

PointNet

ZACH SIEGEL, APEKSHA SINGH, NAT SNYDER

Math 273

Abstract

Object classification and scene segmentation are cornerstones of the computer vision community. While there are an abundance of algorithms and strategies for image-based data, until recently, it was not known how to process point-cloud data in its raw form for the aforementioned objectives. In this paper, we describe PointNet, an novel deep learning architecture capable of object classification and scene segmentation directly from raw point clouds. We discuss PointNet’s architecture design in detail as well as methods researchers and data-scientists alternatively took to circumvent the challenges of raw point-clouds (section 2). We highlight and explore the optimization choices the authors applied to PointNet (section 3) and directly evaluate its performance (and optimization hyperparameters) against a popular 3D convolutional network architecture called VoxelNet (section 4). We compare the performance of PointNet to VoxelNet on the ModelNet40 point-cloud dataset for the task of object classification.

1 Introduction

Researchers in computer vision and machine learning have developed robust algorithms [15], [5], [18] for object classification and decision making [11], [8] through the use of convolutional learning techniques on 2D image inputs. Until the work of PointNet [13] in 2017, the same could not be said about 3-dimensional data involving point-clouds. Point clouds are inherently ‘order-invariant’ data types, meaning the permutation of any or all point-members within a cloud does not affect the information the cloud encompasses. Due to this unique characteristic, point cloud data does not fit the standard image-based convolutional network architecture built for object identification and scene segmentation.

While point clouds and images have inherently different data structures, the objectives for both data types typically surround tasks such as object classification, part segmentation, and scene segmentation. Traditionally, researchers would circumvent the challenges of point clouds by transforming clouds to another data representation. One such technique was to ‘voxelize’ the point-cloud onto a 3D grid representation [21],[22], so that 3D convolutional learning techniques could then be applied. This was not without drawbacks, as depending on the scale of the point cloud, enormous (and sparse) 3d voxel grids would emerge. This also led to the challenge of computing 3d convolution over potentially huge voxel-grids in a reasonable amount of time. Works such as [10],[3] alleviated some of the challenges of sparse 3D convolution, but large amounts of memory were still needed to store 3D voxel-grids.

Another technique described in [12] was to preprocess a point cloud into multiple image views, meaning the 3D point cloud would be transformed into a 2D image from several different perspectives. This technique worked well for simple object classification tasks when a single 3D object point-cloud was isolated from its surrounding environment, but became non-trivial to extend this technique to objectives such as scene or part segmentation. Furthermore, rendering multiple image views from the point-cloud upended any hopes for ‘real-time’ applicability. Hand-crafted features were also proposed [2], [1] in attempts to construct methods which could deal with the various sizes and rotations that basic shapes could take when captured as clouds. These techniques used frequency domain transformations of the points captured in x-y-z space. However, they did not generalize well to scenarios with multiple objects and were not considered for scene or part segmentation. The deep learning community attempted to learn such spectral kernels [16] from data directly to classify manifolds, or meshes. Non-isometric shapes were a large shortcoming to the technique, as well as dealing with undesirable artifacts that were a byproduct of transforming the clouds to a mesh.

A close methodological predecessor of PointNet is [7], which generates point-like samples from image data and transforms them in a way similar to PointNet to achieve rotation-invariant classifications. First, a “differentiable sampling” network learns to extract a grid of point-like pixel clusters from a larger image, which on its own is not significantly different from other convolutional image processing methods. Next, these point-like image samples are passed through a “spatial transformer”, that re-orientates and shears the sampled grid; including a spatial transformer in a classification network is shown to impose a “canonical orientation” on the point-like samples and thus on entire images (see Figure 5 below). A point cloud, on the other hand, is already in point form, and can be directly passed into a spatial transformer; indeed, one of the central innovations of [13] is to do exactly that to achieve rotation-invariant classification.

The motivation behind PointNet was to develop an architecture that could quickly process large point clouds in their raw form, for the tasks of object classification as well as part and scene segmentation. The novel contribution of this work was to both identify and understand why a deep learning architecture could deal with order invariant data directly. In this paper we showcase their network and how it is successfully applied to all the aforementioned objectives. We summarize the novel features of the PointNet architecture in section 2. Furthermore we directly compare the classification accuracy of PointNet against a VoxelNet 3D CNN with voxel-grid input, using the ModelNet40 point-cloud dataset as a train/test source for both network architectures.

2 PointNet Architecture

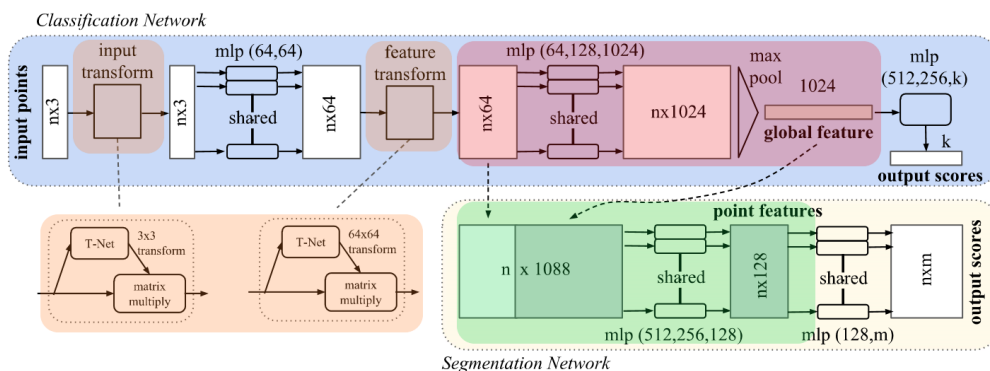


Figure 2. **PointNet Architecture.** The classification network takes n points as input, applies input and feature transformations, and then aggregates point features by max pooling. The output is classification scores for k classes. The segmentation network is an extension to the classification net. It concatenates global and local features and outputs per point scores. “mlp” stands for multi-layer perceptron, numbers in bracket are layer sizes. Batchnorm is used for all layers with ReLU. Dropout layers are used for the last mlp in classification net.

