



**HAL**  
open science

# PoNQ: a Neural QEM-based Mesh Representation

Nissim Maruani, Maks Ovsjanikov, Pierre Alliez, Mathieu Desbrun

► **To cite this version:**

Nissim Maruani, Maks Ovsjanikov, Pierre Alliez, Mathieu Desbrun. PoNQ: a Neural QEM-based Mesh Representation. CVPR 2024 - IEEE/CVF Conference on Computer Vision and Pattern Recognition, Jun 2024, Seattle, United States. hal-04706248

**HAL Id: hal-04706248**

**<https://hal.science/hal-04706248v1>**

Submitted on 2 Dec 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# PoNQ: a Neural QEM-based Mesh Representation

Nissim Maruani

Inria, Université Côte d’Azur  
nissim.maruani@inria.fr

Maks Ovsjanikov

LIX, École Polytechnique, IP Paris  
maks@lix.polytechnique.fr

Pierre Alliez

Inria, Université Côte d’Azur  
pierre.alliez@inria.fr

Mathieu Desbrun

Inria Saclay - Ecole Polytechnique  
mathieu.desbrun@inria.fr

## Abstract

Although polygon meshes have been a standard representation in geometry processing, their irregular and combinatorial nature hinders their suitability for learning-based applications. In this work, we introduce a novel learnable mesh representation through a set of local 3D sample Points and their associated Normals and Quadric error metrics (QEM) w.r.t. the underlying shape, which we denote PoNQ. A global mesh is directly derived from PoNQ by efficiently leveraging the knowledge of the local quadric errors. Besides marking the first use of QEM within a neural shape representation, our contribution guarantees both topological and geometrical properties by ensuring that a PoNQ mesh does not self-intersect and is always the boundary of a volume. Notably, our representation does not rely on a regular grid, is supervised directly by the target surface alone, and also handles open surfaces with boundaries and/or sharp features. We demonstrate the efficacy of PoNQ through a learning-based mesh prediction from SDF grids and show that our method surpasses recent state-of-the-art techniques in terms of both surface and edge-based metrics.

## 1. Introduction

In recent years, learning-based methods have shown great promise as an efficient means to handle ill-posed shape processing tasks with complex priors [20, 23, 59]. Yet, despite the slew of advanced architectures to analyze or process 3D datasets, there is a lack of learnable 3D representations that can capture ridges and corners, while guaranteeing valid output meshes representing real 3D shapes for even the most basic learning tasks.

Early works on shape representation for learning predominantly relied on implicit volumetric representations [59, 60], which, while eventually yielding a mesh through extraction, posed significant challenges. Notably, their training can be costly due to the volumetric nature




Figure 1. **PoNQ representation.** Quadric error metric (QEM) matrices are fitted and visualized on a given shape (left). Note that their aspect ratios (green to blue) capture the underlying surface information: pancakes for flat regions, cigars for sharp edges and balls for corners. Our mesh extraction outputs a watertight (or, optionally, open, see bottom) and non-self-intersecting mesh that preserves salient features of the input shape (center) and is more concise and faithful than current implicit approaches (right).

of these representations. Moreover, the resulting mesh surfaces often lack detail, tending to generate overly smoothed shapes even where sharp features are expected. In response to these limitations, recent advancements have formulated explicit representations that effectively encode shapes through strategically chosen sample points and other geometric entities such as normals, from which more accurate meshes can be derived [35, 41]. While these representations offer higher generalizability and greater control over the generated surfaces, they often cannot guarantee important properties such as *watertightness* (i.e., the mesh must be the boundary of a non-degenerated volume), or the *absence of*

*self-intersections* (i.e., the surface is not just immersed, but embedded in  $\mathbb{R}^3$ ). Only very recently an approach has been proposed that enforces these properties [35]. However, it can suffer from overly complex output meshes (presence of many spurious polygonal facets) and, at times, wrong occupancy guesses create dents on the resulting shapes.

In this work, we propose PoNQ (short for point-normal-QEM), a new learnable 3D representation which encodes a shape through discrete points and other *local* geometric quantities to ensure efficient training and sharp reconstructions. As we demonstrate, a *global* mesh can be extracted from our PoNQ representation through a robust approach that leverages the available geometric data so as to better capture ridges and thin structures. Our key contributions are as follows:

- our neural representation is the first to exploit the quadric error metric (QEM), which has been instrumental in classical geometry processing tasks; consequently, PoNQ excels at capturing sharp features and boundaries, preserving the intricate details that are often lost in existing representations (see Fig. 1);
- PoNQ meshes are guaranteed to be watertight and free of self-intersections, thus broadening their applicability and utility in downstream applications;
- our PoNQ representation can also easily reduce the element count of the output shapes while preserving sharp features due to its reliance on QEM;
- finally, our neural representation outperforms state-of-the-art methods in mesh reconstruction from SDF grids as measured with both surface and edge-based metrics.

## 2. Related Work

We begin with a review of related works, including popular methods in shape learning, but also covering relevant parts of shape reconstruction and shape simplification.

**Shape Reconstruction.** The literature abounds with shape reconstruction methods (typically from input pointsets), each distinguished by their unique priors – necessary to regularize the ill-posed nature of the task – and output representations. In so-called continuous approaches, the reconstructed shape is the fixed-point of a projection operator [2], an algebraic point set surface (APSS [18]), or the isolevel of a 3D distance-like scalar field [26, 27], to mention a few popular approaches. Approaches from the latter category (often referred to as implicit methods) further require a mesh extraction from the scalar field, denoted isosurfacing, typically through Marching Cubes (MC [33], or its improved neural variant NMC [10]), Dual Contouring (DC [25], or its neural variant NDC [8]), or through the recently-proposed Reach-for-the-Spheres (RTS) technique [46] that leverages geometric properties of a signed distance field to improve isosurfacing.

Combinatorial methods, popular in Computational Geometry, rely instead on 3D Delaunay triangulations or their duals, Voronoi diagrams. The popular Crust approach, for instance, proceeds by local filtering of (facets of) 3D Delaunay triangulations [3]. However, it requires dense point samples to correctly capture thin features, and cannot handle sharp edges. Its Powercrust [4] and Tight Cocone [13] variants add filtering of the triangulation to improve various properties, like the capture of holes in the reconstructed surface, but they share the same limitations. A global graph-cut extraction of the faces of the Delaunay triangulation of the input samples [29] was also shown to be more robust to noisy inputs, just like a recent Delaunay-based approach that alternates between filmsticking and sculpting over a 3D triangulation [56].

**Implicit Field Learning.** Most pioneering 3D learning methods [36, 39, 42] represented a shape as a scalar field (very often, a signed distance field in fact), later converted into a mesh via Marching Cubes [33] or directly rendered. These methods need to process the global shape, limiting their generalizability beyond ShapeNet [7]. More recent implicit-based approaches [16, 49] showed that neural networks can be overfitted to single-shape Signed Distance Fields (SDF), and as such, they cannot learn from multiple shapes. Adding point-based convolutions [6], hash tables [37], octree structures [32, 52, 57] or kernel methods [22] have also been proposed to further refine these implicit methods, while the use of unsigned distance fields was leveraged to represent surfaces with boundaries [11, 19, 61]. Although using continuous fields provides topological guarantees, the lack of control of the final locations of mesh vertices yields a distinctively “blobby” aspect that smooths out sharp features and small details. Furthermore, implicit fields are trained over the whole *volume* around and inside the shape, while most applications are only interested in the reconstructed *surface*.

**Explicit Shape Learning.** To overcome the limitations of implicit fields and offer better control over the elements of the output mesh, end-to-end differential isosurfacing approaches were proposed [31, 45, 48]. Yet these approaches still rely on regular grids and often do not guarantee intersection-free output meshes. Methods that best fit a set of canonical geometric primitives [9, 12, 58] or an explicit mesh [38] to an input shape manage to produce concise meshes and preserve sharp features if these primitives (planes, quadric patches, etc) are diverse enough to describe the input mesh. However, they are rapidly compute-intensive even for a limited number of primitives. Deforming template shapes [17, 51] or regular tetrahedral meshes [14, 47] often constrains the output mesh genus and connectivity, and cannot guarantee intersection-free meshes. Another line of contributions approaches the design of learnable 3D representations from a different stand-

point: provided a known reconstruction algorithm, what *geometric estimates* should the neural network predict in order to extract a mesh off of it? Methods relying upon the popular Poisson surface reconstruction algorithm to extract a mesh [41, 43] for instance inherit the topological advantages of an implicit surface representation, but they also suffer from their limitations by exhibiting overly-smoothed sharp features and a large amount of mesh elements. The Neural Marching Cubes (NMC [10]) and Neural Dual Contouring (NDC [8]) methods reduce some of these limitations, at the price of losing topological guarantees. Recently, VoronoiMesh [35], computes a Voronoi diagram so that a subset of the 3D Voronoi facets best fit the input mesh, guaranteeing watertight results; but VoronoiMeshes are littered with small faces with spurious normal orientations, causing both visual artifacts and large mesh element counts. A few learning methods rely on a Delaunay triangulation instead [34, 44, 50, 62], but they are limited to the case of a fixed input pointset. We recap in Tab. 1 some typical properties of explicit learnable 3D representations.

