

PointTriNet: Learned Triangulation of 3D Point Sets

Nicholas Sharp¹[0000-0002-2130-3735] and Maks Ovsjanikov²[0000-0002-5867-4046]

¹ Carnegie Mellon University
² LIX, École Polytechnique, IP Paris

Abstract. This work considers a new task in geometric deep learning: generating a triangulation among a set of points in 3D space. We present PointTriNet, a differentiable and scalable approach enabling point set triangulation as a layer in 3D learning pipelines. The method iteratively applies two neural networks: a classification network predicts whether a candidate triangle should appear in the triangulation, while a proposal network suggests additional candidates. Both networks are structured as PointNets over nearby points and triangles, using a novel triangle-relative input encoding. Since these learning problems operate on local geometric data, our method is efficient and scalable, and generalizes to unseen shape categories. Our networks are trained in an unsupervised manner from a collection of shapes represented as point clouds. We demonstrate the effectiveness of this approach for classical meshing tasks, robustness to outliers, and as a component in end-to-end learning systems.

Keywords: geometric learning, triangulation, geometry processing

1 Introduction

Generating surface meshes is a fundamental problem in visual computing; meshes are essential for tasks ranging from visualization to simulation and shape analysis. In this work, we specifically consider generating surfaces via *point set triangulation*, where a fixed vertex set is given as input and a collection of triangles among those points is returned as output.

Point set triangulation has been widely studied as a classical problem in computational geometry, typically in the context of surface reconstruction from sampled point clouds such as 3D range scan data [29,44,6]. The key property of point set triangulation is that rather than allowing *any* mesh as output, the output is constrained to only those meshes which have the input points as their vertices. When reconstructing surfaces from sensor data this constraint is artificial, and makes the problem needlessly difficult: there simply may not exist any desirable mesh which has the input points as its vertex set. For this reason

Code is available at github.com/nmwsharp/learned-triangulation

most practical methods for surface reconstruction have pivoted to volumetric techniques like Poisson reconstruction [27,28] or learned implicit functions, and extract a mesh only as a post-process [40,13].

However, recent advances in point-based based geometric learning motivate a new setting for point set triangulation, where the points are generated not as the output of a sensor, but rather as the output of a learned procedure [1,42]. If such point sets could be easily triangulated, then any learning-based method capable of generating points could be immediately augmented to generate a mesh. Using point-based representations avoids the implicit smoothing and computational cost that come from working in a volumetric grid, and a differentiable meshing block enables end-to-end training to optimize for the final surface mesh.

Unfortunately, classical techniques from computational geometry are not differentiable, and recent learning-based 3D reconstruction methods are typically either restricted to particular shape classes [13] or too computationally expensive to use as a component of a larger procedure [16]. The focus of our work is to overcome these challenges, and design the first point set triangulation scheme suitable for geometric learning. Several aspects make this problem difficult: first the method must be differentiable, yet the choice of triangles is discrete—we will need to generalize the output to a distribution over possible triangles. The second major challenge is scalability; exhaustive exploration is prohibitively expensive even for moderately-sized point sets. Finally, the method should be generalizable, not tied to specific shape categories.

Our main observation is that while triangulation is a global problem, it is driven largely by local considerations. For example, the local circumscribing-circle test in 2D is sufficient to identify faces in a Delaunay triangulation. Based on this intuition, we propose a novel neural network that predicts the probability that a single triangle should appear in a triangulation, and a second, similarly-structured network to suggest neighboring triangles as new candidates. Iteratively applying these two networks generates a coherent triangle mesh.

Our approach yields differentiable triangle probabilities; gradients can be evaluated via ordinary backpropagation. The use of a proposal network allows our method to efficiently identify candidate triangles, and thus scale to very large inputs. Moreover, since our problems are local and geometric in nature, our method can easily be applied to surface patches across arbitrary shape classes. We validate these properties through extensive experiments, demonstrating that our method is competitive with classical approaches. Finally, we highlight new applications in geometric learning made possible by our differentiable approach.

2 Related work

Surface reconstruction and meshing are among the oldest and most researched areas of computational geometry, computer vision and related fields—their full overview is beyond the scope of this paper. Below we review methods most closely related to ours and refer the interested readers to recent surveys and monographs [29,44,6] for a more in-depth discussion.

A classical approach for surface reconstruction is to estimate a volumetric representation, e.g. a signed distance function, and then extract a mesh [25,15]. Both of these steps have been extended to improve robustness and accuracy, including [46,27,36] among many others (see also related techniques in [38]). These approaches work well on densely sampled surfaces, but both require normal orientation, which is notoriously difficult to estimate in practice, and do not preserve the input point samples, often leading to loss of detail.

Other approaches, based on Delaunay triangulations [9,30,10], alpha shapes [19,7] or Voronoi diagrams [3,4] often come with strong theoretical guarantees and perform well for densely sampled surfaces that satisfy certain conditions (see [18] for an overview of classical techniques). Several of these approaches, e.g. [7,2] also preserve the input point set, but, as we show below can struggle in the presence of poorly sampled data. Interestingly, there are several NP hardness results for surface reconstruction [5,8], which both explain the prevalence of heuristics, and point towards the use of *data-driven* priors. More fundamentally, classical approaches are not differentiable by nature and thus do not allow end-to-end training or backpropagation of the mesh with respect to input positions.