Figure 1: Color coding was added to the figure of the network architecture from the paper to reflect which components of the architecture are inspired by what properties of the point set (red=unorderedness, green=local interactions, orange=invariance to transformations).

The network architecture of the PointNet (Fig. 1) was created with the three main properties of point clouds in mind. The point set in \mathbb{R}^n , which is a subset of points from a Euclidean space, is firstly *unordered*. This is unlike pixels in an image for example. The network therefore must be invariant to permutations of the order in which input data is fed to it. The authors note that there exist three strategies to create such invariance: 1) sort input into canonical order, 2) treat input as a sequence but augment with permutations, or 3) create a symmetric function to aggregate information from each point in a way that is invariant to order, which is the strategy pursued by PointNet.

The authors comment that sorting is often not possible as it essentially requires defining a geometry-preserving bijection between a high-dimension space and a $1d$ line. The sorting should be stable with respect to perturbations of the points, which would require spatial proximity in the high dimension space to be preserved by the bijection. Additionally for point sets, there are often thousands of input elements, so sampling enough permutations of ordered data for training a RNN may also be infeasible. A symmetric function is one that takes n vectors as input and outputs the same vector regardless of the input order of the n vectors. Empirically, the authors demonstrate that these first two strategies do not perform as well as their PointNet’s reliance on a symmetric function. The idea behind the PointNet architecture is to approximate a function f defined on a point set by applying a symmetric function on transformed input elements:

$$f(x_1, \dots, x_n) \approx g(h(x_1), \dots, h(x_n)) \quad (1)$$

where the $x_i \in \mathbb{R}^d$ and $h : \mathbb{R}^d \rightarrow \mathbb{R}^K$, typically for some $K \gg d$. In the PointNet architecture, h is approximated by a multi-layer perceptron (MLP) network and g is approximated by a composition

of a single variable function and a max pooling function (red highlight in Fig. 1). The elementwise maximum, i.e. *max pooling* along each component of a set of vectors, is a simple example of a symmetric function. The authors tried others like elementwise summation and elementwise averaging, but found empirically that max pooling worked best. By using different h , many f can be learned that each capture different properties of the point set. This global signature ($[f_1 \dots f_K]$) is then used as input to train a final MLP for classification. The authors additionally prove that the neural network can approximate set functions (f from above) that are continuous with respect to the Hausdorff $d_H(\cdot, \cdot)$ distance (the greatest of all the distances from a point in one set to the closest point in the other set), formally:

$$\left| f(S) - \gamma\left(\text{MAX}_{x_i \in S}\{h(x_i)\}\right) \right| < \epsilon \quad (2)$$

Here $f : \mathcal{X} \rightarrow \mathbb{R}$, $\mathcal{X} = \{S : S \subseteq [0, 1]^m \text{ and } |S| = n\}$. Additionally x_i are the elements of S ordered arbitrarily, γ is a continuous function, and MAX is the elementwise-maximum of n vectors. The idea behind this proof is to consider the set S' such that $d_H(S, S') < \delta_\epsilon$ so that $|f(S) - f(S')| < \epsilon$ for any such S' given the continuity of f . The interval $[0, 1]$ (taking a one dimensional case for simplicity in notation) is split in K subintervals where $K = \lceil 1/\delta_\epsilon \rceil$ and S' is defined by taking the points in S and mapping them to the left end of the subinterval in which they lie. Next the function $\mathbf{h}(x) : \mathbb{R} \rightarrow \mathbb{R}^K = [h_1(x); \dots; h_K(x)]$ is defined where each h_j is a soft indicator of x occupying interval j (i.e $h_j = e^{-d(x, [\frac{j-1}{K}, \frac{j}{K}])}$ where d is the point to interval distance function). Now the occupancy of interval j by points in S can be found by taking the maximum over $\{h_j(x_1), \dots, h_j(x_n)\}$, denoted by the function $v_j(x_1, \dots, x_n)$. Let $\mathbf{v} = [v_1; \dots; v_K]$, a symmetric function taking in n vectors from \mathbb{R} and returning a value in $\{0, 1\}^K$ which describes the occupancy in each interval. A function τ can then be defined to map this description of the occupancy back to the unique points in S' . Letting $\gamma(\mathbf{v}) = f(\tau(\mathbf{v}))$ from the continuity of f , $|\gamma(\mathbf{v}(x_1, \dots, x_n)) - f(S)| < \epsilon$ and an expression like Equation 2 is obtained by substituting back the elementwise maximum and the function h .

Note in Figure 2 that voxelization can indeed be articulated as the “containment mapping” of points from \mathbb{R}^d to \mathbb{R}^K , where K denotes the number of voxels used to describe the point cloud and $h(x_i)$ is a one-hot encoding of the voxel in which point i lies. In voxelization, the combination of $h(x_1), \dots, h(x_n)$ can be equivalently characterized as a componentwise sum (as in Figure 2), a componentwise average, or a componentwise maximum; the authors of [13] found componentwise maximum to yield the best performance for classification of these three symmetric combinations. Note that the appropriate voxel resolution to represent different point clouds accurately and efficiently will vary with the number and density of points, which is the primary limitation of any voxelization method. Note that in the proof from [13] summarized in the previous paragraph, the simplest case $d = 1$ is addressed, and “voxels” are equal-width subintervals of $[0, 1]$.

