

3D Intracranial Aneurysm Classification and Segmentation via Unsupervised Dual-branch Learning

Di Shao
Deakin University
75 Pigdons Rd, Waurn Ponds, 3216, Australia
shaod@deakin.edu.au

Xuequan Lu* (corresponding author)
Deakin University
75 Pigdons Rd, Waurn Ponds, 3216, Australia
xuequan.lu@deakin.edu.au

Xiao Liu
Deakin University
75 Pigdons Rd, Waurn Ponds, 3216, Australia
xiao.liu@deakin.edu.au

Abstract

Intracranial aneurysms are common nowadays and how to detect them intelligently is of great significance in digital health. While most existing deep learning research focused on medical images in a supervised way, we introduce an unsupervised method for the detection of intracranial aneurysms based on 3D point cloud data. In particular, our method consists of two stages: unsupervised pre-training and downstream tasks. As for the former, the main idea is to pair each point cloud with its jittered counterpart and maximise their correspondence. Then we design a dual-branch contrastive network with an encoder for each branch and a subsequent common projection head. As for the latter, we design simple networks for supervised classification and segmentation training. Experiments on the public dataset (Intra) show that our unsupervised method achieves comparable or even better performance than some state-of-the-art supervised techniques, and it is most prominent in the detection of aneurysmal vessels. Experiments on the ModelNet40 also show that our method achieves the accuracy of 90.79% which outperforms existing state-of-the-art unsupervised models.

Key words: Intracranial Aneurysm Classification, Intracranial Aneurysm Segmentation, 3D Point Cloud, Unsupervised Learning

1. Introduction

Intracranial aneurysms can result in a high rate of mortality, and their classification and segmentation are of great significance. Existing research mainly focused on image data which involve regular pixels [9, 18, 25, 26, 28, 29].

While 3D geometric data such as point clouds can depict more useful information, the research on analysing intracranial aneurysms using point cloud data has been very sparsely exploited to date. Thanks to [38], a point cloud dataset including aneurysmal segments and healthy vessel segments has been published. They have conducted a benchmark using state-of-the-art point-based networks that can directly consume 3D points instead of 2D pixels.

There are many networks available for consuming point cloud data, for example, PointNet [20], PointNet++ [21], SpiderCNN [37], PointCNN [14], SO-Net [13] and DGCNN [32]. PointNet is a seminal method for taking 3D points as input and used for 3D point cloud classification and segmentation. Later on, other point-based methods have been proposed to improve the performance. Since they are all supervised learning methods, annotated data are required for training. However, annotation often requires experts and significant amounts of time, especially for large datasets and medical data.

With the above analysis in mind, we design an unsupervised representation learning method that consumes point clouds of vessel segments. The contrastive learning concept inspires our method. In particular, we first generate a pair of augmented sample of the original point cloud which should have a distinctly difference. Next, we design a dual-branch contrastive network with an encoder for each branch and a follow-up common projection head to facilitate the unsupervised training with a contrastive loss. As for the downstream tasks, we first use the unsupervised trained model to output the representations. Then, we design simple networks and train it by taking the representations as input to classify or segment intracranial aneurysms. Note that we design two unsupervised networks and two correspond-

ing downstream networks to fulfill two different tasks (i.e. classification and segmentation). Supervised methods often need a large scale of labelled data for achieving satisfactory performance. Compared with them, our method does not require labels in unsupervised training, and it can utilise a small scale of labelled data for downstream training. In summary, our contributions in this paper include:

- We propose a simple yet effective method for unsupervised representation learning on 3D point clouds of vessel segments.
- We invent a useful augmentation method for generating pairs of each vessel segment.
- We propose a dual-branch contrastive network with an encoder for each branch.
- We conduct comprehensive experiments and compare with state-of-the-art point-based techniques to demonstrate the superior performance of our method.

2. Related work

2.1. Deep Learning on Intracranial Aneurysms

Intracranial aneurysms are associated with a high mortality rate. Therefore, the detection of intracranial aneurysms is crucial for human health. Traditional methods rely greatly on prior knowledge, which is often inferior to deep learning in terms of capability and accuracy. Due to the excellent performance of deep learning in processing medical images, there are many deep learning methods to detect intracranial aneurysms [24]. [18] proposed a convolutional neural network-based detection system. The system used a 6-layer (CNN) and maximum intensity projection (MIP) algorithm based on the MRA images. This method can achieve almost 100% accuracy for detecting aneurysms greater than 7 mm in diameter. However, it was less sensitive for small vascular aneurysms. To improve this, [28] used the full U-net convolution architecture to predict aneurysm size based on the detection. [29] applied the ResNet-18 network to the MRA images and performed a secondary evaluation on the already detected image data to enhance the detection sensitivity. To better segment the shape of intracranial aneurysms, [26] utilized DeepMedic [10] with 2-pathway architecture and 11-layer convolution to segment intracranial aneurysms from the MRA images on the basis of detection. The above method for detecting intracranial aneurysms used data that are stacked with 2D images. To sum up, nearly all works focused on dealing with medical images rather than 3D geometry like point clouds.