More closely related to ours are recent learning-based methods that aim to exploit the growing collections of geometric objects, which often come with known mesh structure. Similarly to the classical methods, most approaches in this area are based on volumetric representation. This includes, for example, a fully differentiable variant of the marching cubes [31] for computing a mesh via voxel grid occupancy prediction [22,35], and recent generative models for implicit surface prediction [13,20]. Unlike these works, our focus is on directly meshing an input point set. This allows our method to scale better with input complexity, avoid over-smoothing inherent in voxel-based techniques, and differentiate with respect to point positions.

Other surface generation methods include template-based techniques which fit a fixed mesh to the input [34,26,32], or deform a simple template while updating its connectivity [47,39]. More recent representations include fitting parameterized surface patches to points [23,48], or decomposing space in to convex sets [17,12]. Most importantly, none of the methods directly triangulate an arbitrary input point set. More generally, many of these schemes are limited to specific shape categories or topologies, or are too prohibitively expensive to use as a building-block in a larger system.

Perhaps most closely related to ours is the recent Scan2Mesh technique [16], which uses a graph-based formulation to generate triangles in a mesh. Though both generate meshes, there are many differences between this approach and ours, including that Scan2Mesh is applied to a volumetric representation, as opposed to the point-based representation used here, and Scan2Mesh does not attempt to generate watertight meshes. Several of the ideas introduced in this work could perhaps be combined with Scan2Mesh for further benefit, for instance using our proposal network (Section 3.3) to avoid quadratic complexity.

Finally, we note that our *proposal* and *classifier* architecture is inspired by classical approaches for object detection in images [21,43].

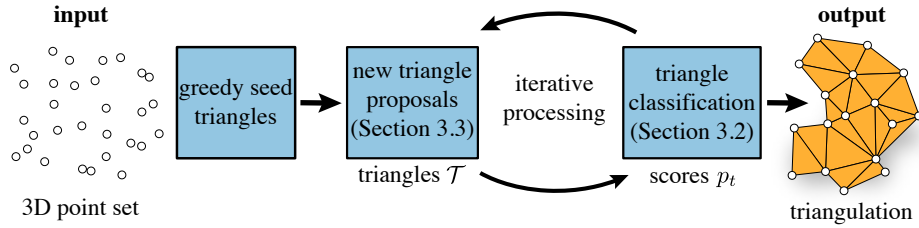


Fig. 1. An overview of our PointTriNet pipeline, which generates a triangulation by alternating between proposing new triangles and classifying them, each with a neural network. The classifier network identifies triangles which should appear in the triangulation, while the proposal network generates new candidates.

3 Method

Our method, called *PointTriNet*, has two essential components, both of which are learned. The first is a *classification* network (Section 3.2), which takes a candidate triangle as input, and outputs a score, which we interpret as the probability of the triangle appearing to the triangulation. The second is a *proposal* network (Section 3.3), which suggests likely neighbor triangles for an existing triangle. PointTriNet alternates between these networks, classifying candidate triangles and proposing new candidates, iteratively generating the output mesh.

3.1 Triangle-relative coordinates

Rather than treating mesh generation as a global learning problem, we pose it as a local problem of predicting a single triangle, and make many such predictions. Both our classification and proposal subproblems are defined with respect to a particular query triangle: either classifying the triangle as a member of the mesh, or proposing candidates adjacent to the triangle. These problems depend on both the triangle’s geometry and nearby neighborhood, but it is not immediately obvious how to encode both of these quantities for a neural network while preserving the expected invariants. Our solution is to design an encoding of points in the neighborhood of a triangle, *relative* to the triangle’s geometry.

For any point p in the neighborhood of a triangle t , our encoding is given by Cartesian coordinates x', y', z' in a frame aligned with the first edge and normal of t , as well as the point’s barycentric coordinates u, v, w after projecting in to the plane of the triangle. This results in the following encoding function, which outputs a 6-dimensional vector for p relative to t

$$\text{encode}(t, p) \rightarrow [x', y', z', u, v, w]. \quad (1)$$

Using both Cartesian and barycentric coordinates is essential to simultaneously encode *both* the shape of the triangle and the location of a neighbor point; either individually would not suffice. This encoding is invariant to scale and rigid

transformations, but not yet invariant to the permutation of the triangle’s vertices. Per-triangle permutation invariance can be easily achieved by averaging results for all 6 possible permutations, though we find this to be unnecessary in practice. Finally, to encode a nearby neighboring triangle using this scheme, we encode each of the neighbor’s vertices, then take a max and min along these encoded values, yielding 12 coordinates. This encoding is easy to evaluate, provides rigid- and (if desired) permutation- invariance, and is well-suited for use in point-based learning architectures.

3.2 Classifier network

The primary tool in our method is a network which classifies whether a single query triangle belongs in the output triangulation. This classification is a function of the nearby points, as well as nearby triangles and their previous classification scores, as shown in Figure 2.