The proof of the ability of PointNet to approximate an arbitrary continuous set function relies on the “worst-case” reduction of voxelization to a symmetric function of a transformation of each point in the point cloud. However, by instead allowing a multi-layer perceptron to map the

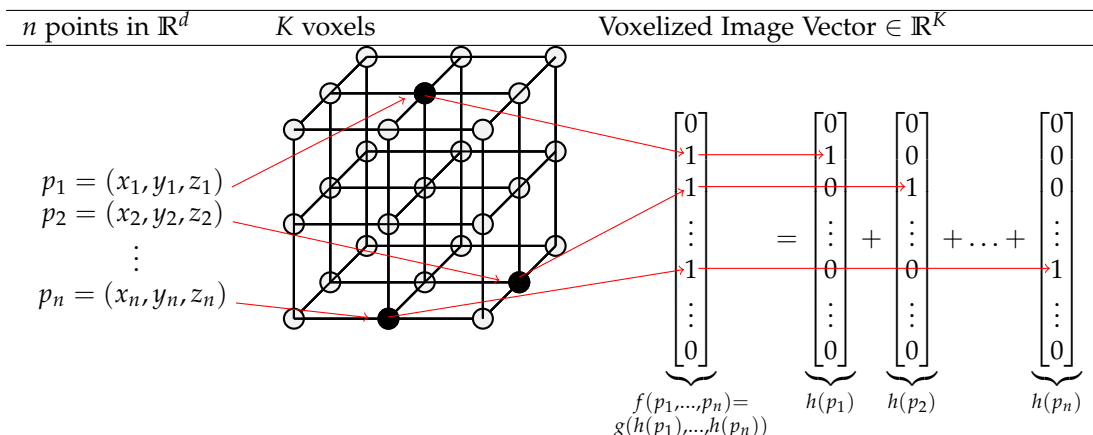


Figure 2: Voxelization is just one way to map $x_1, \dots, x_n \in \mathbb{R}^d$ to $h(x_1), \dots, h(x_n) \in \mathbb{R}^K$ and combine them using a symmetric function.

points from $\mathbb{R}^d \rightarrow \mathbb{R}^K$, backpropagation will ensure that a mapping will be used that yields good classifications. By allowing for a more general mapping to a K -dimensional space, where K no longer has to correspond to voxel resolution, PointNet is able to make classifications as good as (often better than) the best voxelization-based schemes with far lower computational costs.

There are several practical consequences of this theoretical analysis. Since the neural network approximates a continuous function, one take away is that small perturbations in the points of the point cloud should result in similarly small changes in the output. Additionally from the proof it is easy to see that the output of the neural network depends on the value of the K dimensional $\mathbf{v}(x_1, \dots, x_n)$ rather than the n dimensional $\{x_1, \dots, x_n\}$ that comes from the original point cloud. In particular, a critical set of points of size less than or equal to K can be found (Letting $u = \text{MAX}_{x_i \in S} \{\mathbf{h}(x_i)\}$, they can be found by choosing a point, c_j , in \mathcal{X} such that $h_j(c_j) = u_j$ for each element of u) and so long as these critical points are included in the input, the output of the network will not change. This means the network is robust to input data loss and can handle more sparse representations of the point cloud. Conversely adding additional points ($p \in \mathcal{X}$) to the input where $h_j(p) < u_j \forall j$ will also not change the output of the network. The authors call the collection of all points, p , satisfying this condition the upper bound shape. This practically means the network is robust to some input data noise as well. Examples of the critical points and upper bound shape learned by PointNet for some objects is given in Figure 3. As also evident by this analysis, the choice in the value of K , which the authors call the bottleneck size, does significantly impact the performance of the network. They found empirically that increasing K improved classification accuracy on the ModelNet40 dataset. Finally this theoretical analysis unpacks the main idea behind PointNet, that is is to learn a new function h , that performs better than the voxelization described.

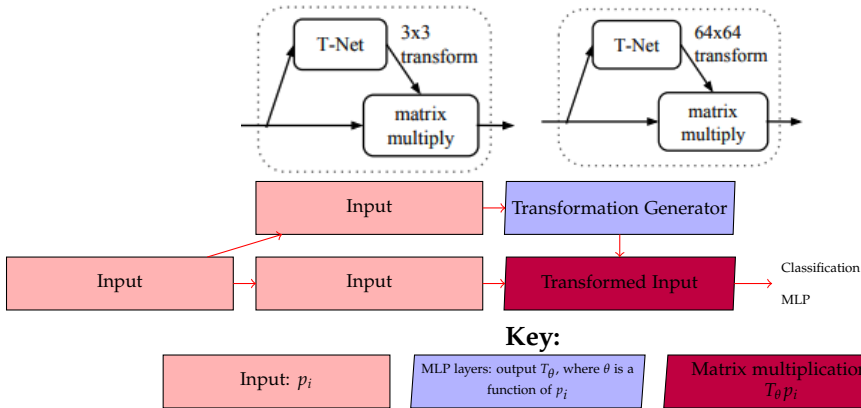


Figure 4: Above: representations from [13] of the action of a spatial transformer network on input. Below: a more verbose representation highlighting that input samples are introduced to the network twice - first, to a MLP that generates a constrained linear transformation; next, the generated linear transformation is applied to those inputs.




