
Polyhedral Complex Derivation from Piecewise Trilinear Networks

Jin-Hwa Kim
NAVER AI Lab & SNU AIIS
Republic of Korea
jinhwa.kim@navercorp.com

Abstract

Recent advancements in visualizing deep neural networks provide insights into their structures and mesh extraction from Continuous Piecewise Affine (CPWA) functions. Meanwhile, developments in neural surface representation learning incorporate non-linear positional encoding, addressing issues like spectral bias; however, this poses challenges in applying mesh extraction techniques based on CPWA functions. Focusing on trilinear interpolating methods as positional encoding, we present theoretical insights and an analytical mesh extraction, showing the transformation of hypersurfaces to flat planes within the trilinear region under the eikonal constraint. Moreover, we introduce a method for approximating intersecting points among three hypersurfaces contributing to broader applications. We empirically validate correctness and parsimony through chamfer distance and efficiency, and angular distance, while examining the correlation between the eikonal loss and the planarity of the hypersurfaces. The code is available at <https://github.com/naver-ai/tropical-nerf.pytorch>.

1 Introduction

Recent advancements in visualizing deep neural networks [1–4] significantly contribute to understanding their intricate structures. This progress provides valuable insights into the expressivity, robustness, training methodologies, and distinctive geometry of neural networks. By leveraging the inherent piecewise linearity in the Continuous Piecewise Affine (CPWA) functions, *e.g.*, ReLU neural networks, each region of the input space is represented as a convex polyhedron. The assembly of these sets constructs a polyhedral complex that delineates the decision boundaries of neural networks [5].

Building upon these breakthroughs, we explore the exciting prospect of analytically extracting a mesh representation from neural implicit surface networks [6–8]. The extraction of a mesh not only offers a precise visualization and characterization of the network’s geometry but also does so efficiently through the utilization of vertices and faces, in contrast to sampling-based methods [9, 10]. The sampling-based methods are limited by their black-box nature, which can obscure the underlying structure and lead to inefficiencies in capturing fine geometric details.

These networks typically learn a signed distance function (SDF) by utilizing ReLU neural networks. However, recent approaches incorporate non-linear positional encoding techniques, such as trigonometric functions [11] or trilinear interpolations using hashing [12] or tensor factorization [13], to mitigate issues like spectral bias [14], ensuring fast convergence, and maintaining high-fidelity. Consequently, applying mesh extraction techniques based on CPWA functions becomes challenging.

Focusing on the successful and widely-adopted trilinear interpolating methods, we present novel theoretical insights and a practical methodology for precise mesh extraction. Following a novel review

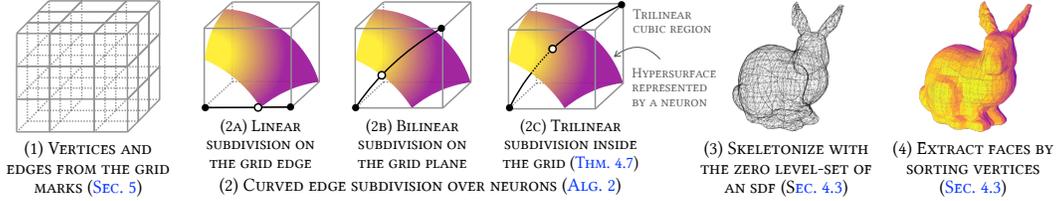


Figure 1: The mesh is analytically extracted from piecewise trilinear networks, comprising both HashGrid [12] and ReLU neural networks, that have been trained to learn a signed distance function with the eikonal loss. We start with initial vertices and edges defined using the grid marks (1), and its linear subdivision (2a); however, intersecting polygons (*ref.* Section 3.2) need bilinear subdivision (2b) and consequentially trilinear subdivision (2c). Note that the linear and bilinear subdivisions are specific instances of trilinear subdivisions with straightforward solutions.

of the *edge subdivision* method [4], an efficient mesh extraction technique for CPWA functions, within the framework of tropical geometry (Section 3), we demonstrate that ReLU neural networks incorporating the trilinear module are piecewise trilinear networks (Definition 4.1). Then, we establish that the hypersurface within the trilinear region becomes a plane under the eikonal [15] constraint (Theorem 4.5), a common practice during the training of SDF.