	Arbitrary connectivity	Sharp features	Watertight, no self-int.	Open surfaces
NDC [8]		✓		✓
DMT [47]		✓		
SAP [41]			✓	
DPF [43]	✓		✓	
VoronoiMesh [35]	✓	✓	✓	
PoNQ	✓	✓	✓	✓

Table 1. Properties of explicit learnable 3D representations

**Quadric Error Metrics.** Finally, we review a workhorse of surface approximation. In the context of mesh decimation, Garland and Heckbert [15] introduced a Quadric Error Metric (QEM), encoded as a  $4 \times 4$  matrix  $Q$ , which evaluates the sum of squared distances to a (possibly large) set of tangent planes. The use of QEM for their application was particularly appropriate: the sum  $Q$  of two matrices  $Q_i$  and  $Q_j$  represent the sum of the squared distances to the union of the tangent planes used for  $Q_i$  and those used for  $Q_j$  – a very economical proxy to all these tangent planes. Consequently, they proposed to use these QEM matrices to perform an efficient mesh decimation: the updated matrices encode the squared distances of all nearby removed facets, thus offering a very fast and low-memory evaluation of the distance to the original mesh throughout the decimation. After a few rounds of simplifications, the  $\epsilon$ -isovalue of QEM error near a vertex  $v_i$  (i.e., the set of 3D positions  $x$  such that  $[x, 1]^t Q_i [x, 1] = \epsilon$ ) thus encodes the local region around  $v_i$  of the initial surface: it will look like a *pancake* if the original region was flat, an elongated ellipsoid if the region was a sharp feature (and the ellipsoid will precisely match the alignment of the sharp feature), or a sphere if the original region was near a corner. While QEM was adapted to deal with colors and texture coordinates [21], spherical distances [53], mesh filtering [30] and even variational mesh reconstruction from 3D pointsets [63], the au-

thors of [54] recently suggested a probabilistic version of QEM as a potential representation for learning-based tasks, while [1] used QEM as a loss function. As we will see next, we will use a QEM matrix per discrete sample point in our learnable *representation* for 3D shapes instead, to better learn the shape of local regions.

### 3. Method

We now delve into our PoNQ representation by first offering a more in-depth introduction to QEM and how we use it in Sec. 3.1. We then describe our representation in Sec. 3.2 along with its use in learning tasks, and finally introduce our PoNQ mesh extraction approach in Sec. 3.3.

#### 3.1. Motivation: Quadric Error Metrics

Given a sample point  $s_k \in \mathbb{R}^3$  and a normal  $n(s_k) \in \mathbb{R}^3$ , the signed distance from a given location  $x$  to the plane passing through  $s_k$  that is normal to  $n(s_k)$  is given simply as  $d_{s_k, n(s_k)}(x) = (x - s_k)^t n(s_k)$ . The squared distance can thus be rewritten as  $d_{s_k, n(s_k)}(x)^2 = x^t A_k x - 2b_k^t x + c_k$ , where  $3 \times 3$  matrix  $A_k$ , vector  $b_k$ , and scalar  $c_k$  are:

$$A_k = n(s_k)n(s_k)^t, \quad b_k = A_k s_k, \quad c_k = s_k^t A_k s_k. \quad (1)$$

The sum of squared distances to a *set of planes* thus yields:

$$\begin{aligned} \text{QEM}(x) &= \sum_k d_{s_k, n(s_k)}^2(x) \\ &= x^t \left( \underbrace{\sum_k A_k}_{:=A} \right) x - 2 \left( \underbrace{\sum_k b_k}_{:=b} \right)^t x + \left( \underbrace{\sum_k c_k}_{:=c} \right) \quad (2) \\ &= [x, 1]^t \left( \underbrace{\begin{array}{c|c} A & -b \\ \hline -b^t & c \end{array}}_{:=Q} \right) [x, 1]. \quad (3) \end{aligned}$$

The summed distances to a fine sampling  $\{s_k\}_k := S$  of a small surface patch can then be concisely encoded by the QEM matrix  $Q \in \mathbb{R}^{4 \times 4}$  from Eq. (3). Moreover,  $Q$  can also indicate the best possible position to represent this surface patch with a single vertex: assuming the submatrix  $A$  to be invertible (which is true in the general case [54]), the quadratic QEM function reaches its unique minimum at a location  $v^* = A^{-1}b$ . We leverage these quantities as part of a powerful 3D *representation* suitable for machine learning.

#### 3.2. PoNQ representation

The PoNQ representation consists in a set  $\mathbf{P} = \{\mathbf{p}_i \in \mathbb{R}^3\}$  of points, augmented with their local normals  $\mathbf{N} = \{\mathbf{n}_i \in \mathbb{R}^3\}$  and quadrics  $\mathbf{Q} = \{\mathbf{Q}_i \in \mathbb{R}^{4 \times 4}\}$  – hence its name. In order to remove possible ambiguities, we will use a bold font to refer to the different quantities involved in the PoNQ representation, i.e., what will be *optimized* or *learned*.

### 3.2.1 QEM-based representation via optimization

In a pure optimization-based setting, we want PoNQ to fit a watertight and non-self-intersecting input shape finely discretized by samples  $s_k \in S$  with their local normals  $n(s_k)$ . We initialize the points  $\mathbf{P}$  on a regular grid around the input shape (other initializations work equally well), and optimize their position to minimize the bi-directional Chamfer Distance CD:

$$\text{CD}(\mathbf{P}, S) = \frac{1}{|\mathbf{P}|} \sum_{\mathbf{p}_i \in \mathbf{P}} \min_{s_k \in S} \|\mathbf{p}_i - s_k\|^2 \quad (4)$$

$$+ \frac{1}{|S|} \sum_{s_k \in S} \min_{\mathbf{p}_i \in \mathbf{P}} \|\mathbf{p}_i - s_k\|^2, \quad (5)$$

The resulting point set will thus lie close to (and spread out over) the target surface. These points define a partition of  $\mathbb{R}^3$  into Voronoi cells  $V(\mathbf{p}_i)$  for which any location  $x \in V(\mathbf{p}_i)$  has  $\mathbf{p}_i$  as its closest point from  $\mathbf{P}$ , see [5].

Once the positions are optimized, we enrich each point  $\mathbf{p}_i \in \mathbf{P}$  with a normal  $\mathbf{n}_i$  and a QEM matrix  $\mathbf{Q}_i$ , where  $\mathbf{n}_i$  represents the average of the sample normals  $n(s_k)$  for all the samples  $s_k$  contained within  $V(\mathbf{p}_i)$  (see Fig. 2), i.e.,

$$\mathbf{n}_i = \frac{1}{|S \cap V(\mathbf{p}_i)|} \sum_{s_k \in S \cap V(\mathbf{p}_i)} n(s_k). \quad (6)$$

Similarly,  $\mathbf{Q}_i$  is assembled using Eqs. (2) and (3) using the tangent planes implied by each sample (and its normal) within  $V(\mathbf{p}_i)$ , i.e.,  $\mathbf{Q}_i$  is the QEM matrix such that

$$[x, 1]^t \mathbf{Q}_i [x, 1] = \sum_{s_k \in S \cap V(\mathbf{p}_i)} d_{s_k, n(s_k)}^2(x). \quad (7)$$

Note that these additional variables are thus proxies for the local geometry around each point  $\mathbf{p}_i$ ; points with no input




Figure 2. 2D illustration of PoNQ: (a) a sampled ground-truth shape  $S$  (blue dots) is represented by PoNQ as points  $\mathbf{p}_i$  (whose Voronoi diagram (dotted lines) partitions the input samples), along with normals  $\mathbf{n}_i$  and quadrics  $\mathbf{Q}_i$  encoding the shape within each Voronoi cell. (b) The PoNQ mesh (black solid lines) is the boundary of the union of labeled tetrahedra from the Delaunay triangulation of the QEM-optimal vertices, providing a better fit than simply interpolating the points (black dotted lines).

samples within their Voronoi cell are simply discarded. As explained in Section 2, each matrix  $\mathbf{Q}_i$  implies, through its sub-matrices  $\mathbf{A}_i$  and  $\mathbf{b}_i$ , an optimal location  $\mathbf{v}^* := \mathbf{A}_i^{-1} \mathbf{b}_i$  with respect to the input surface, hinting at the fact that storing  $\mathbf{v}^*$  instead of  $\mathbf{b}_i$  is a possible alternative, which we will use in the learning context in the next section. We then rely on  $\mathbf{Q}$  and  $\mathbf{N}$  to extract the connectivity of the optimal vertices  $\mathbf{v}^*$ , which will be explained in Sec. 3.3.

### 3.2.2 Learning with PoNQ




Figure 3. Overview of our learning pipeline with PoNQ.

We demonstrate the benefits of PoNQ in a learning context by applying it to reconstruction from Signed Distance Fields (SDF) (see Fig. 3). As shown in several recent works [8, 10, 35], this task is especially interesting as it enables the training of a *local* model that can truly generalize to novel shapes outside the training set.