More precisely, for a given query triangle we gather the n nearest points and m nearest triangles, as measured from triangle barycenters (we use $n = m = 64$ throughout). These neighboring points and triangles are then encoded relative to the query triangle as described in Section 3.1, and for the nearby triangles we additionally concatenate previous classification scores for a total $12 + 1 = 13$ coordinates per triangle. The result is a set of nearby encoded points $\mathcal{N}_{\text{point}} = \{\mathbb{R}^6\}_n$ and the set of nearby encoded triangles $\mathcal{N}_{\text{tri}} = \{\mathbb{R}^{13}\}_m$. We then learn a function

$$f : \mathcal{N}_{\text{point}}, \mathcal{N}_{\text{tri}} \rightarrow [0, 1] \quad (2)$$

which we interpret as the probability that the query triangle appears in the triangulation. We model this function as a PointNet [41], using the multi-layer perceptron (MLP) architecture shown in Figure 2. Here, our careful problem formulation and input encoding enables the use of a small, ordinary PointNet without spatial hierarchies or learned rotations—the inputs are already localized to a small neighborhood and encoded in a rigid-invariant manner.

3.3 Proposal network

The classifier network is effective at identifying good triangles, but requires a set of candidates as input. Naively enumerating all $O(n^3)$ possibilities does not scale, and simple heuristics miss important triangles.

Our approach identifies candidates via second *proposal* network. For each edge of a candidate triangle, the proposal network suggests new neighboring triangles across that edge. These proposals are learned by predicting a scalar function $\mathcal{V} \rightarrow [0, 1]$ on the nearby vertices. Formally, for an edge ij in triangle ijk , we learn the probability that a nearby vertex l forms a neighboring triangle ijl . These probabilities are then sampled to generate new candidates, which are assigned an initial probability as the product of the probability for the triangle that generated them and their learned sampling probability.

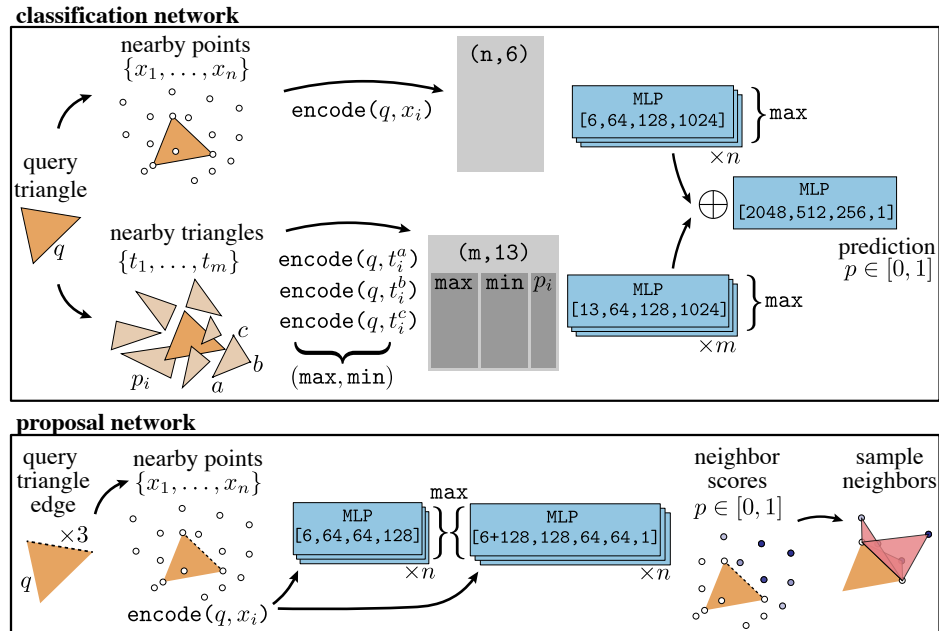


Fig. 2. The core of our technique is a network which classifies whether a given triangle belongs in a triangulation, as a function of nearby points and already-classified candidates. All coordinates are encoded relative to the query triangle. A second, similarly-structured network proposes new triangle candidates for subsequent classification.

The proposal network is again a PointNet over the local neighborhood, though it takes only points as input, and outputs a per-point result. We use the same triangle-relative point encoding (Section 3.1), but here the ordering of query triangle vertices ijk serves as a feature—the network always predicts neighbors across edge ij . Applying the network to the three cyclic permutations of ijk yields three sets of predictions, potential neighbors across each edge.

3.4 Iterative prediction

To triangulate a point set, we alternate between proposing new candidates and classifying them. At each iteration, each candidate triangle t has probability $p_t \in [0, 1]$, the probabilities from each round are features in the next, akin to a recurrent neural network. We initialize with greedy seed triangles among the nearest neighbors of each vertex, then for each iteration we first classify candidates, then generate new candidates, retaining only the highest-probability candidates. If multiple instances of the same triangle arise, we discard lower probability instances. In our experiments, we iterate 5 times, sample 4 neighbors across each edge per iteration, and retain $12|V|$ candidates. If desired, the output can be post-processed to fill small holes (Figure 4,7, right).

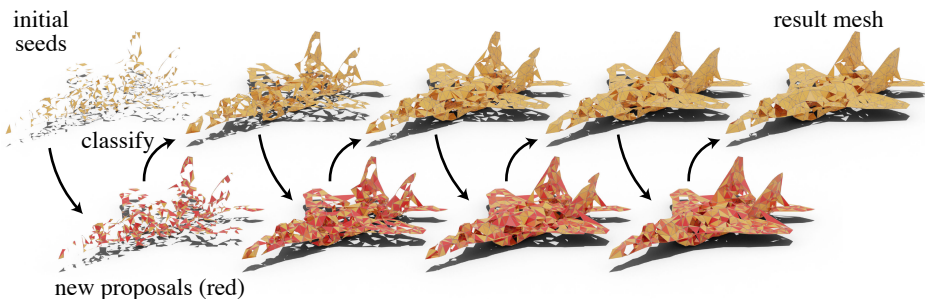


Fig. 3. Our approach iteratively constructs a triangulation, alternating between classifying triangles and proposing new candidates. Classified triangles are drawn in orange, and new proposals are drawn in red, for both only triangles with $p > 0.5$ are shown. The input vertex set is 1k sampled points, not shown here.