2.2. Point-Based Networks

Neural network models for the classification and segmentation of 3D point cloud data have achieved noticeable successes. [20] proposed PointNet to directly process point sets. To obtain permutation invariance and transformation invariance of point clouds, PointNet used the symmetric function and T-net to design the network. It had good results for global features extraction of point clouds. However, it ignored the geometric relationship among points and limited the extraction of local features. To address this issue, [21] proposed PointNet++ using a hierarchical neural network. It used the point sampling and grouping strategy to extract local features of point clouds. However, PointNet++ did not reveal the spatial distribution of the input point cloud. SO-Net [13] constructed the Self-Organizing Map (SOM) [12] to model the spatial distribution of the input point cloud. It allowed SO-Net to adjust the receptive field overlap and performed hierarchical feature extraction. Unlike SO-Net with adjusting the perceptual field of the hierarchical network, PointCNN [14] proposed the χ -transformation to process the point cloud data so that the point cloud data can be weighted or permuted. Thus, it improved the extraction of local features. In addition, SpiderCNN [37] proposed SpiderConv, i.e. parameterized convolutional Filters, to implement convolutional operations on disordered point clouds. DGCNN [32] proposed a convolutional-like operation by constructing local neighbourhood graphs and applying convolutional operations on the edges. It connected adjacent point pairs to exploit the local geometric structure. In addition to common tasks like classification and segmentation, point based networks have also been developed to address other tasks [31,40]. In summary, there has been great progress in analyzing point cloud data in a supervised manner.

2.3. Unsupervised 3D Point Cloud Learning

All the above deep neural networks can classify and segment the point cloud data well. However, considering the complexity of labelling 3D data, it is difficult to get enough data with expert labelling for supervised training in many scenarios. Therefore, it is meaningful for exploiting unsupervised or self-supervised learning for 3D point cloud data. PointContrast [36] proposed an Unsupervised framework with U-net as the backbone network. And it demonstrated the transferability of representation learning to 3D point cloud data and the performance enhancement of pre-training to downstream tasks. Lu et al. [15,16] attempted to address skeleton learning on point cloud sequence data. Jiang et al. [8] introduced a simple yet effective unsupervised learning method on point cloud that only considers rotation as the transformation. Info3D [22] proposed to extend the InfoMax [30] and contrastive learning principles on 3D shapes. It maximized the mutual information between

3D objects and their “chunks” to improve the representation in the aligned dataset. FoldingNet [39] proposed an auto-encoder with graph pooling and MLP layers using the folding operation to deform 2D grids into object surfaces. However, in 3D medical point cloud data, unsupervised methods are still in great demand. We propose an unsupervised representation learning method, which shows excellent performance for the classification and segmentation of point cloud based intracranial aneurysms.

3. Method

Our method consists of two stages which are unsupervised learning and downstream tasks. In stage 1, we first perform augmentation on each point cloud to get a pair of augmented samples which are different in pose (Section 3.1). We then get two representations of a pair of data in a high-dimensional space by means of the dual-branch encoders, which enables each branch of the encoder to extract distinct features. Next, we map representations to a low-dimensional vector [6] with a projection head to improve network training speed (Section 3.2). Last, we employ a contrastive loss to encourage the representations of the pair of point clouds output by the encoders to be similar in the high-dimensional space (Section 3.3). In stage 2, the trained model is used to output unsupervised representations for the downstream task (Section 3.4). The downstream task will evaluate the effectiveness of unsupervised learning. Figure 1 presents the architecture of our method.

3.1. Data Augmentation

We use data augmentation to generate different samples for each point cloud. To generate a pair of data, we consider using data augmentation methods, including jittered, perturbation, and rotation transformations. After experiments, it was found that *jittered* as data augmentation in both branches gave the best results in the downstream tasks, indicating a more discriminative representation learned by the upstream network. Ablation experiments will be presented in Section 4.4.

We take a batch of point clouds with mini-batch size N and input them into the data augmentation module. As shown in the top of Figure 1, for a sample in the mini-batch, we use the jittered function to obtain a pair of samples: one is the jittered point cloud x_i and the other is the jittered point cloud x_j . In this way, we have a batch size of $2N$ in this mini-batch. We randomly select $\{x_i, x_j\}$ as a positive pair, and the other $N - 1$ pairs, which consist of one of positive sample and one of the other samples, are regarded as negative samples in this mini-batch.

3.2. Dual-branch Encoders and Projection Head

As shown in the left bottom of Figure 2, each pair of samples need to be passed through dual-branch encoders

$f(\cdot)$ to obtain two representations h_i and h_j . Features are respectively extracted from a pair of data using two different encoders. Experimentally, we have also compared two different encoders with a common encoder. Ablation experiments will be presented in Section 4.4.

Two encoders are PointNet [20] and PointNet++ [21], respectively. The reason for choosing PointNet and PointNet++ is that PointNet can extract global features while PointNet++ can extract local features. This design highlights distinctions in features and allows for more distinctive representation.

Classification. The first encoder utilizes three consecutive 1D convolutional layers and a max-pooling layer to obtain the representation vector h (1024-dimensional). The second encoder consists of three abstraction levels. Each layer abstracts and processes the point set to create a new point set with fewer elements. The input to the abstraction layer consists of an $n \times (d + c)$ matrix formed by n points with d -dim coordinates and c -dim point features. n is the number of points in a point cloud sample. It outputs a point set group of size $n_1 \times k \times (d + c)$ by sampling n_1 centroids and grouping them, where each group corresponds to a local region. k is the number of points sampled from the centroid point’s neighbourhood. The subsequent pointnet layer outputs a local region feature vector $n_1 \times (d + c_1)$. We take all the sampled points as a group in the last abstraction layer and output the representation vector h (1024-dimensional). We design three linear layers as our projection head $g(\cdot)$ to map each 1024-dimensional representation vector to a 128-dimensional vector z .