We use an architecture similar to NDC [8], consisting of several convolutional neural networks (CNN). Unlike NMC [10] (which uses eight points per voxel) and NDC (which uses only one), we can use an arbitrary number  $P$  of predicted points per cell since our representation does not depend on a regular grid — in practice, we found that  $P = 4$  can represent sub-voxel details nicely without requiring too large networks. We use a shared 5-layer encoder that converts the  $(N+1)^3$  input SDF grid into a  $N^3$ -sized grid of 128 features. These features are processed by five separate 6-layer decoders (which do not share weights) to predict a PoNQ, i.e.,  $P \times N^3$  points  $\mathbf{p}_i$ , along with their associated local normals  $\mathbf{n}_i$  and QEM matrices  $\mathbf{Q}_i$ , to which we add a set  $\mathbf{O}$  of  $N^3$  binary “occupancies”  $\mathbf{o}_i$  to mark the voxels containing the surface at inference time.

To stabilize the learning process and avoid having to perform matrix inversions, we do not store QEM matrices directly but only use the quadratic form  $\mathbf{A}_i$  and the QEM-optimized vertex location  $\mathbf{v}^*$  instead, from which one can reconstruct  $\mathbf{Q}_i = \begin{pmatrix} \mathbf{A}_i & -\mathbf{A}_i \mathbf{v}_i^* \\ -(\mathbf{A}_i \mathbf{v}_i^*)^t & 0 \end{pmatrix}$ . Furthermore, we store instead of  $\mathbf{A}_i$  a  $3 \times 3$  upper triangular matrix  $\mathbf{U}_i$  such that  $\mathbf{A}_i = \mathbf{U}_i \mathbf{U}_i^t$  is the reversed Cholesky decomposition of  $\mathbf{A}_i$ , guaranteeing that  $\mathbf{A}_i$  remains invertible.

We supervise our training with a collection of watertight shapes, each converted into a dense sampling  $S$ . We pre-process the data by computing the ground-truth occupancy set  $O_{gt}$  of the  $m$  voxels containing samples, to which we restrict our loss terms (except for  $L_{occ}$ , which checks how well the trained occupancy set matches the ground-truth one). At

training time, we apply the sum of the following losses and backpropagate their gradients w.r.t. the bold variables:

$$L_{CD} = CD(\mathbf{P}, S) \quad (\text{see Eq. (5)}) \quad (8)$$

$$L_{\mathbf{n}} = \frac{1}{m} \sum_{i=1}^m \sum_{s_k \in S \cap V(\mathbf{p}_i)} \|\mathbf{n}_i - n(s_k)\|^2 \quad (9)$$

$$L_{\mathbf{A}} = \frac{1}{m} \sum_{i=1}^m \sum_{s_k \in S \cap V(\mathbf{p}_i)} \|\mathbf{U}_i \mathbf{U}_i^t - n(s_k) n(s_k)^t\|^2 \quad (10)$$

$$L_{\mathbf{v}^*} = \frac{1}{m} \sum_{i=1}^m \sum_{s_k \in S \cap V(\mathbf{p}_i)} (n(s_k)^t (\mathbf{v}_i^* - s_k))^2 \quad (11)$$

$$L_{\text{reg}} = \frac{1}{m} \sum_{i=1}^m \|\mathbf{v}_i^* - \mathbf{p}_i\|^2 \quad (\mathbf{p}_i \text{ assumed fixed here}) \quad (12)$$

$$L_{\text{occ}} = \|\mathbf{O} - O_{\text{gt}}\|^2 \quad (13)$$

In the order in which they are listed above, these losses were designed to: help spread around the pointset  $\mathbf{P}$  and best fit the inputs; enforce that each normal is the mean normal of the local sample normals; enforce that the quadratic forms  $\mathbf{A}_i$  correspond to the proper submatrix of the QEM of the local samples and normals around point  $\mathbf{p}_i$ ; enforce, similarly, that  $\mathbf{v}_i^*$  minimizes the sum of squared distance to the local tangent planes around point  $\mathbf{p}_i$ ; regularize (with a very small coefficient) the positions of the optimal points  $\mathbf{v}_i^*$  in flat regions — as in this case, any point on this flat region is optimal in theory, so we force it to stay close to  $\mathbf{p}_i$ ; and make sure that we match the ground-truth occupancy set. See §2.3 of the supplementary material for ablation studies.

### 3.3. Meshing our representation

Given a PoNQ representation (either optimized or produced by a trained network), one can easily extract a mesh by combining two approaches from computational geometry [3, 29] to ensure robustness. Note that we do not even use the pointset  $\mathbf{P}$  which only served to build a shape-adapted partition: we construct a mesh whose vertices are the QEM-optimal positions  $\mathbf{v}_i^*$  as they capture features best. We present here a concise overview of our meshing method; see our supplemental material for details.

**Pre-processing** We first compute the Delaunay tetrahedralization of the optimal vertices  $\mathbf{v}_i^*$  deriving from the QEM matrices  $\mathbf{Q}$ . The next two steps will tag each tetrahedron as either *inside* or *outside* based on local information, so that *our final PoNQ mesh will be simply the triangle mesh forming the inside/outside boundary* — ensuring watertightness and no self-intersections by design.

**Tagging obvious inside/outside tetrahedra.** We leverage the circumcenter criterion put forth in the Crust algorithm [3]. In our case, each vertex  $\mathbf{v}_i^*$  and its assigned normal  $\mathbf{n}_i$  define an oriented plane: we tag a tetrahedron as *outside* (resp., *inside*) if both its circumcenter and barycenter



Figure 4. Optimization-based results ( $32^3$ ).

are determined to be in the *outside* (resp., *inside*) half-space of each of the four vertices. Considering that the shape is contained within the convex hull of the  $\mathbf{v}_i^*$  and that each vertex of the Delaunay tetrahedralization must be part of the final surface allows us to further tag a series of tetrahedra where there is no ambiguity; see our supplemental material for a detailed rationale.

**Tagging remaining tetrahedra.** Delaunay-based meshing approach (like *Crust* [3]) require a dense point sampling (formally, an  $\epsilon$ -sampling), which is not compatible with our desire to deal with thin structures, sharp features and corners — and this is the main reason why our earlier phase can end up not providing a tag for *every* tetrahedra. To finish our tetrahedron tagging based on the ones we already have, we use a graph cut approach, inspired by (but simpler than) an existing spectral graph partitioning [29]. For each Delaunay triangle  $T$ , we compute a likelihood score  $S(T)$  that evaluates how confident we are that this triangle is to appear on the final output mesh: this triangle score evaluates the fitness of  $T$  based on the local normals  $\mathbf{N}$  and quadric matrices  $\mathbf{Q}$  as explained in the supplemental material. We now tag the remaining undetermined tetrahedra through a *minimum cut* of the Voronoi graph (in which each dual of a tetrahedron is a node, and each dual of a Delaunay triangle  $T$  is an edge with weight  $S(T)$ ) using the already-tagged “inside” ones as a source and the “outside” one as a drain.

**Surface extraction.** We extract the final *PoNQ mesh* as the triangle mesh forming the boundary between the inside and outside tetrahedra (see Fig. 2). As mentioned above, this automatically enforces by design the fact that our mesh is watertight and intersection-free.

## 4. Experimental Results

We implemented PoNQ for both optimization and learning tasks in Python, using PyTorch [40], SciPy [55] and libigl [24] — our code is available on [our project page](#). All timings were computed on a single workstation with 54 cores, a NVidia A6000 GPU and 512 GB of RAM.

## 4.1. Optimization-based 3D Reconstruction

For our tests in optimization-based 3D reconstruction, we use the 30 watertight shapes from the Thingi10k [64] dataset chosen in VoroMesh [35] as it is arguably the most related and most recent neural representation with which to compare. We consider three grid resolutions  $\text{res} \in [32^3, 64^3, 128^3]$ , and sample  $1024 \times \text{res}^{2/3}$  surface samples. As described in Sec. 3.2.1, we use the chamfer distance as a loss to optimize the points for 400 epochs with the Adam optimizer before computing the mean normals and quadrics with our GPU-based implementation. We then extract the PoNQ representation as explained in Sec. 3.3.

**Baselines.** Besides VoroMesh, we also compare our results with Shape As Points (SAP) [41] and Dynamic Point Field (DPF) [43] as these three point-based neural representations all guarantee watertight outputs. We add a variant of DPF, trained with the Chamfer distance only (i.e., without the image-based loss), which we denote  $\text{DPF}_{\text{chamfer}}$ . For each method, we use the same number of optimized points.

**Metrics.** We use the most common surface-based metrics, i.e., chamfer distance (CD), F-score (F1, with a threshold of 0.003), and normal consistency (NC). In order to assess the reconstruction quality of sharp edges, we sample  $10^5$  points on the edges featuring a dihedral angle larger than  $\frac{\pi}{6}$ , and compute the edge-chamfer distance (ECD) and edge-F-score (EF1, with a threshold of 0.005) between the ground truth and the reconstruction samples (see supplemental material for details). We also report the number of triangles and faces of the extracted 3D models. Finally, we provide timings of the optimization step for all methods, as it is systematically the most time-consuming phase; see supplemental material for additional timings.