Building upon this observation, we propose a method for approximating the determination of intersecting points among three hypersurfaces. Given the inherent curvature of these surfaces, achieving an exact solution is infeasible. In our approach, we utilize two hypersurfaces and a diagonal plane for a more feasible and precise approximation (Theorem 4.7). Additionally, we introduce a methodology for arranging vertices composed of faces to implicitly represent normal vectors. We experimentally confirm accuracy and simplicity using chamfer distance, efficiency, and angular distance, simultaneously exploring the relationship between the eikonal loss and the planarity of the hypersurface. In Figure 1, we present the analytically generated skeleton and normal map produced by our method using piecewise trilinear networks composed of HashGrid and ReLU neural networks. For a more detailed view, please refer to Figure 10 in Appendix.

Our contributions are summarized as follows:

1. We present novel theoretical insights and a practical methodology for precise mesh extraction, from the piecewise trilinear networks using the HashGrid.
2. We provide a theoretical analysis of the piecewise trilinear networks under the eikonal constraint revealing that, within a trilinear region, a hypersurface transforms into a plane.
3. We validate its correctness and parsimony through chamfer distance and efficiency, and angular distance, showing the correlation between the eikonal loss and hypersurface planarity.

2 Related work

Mesh extraction in deep neural networks. A conventional black-box approach is to assess a given function at sampled points to derive the complex of the function, commonly through marching cubes [9] or marching tetrahedra [10, 16]. Despite the inclusion of mesh optimization techniques [17] in modern mesh packages, this exhaustive method often generates redundant complexities due to discretization, increasing the computational load by using unnecessary mesh elements (*e.g.*, fragmented planars). Aiming for exact extraction, initial efforts focused on identifying linear regions [18] and employing *Analytical Marching* [2] to exactly reconstruct the zero-level isosurface. Contemporary approaches involve region subdivision, wherein the regions of the complex are progressively subdivided from neuron to neuron and layer to layer, allowing for the efficient calculation of the exponentially increasing linear regions [4, 19–21]. Nevertheless, many previous studies have explored CPWA functions consisting of fully connected layers and ReLU activation, primarily owing to the mathematical simplicity of hyperplanes defining linear regions. However, this restricts their applications in emerging domains of deep implicit surface learning that diverge from CPWA functions.

Positional encoding. The spectral bias of a multi-layer perceptron (MLP) hinders effectively learning high frequencies, both theoretically and empirically [14]. To address this limitation, Fourier feature mapping, employing trigonometric projections similar to those used in the Transformer architecture [22], proves successful in representing complex 3D objects and scenes. Furthermore, to address its slow convergence and enhance rendering quality, novel positional encoding methods using trilinear interpolations, *e.g.*, HashGrid [12] or TensorRF [13], are introduced. Both techniques generate positional features through trilinear interpolation among the eight nearest corner features derived from pre-defined 3D grids. These features are obtained by hashing from multi-resolution hash tables or factorized representations of a feature space, respectively. However, these approaches cause the function represented by neural networks to deviate from CPWA functions, as trilinear interpolation is not an affine transformation. Consequently, the complex regions are divided by intricate hypersurfaces.

Eikonal equation and SDF. An eikonal¹ equation is a non-linear first-order partial differential equation governing the propagation of wavefronts. Notably, this finds application in the regularization of a signed distance function (SDF) in the context of neural surface modeling [6, 8, 23]. If Ω represents a subset of the space \mathcal{X} equipped with the metric d , the SDF $f(\mathbf{x})$ is defined as follows:

$$f(\mathbf{x}) = \begin{cases} -d(\mathbf{x}, \partial\Omega) & \text{if } \mathbf{x} \in \Omega \\ d(\mathbf{x}, \partial\Omega) & \text{if } \mathbf{x} \in \Omega^c \end{cases} \quad (1)$$