Segmentation. The segmentation encoders are based on the encoders for classification. The 1024-dimensional representation vector output by encoder 1 is copied n times to form an n tensor. Concatenating it with the $n \times 1024$ tensor obtained from the last convolutional layer gives a $n \times 2048$ tensor. As such, this tensor contains both global features and features for each point. The encoder 2 extends the classification’s encoder 2 by adding three propagation layers. We adopt distance-based interpolation and a skip-link across levels propagation strategy. In a propagation level, we propagate point features from $n_1 \times (d + c_1)$ points to n points. We achieve feature propagation by concatenating interpolating feature values c_1 of n_1 points with skip linked point features from the set abstraction level c of the n points. It outputs a $n \times (d + c_1 + c)$ vector, which is then passed through the unit pointnet to obtain a $n \times 1024$ tensor. The 1024-dimensional vector output from the abstraction layer is copied n times and concatenated with the $n \times 1024$ -dimensional tensor output from the propagation layer to obtain the final $n \times 2048$ representation tensor. The tensor obtained by encoder 1 and encoder 2 are max-pooled separately to obtain a 2048-dimensional representation vector. These two vectors are used as the feature

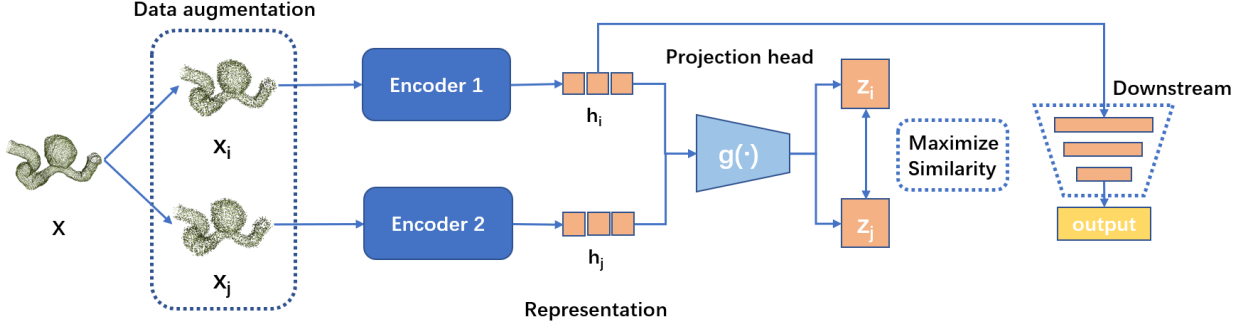


Figure 1. The architecture of our method including data augmentation, encoder, projection head, loss function and downstream tasks. We first jitter a point cloud x to construct a pair (x_i, x_j) . The representation vectors z_i and z_j to the pair of point clouds are then extracted via the dual-branch encoders and mapping head, and network optimisation is performed by a contrastive loss. The representation h obtained by dual-branch encoders will be used for downstream tasks.

representation h for the downstream segmentation network. We design two linear layers as our projection head $g(\cdot)$ to map each 2048-dimensional representation vector to a 512-dimensional vector z .

3.3. Contrastive Loss

We use a contrastive loss function similar to [4]. With this loss function, unsupervised learning can effectively learn separable features for point clouds. After the projection head $g(\cdot)$, for each sample in the mini-batch, we obtain the projection representation z . For the pair x_i and x_j , we use their projection representations z_i and z_j to measure the cosine similarity between the two samples, as follows:

$$s_{i,j} = \frac{z_i^\top z_j}{(\|z_i\| \|z_j\|)} \quad (1)$$

Intuitively, the similarity for a positive sample pair should be high. A combination pair of a positive and a negative sample should be low. Then, we use to get a similar probability of each positive sample pair in a mini-batch. $\mathbb{1}_{[k \neq i]} \in \{0, 1\}$ is an indicator function evaluating to 1 iff $k \neq i$. The equation for calculating the probability of similarity is as follows:

$$S(i, j) = \frac{\exp(s_{i,j})}{\sum_{k=1}^{2N} \mathbb{1}_{k \neq i} \exp(s_{i,k})} \quad (2)$$

We use the negative logarithm to calculate the loss of the sample pair. This loss has been used in previous works [19, 27, 35]. τ denotes a temperature parameter which scales the input and expands the range of cosine similarity. This loss is known as the normalized temperature-scaled cross-entropy loss [2, 19] as follows:

$$l(i, j) = -\log \frac{\exp(s_{i,j}/\tau)}{\sum_{k=1}^{2N} \mathbb{1}_{k \neq i} \exp(s_{i,k}/\tau)} \quad (3)$$