**Results.** As SAP and DPF both rely on Poisson surface reconstruction, they smooth out sharp edges and small details. DPF does not optimize the extracted field, but rather leverages local information from which the mesh is assembled without supervision, leading to faster convergence and better scores than SAP for all metrics. As the evaluation

Method	Grid Size	CD ↓ ( $\times 10^{-5}$ )	F1 ↑	NC ↑	ECD ↓ ( $\times 10^3$ )	EF1 ↑ ( $\times 10^3$ )	# V ( $\times 10^3$ )	# F ( $\times 10^3$ )	Time (s)
SAP [41],	$32^3$	6.475	0.589	0.894	0.235	0.058	1.6	3.2	51.2
DPF [43]	$32^3$	2.256	0.724	0.935	0.147	0.115	5.8	11.6	45.5
$\text{DPF}_{\text{chamfer}}$ [43]	$32^3$	2.077	0.717	0.933	0.160	0.104	5.7	11.3	2.5
VoroMesh [35]	$32^3$	<b>0.802</b>	<b>0.919</b>	0.957	0.257	0.242	5.9	11.8	2.0
PoNQ	$32^3$	0.972	0.892	<b>0.961</b>	<b>0.106</b>	<b>0.447</b>	2.3	4.6	2.6
SAP [41],	$64^3$	1.912	0.858	0.949	0.119	0.267	7.0	13.9	106.3
DPF [43]	$64^3$	0.795	0.909	0.971	0.104	0.435	22.6	45.1	51.5
$\text{DPF}_{\text{chamfer}}$ [43]	$64^3$	0.797	0.907	0.970	0.094	0.415	24.1	48.3	4.0
VoroMesh [35]	$64^3$	<b>0.645</b>	<b>0.938</b>	0.975	0.251	0.249	23.4	46.9	4.1
PoNQ	$64^3$	0.655	0.936	<b>0.981</b>	<b>0.061</b>	<b>0.645</b>	10.1	20.1	4.1
SAP [41],	$128^3$	0.671	0.934	0.978	0.060	0.619	28.8	57.7	191.3
DPF [43]	$128^3$	0.644	0.938	0.986	0.089	0.665	90.0	180.0	71.0
$\text{DPF}_{\text{chamfer}}$ [43]	$128^3$	0.644	0.938	0.986	0.086	0.664	97.3	194.5	17.7
VoroMesh [35]	$128^3$	<b>0.634</b>	<b>0.939</b>	0.982	0.264	0.213	91.3	182.6	36.3
PoNQ	$128^3$	0.637	<b>0.939</b>	<b>0.988</b>	<b>0.039</b>	<b>0.795</b>	42.3	84.6	17.9

Table 2. Optimization-based results. Quantitative comparisons of Chamfer distance (CD), F1 score, and normal consistency (NC) on the Thingi30 dataset for three different grid resolutions.

point cloud is sampled from the reconstructed Poisson mesh (rather than the predicted points), the image-based loss does not significantly improve results, and  $\text{DPF}_{\text{chamfer}}$  is faster than DPF with comparable performance.

VoroMesh exhibits finer details and better sharpness overall. With the same number of points as SAP, DPF and PoNQ, it has the best overall surface fitting scores due to the large number of Voronoi vertices and faces. However many of these faces have spurious normals and create undesired sharp edges, thus impacting the NC, ECD, and EF1 scores.

Qualitatively, PoNQ is better at dealing with sharp features than SAP and DPF and does not generate surface artifacts like VoroMesh (see Figure 4). Quantitatively, our method yields the best normal consistency and sharp edge fitting scores (see Tab. 2). In terms of surface fitting, it matches the considered baselines at resolution  $128^3$ , and comes close second behind VoroMesh at resolutions  $32^3$  and  $64^3$ , but with significantly lower face counts. We will discuss in Sec. 4.3 how one can further lower the face count of PoNQ meshes at very little cost on the scores (see Fig. 7), setting PoNQ apart even more prominently.

## 4.2. Learning-based 3D reconstruction

### 4.2.1 Reconstruction from SDF

We now assess the behavior of our PoNQ representation in the learning-based task of 3D shape reconstruction from SDF grids. We train and evaluate our method on the CAD shapes of the ABC dataset [28]. We also assess the generalizability of our method on the free-form shapes of the Thingi10k [64] dataset, without any fine-tuning. For fairness, we use the train/test split provided in VoroMesh [35]: 3,843/962 in the training and testing set for ABC, with 30 watertight validation shapes for Thingi10k. We train our network for 600 epochs while increasing the number of sampled points and decreasing the learning rate and regularization – see supplemental material for additional details.

**Baselines.** We compare our method against the three most closely related baselines: NMC [10], NDC [8] and VoroMesh [35] using the authors’ code and their best pre-trained model (we also discuss comparisons with RTS [46] and DMTet [47] in the supplementary material).

**Metrics.** We use the same evaluation metrics that were already mentioned in Sec. 4.1.

**Results.** Our key results are reported in Tabs. 3 and 4, where PoNQ outperforms state-of-the-art methods *on every resolution, dataset, and metric while guaranteeing watertight output meshes that are devoid of self-intersections*. We note that NMC and NDC both rely on a regular-grid based meshing. As a result, they fail to capture thin structures and exhibit aliasing artifacts (see Figs. 5 and 6), which impacts their surface, normal and sharp-edge fitting scores.

VoroMesh does not rely on a regular grid, so its ability to capture thin surfaces leads to better results than NMC




Figure 5. Learning results (top:  $32^3$ ; bottom:  $64^3$ ) on ABC.

Method	Grid Size	CD ↓ ( $\times 10^{-5}$ )	F1 ↑	NC ↑	ECD ↓	EF1 ↑	Watertight ↑ no self-int.	# V ( $\times 10^3$ )	# F ( $\times 10^3$ )
NDC [8]	$32^3$	66.004	0.787	0.941	0.445	0.658	44%	1.3	2.6
NMC [10]	$32^3$	60.755	0.833	0.954	0.350	0.693	26%	9.7	19.3
VoroMesh [35]	$32^3$	2.228	0.835	0.941	0.802	0.232	100%	10.0	20.0
PoNQ-lite	$32^3$	3.539	0.810	0.953	0.296	0.658	100%	1.3	2.6
PoNQ	$32^3$	<b>1.514</b>	<b>0.852</b>	<b>0.964</b>	<b>0.184</b>	<b>0.713</b>	100%	5.1	10.2
NDC [8]	$64^3$	2.211	0.882	0.975	0.223	0.855	23%	5.5	11.0
NMC [10]	$64^3$	2.138	0.891	<b>0.980</b>	0.254	0.854	18%	42.8	85.5
VoroMesh [35]	$64^3$	1.219	0.886	0.966	0.796	0.207	100%	38.4	76.9
PoNQ-lite	$64^3$	1.074	0.888	0.978	0.128	0.858	100%	5.5	10.9
PoNQ	$64^3$	<b>0.886</b>	<b>0.892</b>	<b>0.980</b>	<b>0.109</b>	<b>0.866</b>	100%	21.2	42.3
NDC [8]	$128^3$	1.889	<b>0.896</b>	0.983	0.095	<b>0.947</b>	14%	22.1	44.2
NMC [10]	$128^3$	1.888	<b>0.896</b>	<b>0.984</b>	0.349	0.859	12%	175.9	351.9
VoroMesh [35]	$128^3$	1.069	0.894	0.974	0.792	0.189	100%	149.1	298.2
PoNQ-lite	$128^3$	1.007	<b>0.896</b>	<b>0.984</b>	<b>0.043</b>	0.933	100%	21.9	43.8
PoNQ	$128^3$	<b>0.920</b>	<b>0.896</b>	<b>0.984</b>	0.191	0.878	100%	85.7	171.2

Table 3. Results on ABC with our network trained on ABC.

and NDC for the complex ABC dataset or on low-resolution models of Thingi30. However, as noticed in Tabs. 3 and 4, it suffers from local surface artifacts leading to sharp edges and spurious normals, resulting in worst NC, ECD and EF1. Moreover, its two-stage training implies that its encoder is not trained for occupancy prediction; as a result, it can mis-label Voronoi generators, leading to either missing parts or floating volumes around the shapes (see Figs. 5 and 6).

In contrast, our method is able to capture fine details due to our use of QEM, which helps to capture surface characteristics, and it does not visually exhibit any of the artifacts of other approaches due to our PoNQ mesh being extracted from a Delaunay tetrahedralization. Ultimately, PoNQ leads to more faithful reconstructions while guaranteeing 100% watertight and intersection-free results.

### 4.3. Additional extensions

We conclude this section with other examples leveraging the unique nature of our neural representation.

**Surfaces with boundary.** With minor modifications, PoNQ can output surfaces with boundaries as well (Fig. 1 and supplementary material). To optimize or train our representation for this case, we simply compute a *boundary sampling*  $S_b$  of the input surface boundaries, and duplicate it




Figure 6. Learning results (top:  $32^3$ ; bottom:  $128^3$ ) on Thingi30.