Figure 3: Figure from the paper demonstrating the critical point set and upper bound shapes learned by PointNet for ModelNet40 objects.

The next property of point sets is that there is *interaction between points*; they are not isolated and their local structure matters. Considering local structure, in addition to global structure, is especially important in segmentation tasks. The global features previously described are concatenated with each of the points' features (green highlight in Fig. 1). New point features are extracted from this modified input and then these point features are used to classify each point (both steps are implemented with a MLP).

The final property of point sets is that they possess *invariance under transformations*. Rotations or translations of all the points should not affect the results of classification or segmentation, since the point clouds represent physical objects. One solution is to align all input to a canonical space before extracting features. In PointNet an affine transformation (3×3) matrix is predicted by a mini-network called T-net from the input data and this transformation is then applied to the input data (orange highlight in Fig. 1). The computational pattern by which the T-nets process

inputs is visualized in Figure 4. The T-net mimics the larger PointNet architecture. It is composed of a shared MLP on each point with layer output sizes 64, 128, and 1024, then max pooling across points, and finally two fully connected layers with output sizes 512 and 256. PointNet also predicts another affine transformation matrix to align the point features: the same T-net strategy predicts a 64×64 matrix and adds the following regularization term to the loss function to constrain the matrix predicted, A to be close to orthogonal and to stabilize optimization

$$L_{reg} = \|I - AA^T\|_F^2 \quad (3)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. Note that the 3×3 (affine) transformation matrix returned by the first T-net corresponds to a shearing of the entire point cloud (a rotation and linear deformation in 3 dimensions), which is then applied to each point in the point cloud. The latter T-net has the exact same structure mathematically, but it acts on the output of a layer of the network that happens to be 64-dimensional, and so it has a less geometric interpretation. The geometrically interpretable effect of a low-dimensional spatial transformer is visualized in Figure 5.

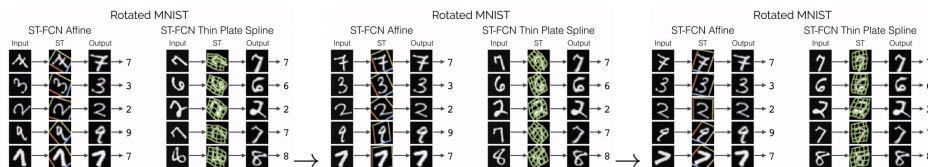


Figure 5: Three still frames of an animation depicting a spatial transformer learning over several epochs to orient handwritten digits from the “rotated MNIST” dataset from each class in a similar way (a “canonical orientation” - not necessarily upright) to yield better classifications by subsequent multi-layer perceptron. Images from [7].

The STN idea was first introduced in [7]. In section 3.2 they discuss at length the family of transformations they consider that can be learned during backpropagation and applied to two-dimensional points sampled from images as

$$T_{\theta} \mathbf{x}_i = \mathbf{A} \cdot \begin{bmatrix} x_i \\ y_i \\ 1 \end{bmatrix}.$$

For example, an “attention” transformation $T_{\theta}x$ defined by

$$A_{\theta} = \begin{bmatrix} s & 0 & t_x \\ 0 & s & t_y \end{bmatrix},$$

which corresponds to “cropping, translation, and isotropic scaling”, i.e. giving “attention” to a non-rectangular, non-upright region of an image. The broader class of transformations they consider in their model for image tasks utilizes all six entries of A_{θ} and allows for “cropping, translation, rotation, scale, and skew to be applied to the input feature map.” They note that T_{θ}

can in principle be more general than a linear transformation as long as it is differentiable with respect to the parameters. Note that the “differentiable image sampling” technique used in [7] finds a 2-dimensional point cloud-like representation of an image before transforming it; this idea is the closest intellectual predecessor of PointNet [13].

Section 4 presents our empirical results using PointNet versus a voxel-based classification network. The authors of [13] compared their best-performing network to several variants that excluded some or all of: 1) the 3×3 spatial-transformer, 2) the 64×64 spatial transformer, and 3) the regularizing loss-term L_{reg} to ensure a near-unitary transformation matrix. Their results in Figure 6 demonstrate the relative value of these network features.

Transform	accuracy
none	87.1
input (3x3)	87.9
feature (64x64)	86.9
feature (64x64) + reg.	87.4
both	89.2

Figure 6: In [13], four network variants are compared including and excluding the two spatial transformers and the regularization term L_{reg} for the 64×64 spatial transformer. For this paper, an expanded comparison of methods is provided in Figure 7.

3 Optimization Methods in PointNet Training

Training of the PointNet network requires solving an optimization problem, in particular minimizing a loss function. Here the loss function is defined as the sum of a softmax training loss and the regularization term mentioned previously. The softmax function is a generalization of the logistic function to multiple dimensions. In the final layer of PointNet the softmax function generates the probability that an instance belongs to each class. The cross-entropy or log-loss, which is the objective to be minimized, is then calculated between the actual class labels and the predicted probabilities and summed over all instances in the training dataset. This loss function is a function of the weights in the network architecture. Like most deep learning models, PointNet uses a gradient-descent based algorithm to optimize its loss function and these gradients are computed with the backpropagation algorithm, which takes advantage of the structure of neural networks, calculating the gradient of the loss function with respect to each weight one layer at a time using the chain rule. In particular, PointNet uses the Adam (adaptive moment estimation) optimizer [9]. Adam is a variant of stochastic gradient descent, which are advantageous when gradient calculations can be computationally intensive and some can potentially avoid local minima.