3.5 Losses and training

Our method outputs a probabilistic surface, where each triangle has an associated probability. This is a necessary generalization: the choice of triangles is a discrete set, and thus does not have traditional smooth derivatives. We thus make use of probabilistic losses detailed below, which are essentially the expected value of common geometric losses, like expected Chamfer distance.

For training, we assume a known, ground-truth surface \mathcal{S} , and the probabilistic surface \mathcal{T} with probabilities $p : \mathcal{T} \rightarrow [0, 1]$ predicted by our algorithm. Note that we do not assume to be given a ground truth triangulation of \mathcal{S} , only some representation of the underlying surface sufficient to measure distance. In the losses below we assume \mathcal{S} is represented by a sampled point cloud, but other representations like an implicit function could easily be substituted as training data. Notice that our method is thus unsupervised, in the sense that it learns to generate triangulations given un-annotated point clouds as input—it is not trained to match existing triangulations.

Expected Chamfer distance (forward) The distance from the ground truth surface to our predicted surface is measured as

$$\begin{aligned} \mathcal{L}_{\vec{\mathcal{C}}} &:= E\left[\int_{x \in \mathcal{S}} \min_{y \in \mathcal{T}} d(x, y)\right] = \int_{x \in \mathcal{S}} E\left[\min_{y \in \mathcal{T}} d(x, y)\right] \\ &= \int_{x \in \mathcal{S}} \int_{y \in \mathcal{T}} \gamma(x, y) d(x, y) \end{aligned} \quad (3)$$

where $\gamma(x, y)$ denotes the probability that y is the closest point in \mathcal{T} to x . Discretely, we evaluate this by sampling points x on \mathcal{S} , then for each point sorting all triangles in \mathcal{T} by distance from x . A cumulative product of probabilities for these sorted triangles gives the probability that each triangle is the closest to x , and taking an expectation of distance under these probabilities yields the desired

expected distance from the sample to our surface. In practice we truncate to the $k = 32$ nearest triangles.

Expected Chamfer distance (reverse) The distance from our predicted surface to the ground truth surface is measured as

$$\mathcal{L}_{\mathcal{C}}^{\leftarrow} := E\left[\int_{y \in \mathcal{T}} \min_{x \in \mathcal{S}} d(x, y)\right] = \int_{y \in \mathcal{T}} p(y) \min_{x \in \mathcal{S}} d(x, y) \quad (4)$$

Discretely, we evaluate this by sampling points y on \mathcal{T} , then for each point we measure the distance to \mathcal{S} and sum, weighting by the probability $p(y)$. Similar losses have appeared in 3D reconstruction, e.g. [31].

Overlap kernel loss To discourage triangles from overlapping in space, we define a spatial kernel around each triangle

$$g_t(x) = p_t \max\left(0, 1 - \frac{d_n(x)}{d_e(x)}\right) \quad (5)$$

where x is an arbitrary point in space, p_t is the probability of triangle t , d_n is the distance in the normal direction from the triangle, and d_e is the signed perpendicular distance from the triangle’s edge (see supplement for details). In a good triangulation, for any point x on the surface there will be exactly one triangle t for which $g_t(x) \approx 1$ while $g_{t'}(x) \approx 0$ for all other triangles t' . This is modeled by the loss

$$\mathcal{L}_{\mathcal{O}} := \int_{x \in \mathcal{T}} \left(-1 + \sum_{t \in \mathcal{T}} g_t(x)\right)^2 + \left(-1 + \max_{t \in \mathcal{T}} g_t(x)\right)^2 \quad (6)$$

which is minimized when exactly one kernel contributes a value of 1 at x . Discretely, we sample points x on the surface of the generated triangulation \mathcal{T} to evaluate the loss. One could instead sample x on the ground truth surface, but using the generated triangulation \mathcal{T} makes the loss a regularizer, applicable in generative settings.

Watertight loss To encourage a watertight mesh, we explicitly maximize the probability that each edge in the triangulation is watertight, via

$$\mathcal{L}_{\mathcal{W}} := \sum_{ij \in \mathcal{E}} p_{ij}(1 - p_{ij}^{\text{water}}) \quad (7)$$

where \mathcal{E} is the set of edges, p_{ij} denotes the probability that ij appears in the triangulation, and p_{ij}^{water} denotes the probability that ij is watertight. An expression for evaluating this loss from triangle probabilities is given in the supplementary material. Notice that this loss does not directly penalize vertex-manifoldness—we observe that in practice almost all watertight configurations are also vertex-manifold, and thus simply encourage watertightness.