Method	Grid Size	CD ↓ ( $\times 10^{-5}$ )	F1 ↑	NC ↑	ECD ↓	EF1 ↑	Watertight ↑ no self-int.	# V ( $\times 10^3$ )	# F ( $\times 10^3$ )
NDC [8]	$32^3$	6.390	0.745	0.920	0.172	0.245	40%	1.3	2.6
NMC [10]	$32^3$	5.188	0.796	0.936	0.148	0.271	0%	8.7	17.3
VoroMesh [35]	$32^3$	2.825	0.758	0.902	0.263	0.156	100%	9.9	19.9
PoNQ-lite	$32^3$	1.705	0.754	0.934	0.154	0.270	100%	1.3	2.6
PoNQ	$32^3$	<b>1.344</b>	<b>0.810</b>	<b>0.942</b>	<b>0.137</b>	<b>0.314</b>	100%	4.8	9.7
NDC [8]	$64^3$	0.849	0.908	0.961	0.106	0.441	3%	5.4	10.8
NMC [10]	$64^3$	0.776	0.923	0.969	0.115	0.467	0%	36.8	73.6
VoroMesh [35]	$64^3$	1.021	0.906	0.939	0.259	0.192	100%	39.4	78.9
PoNQ-lite	$64^3$	0.769	0.914	0.968	<b>0.090</b>	0.495	100%	5.3	10.6
PoNQ	$64^3$	<b>0.758</b>	<b>0.924</b>	<b>0.971</b>	0.100	<b>0.511</b>	100%	19.9	39.9
NDC [8]	$128^3$	0.650	0.937	0.980	0.065	0.644	0%	22.0	44.1
NMC [10]	$128^3$	0.642	<b>0.939</b>	<b>0.984</b>	0.131	0.574	0%	151.7	303.3
VoroMesh [35]	$128^3$	0.731	0.932	0.959	0.260	0.198	100%	157.2	314.5
PoNQ-lite	$128^3$	0.644	0.938	<b>0.984</b>	<b>0.055</b>	<b>0.699</b>	100%	21.7	43.3
PoNQ	$128^3$	<b>0.641</b>	<b>0.939</b>	<b>0.984</b>	0.123	0.592	100%	80.8	161.8

Table 4. Results on Thingi30 with our network trained on ABC.

to create a sampling  $S'_b$  where we just rotate all the normals by  $\pi/2$  around the boundary. Changing  $S$  into  $S \cup S_b \cup S'_b$  will thus enforce that all the QEM matrices on the boundaries will generate elongated ellipsoids aligned with the local boundary. Meshing can proceed as before; but the output mesh will automatically close the holes. We further cull any triangle  $T$  of the extracted mesh for which the anisotropy of the QEM of each of its vertices is above 40%. In practice, we measure the anisotropy through the ratio  $r_i = \lambda_2/\lambda_1$  of the two largest eigenvalues of the matrix  $A_i$ . Obviously, the surface no longer bounds a volume since we cut holes in the original PoNQ mesh extraction, but it remains devoid of self-intersections. Note that this PoNQ variant cannot handle special cases such as a single square sheet.



**PoNQ-lite.** While using  $P=4$  predicted points per voxel yielded the best performances in our learning-based tests as mentioned in Sec. 3.2.2, we can trivially provide a “lite” version of our output PoNQ with a single point per cell using the same network: due to our reliance on QEM matrices, one can simply sum the  $P$  matrices  $\mathbf{Q}_i$  within a cell to directly create a single QEM matrix per cell, from which is derived a new optimal position  $\mathbf{v}_i^*$  for each cell. The normals  $\mathbf{n}_i$  are also trivially averaged into the new one. Sharp features are still well preserved due to our use of quadric error metrics (see Fig. 7), and these PoNQ-lite meshes in fact outperform NDC for an equal level of element count (see Tabs. 3 and 4), proving the superiority of grid-free methods. If even coarser meshes are desired, one can also construct, at nearly no cost, a whole hierarchy of PoNQ meshes by first applying PoNQ-lite, and then merging each group of eight cells into a single larger one to form a twice-coarser grid ( $2 \times 2$  average pool; see Fig. 8).




Figure 7. For a network trained on ABC with  $P=4$  predicted QEM matrices per cell (left), one can sum these matrices – i.e., via average pooling – per cell to produce a more compact mesh (right) that still naturally preserves the detected sharp features.

Although PoNQ outperforms PoNQ-lite on almost all metrics, the latter slightly pulls ahead on ECD and EF1 on Thing30 at  $128^3$  resolution. This opens the door to a series of further investigations, out of scope for this paper: one could potentially simplify a PoNQ output *adaptively* depending on the contents of the cell.

## 5. Discussion

The QEM matrices, encoded via points  $v_i^*$  and SPD matrices  $A_i$  are essential to our work: besides their role in capturing sharp features and second-order shape properties, they (a) disambiguate meshing compared to just point+normals to achieve state-of-the-art results (see supplementary material), (b) allow direct simplification through quadric collapses (PoNQ-lite, Fig. 7 & 8), and (c) allow open boundaries (Fig. 1). Yet, compared to SAP or VoroMesh, one may wonder if the added quadric information are worth a higher network size. In fact, producing QEM information does not

significantly affect network size since 80% of the weights are concentrated in the shared encoder. Moreover, the task of fitting points, QEM and normals is quite straightforward and does not require overly large networks (e.g., VoroMesh requires 8.4M parameters while PoNQ only requires 2.7M).

Another possible perceived limitation is that our PoNQ mesh extraction may create non-manifold vertices or edges despite being watertight – for instance, a 1-ring of a vertex may contain two *non-adjacent* tetrahedra both labeled as *inside*. However, one can simply duplicate these few non-manifold elements and connect them to their neighboring mesh elements to enforce manifoldness [35]. In addition, final results can be affected by two types of flaws: wrong estimates of geometric quantities (CNN failure, noise in the motor plate in Fig. 5), or tetrahedra mislabelling (reconstruction failure, unwanted link elbow-body Fig 6).

There are also exciting aspects of PoNQ we have not explored yet. For example, our meshing of open boundaries can potentially be further refined to extract directly the surface through a tagging of boundary edges, instead of relying on a potentially brittle final score-based filtering. A nice extension to our work would be to integrate the PoNQ representation in a differentiable rendering pipeline. Finally, the naturally multiscale nature of PoNQ through average pooling is also bound to be exploitable in a number of contexts.




Figure 8. Learning-based results. From PoNQ-lite ( $128^3$ ), to power-of-two simplifications down to  $16^3$  via average pooling.

## 6. Conclusion

We proposed a novel learnable 3D shape representation, coined PoNQ which combines the power of the quadric error metric (QEM) originally devised for mesh decimation and insights from computational geometry. PoNQ relies on points, normals, and QEM matrices to represent local geometric information, which are later leveraged to construct a triangle mesh that is guaranteed to be the boundary of a volume and devoid of self-intersections. We demonstrated the representation power of PoNQ through optimization-based tasks and learning-based reconstruction experiments, showing significant improvement upon previous 3D representations. We thus believe that PoNQ is poised to find many applications and extensions in neural shape processing.

**Acknowledgments.** Work supported by 3IA Côte d’Azur (ANR-19-P3IA-0002), ERC Starting Grant 758800 (EX-PROTEA), ERC Consolidator Grant 101087347 (VEGA), ANR AI Chair AIGRETTE, Ansys, Adobe Research, and a Choose France Inria chair.