where $\partial\Omega$ denotes the boundary of Ω . The metric d is defined (with a slight notational abuse) as:

$$d(\mathbf{x}, \partial\Omega) := \inf_{\mathbf{y} \in \partial\Omega} d(\mathbf{x}, \mathbf{y}) \quad (2)$$

where d represents the shortest distance from \mathbf{x} to $\partial\Omega$ in the Euclidean space of \mathcal{X} . In the Euclidean space with a piecewise smooth boundary, the SDF exhibits differentiability almost everywhere, and its gradient adheres to *the eikonal equation*. Specifically, for points \mathbf{x} on the boundary, $\nabla f(\mathbf{x}) = N(\mathbf{x})$ where the Jacobian of f represents the outward normal vector, or equivalently, $|\nabla f(\mathbf{x})| = 1$.

3 Preliminaries

We present an overview of tropical geometry (Section 3.1) to provide a formal definition of the tropical hypersurface and the tropical algebra of neural networks. Then, we discuss the edge subdivision algorithm [4] for extracting the tropical hypersurface (Section 3.2). We supplement Appendix A.1 for less familiar readers. Later, we extend this idea to include considerations for trilinear interpolation in Section 4. You may choose to skip these sections if you are already acquainted with these concepts.

3.1 Tropical geometry

Tropical² geometry [25, 26] describes polynomials and their geometric characteristics. As a skeletonized version of algebraic geometry, tropical geometry represents piecewise linear meshes using *tropical semiring* and *tropical polynomial* (ref. Appendix A.1). The set of points where a tropical polynomial f is non-differentiable, due to tropical sum $\oplus := \max(\cdot)$, is called *tropical hypersurface*:

Definition 3.1 (Tropical hypersurface). The tropical hypersurface of a tropical polynomial $f(\mathbf{x}) = c_1 \mathbf{x}^{\alpha_1} \oplus \dots \oplus c_r \mathbf{x}^{\alpha_r}$ is defined as:

$$V(f) = \{\mathbf{x} \in \mathbb{R}^d : c_i \mathbf{x}^{\alpha_i} = c_j \mathbf{x}^{\alpha_j} = f(\mathbf{x}) \text{ for some } \alpha_i \neq \alpha_j\}$$

where the tropical monomial $c_i \mathbf{x}^{\alpha_i}$ with d -variate \mathbf{x} is defined as $c_i \odot x_1^{\alpha_{i,1}} \odot \dots \odot x_d^{\alpha_{i,d}}$ with d -variate $\alpha_i = (a_{i,1}, \dots, a_{i,d})$ and *tropical product* $\odot := +$.

Definition 3.2 (ReLU neural networks). Assuming input and output dimensions are consistent, L -layer neural networks with ReLU activation are a composition of functions defined as:

$$\nu^{(L)} = \rho^{(L)} \circ \sigma^{(L-1)} \circ \rho^{(L-1)} \dots \sigma^{(1)} \circ \rho^{(1)} \quad (3)$$

where $\sigma(\mathbf{x}) = \max(\mathbf{x}, 0)$ and $\rho^{(i)}(\mathbf{x}) = \mathbf{W}^{(i)}\mathbf{x} + \mathbf{b}^{(i)}$.

¹The *eikonal* is derived from the Greek, meaning image or icon.

²The *tropical* is named in honor of Hungarian-born Brazilian computer scientist Imre Simon [24].

Although the neural networks ν is generally non-convex, it is known that the difference of two tropical *signomials* (allowing a real-valued α in a tropical polynomial) can represent ν [27] (ref. Appendix A.3). Given this observation, we proceed to examine the following proposition:

Proposition 3.3 (Decision boundary of neural networks). *The decision boundary for L -layer neural networks with ReLU activation $\nu^{(L)}(\mathbf{x}) = 0$ is a subset of or equals with the region where two tropical monomials, or two arguments of max, equal (Definition 3.1), as follows:*

$$\mathcal{B} \subseteq \{ \mathbf{x} : \nu_j^{(i)}(\mathbf{x}) = 0, i \in (1, \dots, L), j \in (1, \dots, H) \text{ if } i \neq L \text{ else } (1) \}$$

where the subscript j denotes the j -th neuron, assuming the hidden size of neural networks is H , while the final L -th layer has a single output for rather discussions as an SDF (Definition 4.3). Remind that tropical operations satisfy the usual laws of arithmetic for recursive consideration.