We calculate the average loss of both (i, j) and (j, i) in the mini-batch. Based on this loss, the representation of the encoder and projection head improves over time, and the trained network places similar samples closer in the representation space. Specifically, the loss function is given by:

$$L = \frac{1}{2N} \sum_{k=1}^N [l(2k-1, 2k) + l(2k, 2k-1)] \quad (4)$$

3.4. Downstream Tasks

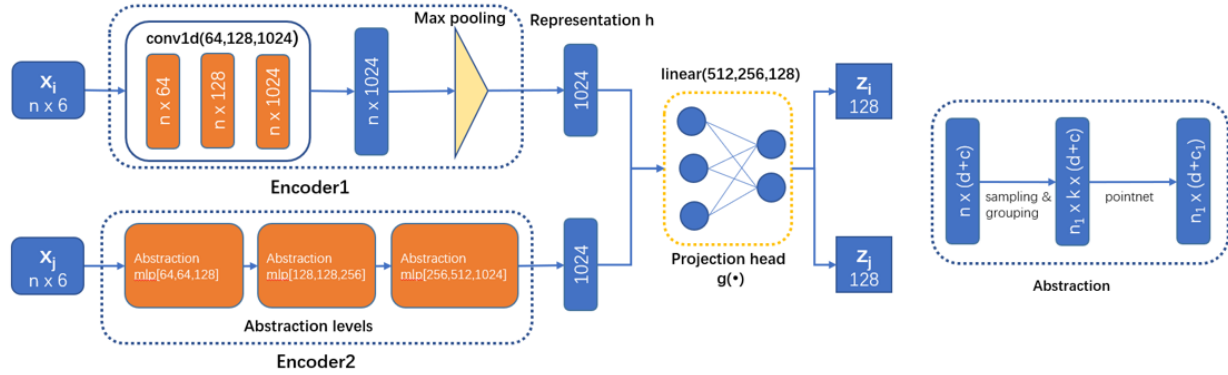
We design two simple downstream networks to evaluate the unsupervised learned representations for classification and segmentation, respectively. Each point cloud of a vessel segment is fed into the unsupervised dual-branch encoders to obtain two representations. We then concatenate the two representations into one and use this representation as input to train the downstream network. As for the binary classification task, we use four linear layers (512, 256, 128, 2) as the downstream network. Regarding the segmentation task, we employ four 1D convolutional layers (1024, 512, 256, m), where m is the number of segmentation labels.

4. Evaluation

4.1. Datasets

Intra [38] consists of complete models of aneurysms, generated vessel segments and annotated aneurysm segments. Intra collected 103 3D models of the entire cerebral vasculature by reconstructing 2D MRA images scanned for patients. Intra generated 1,909 vessel segments from the complete model, including 1,694 healthy vessel segments and 215 aneurysmal segments. Additionally, 116 aneurysm segments were manually annotated for each point. In Intra, each sample was represented as a 3D point cloud. Each

Classification Network



Segmentation Network

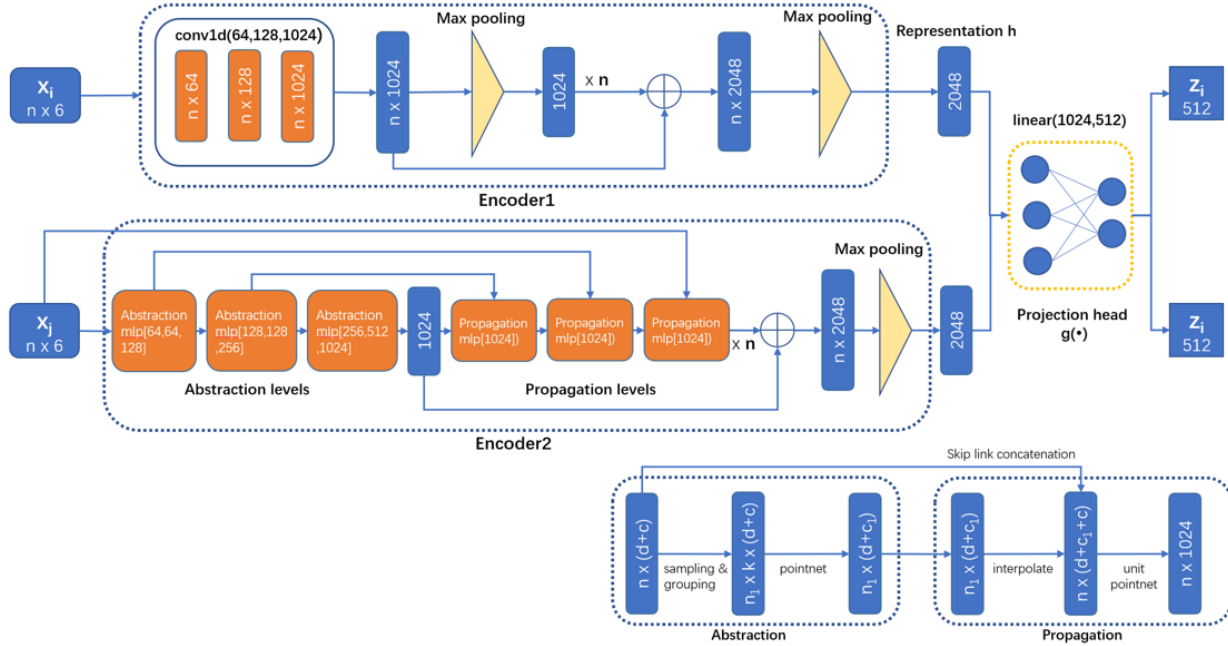


Figure 2. Dual-branch Encoders and projection head. “conv1d”: 1D convolution, “linear”: fully connected layers, “mlp” stand for multi-layer perceptron. The numbers in brackets represent the layer sizes. All convolutions and fully connected layers include batchnorm and ELU.

point p is a 6D vector composing of its coordinates and normal vector. Following Intra, we combined the generated vessel segments and manually annotated aneurysms to achieve a total of 2,025 samples. These 2,025 vessel segments will be used as the dataset for our unsupervised training. All 2,025 vessel segments will be used for the downstream classification task. 116 annotated aneurysm segments will be used for the downstream segmentation task.