## References

- [1] Nitin Agarwal, Sung-eui Yoon, and M. Gopi. Learning Embedding of 3D models with Quadric Loss, July 2019. arXiv:1907.10250 [cs].
- [2] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C.T. Silva. Point set surfaces. In *Proceedings Visualization, 2001. VIS '01.*, pages 21–537, San Diego, CA, USA, 2001. IEEE.
- [3] Nina Amenta, Marshall Bern, and Manolis Kamvyselis. A new Voronoi-based surface reconstruction algorithm. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques - SIGGRAPH '98*, pages 415–421, Not Known, 1998. ACM Press.
- [4] Nina Amenta, Sunghye Choi, and Ravi Krishna Kolluri. The power crust. In *Proceedings of the sixth ACM symposium on Solid modeling and applications, SMA '01*, pages 249–266, New York, NY, USA, May 2001. Association for Computing Machinery.
- [5] Franz Aurenhammer. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3):345–405, Sept. 1991.
- [6] Alexandre Boulch and Renaud Marlet. POCO: Point Convolution for Surface Reconstruction. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6292–6304, New Orleans, LA, USA, June 2022. IEEE.
- [7] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiang Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository, Dec. 2015. arXiv:1512.03012 [cs].
- [8] Zhiqin Chen, Andrea Tagliasacchi, Thomas Funkhouser, and Hao Zhang. Neural dual contouring. *ACM Transactions on Graphics*, 41(4):1–13, July 2022.
- [9] Zhiqin Chen, Andrea Tagliasacchi, and Hao Zhang. BSP-Net: Generating Compact Meshes via Binary Space Partitioning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 42–51, Seattle, WA, USA, June 2020. IEEE.
- [10] Zhiqin Chen and Hao Zhang. Neural marching cubes. *ACM Transactions on Graphics*, 40(6):1–15, Dec. 2021.
- [11] Thor V. Christiansen, Jakob Andreas Bærentzen, Rasmus R. Paulsen, and Morten R. Hannemose. Neural Representation of Open Surfaces. *Computer Graphics Forum*, 42(5):e14916, 2023. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14916>.
- [12] Boyang Deng, Kyle Genova, Soroosh Yazdani, Sofien Bouaziz, Geoffrey Hinton, and Andrea Tagliasacchi. CvxNet: Learnable Convex Decomposition. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 31–41, Seattle, WA, USA, June 2020. IEEE.
- [13] Tamal K. Dey and Samrat Goswami. Tight cocone: a watertight surface reconstructor. In *Proceedings of the eighth ACM symposium on Solid modeling and applications*, pages 127–134, Seattle Washington USA, June 2003. ACM.
- [14] Jun Gao, Wenzheng Chen, Tommy Xiang, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Learning Deformable Tetrahedral Meshes for 3D Reconstruction. In *Advances in Neural Information Processing Systems*, volume 33, pages 9936–9947. Curran Associates, Inc., 2020.
- [15] Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques - SIGGRAPH '97*, pages 209–216, Not Known, 1997. ACM Press.
- [16] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit Geometric Regularization for Learning Shapes. In *Proceedings of the 37th International Conference on Machine Learning*, pages 3789–3799. PMLR, Nov. 2020. ISSN: 2640-3498.
- [17] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. A Papier-Mache Approach to Learning 3D Surface Generation. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 216–224, Salt Lake City, UT, USA, June 2018. IEEE.
- [18] Gaël Guennebaud and Markus Gross. Algebraic point set surfaces. In *ACM SIGGRAPH 2007 papers*, page 23, San Diego California, July 2007. ACM.
- [19] Benoît Guillard, Federico Stella, and Pascal Fua. MeshUDF: Fast and Differentiable Meshing of Unsigned Distance Field Networks. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, Lecture Notes in Computer Science, pages 576–592, Cham, 2022. Springer Nature Switzerland.
- [20] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennamoun. Deep Learning for 3D Point Clouds: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(12):4338–4364, Dec. 2021.
- [21] Hugues Hoppe. New quadric metric for simplifying meshes with appearance attributes. In *Proceedings Visualization '99 (Cat. No.99CB37067)*, pages 59–510, San Francisco, CA, USA, 1999. IEEE.
- [22] Jiahui Huang, Zan Gojcic, Matan Atzmon, Or Litany, Sanja Fidler, and Francis Williams. Neural Kernel Surface Reconstruction. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4369–4379, Vancouver, BC, Canada, June 2023. IEEE.
- [23] Zhangjin Huang, Yuxin Wen, Zihao Wang, Jinjuan Ren, and Kui Jia. Surface Reconstruction from Point Clouds: A Survey and a Benchmark, May 2022. arXiv:2205.02413 [cs].
- [24] Alec Jacobson and Daniele Panozzo. libigl.
- [25] Tao Ju, Frank Losasso, Scott Schaefer, and Joe Warren. Dual contouring of hermite data. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques, SIGGRAPH '02*, pages 339–346, New York, NY, USA, July 2002. Association for Computing Machinery.
- [26] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. *Poisson Surface Reconstruction*. The Eurographics Association, 2006. Accepted: 2014-01-29T08:14:02Z ISSN: 1727-8384.
- [27] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics*, 32(3):1–13, June 2013.
- [28] Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis

- Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. ABC: A Big CAD Model Dataset For Geometric Deep Learning, Apr. 2019. arXiv:1812.06216 [cs].
- [29] Ravikrishna Kolluri, Jonathan Richard Shewchuk, and James F. O’Brien. Spectral surface reconstruction from noisy point clouds. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 11–21, Nice France, July 2004. ACM.
- [30] H el ene Legrand, Jean-Marc Thiery, and Tamy Boubekeur. Filtered Quadrics for High-Speed Geometry Smoothing and Clustering. *Computer Graphics Forum*, 38(1):663–677, Feb. 2019.
- [31] Yiyi Liao, Simon Donne, and Andreas Geiger. Deep Marching Cubes: Learning Explicit Surface Representations. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2916–2925, Salt Lake City, UT, June 2018. IEEE.
- [32] Shi-Lin Liu, Hao-Xiang Guo, Hao Pan, Peng-Shuai Wang, Xin Tong, and Yang Liu. Deep Implicit Moving Least-Squares Functions for 3D Reconstruction. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1788–1797, Nashville, TN, USA, June 2021. IEEE.
- [33] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, Aug. 1987.
- [34] Yiming Luo, Zhenxing Mi, and Wenbing Tao. DeepDT: Learning Geometry From Delaunay Triangulation for Surface Reconstruction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(3):2277–2285, May 2021. Number: 3.
- [35] Nissim Maruani, Roman Klokov, Maks Ovsjanikov, Pierre Alliez, and Mathieu Desbrun. VoroMesh: Learning Watertight Surface Meshes with Voronoi Diagrams. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14565–14574, 2023.
- [36] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy Networks: Learning 3D Reconstruction in Function Space, Apr. 2019. arXiv:1812.03828 [cs].
- [37] Thomas M uller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4):1–15, July 2022.
- [38] Charlie Nash, Yaroslav Ganin, S. M. Ali Eslami, and Peter Battaglia. PolyGen: An Autoregressive Generative Model of 3D Meshes. In *Proceedings of the 37th International Conference on Machine Learning*, pages 7220–7229. PMLR, Nov. 2020. ISSN: 2640-3498.
- [39] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, Long Beach, CA, USA, June 2019. IEEE.
- [40] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas K opf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library, Dec. 2019. arXiv:1912.01703 [cs, stat].
- [41] Songyou Peng, Chiyu Jiang, Yiyi Liao, Michael Niemeyer, Marc Pollefeys, and Andreas Geiger. Shape As Points: A Differentiable Poisson Solver. In *Advances in Neural Information Processing Systems*, volume 34, pages 13032–13044. Curran Associates, Inc., 2021.
- [42] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional Occupancy Networks, Aug. 2020. arXiv:2003.04618 [cs].
- [43] Sergey Prokudin, Qianli Ma, Maxime Raafat, Julien Valentin, and Siyu Tang. Dynamic Point Fields. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7964–7976, 2023.
- [44] Marie-Julie Rakotosaona, Paul Guerrero, Noam Aigerman, Niloy Mitra, and Maks Ovsjanikov. Learning Delaunay Surface Elements for Mesh Reconstruction. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22–31, Nashville, TN, USA, June 2021. IEEE.
- [45] Edoardo Remelli, Artem Lukoianov, Stephan Richter, Benoit Guillard, Timur Bagautdinov, Pierre Baque, and Pascal Fua. MeshSDF: Differentiable Iso-Surface Extraction. In *Advances in Neural Information Processing Systems*, volume 33, pages 22468–22478. Curran Associates, Inc., 2020.
- [46] Silvia Sell an. Reach For the Spheres: Tangency-Aware Surface Reconstruction of SDFs. *ACM Transactions on Graphics*, 2023.
- [47] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep Marching Tetrahedra: a Hybrid Representation for High-Resolution 3D Shape Synthesis. In *Advances in Neural Information Processing Systems*, volume 34, pages 6087–6101. Curran Associates, Inc., 2021.
- [48] Tianchang Shen, Jacob Munkberg, Jon Hasselgren, Kangxue Yin, Zian Wang, Wenzheng Chen, Zan Gojcic, Sanja Fidler, Nicholas Sharp, and Jun Gao. Flexible Isosurface Extraction for Gradient-Based Mesh Optimization. *ACM Transactions on Graphics*, 42(4):1–16, Aug. 2023.
- [49] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit Neural Representations with Periodic Activation Functions. In *Advances in Neural Information Processing Systems*, volume 33, pages 7462–7473. Curran Associates, Inc., 2020.
- [50] Raphael Sulzer, Loic Landrieu, Renaud Marlet, and Bruno Vallet. Scalable Surface Reconstruction with Delaunay-Graph Neural Networks. *Computer Graphics Forum*, 40(5):157–167, 2021. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14364>.
- [51] Qingyang Tan, Lin Gao, Yu-Kun Lai, and Shihong Xia. Variational Autoencoders for Deforming 3D Mesh Models. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5841–5850, 2018.
- [52] Jia-Heng Tang, Weikai Chen, Jie Yang, Bo Wang, Songrun Liu, Bo Yang, and Lin Gao. OctField: Hierarchical Implicit Functions for 3D Modeling, Nov. 2021. arXiv:2111.01067