The piecewise hyperplanes corresponding to $\nu_j^{(i)}$ are the *candidates* shaping the decision boundary. It enables for-loop iterations over the neurons across layers, as in Algorithm 2, not visiting every exponentially growing linear region [18, 27]. The equality in Proposition 3.3 generally does not hold, but we can filter them by checking if $\nu_1^{(L)}(\mathbf{x}) = 0$. For more details, please refer to Berzins [4]. We note that our formal explanation through tropical geometry is unprecedented in the previous work [4].

3.2 Edge subdivision for tropical geometry

Polyhedral³ complex derivation from a multitude of linear regions is inefficient and may be infeasible in some cases. Rather, Berzins [4] argue that tracking vertices and edges of the decision boundary sequentially considering the polynomial number of hyperplanes is particularly efficient, even enabling parallel tensor computations for edge subdivisions.

Initially, we start with a unit cube of eight vertices and a dozen edges, denoted by \mathcal{V} and \mathcal{E} , respectively. For each piecewise linear, or *folded* in their term, hyperplane, we find the intersection points of each one of the edges and the hyperplane (if any) to add to \mathcal{V} . The divided edges by the intersection are also added to \mathcal{E} . Note that we also find new edges of polygons, the intersections of convex polyhedra representing corresponding linear regions [5], and the folded hyperplane. Repeating these for all folded hyperplanes, before selecting all vertices \mathbf{x}^* where $\nu(\mathbf{x}^*) = 0$. Then, the valid edges connecting those vertices would represent the decision boundary.

To elaborate with details, the order of subdividing is invariant as stated in Proposition 3.3; however, we must perform the subdivision with all hyperplanes in the previous layers in advance if we want to keep the current set of edges not crossing linear regions. One critical advantage of *this principle* is that we can find the intersection point by evaluating the output ratio of the two vertices of an edge, $\nu(\mathbf{x}_0)$ and $\nu(\mathbf{x}_1)$, which are proportional to the distances to the hyperplane by Thales's theorem [28]. The intersection point would simply be $\hat{\mathbf{x}}_{0,1} = (1-w)\mathbf{x}_0 + w\mathbf{x}_1$ where $w = |d_0|/|d_0 - d_1|$ and $d_k = \nu_j^{(i)}(\mathbf{x}_k)$, upon the fact that $d_0 \cdot d_1 < 0$ if the hyperplane divides the edge.

Updating edges needs two steps: dividing the edge into two new edges and adding new edges of intersectional polygons by a hyperplane for convex polyhedra representing linear regions. The latter is a tricky part that we need to find every pair of two vertices among $\{\hat{\mathbf{x}}_{\cdot,\cdot}\}$, both within the same linear region *and* on the two common hyperplanes, obviously including the current hyperplane. They argue that the *sign-vectors* for the preactivations provide an efficient way to find them.

Definition 3.4 (Sign-vectors). For the current hyperplane specified by the i^* -th layer and its j^* -th preactivation, we define the sign-vectors with a small positive constant ϵ :

$$\gamma(\mathbf{x}_k)_{\phi(i,j)} = \begin{cases} +1 & \text{if } \nu_j^{(i)}(\mathbf{x}_k) > +\epsilon \\ 0 & \text{if } |\nu_j^{(i)}(\mathbf{x}_k)| \leq +\epsilon \\ -1 & \text{if } \nu_j^{(i)}(\mathbf{x}_k) < -\epsilon \end{cases} \quad (4)$$

for all i and j where $\phi(i,j) \leq \phi(i^*,j^*)$, and $\phi(i,j) \in \{0\} \cup \mathbb{N}$ is an ascending indexing function from lower layers, handling arbitrary hidden sizes of networks to vectorize a signed matrix.