ModelNet-40 [34] is a collection of 40 categories and 12,311 models culled from the mesh surfaces of CAD models. Following previous practice, 9,843 models are employed for training, and 2,468 for testing. Each point cloud

has 2,048 points, and all of the points’ coordinates are normalised to the unit sphere. Each point is a 6D vector made up of its coordinates and normal vector. We take 1,024 points from each item and augment the data by jittered. This dataset is used for comparisons of our method with other unsupervised methods.

4.2. Experimental Setting

For unsupervised training, we use the Adam optimizer with a weight decay of 10^{-6} . The mini-batch size is set to 32. The number of epochs is 200. The initial learning rate is 10^{-3} . The learning rate is scheduled to be multiplied by 0.5

	Network	#.Points	V.(%)	A.(%)	F1-Score
Supervised	SpiderCNN [37]	512	98.05	84.58	0.8692
		1024	97.28	87.9	0.8722
		2048	97.82	84.89	0.8662
	SO-Net [13]	512	98.76	84.24	0.8840
		1024	98.88	81.21	0.8684
		2048	98.88	83.94	0.8850
	PointCNN [14]	512	98.38	78.25	0.8494
		1024	98.79	81.28	0.8748
		2048	98.95	85.81	0.9044
	DGCNN [32]	512/10	95.22	60.73	0.6578
		1024/20	95.34	72.21	0.7376
		2048/40	97.93	83.40	0.8594
	PointNet++ [21]	512	98.52	86.69	0.8928
		1024	98.52	88.51	0.9029
		2048	98.76	87.31	0.9016
PointNet [20]	512	94.45	67.66	0.6909	
	1024	94.98	64.96	0.6835	
	2048	93.74	69.50	0.6916	
Unsupervised	FoldingNet [39]	512	91.37	77.41	0.6159
		1024	91.83	78.28	0.6241
		2048	91.64	79.54	0.6316
	Our(single_PN)	512	94.33	75.55	0.7233
		1024	94.21	78.33	0.7424
	Our(single_PN++)	2048	94.84	77.05	0.7408
		512	95.33	80.55	0.7679
	Our(single_PN++)	1024	95.63	83.60	0.7968
		2048	95.74	83.41	0.7988
	Our(dual)	512	96.74	82.35	0.8296
1024		97.45	84.28	0.8613	
		2048	95.41	89.47	0.8226

Table 1. Classification results of each method. The additional input K is required for DGCNN. PN: PointNet, PN++: PointNet++.

in every 10 epochs. We use jittered as the data augmentation method, which directly adds Gaussian noise to every coordinate and normal information of input point clouds. In the encoder 2, the number of points sampled from the centroid point’s neighbourhood k is set to [32, 64, “None”]. “None” means that all points are sampled. The projection head outputs a feature z , the dimension of which is set to 128 for the classification task and 512 for the segmentation task. In the loss function, the temperature parameter τ is set to 0.5.

For the downstream network, the optimizer, the number of epochs, representation dimensions, initial learning rate, learning rate decay schedule and mini-batch size are the same as those in unsupervised training. We sample 512,1024 and 2048 points separately in each point cloud for both experiments. For classification task weight decay is set to 10^{-6} and size of linear is set to [512, 256, 128, a]. a is the number of categories in the point cloud. For seg-

mentation task, weight decay is set to 1.0 and size of MLPs is set to [1024, 512, 256, b]. b is the number of categories of points in the point cloud.

Experiments were implemented using PyTorch on a GeForce GTX 1080 GPU. For IntrA dataset, the time for both unsupervised training on classification and segmentation is approximately 1 hour. The downstream classification and segmentation training are approximately 40 minutes and 50 minutes, respectively.

4.3. Experimental Results

We evaluate the classification task and the segmentation task separately on **IntrA** [38]. To demonstrate the generalisation of our method, we also perform the classification task on **ModelNet-40** [34], and then compare our method with start-of-the-art unsupervised methods to verify the effectiveness of our method.

Classification task. On IntraA, we evaluate the performance using three metrics: (1) V. Accuracy, measuring the percentage of correctly predicted healthy vessels’ samples over all healthy vessels’ samples, (2) A. Accuracy, indicating the percentage of correctly predicted aneurysm vessels’ samples over all aneurysm vessels’ samples, (3) F1 score, representing the harmonic average of precision and recall and evaluating the quality of the model. On ModelNet-40, we evaluate the performance using overall accuracy.