Stochastic gradient descent can be applied to minimizing any objective function of the form $Q(w) = \frac{1}{n} \sum_{i=1}^n Q_i(w)$ where the parameter w are the weights in the case of training a neural network and the Q_i correspond to loss function for the i th instance in the training dataset. In typical stochastic gradient descent, the gradient of the loss function is approximated at each iteration by computing the gradient based on a sample of the instances in the training data. The size of this sample is referred to as the batch size and the authors of PointNet use a batch size of

32 for training. At each iteration the samples are randomly shuffled and a new batch is selected for approximating the gradient. In vanilla stochastic gradient descent the step-size (α) or "learning rate" is fixed. This can be problematic because too large of a step size can cause instability and too small of a step size can lead to slow convergence. There have been several extensions and variants proposed based on stochastic gradient descent to address this. These include implicit stochastic gradient descent (ISGD) [20] where the stochastic gradient is essentially evaluated at the next iterate rather than the current one (which can be interpreted as a stochastic version of the proximal algorithms we saw in class). Additionally there are variants that consider a momentum term where the change in weights at each iteration is a linear combination of the stochastic gradient calculated at that iteration and the change in weights from the previous iteration (Δw), where the coefficient in front of Δw is typically an exponential decay factor between 0 and 1. Finally there are stochastic gradient methods where the step size is adaptive, such as AdaGrad, RMSProp, and Adam, in which we will focus here on the last one since it is utilized in training PointNet.

In addition to calculating the stochastic gradient at each iteration, the Adam algorithm updates exponential moving average of the gradient (m) and the squared gradient (v) (i.e. the first and second moments) at each iteration (the decay rates (β_1, β_2) of these moments are hyperparameters of the algorithm for which the authors of Adam provide suggested good default values). The updated weights at each iteration are then given by $w_{t+1} = w_t - \alpha \frac{\hat{m}}{\sqrt{\hat{v} + \epsilon}}$ where ϵ is another hyperparameter for which the authors of Adam suggest a small scalar value. The authors of PointNet utilized these default values in their TensorFlow implementation ($\beta_1 = 0.9, \beta_2 = 0.99, \epsilon = 10^{-8}$) They utilize an initial learning rate, α , of 0.001 (the value suggested by the authors of Adam), but additionally divide the learning rate by 2 every 20 epochs. Since the moving averages are initialized at zero, \hat{m} and \hat{v} are the values of m and v with bias correction. Taking ϵ to zero, it can be shown that the effective step size at each iteration is bounded above by either $\alpha(1 - \beta_1) / \sqrt{(1 - \beta_2)}$ (depending on the values of β_1 and β_2) or α . The first case only emerges with extreme sparsity, i.e. the gradient with respect to a parameter is now nonzero but has been zero at all prior timesteps. In less sparse cases the bound will be smaller, demonstrating the adaptive behavior of Adam.

The authors of Adam provide an analysis of its convergence using the online learning framework [23]. In this framework, algorithm performance is measured by minimization of the regret defined as $R(T) = \sum_{t=1}^T [Q_t(w_t) - Q_t(w^*)]$ where w^* are the optimized weights. They provide a proof that Adam has a $O(\sqrt{T})$ regret bound when the gradients are bounded, $\frac{\beta_1}{\sqrt{\beta_2}} < 1$, and α is decaying at a rate of $t^{-1/2}$ for convex functions Q_t, f_t . The idea behind their proof comes from the convexity of f_t ; since the f_t are convex, $f_t(w_t) - f_t(w^*) \leq g_t^T(w_t - w^*)$, where g_t is the gradient at iteration t . Therefore finding an upper bound on the regret is now a problem of finding an upper bound on the quantity involving the gradient from the previous expression. This relies on showing that $\sum_{t=1}^T \sqrt{\frac{g_{t,i}^2}{t}}$ where $g_{t,i}$ is the i th element of the gradient at iteration t and $\sum_{t=1}^T \frac{\hat{m}_{t,i}^2}{\sqrt{t\hat{\sigma}_{t,i}^2}}$ are bounded in their proof. Some recent works (such as in [14]) have provided examples of where Adam can fail because the average regret does not go to zero as $T \rightarrow \infty$, demonstrating the importance in hyperparameter choices and suggesting modifications to Adam, such as placing

essentially a longer term memory on past gradients (AMSGrad).

Utilizing a gradient-based optimization method requires calculation of a gradient. This warrants some discussion since all the layers of PointNet include ReLU, the rectified linear activation function. ReLU returns its input if it is greater than zero and returns zero if not (i.e. $f(x) = \max(x, 0)$). ReLU is thus a simple function that is not computationally intensive to deal with, which can contribute to faster model training times, and can give sparser solutions. Additionally traditional activation functions such as the sigmoid and tanh function can suffer from the vanishing gradient problem in gradient-descent based optimization methods. Since the sigmoid and tanh functions plateau over most values, the value of the gradient can often be near zero. This means the gradient of the front layers of neural network will be a multiplication of many small values through the backpropagation algorithm, and a vanishing gradient means the weights will not update at each iteration. Of note though, ReLU can suffer from the opposite problem of exploding gradients, which can make optimization unstable. In terms of calculating the gradient, ReLU is not differentiable at zero, however from our discussion of subdifferentials in class, gradient-based methods can simply pick a value between 0 and 1 at this point.