Network training The full loss for training the classification network is given by

$$\mathcal{L} := \mathcal{L}_{\vec{C}} + \mathcal{L}_{\bar{C}} + \lambda_{\text{O}}\mathcal{L}_{\text{O}} + \lambda_{\text{W}}\mathcal{L}_{\text{W}}, \quad (8)$$

where we use $\lambda_{\text{O}} = 0.01$ and $\lambda_{\text{W}} = 1$ for all experiments. All losses are normalized by the surface area or number of elements as appropriate, for scale invariance. During training, we backpropagate gradients only through the last iteration of classification. This strategy proves effective here because the network must be able to classify the most useful triangles regardless of the quality of the current candidate set.

We train the proposal network simultaneously with the classifier by encouraging its suggestions to match the classification scores. Intuitively, this means that the network attempts to propose triangles which will receive a high classification score on the next iteration. During training, we perform one last round of proposal to generate a suggestion set P with probabilities u_t . Rather than merging these triangles into the candidate set T as usual, we instead evaluate the loss

$$\mathcal{L}_{\text{M}} := \frac{1}{|P \cap T|} \sum_{t \in P \cap T} (u_t - p_t)^2 \quad (9)$$

and backpropagate through this loss only with respect to the proposal scores u_t . During training we also augment the proposal network to predict a random neighbor with 25% probability, to encourage diversity in the proposal set.

4 Experiments

Architecture and training The layer sizes of our networks are given in Figure 2; we use ReLU activation functions throughout [37], except for a final single sigmoid activation to generate predictions. We use dropout with $p = 0.5$ on only the hidden layers of the prediction MLPs [45]. Neither batch nor layer normalization are used, as they were observed to negatively impact both result quality and the rate of convergence.

We train using the ADAM optimizer with a constant learning rate of 0.0001, and batch size of 8. Our networks are trained for 3 epochs over 20k training samples, which amounts to ≈ 5 hours on a single RTX 2070 GPU.

Datasets We validate our approach on the ShapeNetCore V2 dataset [11], which consists of about 50,000 3D models. We use the dataset-recommended 80%/20% train/test split, and further reserve 20% of the training models as a validation set. Our method is geometric in nature, so unlike semantic networks which operate per-category, we train and test simultaneously on all of ShapeNet, and do not need to subdivide by category.

We uniformly sample 1k points on the surface of each model, which will serve as the vertex set to be triangulated, and then separately sample another 10k points from each model, which serve as the representation of the surface when evaluating loss functions. Note that we *do not* use the mesh structure of the dataset during our training procedure.

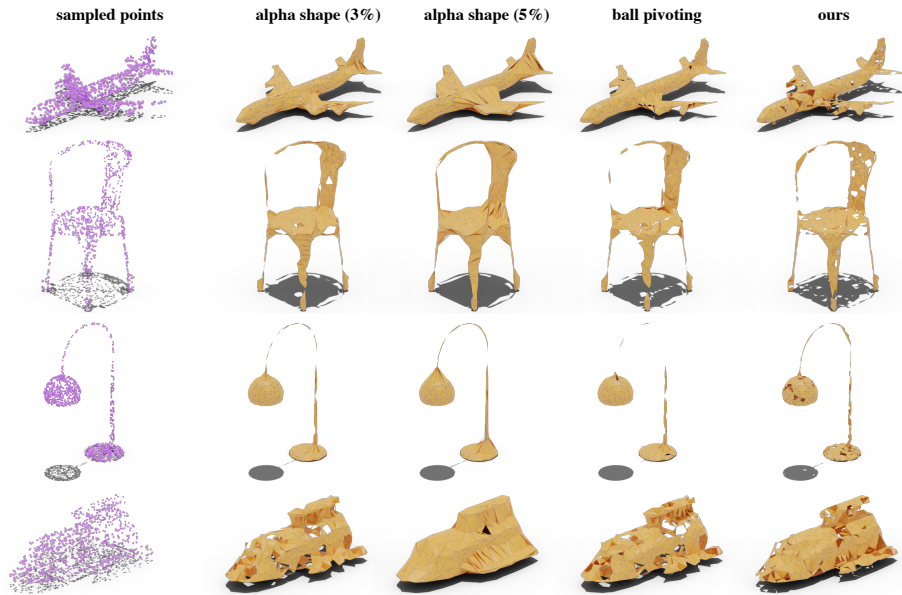


Fig. 4. A selection of outputs from PointTriNet and baselines on ShapeNet. The input is 1k points sampled on the surface of the shape, and the output is a mesh using those points as vertices. Constructing point set triangulations which are both accurate and watertight is still an open problem, but our method shows that a differentiable, learning-based approach can yield results competitive with classical computational geometry schemes.




Fig. 5. Our method is data-driven, but its geometric nature leads to excellent generalization. *Left:* reconstructing a synthetic model with vertices of varying regularity and density. *Right:* reconstructing a 3D range scan of a cathedral (from [24]). Both use our networks trained on uniformly-sampled ShapeNet in Section 4.

To generate training examples in small neighborhoods, we choose a random center point on some shape, and gather the 256 nearest vertex samples (among the 1k), along with all surface samples (among the 10k) in the same radius. We use 20k such neighborhoods as a training set, with another 5k for validation. These neighborhoods are used only for efficient training—to evaluate our method on the test set we triangulate the full 1k sampled vertices for each model.

Baselines We compare our learned approach against two classical methods for point set triangulation: ball pivoting and α -shapes [19,7]. For ball-pivoting, the ball radius is automatically guessed as the bounding box diagonal divided by the square root of the vertex count. For α -shapes, we report two different choices of the radius parameter α , as 3% and 5% of the bounding box diagonal. No prior learning-based methods can directly serve as a baseline, but in the supplement we adapt a variant of Scan2Mesh [16] for further comparison.