³In geometry, a polyhedron is a three-dimensional shape with polygonal faces, having straight edges and sharp vertices.

Notice that a hyperplane divides a space into two half-spaces, two linear regions. Collectively, if the sign-vectors for the i -th layer and its j -th preactivation of two vertices have the same values except for zeros, which are wild cards for matching, we can say that the two vertices are within the same linear region in the current step (i^*, j^*) . For the second, the sign-vector has at least three zeros specifying a point by at least three hyperplanes. If two zeros are at the same indices in the two sign-vectors, the two vertices are on the same two hyperplanes, forming an edge. Notice that our principle asserts edges should be within linear regions.

Berzins [4] argue that the edge subdivision provides the optimal time and memory complexities of $\mathcal{O}(|\mathcal{V}|)$, linear in the number of vertices, while it can be efficiently computed in parallel. According to Hanin and Rolnick [20], the number of linear regions is known for $o(N^L/L!)$ where N is the total number of neurons, or $\phi(L, 1)$ in our notation.

4 Method

4.1 Piecewise trilinear networks

Definition 4.1 (Trilinear interpolation). Let a unit cube be in the D -dimensional space, where its corners are represented by $\mathbf{H} \in \mathbb{R}^{2^D \times F}$ where F is a corner-feature dimension. Given weights $\mathbf{w} \in [0, 1]^D$, the trilinear interpolation function ψ is defined as:

$$\psi(\mathbf{w}; \mathbf{H}) = \sum_{i=0}^{2^D-1} \omega(i, \mathbf{w}) \cdot \mathbf{H}_i \in \mathbb{R}^F \quad (5)$$

where the interpolating weight $\omega(i, \mathbf{w}) \in \mathbb{R}$ is defined using a left-aligned and zero-trailing binarization function $\mathbf{B} \in \{0, 1\}^D$ (ref. Appendix B.1) as follows:

$$\omega(i, \mathbf{w}) = \prod_{j=1}^D (1 - \mathbf{B}(i)_j)(1 - \mathbf{w}_j) + \mathbf{B}(i)_j \mathbf{w}_j \quad (6)$$

which is the volume of the opposite subsection of the hypercube divided by the weight point \mathbf{w} . For a general hypercube, scaling the weight for a unit hypercube gets the same result. Going forward, we assume $D=3$ for trilinear interpolation without explicitly stating it otherwise.

Lemma 4.2 (Nested trilinear interpolation). *Let a nested cube be inside of a unit cube, where its positions of the eight corners deviate $-\mathbf{a}_j$ or $+\mathbf{b}_j$, such that $\mathbf{a}_j, \mathbf{b}_j \geq 0$, from \mathbf{w}_j for each dimension j , using the notations of Definition 4.1. The eight corners of the nested cube are the trilinear interpolations of the eight corners of the unit cube. Then, the trilinear interpolation with the unit cube and the nested cube for \mathbf{w} is identical.*

Proof. Notice that trilinear interpolation is linear for each dimension to prove it. The detailed proof can be found in Lemma D.1. \square

Definition 4.3 (Piecewise trilinear networks). Let a positional encoding module be $\tau : \mathbb{R}^D \rightarrow \mathbb{R}^F$ using the trilinear interpolation of spatially nearby learnable vectors on a three-dimensional grid, e.g., HashGrid [12] or TensoRF [13], where τ is an instance of ψ in Definition 4.1. Usually, they transform the input coordinate \mathbf{x} to \mathbf{w} as a relative position inside a unit grid in a corresponding resolution. \mathbf{H} is a learnable parameter specified in the corresponding method. Then, we define trilinear neural networks $\tilde{\nu}$, by prepending τ to the neural networks ν from Definition 3.2. Here, the input and output dimensions of $\nu^{(L)}$ are F and 1 (an SDF distance), respectively.

$$\tilde{\nu}^{(L)} = \nu^{(L)} \circ \tau \quad (7)$$

Note that τ is *gridwise* trilinear. Using Lemma 4.2, $\tilde{\nu}$ is trilinear within the composite intersections of linear regions of ν and trilinear regions of τ . (We discuss how to access *virtual* cubic corner features where trilinear regions are not cubic in Section 5.) So, we regard $\tilde{\nu}$ as *piecewise* trilinear.