The PointNet network also includes a batch normalization in all layers except the last one. In batch normalization [6] a standard normalization (subtraction of mean and dividing by variance) is applied to the batch output of each layer before it is fed into the next layer. Empirically, batch normalization has led to faster model training with more robustness to hyperparameter choices, but it's worth trying to understand why these advantages occur. Viewing each layer of a neural network as a separate optimization problem, the inputs to this problem change at each step depending on the output from previous layers. Hence the optimization problem is not fixed, which the authors of the batch normalization paper term, internal covariate shift. The authors hypothesize that this constant shifting of the optimization problem hampers convergence, thus they propose batch normalization to alleviate this effect and contribute to faster convergence. More recent analysis [19] has revealed that in many settings, batch normalization actually helps to reduce the value of the gradient with respect to a layer's output and smooth the objective function. When optimizing a neural network using a gradient-based method, the internal covariate shift is related to the difference in the values of the gradient before and after updating the weights, so these two intuitions behind batch normalization's performance are related. Additionally, there are momentum versions of batch normalization, which is what is used by the authors in training PointNet. In these variations, an average of past batch mean and variances is maintained and used in computing the batch normalization, which is intended to make results more robust to batch size. For the training of PointNet, the authors actually implement a rate that starts at 0.5 and then gradually increases to 0.99.

Finally of note in the discussion of the optimization in training PointNet is the spatial transformer networks. These components do cause PointNet to be different than more standard feed-forward architectures since the input essentially enters the network twice, once for learning the affine transformation, and then again when the learned affine transformation is applied to the input (a matrix multiplication). At first glance, it might appear that such a departure from

standard architecture would cause problems for backpropagation, however this is not the case. The application of the affine transformation can be simply thought of as another layer within PointNet. The input to this layer is the affine transformation and the output to this layer is the affine transformation applied to the fixed input data. Also it is obvious that the output of this layer can be easily differentiated with respect to its input.

To conclude this section, we provide some demonstrations of the optimization methods and network components utilized in the training of PointNet by performing classification of the MNIST handwritten digit database. We implemented a simple network of three fully connected layers for this example in TensorFlow with softmax classification loss. The layers contain batch normalization and ReLU activation functions to match the PointNet implementation. First we wanted to explore the convergence behavior of the Adam optimizer. From the convergence analysis presented by the authors of Adam, we should be able to predict how changing the hyperparameters of the optimization algorithm affects convergence. In 7A we plot the training accuracy (inversely related to the loss function) versus the training epoch. We plot the results for Adam with the default parameter values mentioned previously, and also utilized by the authors PointNet. In addition we plot the results with an increased β_1 , decreased β_2 , and increased α . As predicted by the theoretical analysis, these modifications all negatively impact the convergence of the Adam optimizer. Next in 7B, we explored the effect of swapping Adam for a vanilla SGD with the same learning rate, but fixed. The results demonstrate the clear advantage of a adaptive step size, with Adam enjoying much faster convergence than Vanilla SGD (the authors of Adam have also demonstrated the superior performance of Adam compared to other adaptive methods on MNIST classification [9]). Additionally in 7B, keeping the default Adam optimizer, we explored the effect of eliminating batch normalization and switching the ReLU activation function for a sigmoid one. It should be explicitly stated here that the underlying optimization has changed with these modifications, so the values for training accuracy obtained are not directly comparable. However we can still see in this example that these modifications to the network architecture do not greatly affect the stability nor speed of convergence of the optimization process here. Of note though, the default Adam implementation does maintain the best accuracy on a separate testing dataset, suggesting it might have learned a sparser solution providing for more generalizability. The advantages of batch normalization and the ReLU activation function become more apparent in more complex problems like PointNet, where the example used here is more trivial. Finally in 7C, we utilized an implementation of the spatial transformer network with a CNN for classification of the MNIST dataset [7] as a starting point to compare results with and without the spatial transformer network. The spatial transformer creates a more complex optimization problem, which is reflected by the more oscillatory plot of loss over epochs, but ultimately the network with the spatial transformer performs slightly better on unseen data. Like mentioned previously, PointNet stabilizes the optimization by adding a regularization term and the addition of the spatial transformation network provides slightly improved classification results.

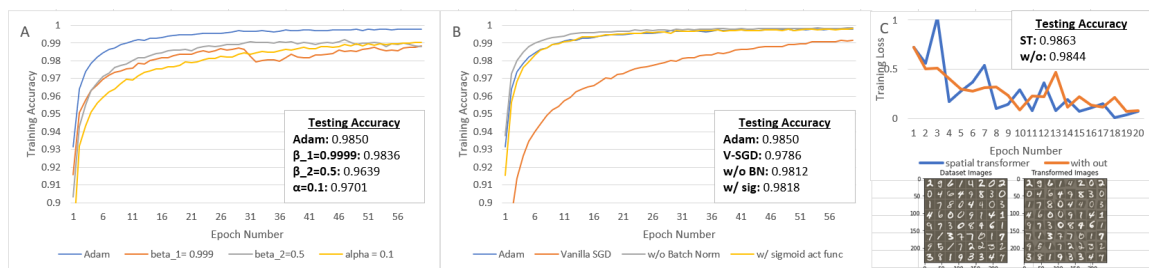
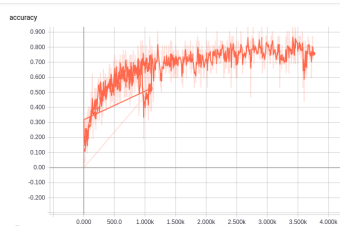


Figure 7: The PointNet methods were reproduced for this paper in several variants. Here we compare the performance of several PointNet variants including different learning rates, different SGD methods, and both including and excluding the spatial transformer networks. This is an expansion on the comparison by the authors of [13] in Figure 6.