Metrics Geometric accuracy is assessed with a bi-directional Chamfer distance

$$\frac{1}{|\mathcal{A}|} \int_{x \in \mathcal{A}} \min_{y \in \mathcal{B}} d(x, y) + \frac{1}{|\mathcal{B}|} \int_{y \in \mathcal{B}} \min_{x \in \mathcal{A}} d(x, y) \quad (10)$$

and evaluated discretely by sampling 10k points on both meshes. Reported Chamfer distances are scaled by $100\times$ for readability. We measure mesh connectivity by *watertightness*, the percentage of edges which have exactly two incident triangles, and (edge-) *manifoldness*, the percentage of edges with one or two incident triangles—with this terminology mesh boundaries are manifold but not watertight. For our method, we take the triangles with $p > 0.9$ as the output.

4.1 Results

The most basic task for PointTriNet is to generate a mesh for a point set sampled from some underlying shape. Table 1 shows the effectiveness of our method on the sampled ShapeNet dataset, compared to the baseline approaches. Our method is geometrically more accurate than all of the classical schemes, though it achieves somewhat less regular connectivity than ball pivoting. We note that it may not be possible for any algorithm to significantly improve on these metrics: an imperfect

Table 1. When evaluating on sampled ShapeNet, PointTriNet outperforms classical methods in geometric accuracy, though it has moderately less watertight connectivity.

	Chamfer	watertight	manifold
ours	0.7417	77.0%	97.4%
ballpivot	1.3440	84.1%	100.0%
alpha-3	1.1099	49.8%	61.3%
alpha-5	0.9195	47.6%	53.0%

(e.g., random) vertex sampling likely does not admit any triangulation which has both perfect geometric accuracy and perfect manifoldness. The supplement contains an ablation study, a breakdown by class, and triangle quality statistics.

Performance Triangulating a point set with PointTriNet amounts to evaluating standard MLPs, in contrast to methods which e.g. solve an optimization problem for each input (Section 2). In our unoptimized implementation, triangulating 1000 points takes about 1 second on an RTX 2070 GPU. To reduce the memory footprint for large point sets (Figure 4), one can evaluate the MLPs independently over many small patches of the input.

Outliers An immediate benefit of a learned approach is that it can adapt to deficiencies in the input data. We add synthetic noise to our sampled ShapeNet dataset by perturbing 25% of the points according to a Gaussian with 2%-bounding-box deviation, and train our method from scratch on this data (Table 2). The classical baselines are not designed to handle noise, so our approach yields an even larger improvement in geometric accuracy in this setting.

Table 2. Evaluation on sampled ShapeNet, with noise added to a subset of the samples. Our method can adapt to the noise, further improving geometric accuracy.

	Chamfer	watertight	manifold
ours	1.0257	74.3%	98.0%
ballpivot	1.5418	83.9%	100%
alpha-3	1.3260	50.8%	61.5%
alpha-5	1.2309	47.5%	52.7%

Geometry processing A key benefit of our strategy is that PointTriNet directly generates a standard triangle mesh. Although these meshes are not necessarily manifold (a property shared with classical methods like α -shapes), they are still suitable for many standard algorithms in geometry processing (Figure 6).

4.2 PointTriNet in geometric learning

Our method is designed to be a new building block in geometric learning.

Learned vertex improvement We design a network which additionally learns a position update at each of the input vertices to further improve the quality of the output mesh, inspired by methods like [42]. More precisely, we introduce an additional MLP per-triangle, structured identically to the prediction MLP except the last layer has dimension $\mathbb{R}^{3 \times 2}$: an offset in the triangle’s tangent basis for each of its three vertices, accumulated to yield a per-vertex update in the output mesh. This network, coupled with the base PointTriNet, is trained end to

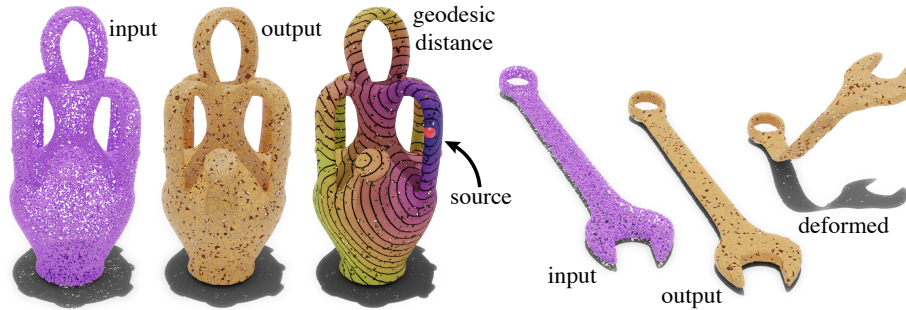


Fig. 6. PointTriNet directly generates a triangle mesh, opening the door to many standard geometric algorithms. *Left:* geodesic distance is computed from a source point [14]. *Right:* a reconstructed object is deformed via an anchor at one endpoint [33].

end as described above, including an additional loss term which encourages mesh quality by penalizing the deviation of the edge lengths. Figure 7 shows how these learned offsets further improve mesh quality with a more natural distribution of vertices.