As shown in Table 1, as for our dual-branch encoders method with PointNet and PointNet++ backbones (i.e. PN, PN++), 1,024 sample points have the best results in terms of the F1 score and V. Accuracy compared with other numbers of sample points. The results for 512 sample points are still impressive, though the number of points in each point cloud is much smaller. Although the 2,048 sample point result is not the best in terms of F1 score and V. accuracy. *Notice that our method in 2,048 sample point has the best A. Accuracy results compared with all mentioned methods. The ability to identify aneurysms is essential in this case.* Furthermore, we find that the A. accuracy increase with more sample points. Compared with other supervised methods, our results are better than the supervised PointNet in all metrics. Besides, it also outperforms more advanced supervised networks such as DGCNN in general. Result of 1,024 sample points is very close to SO-Net and SpiderCNN in terms of F1-Score. We also compare our method with FoldingNet, one of the most representative unsupervised methods. Obviously, our method performs better on all metrics. The effectiveness of unsupervised learning is inherently limited due to the unsupervised nature, and the tubular structure of intracranial aneurysms is much less prominent compared with other data. It causes our method (dual) to perform less well than supervised PN++.

Method	ModelNet40(%)
SPH [11]	68.2%
LFD [3]	75.5%
T-L Network [5]	74.4%
VConv-DAE [23]	75.5%
3D-GAN [33]	83.3%
Latent-GAN [1]	85.7%
FoldingNet [39]	88.4%
PointCapsNet [41]	88.9%
MultiTask [7]	89.1%
Our(dual)	90.79%

Table 2. Classification accuracy of unsupervised learning on ModelNet40.

As shown in Table 2, we compare the performance of our model with other unsupervised methods on ModelNet40 [34]. we can see that our method outperforms all other un-

supervised methods, which again confirms its effectiveness in unsupervised representation learning.

Network	#.Points	IoU_V.(%)	IoU_A.(%)
SO-Net	512	94.22	80.14
	1024	94.42	80.99
	2048	94.46	81.40
PN++	512	93.42	76.22
	1024	93.35	76.38
	2048	93.24	76.21
PointCNN	512	92.49	70.65
	1024	93.47	74.11
	2048	93.59	73.58
SpiderCNN	512	90.16	67.25
	1024	87.95	61.60
	2048	87.02	58.32
PointGrid	16/2	78.32	35.82
	16/4	79.49	38.23
	32/2	80.11	42.42
PointNet	512	73.99	37.30
	1024	75.23	37.07
	2048	74.22	37.75
Our(PN)	512	80.05	44.66
	1024	82.54	46.55
	2048	81.65	48.45
Our(PN++)	512	80.05	40.66
	1024	82.05	41.42
	2048	82.65	42.45
Our(dual)	512	82.25	48.66
	1024	84.35	50.92
	2048	82.65	51.45

Table 3. Segmentation results of each network.

Segmentation task. Following [38], we evaluate the segmentation performance using two metrics: (1) V. IoU, indicating the IoU of healthy vessel, and (2) A. IoU, indicating the IoU of aneurysm vessel.

As shown in Table 3, as for our method (dual), 1,024 sample points have the best results in terms of V. IoU. But 2,048 sample points have the best results in terms of A. IoU. In comparison, our method outperforms the supervised PointNet on both V. IoU and A. IoU. *Notice that our method is better than more advanced supervised networks like PointGrid.* Compared to the supervised PointNet, which is trained with only 116 labelled samples, our method is able to learn unsupervised features from a much wider range of data, thus facilitating downstream network training. Our method generally generates better results with increasing the point number, and produces better results than the supervised PointNet in both metrics. Our method (PN++) is still inferior to the supervised PointNet++, which

is considered to be limited by the unsupervised nature.

4.4. Ablation Studies

We explore the factors that make our method effective through ablation experiments. We conduct two ablation experiments to further understand which data augmentation is more effective, and the effect of dual-branch encoders. We also analyse the effectiveness of our method in the case of sparse labels. The ablation experiments sample 1,024 points in each point cloud.

Augmentation. We try to find the best data augmentation method for our unsupervised method, by considering three different augmentation methods including *rotation*, *jittered* and *perturbation*.

Augmentation	V.(%)	A.(%)	F1-Score
rotation	95.23	75.52	0.7637
perturbation	95.35	81.87	0.8121
jittered & perturbation	95.63	81.76	0.8240
jittered	97.45	84.28	0.8613

Table 4. Ablation study on augmentation.

Rotation means randomly rotating the point cloud along the *Y*-axis. Perturbation means randomly rotating the point cloud by a small angle along the *XYZ*-axis. Jittered is the addition of Gaussian noise to the *XYZ* coordinates and normal information of the point cloud. As shown in Table 4, jittered is the best data augmentation for both branches, and the method with both jittered and perturbation is the second and generally better than the perturbation for both branches. The data augmentation with both branches rotation is the least effective. Based on the results, we can find that the data augmentation using jittered allows the encoder to learn the distinctive features of the point cloud more effectively, thereby giving better results.




Figure 3. Ablation study on dual-branch encoders.

Dual-branch encoders. In order to investigate the effectiveness of dual-branch encoders, experiments are designed to compare it with the traditional single encoder method. As

shown in Figure 3, the dual-branch encoders method has the best performance for both classification and segmentation tasks, in particular for the classification task. Besides, the single encoder method based on the more advanced PN++ is inferior to the PN-based method. This is probably due to contrastive learning depends largely on global information.