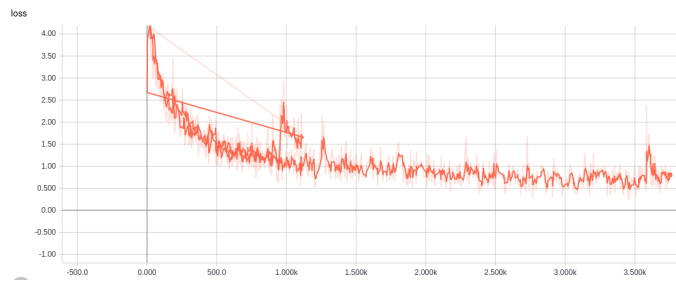
4 PointNet and Experimental Results

In this section we highlight our findings of testing the point-cloud architecture against voxelization. Particularly, we test the PointNet architecture against an architecture called ‘Voxel-Net’ which is a 3D convolutional neural network capable of performing classification over a 3D voxel-grid input. Here we compare PointNet to Voxel-Net using the ModelNet40 dataset, which consists of 40 different classes of point-cloud objects.

As the PointNet architecture has been discussed at length in the previous sections, we summarize the hyperparameters chosen to run the PointNet. An input size of 1024 points is subsampled from an object’s point-cloud, standardizing the input dimension of the cloud to 1024×3 . Standard stochastic gradient descent (SGD) with mini-batching of size 32 is chosen for use in conjunction with the Adam optimizer. A learning rate of .001 is initially set which is set to decay at a rate of .7 after every epoch. An epoch is declared as a single pass over all the training data divided by the size of a mini-batch. Before a point cloud is fed to the network for training, the cloud is rotated and jittered as well to augment the learning process. We see in Figure 8a an accuracy of 80% is achieved after 3800 mini-batches (18 epochs). The authors of PointNet note the testing accuracy can achieve 89.2% on this ModelNet40. Figure 8b depicts the training loss. Note that after 3000 batches the loss is seen to become very flat, and it is unclear whether more hyperparameter tuning would aid in further convergence. Also note that swapping the adam optimizer for a momentum based optimizer was (empirically) not seen to have any large effects on training accuracy or model loss. It should lastly be noted that model training was ended after 30 minutes to prevent any overheating of our deep learning machinery, as temperatures reached near-critical system tolerances. More elegant methods of saving and reloading parameters could have been incorporated to overcome this problem but we leave this for future work.

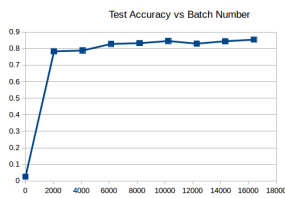


(a) *PointNet*: Test accuracy vs. number of training mini-batches. 80% accuracy was achieved after approximately 3800 minibatches (approximately 30 minutes of training on a NVIDIA GeForce 1060x GPU). Graphic from Tensorboard.

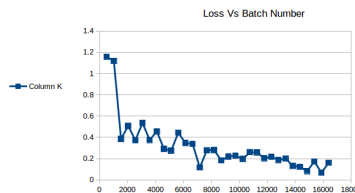


(b) *PointNet*: training loss vs number of training mini-batches.

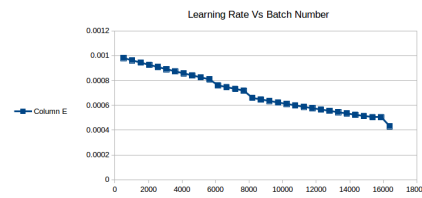
The performance of the PointNet architecture was compared to a 3D convolutional neural network which classified voxelized representations of the point-clouds. This type of network is typically dubbed a *VoxelNet*. First, all point clouds of the ModelNet40 dataset were transformed into 3D voxel-grids. (It should be noted the memory requirements to do so is quite large. 400MB of point-clouds were converted to greater than 5GB of voxel-grids). A coarse voxelgrid of $32 \times 32 \times 32$ voxels in the (x,y,z) was then chosen to uniformly discretize the point-cloud space for 3D convolution. Note this would take much more preprocessing if the ModelNet40 had not formatted the point-clouds appropriately to begin. Four layers of 3D convolution were used, which had kernel sizes of $(64,5,2)$, $(64,3,1)$, $(128,2,2)$, $(128,2,2)$, respectively. For example, layer one's $(64,5,2)$ can be read as 64 5×2 kernels, which are then moved across the 3D voxel-grid with a stride of 2 (our chosen hyper-parameter). 3D max-pooling was used after the fourth layer to aggregate latent information and reduce dimensionality. Lastly, two fully connected layers are applied to transform the (flattened) latent information after layer 4. Finally a softmax is applied to classify the 40 different objects appropriately. Note that no batch normalization or ReLu activation was applied within the last fully connected layer, but both were applied in the first fully connected layer.



(a) *VoxelNet*: Test accuracy vs. mini-batch number



(b) *VoxelNet*: loss vs. mini-batch number



(c) *VoxelNet*: learning rate vs. mini-batch number

As seen in Figure 9a, a testing accuracy of 85% was achieved using this strategy. We observe a large increase in accuracy over the first 2000 mini-batches, which eventually levels out around the 16000th mini-batch. We should note that the voxelnet was trained for more batches than pointnet only because computer overheating was (somewhat strangely) not observed in this case. The learning rate was set to decrease at a near linear rate (Figure 9c) after every 512 mini-batches,

resulting in a reasonably decreasing loss (Figure 9b).

We note that the main advantage of the PointNet over other methods (such as voxelization) is that it does not require the initial data transformation the VoxelNet requires. This makes the PointNet architecture a serious candidate for real time object classification using raw point-cloud inputs, as it is comparably as accurate but sufficiently faster with respect to the alternate techniques mentioned.