- [cs].
- [53] Jean-Marc Thiery, Émilie Guy, and Tamy Boubekeur. Sphere-Meshes: shape approximation using spherical quadric error metrics. *ACM Transactions on Graphics*, 32(6):1–12, Nov. 2013.
  - [54] Philip Trettner and Leif Kobbelt. Fast and Robust QEF Minimization using Probabilistic Quadrics. *Computer Graphics Forum*, 39(2):325–334, May 2020.
  - [55] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C. J. Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, and Paul van Mulbregt. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3):261–272, Mar. 2020. Number: 3 Publisher: Nature Publishing Group.
  - [56] Pengfei Wang, Zixiong Wang, Shiqing Xin, Xifeng Gao, Wenping Wang, and Changhe Tu. Restricted Delaunay Triangulation for Explicit Surface Reconstruction. *ACM Transactions on Graphics*, 41(5):1–20, Oct. 2022.
  - [57] Peng-Shuai Wang, Yang Liu, and Xin Tong. Dual Octree Graph Networks for Learning Adaptive Volumetric Shape Representations. *ACM Transactions on Graphics*, 41(4):1–15, July 2022. arXiv:2205.02825 [cs].
  - [58] Ji Wu, Huai Yu, Wen Yang, and Gui-Song Xia. QuadricsNet: Learning Concise Representation for Geometric Primitives in Point Clouds, Sept. 2023. arXiv:2309.14211 [cs].
  - [59] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural Fields in Visual Computing and Beyond. *Computer Graphics Forum*, 41(2):641–676, 2022. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14505>.
  - [60] Guandao Yang, Serge Belongie, Bharath Hariharan, and Vladlen Koltun. Geometry Processing with Neural Fields. In *Advances in Neural Information Processing Systems*, volume 34, pages 22483–22497. Curran Associates, Inc., 2021.
  - [61] Congyi Zhang, Guying Lin, Lei Yang, Xin Li, Taku Komura, Scott Schaefer, John Keyser, and Wenping Wang. Surface Extraction from Neural Unsigned Distance Fields. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22531–22540, 2023.
  - [62] Chen Zhang, Ganzhangqin Yuan, and Wenbing Tao. DMNet: Delaunay Meshing Network for 3D Shape Representation. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14418–14428, 2023.
  - [63] Tong Zhao, Laurent Busé, David Cohen-Steiner, Tamy Boubekeur, Jean-Marc Thiery, and Pierre Alliez. Variational Shape Reconstruction via Quadric Error Metrics. In *ACM SIGGRAPH 2023 Conference Proceedings*, Aug. 2023.
  - [64] Qingnan Zhou and Alec Jacobson. Thingi10K: A Dataset of 10,000 3D-Printing Models, July 2016. arXiv:1605.04797 [cs].

# PoNQ: a Neural QEM-based Mesh Representation

## Supplementary Material

In this supplemental material, we discuss further evaluations of our PoNQ representation by comparing it with [46] and [47] in Sec. 7, before providing a number of implementation details in Sec. 8 which were not thoroughly covered in the main paper.

### 7. Discussing Other Related Works

To complement the series of comparisons presented in the main paper, we also discuss how PoNQ differs from two other (less) related works.

#### 7.1. Reach for The Spheres [46]



Figure 9. Learning-based results for PoNQ. Optimization-based results for Reach For the Spheres [46]. Note that both methods rely on the same input, here a  $32^3$  SDF grid. RTS can either miss thin structures (bottom rows) or fill thin voids (top rows).

A recent work on surface reconstruction of SDF, called Reach for The Spheres (RTS) [46], relies on the deformation of an initial mesh, thus requiring that its genus matches

the genus of the underlying SDF-encoded shape. In practice, several shapes of the Thing30 have a genus strictly higher than one. Consequently, and despite our best efforts, we could not run the code provided by the authors of Reach for The Spheres (RTS) [46] on the Thing30 dataset: segmentation faults and inverted elements systematically appeared in the optimization process when starting from a unit sphere — and using the output of a marching-cube (MC) extraction instead did not solve this problem, since the resulting genus depends on the MC table being used. Nevertheless, we provide for completeness a visual comparison of the two methods in Fig. 9 for two shapes that RTS could handle. While RTS exhibits many artifacts and PoNQ is far more robust to the genus of shapes, we note that the RTS empty sphere constraint could potentially be added to our loss function.

#### 7.2. Deep Marching Tetrahedra [47]

Regarding Deep Marching Tetrahedra (DMTet) [47], which relies on the deformation of a regular tetrahedra grid, two critical differences exist between their work and ours:

- strictly speaking, DMTet is not a point-based approach: their hybrid representation requires a neural network to model a continuous field of SDF and displacements;
- unless one restricts the amplitude of the deformation, it cannot guarantee intersection-free output meshes.

DMTet is thus quite distinct from PoNQ, making it difficult to compare them fairly. Nonetheless, we tried to provide a comparison; but given the absence of DMTet implementation for surface reconstruction from SDF grids, we found ourselves unable to provide a meaningful visual comparison — just like VoroMesh [35] was unable to compare with DMTet for the simpler case of the optimization-based task.

### 8. Implementation Details

We now cover in detail a series of aspects of PoNQ that were not fully detailed in the main paper.

#### 8.1. Training

For all the examples we show in our paper, we trained our PoNQ network with the AdamW optimizer and a linear combination of the losses, i.e.,

$$L = \alpha_{CD}L_{CD} + \alpha_nL_n + \alpha_A L_A + \alpha_{v^*}L_{v^*} + \alpha_{reg}L_{reg} + \alpha_{occ}L_{occ}$$

We rely on three training phases of 200 epochs each, with batches of size 16 with the following learning rates  $\gamma$ , weights  $\alpha = (\alpha_{CD}, \alpha_n, \alpha_A, \alpha_{v^*}, \alpha_{reg}, \alpha_{occ})$  and point samples  $S$ :

- $\gamma = 6.4 \cdot 10^{-5}$ ,  $\alpha = (100, .1, .1, 100, 100, .1)$ ,  $S = 5 \cdot 10^5$
- $\gamma = 3.2 \cdot 10^{-5}$ ,  $\alpha = (100, .1, .1, 100, 100, .1)$ ,  $S = 7 \cdot 10^5$
- $\gamma = 3.2 \cdot 10^{-5}$ ,  $\alpha = (100, .1, .1, 100, 1, .1)$ ,  $S = 7 \cdot 10^5$ .

## 8.2. Meshing

As explained in the main paper, the choices we made to design our QEM-based PoNQ representation allow for a simple meshing approach, inspired by computational geometry to ensure robustness: due to the local optimality of the positions  $\mathbf{v}_i^*$  (in particular, their ability to capture corners and sharp features), our output mesh is produced by filtering the triangle facets of a 3D Delaunay triangulation of the optimal positions  $\mathbf{v}_i^*$ . Given the PoNQ data (produced by a trained network or through optimization), we now describe how a PoNQ mesh is extracted in full detail.

**0. Pre-processing** We first normalize the quadrics by dividing them by their largest eigenvalue. While this technically changes their meaning (they now measure Euclidean distances *up to a multiplicative constant*), it also removes the possible bias due to variable sampling density. Additionally, a bounding box  $B$  of all the points  $\mathbf{v}_i^*$  is computed.

**1. Triangulation of QEM optimal positions.** We then compute the Delaunay tetrahedralization of all points  $\mathbf{v}_i^*$ , to which we add eight “protective” points defined as the corners of the bounding box  $B$ : all the adjacent tetrahedra to these protective points will be guaranteed to be outside of the shape we wish to reconstruct. The next two steps will tag each of the remaining tetrahedra as either *inside* or *outside* based on local geometric information, so that *our final PoNQ mesh will simply be the triangle mesh forming the boundary between the inside and outside regions*, ensuring watertightness and no self-intersections by design.

**2. Tagging obvious inside/outside tetrahedra.** We first tag all the tetrahedra adjacent to the eight protective point as *outside*. Tagging the rest of the tetrahedra seems daunting, but by leveraging the ideas put forth in the Crust algorithm [3], we note that the location of the circumcenter of a Delaunay tetrahedron whose four vertices are on the surface to mesh (which is the case we are in) can often determine the insideness or outsideness of this tetrahedron — corresponding to whether this circumcenter is on the inside (resp. outside) medial axis. We use a similar approach here, except that our Delaunay vertices have additional information to help us: we also know a local normal  $\mathbf{n}_i$  in the vicinity of each vertex  $\mathbf{v}_i^*$ . Thus, each vertex and its assigned normal defines an oriented plane, forming the boundary between the outside half-space (the one pointed by the normal) and the inside half-plane. Using this local test to determine inside/outside, we tag a tetrahedron as *outside* (resp., *inside*) if both its circumcenter and barycenter are determined to be in the outside (resp., inside) half-space of each of its four vertices. We then go through every (non-protective) vertex

that already has at least one tagged adjacent tetrahedron. If such a vertex does not have any adjacent *outside* (resp., *inside*) tetrahedron, we pick its untagged adjacent tetrahedron with the smallest edge; if the size of this smallest edge is below a certain large threshold (i.e., we are not in a very sparse region of the domain), we then tag this selected tetrahedron as *outside* (resp., *inside*). The rationale behind this last round of tagging is that we know that all vertices of our 3D Delaunay triangulation are *on* the surface, so unless the smallest tetrahedron is too big (in this case, there is clearly a large uncertainty), we can safely tag it to be on the opposite side of the surface than what the other tags had already determined. While this tagging procedure can, on rare occasions, tag *all* tetrahedra, there are often a few remaining untagged tetrahedra (typically caused by the presence of near sharp features or thin structures) that are too ambiguous to tag. We now need to lift the remaining uncertainty based on additional PoNQ data.