Based on these results, we have the following findings:

- Dual-branch encoders are able to extract more discriminative features. In particular, the two encoders in our design are PN and PN++ respectively where PN focuses on global features and PN++ on local features.
- Contrastive learning can better understand the distinctions between the features extracted by the two encoders. Therefore it is more effective than a single encoder. In our design, PN and PN++ as encoders can better highlight the distinctions between the local and global features of a point cloud sample, allowing the contrastive loss to optimise the network more effectively.
- Contrastive learning is excellent at describing objects as a whole, but is weak at describing them at the point scale. Our method achieves outstanding results in classification tasks, but is moderately effective in segmentation tasks. This is because our contrastive learning is not a comparison between points but between point clouds as a whole.

Network	Label(%)	V.(%)	A.(%)	F1-Score
PointNet	10	87.43	53.33	0.3298
	5	86.53	42.85	0.3012
	1	84.86	30.58	0.2485
PointNet++	10	94.37	70.58	0.7111
	5	92.89	63.23	0.6370
	1	89.55	45.07	0.4637
Our(dual)	10	95.34	71.19	0.7294
	5	94.39	67.27	0.6935
	1	90.84	58.31	0.5712

Table 5. Ablation study for limited labeled Data.

Limited Labeled data. In real-world situations, we frequently lack sufficient labeled data. We divide the original dataset into two parts, A and B, assume A to be the unlabeled data and B to be the labeled data, to represent such circumstance. The percentages of labeled data are set to 10%, 5% and 1%, respectively. In unsupervised learning, we use A+B to pre-train the model, and then use B for training the downstream tasks. Because of the nature of supervised learning, only B is used for other supervised training. The experiments are set up with a classification

task and performed on the IntrA dataset. As shown in Table 5, the accuracy of the classification gradually decreases as the amount of annotated data decreased. However, the accuracy of our model consistently outperforms that of the supervised models. This suggests that our method is more robust by making use of the unlabeled data for unsupervised learning. An interesting point is to combine evolutionary optimization with the proposed method to enhance the performance on limited labeled data [17].

5. Conclusion

In this work, we have presented an unsupervised representation learning method for the classification and segmentation of 3D intracranial aneurysms. It first augments a point cloud into two samples, and pairs them up for going through the dual-branch encoders and a subsequent common projection head. Distinctive features are learned by maximising the correspondence for a pair. The representations learned by the unsupervised trained encoders are used as input for the downstream tasks. Experiments demonstrated that our method is effective in learning unsupervised representations and can achieve better or comparable performance than state-of-the-art supervised and unsupervised learning methods.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Representation learning and adversarial generation of 3d point clouds. *arXiv preprint arXiv:1707.02392*, 2(3):4, 2017. 7
- [2] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *arXiv preprint arXiv:1906.00910*, 2019. 4
- [3] Ding-Yun Chen, Xiao-Pei Tian, Yu-Te Shen, and Ming Ouhyoung. On visual similarity based 3d model retrieval. In *Computer graphics forum*, volume 22, pages 223–232. Wiley Online Library, 2003. 7
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 4
- [5] Rohit Girdhar, David F Fouhey, Mikel Rodriguez, and Abhinav Gupta. Learning a predictable and generative vector representation for objects. In *European Conference on Computer Vision*, pages 484–499. Springer, 2016. 7
- [6] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742. IEEE, 2006. 3
- [7] Kaveh Hassani and Mike Haley. Unsupervised multi-task feature learning on point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8160–8171, 2019. 7
- [8] Jincen Jiang, Xuequan Lu, Wanli Ouyang, and Meili Wang. Unsupervised representation learning for 3d point cloud data. *arXiv preprint arXiv:2110.06632*, 2021. 2
- [9] Bio Joo, Sung Soo Ahn, Pyeong Ho Yoon, Sohi Bae, Beom-seok Sohn, Yong Eun Lee, Jun Ho Bae, Moo Sung Park, Hyun Seok Choi, and Seung-Koo Lee. A deep learning algorithm may automate intracranial aneurysm detection on mr angiography with high diagnostic performance. *European Radiology*, 30:5785–5793, 2020. 1
- [10] Konstantinos Kamnitsas, Enzo Ferrante, Sarah Parisot, Christian Ledig, Aditya V Nori, Antonio Criminisi, Daniel Rueckert, and Ben Glocker. Deepmedic for brain tumor segmentation. In *International workshop on Brainlesion: Glioma, multiple sclerosis, stroke and traumatic brain injuries*, pages 138–149. Springer, 2016. 2
- [11] Michael Kazhdan, Thomas Funkhouser, and Szymon Rusinkiewicz. Rotation invariant spherical harmonic representation of 3 d shape descriptors. In *Symposium on geometry processing*, volume 6, pages 156–164, 2003. 7
- [12] Teuvo Kohonen. The self-organizing map. *Proceedings of the IEEE*, 78(9):1464–1480, 1990. 2
- [13] Jiaxin Li, Ben M Chen, and Gim Hee Lee. So-net: Self-organizing network for point cloud analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9397–9406, 2018. 1, 2, 6
- [14] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on χ -transformed points. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 828–838, 2018. 1, 2, 6
- [15] Xuequan Lu, Honghua Chen, Sai-Kit Yeung, Zhigang Deng, and Wenzhi Chen. Unsupervised articulated skeleton extraction from point set sequences captured by a single depth camera. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. 2
- [16] Xuequan Lu, Zhigang Deng, Jun Luo, Wenzhi Chen, Sai-Kit Yeung, and Ying He. 3d articulated skeleton extraction using a single consumer-grade depth camera. *Computer Vision and Image Understanding*, 188:102792, 2019. 2
- [17] T. Nakane, N. Bold, H. Sun, X. Lu, T. Akashi, and C. Zhang. Application of evolutionary and swarm optimization in computer vision: a literature survey. *IPSJ Transactions on Computer Vision and Applications*, 12(1):1–34, 2020. 9
- [18] Takahiro Nakao, Shouhei Hanaoka, Yukihiro Nomura, Issei Sato, Mitsutaka Nemoto, Soichiro Miki, Eriko Maeda, Takeharu Yoshikawa, Naoto Hayashi, and Osamu Abe. Deep neural network-based computer-assisted detection of cerebral aneurysms in mr angiography. *Journal of Magnetic Resonance Imaging*, 47(4):948–953, 2018. 1, 2
- [19] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 4
- [20] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas 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*, pages 652–660, 2017. 1, 2, 3, 6