5 Conclusion

The effort in machine learning to classify image and video data produced the essential components of the PointNet architecture: fully-connected multi-layer perceptrons to approximate classification functions and convolutional layers to capture geometric relationships between features. These fundamental techniques are known to produce good results both empirically and with a theoretical basis in several “universal approximation” theorems.

Non-uniformity in input data has prompted a wide range of innovations. For image and video data of different dimensions, replacing downsampling and other pre-processing with network architectures that are *scale-invariant* has been an effort noted in several of the works cited throughout. However, PointNet was the method proposed to process point data in a way that is *permutation-invariant* regarding the inputs constituting a given point cloud.

PointNet builds directly on [7], which attempts to classify images in a way that is rotation-invariant, and originates the “spatial transformer network” idea. Before this network learns and applies an affine transformation to achieve transformation-invariance, it extracts point-like pixel clusters sampled from the larger image, which are highly reminiscent of point data. This is achieved via a “differentiable sampling” technique that is in practice simply a special case of a convolutional layer, but it is clear that their work was an important precursor to the PointNet method.

On one hand, the PointNet method is a natural extension of other rotation- and scale-invariant classification methods to permutation-invariance for point cloud data, which has led to its benchmark-setting on point cloud classification tasks both in accuracy and efficiency. Furthermore, the reduction of voxelization to max-pooling composed with a learned higher-dimensional embedding gives a lower-bound on the performance of this technique for any point data. On the other hand, for many real-time point-cloud classification tasks, there may be considerably more structure than a general point cloud, and PointNet may not be appropriate. For example, a LIDAR sensor on an autonomous vehicle or a facial identification module may generate a point cloud, but one that is oriented somewhat uniformly, especially if many readings are taken over time. For such a “mesh”, some of the features of the PointNet architecture may be irrelevant or even disadvantageous. For example, the symmetric max pooling that makes PointNet invariant to input permutation is technically a lossy function, and for a more regular data format and a requirement of high precision (such as with facial ID) or implementability on extremely limited hardware, it’s possible that PointNet is not appropriate.

Future work will likely find ways to take advantage of the innovations of PointNet. The empirically observed classification improvement after applying a regularized spatial transformer to abstract, 64-dimensional feature vectors demonstrates that a potentially wide range of natural data (image, video, language) might benefit from this type of stabilization and regularization. It's also possible that other data that are generated or stored as a sequence but are intrinsically unordered could benefit from the permutation-invariant methods of PointNet.

References

- [1] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 1626–1633, 2011.
- [2] Alexander M. Bronstein. Spectral descriptors for deformable shapes. *arXiv e-prints*, page arXiv:1110.5015, October 2011.
- [3] Martin Engelcke, Dushyant Rao, Dominic Zeng Wang, Chi Hay Tong, and Ingmar Posner. Vote3Deep: Fast Object Detection in 3D Point Clouds Using Efficient Convolutional Neural Networks. *arXiv e-prints*, page arXiv:1609.06666, September 2016.
- [4] Dumitru Erhan, Christian Szegedy, Alexander Toshev, and Dragomir Anguelov. Scalable object detection using deep neural networks. *CoRR*, abs/1312.2249, 2013.
- [5] Ross Girshick. Fast R-CNN. *arXiv e-prints*, page arXiv:1504.08083, April 2015.
- [6] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift, 2015.
- [7] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. *CoRR*, abs/1506.02025, 2015.
- [8] Mattis Manfred Kämmerer. On Policy Gradients. *arXiv e-prints*, page arXiv:1911.04817, November 2019.
- [9] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [10] Yangyan Li, Soeren Pirk, Hao Su, Charles R. Qi, and Leonidas J. Guibas. FPNN: Field Probing Neural Networks for 3D Data. *arXiv e-prints*, page arXiv:1605.06240, May 2016.
- [11] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing Atari with Deep Reinforcement Learning. *arXiv e-prints*, page arXiv:1312.5602, December 2013.
- [12] Charles R. Qi, Hao Su, Matthias Niessner, Angela Dai, Mengyuan Yan, and Leonidas J. Guibas. Volumetric and Multi-View CNNs for Object Classification on 3D Data. *arXiv e-prints*, page arXiv:1604.03265, April 2016.

- [13] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *CoRR*, abs/1612.00593, 2016.
- [14] Sashank J. Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond, 2019.
- [15] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You Only Look Once: Unified, Real-Time Object Detection. *arXiv e-prints*, page arXiv:1506.02640, June 2015.
- [16] Oren Rippel, Jasper Snoek, and Ryan P. Adams. Spectral Representations for Convolutional Neural Networks. *arXiv e-prints*, page arXiv:1506.03767, June 2015.
- [17] Sebastian Ruder. An overview of gradient descent optimization algorithms. *CoRR*, abs/1609.04747, 2016.
- [18] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *arXiv e-prints*, page arXiv:1409.0575, September 2014.
- [19] Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. How does batch normalization help optimization?, 2019.
- [20] Panos Toulis and Edoardo M. Airoidi. Asymptotic and finite-sample properties of estimators based on stochastic gradients, 2016.
- [21] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D ShapeNets: A Deep Representation for Volumetric Shapes. *arXiv e-prints*, page arXiv:1406.5670, June 2014.
- [22] Yin Zhou and Oncel Tuzel. VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection. *arXiv e-prints*, page arXiv:1711.06396, November 2017.
- [23] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, ICML'03, page 928–935. AAAI Press, 2003.