**3. Finishing up triangle selection.** Delaunay-based meshing approaches (like *Crust*) require a dense point sampling (formally, an  $\epsilon$ -sampling) to offer topological and geometric guarantees, which is not compatible with our desire to deal with thin structures, sharp features and corners — and this is the main reason why our earlier phase often ends up not providing a tag for *every* tetrahedron. To finish our tetrahedron tagging based on the ones we already have, we propose to use a graph cut approach, inspired by existing spectral graph partitioning [29]. For each Delaunay triangle  $T$  between two adjacent tetrahedra for which *at least* one of them is still undetermined, we compute a likelihood score  $S(T)$  that evaluates how confident we are that this triangle is to appear on the final output mesh. We propose a score that evaluates the fitness of  $T$  based on the local PoNQ normals and the local PoNQ quadrics matrices:

$$S(T) = S_{\mathbf{n}}(T) + hS_{\mathbf{Q}}(T)$$

(with  $h$  set as the squared inverse of the edge length of the SDF grid), where:

$$S_{\mathbf{n}}(T) = \left( \frac{2}{\pi} \sum_{\mathbf{v}_i^* \in T} \arccos(n_T^t \mathbf{n}_i) \right)^2,$$

$$S_{\mathbf{Q}}(T) = \sum_{\mathbf{v}_i^* \in T} \sum_{\substack{\mathbf{v}_j^* \in T \\ i \neq j}} [\mathbf{v}_j^*, 1]^t \mathbf{Q}_i [\mathbf{v}_j^*, 1],$$

where  $n_T$  denotes the normal of triangle  $T$ . We can now tag the remaining undetermined tetrahedra with a definite *inside* or *outside* label: we compute a minimum cut of the Voronoi graph (in which each dual of a tetrahedron is a node, and each dual of a Delaunay triangle face is an edge) using the already-tagged “inside” ones as a source and the “outside” ones as a drain, and each edge weight between two tetrahedra (with a common face  $T$ ) set to the score  $S(T)$ . Since most of tetrahedra are already tagged, we can merge Voronoi edges between marked tetrahedra to reduce the graph size and thus accelerate computations. We now

just extract the final *PoNQ mesh* as the triangle mesh forming the boundary between the inside and outside regions.

### 8.3. Ablation studies




Figure 10. A network trained without QEM (b) fails to recover sharp edges. The second-order QEM information provided by  $S_Q$  helps to disambiguate tetrahedra labeling (c).

To further justify the need for QEM, we provide two ablation studies: we re-generated models with a simpler version of our meshing algorithm that does not rely on quadrics (thus removing  $S_Q$ ), and we re-trained a network producing only points, normals, and voxel occupancies (thus removing the QEM part and its associated losses  $L_A$ ,  $L_{v^*}$ , and  $L_{reg}$ ). Figure 10 shows that the QEM-optimal vertex placement is essential to fit sharp features, and that using only  $S_n$  can mess up labeling in intricate/thin regions. Quantitatively, non-QEM versions fall behind our competitors in both surface and sharp-edge fitting scores, see Table 5.

Method	Grid size	CD ↓ ( $\times 10^{-5}$ )	F1 ↑	NC ↑	ECD ↓	EF1 ↑
retrained w/o QEM	$32^3$	<b>1.327</b>	0.840	0.960	<b>0.184</b>	0.598
PoNQ w/o $S_Q$	$32^3$	1.801	0.851	<b>0.964</b>	0.191	<b>0.715</b>
PoNQ	$32^3$	1.514	<b>0.852</b>	<b>0.964</b>	<b>0.184</b>	0.713
retrained w/o QEM	$64^3$	0.921	0.891	0.979	0.115	0.837
PoNQ w/o $S_Q$	$64^3$	0.931	<b>0.892</b>	<b>0.980</b>	<b>0.103</b>	0.863
PoNQ	$64^3$	<b>0.886</b>	<b>0.892</b>	<b>0.980</b>	0.109	<b>0.866</b>
retrained w/o QEM	$32^3$	1.387	0.803	0.942	0.139	0.286
PoNQ w/o $S_Q$	$32^3$	1.475	<b>0.810</b>	<b>0.943</b>	<b>0.137</b>	<b>0.315</b>
PoNQ	$32^3$	<b>1.344</b>	<b>0.810</b>	0.942	<b>0.137</b>	0.314
retrained w/o QEM	$64^3$	0.784	0.922	<b>0.971</b>	0.102	0.489
PoNQ w/o $S_Q$	$64^3$	0.779	<b>0.924</b>	<b>0.971</b>	0.102	<b>0.511</b>
PoNQ	$64^3$	<b>0.758</b>	<b>0.924</b>	<b>0.971</b>	<b>0.100</b>	<b>0.511</b>
retrained w/o QEM	$128^3$	0.645	<b>0.939</b>	<b>0.984</b>	0.135	0.556
PoNQ w/o $S_Q$	$128^3$	0.671	0.938	<b>0.984</b>	0.126	0.584
PoNQ	$128^3$	<b>0.641</b>	<b>0.939</b>	<b>0.984</b>	<b>0.123</b>	<b>0.592</b>

Table 5. Ablation studies on ABC (top) and Thingi30 (bottom).

### 8.4. Open surfaces

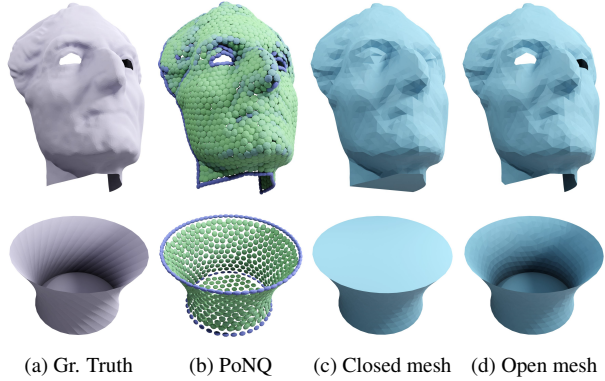


Figure 11. Optimization-based results for PoNQ when open boundaries are used.

We provide more results in Fig. 11, this time for open surfaces. Note that the QEM-optimized vertices  $v^*$  snap to sharp and boundary features, allowing for clean reconstructions. As explained in the main paper, we currently transit via a closed mesh (see Fig. 11c) that is then filtered to create holes between open boundaries, thus limiting our representation of open surfaces. However, since the PoNQ representation provides a precise fit of the target surface (see Fig. 11b), we believe a more sophisticated meshing could allow for arbitrary open surfaces.

### 9. Edge-based CD: Discussion




Figure 12. Comparison between two different ECD evaluation methods: sampling surface points and identifying right-angle normal changes between neighbors (left image) is less accurate than sampling sharp edges (right image), and fails to capture sharp edges with angles different than 90 degrees.

The Edge Chamfer Distance was introduced in BSP-Net [9] to evaluate sharp reconstructions on ShapeNet [7]. It is based on a sampling of the surface, and only captures sharp edges featuring an angle close to 90 degrees, see Fig. 12a. While this metric might be sufficient for ShapeNet, in which most sharp angles lie on the edges of boxy objects (chairs, televisions...), it is no longer true for

ABC [28] where sharp edges can have a variety of dihedral angles. As a result, we observed a large variance in this metric evaluation depending on the initial random surface sampling; moreover, it also failed to capture the spurious sharp edges of VoronoiMesh [35] due to faces of small area not being sampled. For these two reasons, and since most meshes of the ABC dataset have clean and well-defined sharp edges, we found that sampling the sharp edges (defined as those with a dihedral angle larger than  $\frac{\pi}{6}$ ) directly provided a more robust and faithful metric (see cog teeth in Fig. 12b).

### 9.1. Additional timings

VoronoiMesh [35] provides two implementations for their mesh extraction, based on SciPy and CGAL. Since our mesh extraction is also based on SciPy, we use their SciPy one out of fairness in Tab. 6 to evaluate and compare timings; we also added the timings of the meshing phase of NMC [10] and NDC [8] for comparison purposes. While our code is not yet optimized, an implementation of PoNQ in CGAL would undoubtedly allow for an even faster extraction.

Method	Grid	Optimization (s)	Meshing (s)
NDC [8]	$32^3$	-	0.001
NMC [10]	$32^3$	-	0.1
VoronoiMesh [35]	$32^3$	2.0	0.3
PoNQ	$32^3$	2.6	0.3
NDC [8]	$64^3$	-	0.02
NMC [10]	$64^3$	-	0.9
VoronoiMesh [35]	$64^3$	4.2	1.1
PoNQ	$64^3$	4.1	1.3
NDC [8]	$128^3$	-	0.2
NMC [10]	$128^3$	-	7.3
VoronoiMesh [35]	$128^3$	36.2	4.8
PoNQ	$128^3$	17.9	7.8

Table 6. Timings for optimization-based experiments.

## 10. Additional Renders

Finally, Figs. 13-17 exhibit further results comparing PoNQ to previous works.





Figure 13. Optimization-based results (top to bottom:  $32^3$ ,  $64^3$ ,  $128^3$ ) on Thingy30.




Figure 14. Learning results (top:  $32^3$ ; bottom:  $64^3$ ) on ABC.




Figure 15. Learning results (top:  $32^3$ ; bottom:  $64^3$ ) on ABC.




Figure 16. Learning results (top to bottom:  $32^3$ ,  $64^3$ ,  $128^3$ ) on Thingi30. Networks trained on ABC.




Figure 17. Learning results (top to bottom:  $32^3$ ,  $64^3$ ,  $128^3$ ) on Thingi30. Networks trained on ABC.