- [21] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017. [1](#), [2](#), [3](#), [6](#)
- [22] Aditya Sanghi. Info3d: Representation learning on 3d objects using mutual information maximization and contrastive learning. In *European Conference on Computer Vision*, pages 626–642. Springer, 2020. [2](#)
- [23] Abhishek Sharma, Oliver Grau, and Mario Fritz. Vconv-dae: Deep volumetric shape learning without object labels. In *European Conference on Computer Vision*, pages 236–250. Springer, 2016. [7](#)
- [24] Z Shi, B Hu, UJ Schoepf, RH Savage, DM Dargis, CW Pan, XL Li, QQ Ni, GM Lu, and LJ Zhang. Artificial intelligence in the management of intracranial aneurysms: current status and future perspectives. *American Journal of Neuroradiology*, 41(3):373–379, 2020. [2](#)
- [25] Zhao Shi, Chongchang Miao, U Joseph Schoepf, Rock H Savage, Danielle M Dargis, Chengwei Pan, Xue Chai, Xiu Li Li, Shuang Xia, Xin Zhang, et al. A clinically applicable deep-learning model for detecting intracranial aneurysm in computed tomography angiography images. *Nature communications*, 11(1):1–11, 2020. [1](#)
- [26] T Sichter, A Faron, R Sijben, N Teichert, J Freiherr, and M Wiesmann. Deep learning-based detection of intracranial aneurysms in 3d tof-mra. *American Journal of Neuroradiology*, 40(1):25–32, 2019. [1](#), [2](#)
- [27] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 1857–1865, 2016. [4](#)
- [28] Joseph N Stember, Peter Chang, Danielle M Stember, Michael Liu, Jack Grinband, Christopher G Filippi, Philip Meyers, and Sachin Jambawalikar. Convolutional neural networks for the detection and measurement of cerebral aneurysms on magnetic resonance angiography. *Journal of digital imaging*, 32(5):808–815, 2019. [1](#), [2](#)
- [29] Daiju Ueda, Akira Yamamoto, Masataka Nishimori, Taro Shimono, Satoshi Doishita, Akitoshi Shimazaki, Yutaka Katayama, Shinya Fukumoto, Antoine Choppin, Yuki Shimahara, et al. Deep learning for mr angiography: automated detection of cerebral aneurysms. *Radiology*, 290(1):187–194, 2019. [1](#), [2](#)
- [30] Petar Veličković, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. Deep graph info-max. *arXiv preprint arXiv:1809.10341*, 2018. [2](#)
- [31] Weijia Wang, Xuequan Lu, Dasith de Silva Edirimuni, Xiao Liu, and Antonio Robles-Kelly. Deep point cloud normal estimation via triplet learning. *arXiv preprint arXiv:2110.10494*, 2021. [2](#)
- [32] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019. [1](#), [2](#), [6](#)
- [33] Jiajun Wu, Chengkai Zhang, Tianfan Xue, William T Freeman, and Joshua B Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *arXiv preprint arXiv:1610.07584*, 2016. [7](#)
- [34] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. [5](#), [6](#), [7](#)
- [35] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3733–3742, 2018. [4](#)
- [36] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *European Conference on Computer Vision*, pages 574–591. Springer, 2020. [2](#)
- [37] Yifan Xu, Tianqi Fan, Mingye Xu, Long Zeng, and Yu Qiao. Spidercnn: Deep learning on point sets with parameterized convolutional filters. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 87–102, 2018. [1](#), [2](#), [6](#)
- [38] Xi Yang, Ding Xia, Taichi Kin, and Takeo Igarashi. Intra: 3d intracranial aneurysm dataset for deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2656–2666, 2020. [1](#), [4](#), [6](#), [7](#)
- [39] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 206–215, 2018. [3](#), [6](#), [7](#)
- [40] Dongbo Zhang, Xuequan Lu, Hong Qin, and Ying He. Point-filter: Point cloud filtering via encoder-decoder modeling. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2015–2027, 2020. [2](#)
- [41] Yongheng Zhao, Tolga Birdal, Haowen Deng, and Federico Tombari. 3d point capsule networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1009–1018, 2019. [7](#)