–3152.
- [59] Goyal, A, Law, H, Liu, B, Newell, A, Deng, J. Revisiting point cloud classification with a simple and effective baseline. In: ICML. 2020,.
- [60] Qiu, S, Anwar, S, Barnes, N. Dense-resolution network for point cloud classification and segmentation. In: Proc. of the IEEE Winter Conference on Applications of Computer Vision (WACV). 2021, p. 3813–3822.
- [61] Qiu, S, Anwar, S, Barnes, N. Geometric back-projection network for point cloud classification. IEEE Transactions on Multimedia 2021;doi:10.1109/TMM.2021.3074240.
- [62] Wiersma, R, Nasikun, A, Eisemann, E, Hildebrandt, K. Deltaconv: Anisotropic point cloud learning with exterior calculus. arXiv preprint arXiv:211108799 2021;.
- [63] Xu, M, Ding, R, Zhao, H, Qi, X. 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, p. 3173–3182.
- [64] Zhao, H, Jiang, L, Jia, J, Torr, PH, Koltun, V. Point transformer. In: ICCV. 2021, p. 16259–16268.
- [65] Ran, H, Zhuo, W, Liu, J, Lu, L. Learning inner-group relations on

point clouds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. 2021, p. 15477–15487.

[66] Han, Z, Shang, M, Liu, Z, Vong, CM, Liu, YS, Zwicker, M, et al. SeqViews2SeqLabels: Learning 3d global features via aggregating sequential views by rnn with attention. IEEE Transactions on Image Processing 2018;28:658–672.