Generative triangulation PointTriNet enables triangulation as a component in end-to-end systems—algorithms designed to output points can be augmented to generate meshes useful for downstream applications. As a proof-of-concept, we construct a shape autoencoder which directly outputs a mesh using our technique. The encoder is a PointNet mapping input points to a 128 dimensional latent space, then decoding to 512 vertices with the architecture of [1]; these vertices are then triangulated with PointTriNet. We evaluate this scheme on the planes class of ShapeNet, as shown in Figure 8. The point encoder/decoder and PointTriNet are pretrained as in [1] and Table 1, respectively, and the full system is then trained end-to-end for 3 epochs with the losses in Section 3.5.




Fig. 7. Our method can be incorporated with other learning-based techniques. Here, we additionally learn a position offset at points to further improve output mesh quality.

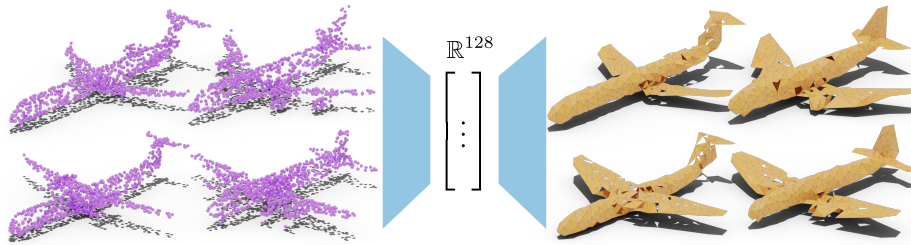


Fig. 8. A preliminary example of an autoencoder, which uses our method to directly output a mesh without any intermediate volumetric representations. Input point clouds (*left*) are decoded to meshes with 512 vertices (*right*).

This autoencoder is intentionally simplistic, and we leave a full investigation of its effectiveness to future work. However, it serves to demonstrate that PointTriNet can indeed directly generate mesh outputs from point-based geometric learning, while nearly all past approaches instead relied on volumetric representations or template deformations.

5 Limitations and future work

We introduce PointTriNet, a data-driven approach to address the classical problem of triangulating a given set of points. Our scheme uses point-based neural network architectures, and is shown to be accurate, generalizable, and scalable.

The meshes resulting from our method are competitive in accuracy with state-of-the-art point set triangulation methods from computational geometry. Crucially, however, our approach is fully differentiable, enabling novel applications such as building an autoencoder network with a triangle mesh output.

Our meshes are still noticeably less smooth than those extracted from volumetric representations—this is an inherent challenge of the strictly-harder point set triangulation problem, as opposed to general surface reconstruction. By incorporating topological priors, hole-filling, and building upon the differentiability of the meshing network, future work could further improve mesh quality.

We believe that this new approach to a classical problem will spur further research in geometric learning, and prove to be a useful component in a wide variety of applications.

Acknowledgements The authors are grateful to Marie-Julie Rakotosaona and Keenan Crane for fruitful initial discussions, and to Angela Dai for assistance comparing with Scan2Mesh. Parts of this work were supported by an NSF Graduate Research Fellowship, the KAUST OSR Award No. CRG-2017-3426 and the ERC Starting Grant No. 758800 (EXPROTEA).

References

1. Achlioptas, P., Diamanti, O., Mitliagkas, I., Guibas, L.: Learning representations and generative models for 3d point clouds. arXiv preprint arXiv:1707.02392 (2017)
2. Amenta, N., Bern, M.: Surface reconstruction by voronoi filtering. *Discrete & Computational Geometry* **22**(4), 481–504 (1999)
3. Amenta, N., Bern, M., Kamvysselis, M.: A new voronoi-based surface reconstruction algorithm. *Proc. SIGGRAPH* pp. 415–421 (1998)
4. Amenta, N., Choi, S., Kolluri, R.K.: The power crust, unions of balls, and the medial axis transform. *Computational Geometry* **19**(2-3), 127–153 (2001)
5. Barequet, G., Dickerson, M., Eppstein, D.: On triangulating three-dimensional polygons. *Computational Geometry* **10**(3), 155–170 (1998)
6. Berger, M., Tagliasacchi, A., Seversky, L.M., Alliez, P., Guennebaud, G., Levine, J.A., Sharf, A., Silva, C.T.: A survey of surface reconstruction from point clouds. In: *Computer Graphics Forum*. vol. 36, pp. 301–329. Wiley Online Library (2017)
7. Bernardini, F., Mittleman, J., Rushmeier, H., Silva, C., Taubin, G.: The ball-pivoting algorithm for surface reconstruction. *IEEE transactions on visualization and computer graphics* **5**(4), 349–359 (1999)
8. Biedl, T., Durocher, S., Snoeyink, J.: Reconstructing polygons from scanner data. *Theoretical Computer Science* **412**(32), 4161–4172 (2011)
9. Boissonnat, J.D.: Geometric structures for three-dimensional shape representation. *ACM Transactions on Graphics (TOG)* **3**(4), 266–286 (1984)
10. Boissonnat, J.D., Oudot, S.: Provably good sampling and meshing of surfaces. *Graphical Models* **67**(5), 405–451 (2005)
11. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
12. Chen, Z., Tagliasacchi, A., Zhang, H.: Bsp-net: Generating compact meshes via binary space partitioning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 45–54 (2020)
13. Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5939–5948 (2019)
14. Crane, K., Weischedel, C., Wardetzky, M.: The heat method for distance computation. *Commun. ACM* **60**(11), 90–99 (Oct 2017)
15. Curless, B., Levoy, M.: A volumetric method for building complex models from range images. In: *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*. pp. 303–312. *Proc. SIGGRAPH* (1996)
16. Dai, A., Nießner, M.: Scan2mesh: From unstructured range scans to 3d meshes. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5574–5583 (2019)
17. Deng, B., Genova, K., Yazdani, S., Bouaziz, S., Hinton, G., Tagliasacchi, A.: Cvxnet: Learnable convex decomposition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 31–44 (2020)
18. Dey, T.K.: *Curve and surface reconstruction: algorithms with mathematical analysis*, vol. 23. Cambridge University Press (2006)
19. Edelsbrunner, H., Mücke, E.P.: Three-dimensional alpha shapes. *ACM Transactions on Graphics (TOG)* **13**(1), 43–72 (1994)
20. Genova, K., Cole, F., Vlastic, D., Sarna, A., Freeman, W.T., Funkhouser, T.: Learning shape templates with structured implicit functions. arXiv preprint arXiv:1904.06447 (2019)

21. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 580–587 (2014)
22. Gkioxari, G., Malik, J., Johnson, J.: Mesh r-cnn. arXiv preprint arXiv:1906.02739 (2019)
23. Groueix, T., Fisher, M., Kim, V.G., Russell, B.C., Aubry, M.: A papier-mâché approach to learning 3d surface generation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 216–224 (2018)
24. Hackel, T., Savinov, N., Ladicky, L., Wegner, J.D., Schindler, K., Pollefeys, M.: Semantic3d. net: A new large-scale point cloud classification benchmark. arXiv preprint arXiv:1704.03847 (2017)
25. Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W.: Surface reconstruction from unorganized points, vol. 26. ACM (1992)
26. Kanazawa, A., Tulsiani, S., Efros, A.A., Malik, J.: Learning category-specific mesh reconstruction from image collections. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 371–386 (2018)
27. Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. In: Proceedings of the fourth Eurographics symposium on Geometry processing. vol. 7 (2006)
28. Kazhdan, M., Hoppe, H.: Screened poisson surface reconstruction. *ACM Transactions on Graphics (ToG)* **32**(3), 1–13 (2013)
29. Khatamian, A., Arabnia, H.R.: Survey on 3d surface reconstruction. *Journal of Information Processing Systems* **12**(3) (2016)
30. Kolluri, R., Shewchuk, J.R., O’Brien, J.F.: Spectral surface reconstruction from noisy point clouds. In: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing. pp. 11–21. ACM (2004)
31. Liao, Y., Donne, S., Geiger, A.: Deep marching cubes: Learning explicit surface representations. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2916–2925 (2018)
32. Lin, C.H., Wang, O., Russell, B.C., Shechtman, E., Kim, V.G., Fisher, M., Lucey, S.: Photometric mesh optimization for video-aligned 3d object reconstruction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 969–978 (2019)
33. Lipman, Y., Sorkine, O., Cohen-Or, D., Levin, D., Rossi, C., Seidel, H.P.: Differential coordinates for interactive mesh editing. In: Proceedings Shape Modeling Applications, 2004. pp. 181–190. IEEE (2004)
34. Litany, O., Bronstein, A., Bronstein, M., Makadia, A.: Deformable shape completion with graph convolutional autoencoders. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1886–1895 (2018)
35. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4460–4470 (2019)
36. Mullen, P., De Goes, F., Desbrun, M., Cohen-Steiner, D., Alliez, P.: Signing the unsigned: Robust surface reconstruction from raw pointsets. In: *Computer Graphics Forum*. vol. 29, pp. 1733–1741. Wiley Online Library (2010)
37. Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. In: Proceedings of the 27th international conference on machine learning (ICML-10). pp. 807–814 (2010)
38. Newman, T.S., Yi, H.: A survey of the marching cubes algorithm. *Computers & Graphics* **30**(5), 854–879 (2006)

39. Pan, J., Han, X., Chen, W., Tang, J., Jia, K.: Deep mesh reconstruction from single rgb images via topology modification networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 9964–9973 (2019)
40. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 165–174 (2019)
41. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 652–660 (2017)
42. Rakotosaona, M.J., La Barbera, V., Guerrero, P., Mitra, N.J., Ovsjanikov, M.: Pointcleannet: Learning to denoise and remove outliers from dense point clouds. In: Computer Graphics Forum. Wiley Online Library (2019)
43. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in neural information processing systems. pp. 91–99 (2015)
44. Shewchuk, J., Dey, T.K., Cheng, S.W.: Delaunay mesh generation. Chapman and Hall/CRC (2016)
45. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* **15**(1), 1929–1958 (2014)
46. Treece, G.M., Prager, R.W., Gee, A.H.: Regularised marching tetrahedra: improved iso-surface extraction. *Computers & Graphics* **23**(4), 583–598 (1999)
47. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2mesh: Generating 3d mesh models from single rgb images. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 52–67 (2018)
48. Williams, F., Schneider, T., Silva, C., Zorin, D., Bruna, J., Panozzo, D.: Deep geometric prior for surface reconstruction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10130–10139 (2019)