4.2 Curved edge subdivision in a trilinear space

In a trilinear space, τ projects a line to a curve except the lines on the grid. Notably, the diagonal line $\mathbf{t} = (t, t, t)^\top$ where $t \in [0, 1]$ is projected to a cubic Bézier curve (ref. Proposition D.2). We aim to generalize for the curved edge subdivision for hypersurfaces.

Lemma 4.4 (Curved edge of two hypersurfaces). *In a piecewise trilinear region, let an edge $(\mathbf{x}_0, \mathbf{x}_7)$ be the intersection of two hypersurfaces $\tilde{\nu}_i(\mathbf{x}) = \tilde{\nu}_j(\mathbf{x}) = 0$, while \mathbf{x}_0 and \mathbf{x}_7 are on the two hypersurfaces. Then, the edge is defined as:*

$$\{\mathbf{x} : \tau(\mathbf{x}) = (1-t)\tau(\mathbf{x}_0) + t\tau(\mathbf{x}_7) \text{ where } \mathbf{x} \in \mathbb{R}^D \text{ and } t \in \mathbb{R}\}.$$

Proof. The detailed proof using the piecewise linearity of ν is provided in Lemma D.4. \square

Notice that it does not decrease the number of variables to find a line solution since the trilinear interpolation τ still forms hypersurfaces making complex cases. Yet, we theoretically demonstrate that the eikonal constraints on $\tilde{\nu}$ render them hyperplanes with linear solutions.

Theorem 4.5 (Hypersurface and eikonal constraint). *A hypersurface $\tau(\mathbf{x}) = 0$ intersects two points $\tau(\mathbf{x}_0) = \tau(\mathbf{x}_7) = 0$ while $\tau(\mathbf{x}_{1\dots 6}) \neq 0$ for the remaining six points. These points form a cube, with \mathbf{x}_0 and \mathbf{x}_7 positioned on the diagonal of the cube. The hypersurface satisfies the eikonal constraint $\|\nabla\tau(\mathbf{x})\|_2^2 = 1$ for all $\mathbf{x} \in [0, 1]^3$. Then, the hypersurface of $\tau(\mathbf{x}) = 0$ is a plane.*

Corollary 4.6 (Affine-transformed hypersurface and eikonal constraint). *Let $\nu_0(\mathbf{x})$ and $\nu_1(\mathbf{x})$ be two affine transformations defined by $\mathbf{W}_i^\top \mathbf{x} + \mathbf{b}_i$, $i \in \{0, 1\}$. A hypersurface $\nu_0(\tau(\mathbf{x})) = 0$ passing two points $\nu_0(\tau(\mathbf{x}_0)) = \nu_0(\tau(\mathbf{x}_7)) = 0$ while $\nu_0(\tau(\mathbf{x}_{1\dots 6})) \neq 0$ satisfies the eikonal constraint $\|\nabla\nu_1(\tau(\mathbf{x}))\|_2^2 = 1$ for all $\mathbf{x} \in [0, 1]^3$. Then, the hypersurface of $\nu_0(\tau(\mathbf{x})) = 0$ is a plane.*

The proofs using its partial derivatives can be found in Theorem D.5 and Corollary D.6.

Since we aim for a practical polyhedral complex derivation that piecewise trilinear networks represent, we replace one of the hypersurfaces forming the curved edge with a diagonal plane in a piecewise trilinear region. We choose this option not only for its mathematical simplicity but also due to the characteristics of trilinear hypersurfaces, as illustrated in Figure 9, Appendix G. Thus, the new vertices lie on at least two hypersurfaces, and the new edges exist on the same hypersurface, while the eikonal constraint minimizes any associated error from this approximation. This error is tolerated by the hyperparameter $\epsilon = 1e-4$ as specified in Definition 3.4 and Section 6. More discussions on this matter can be found in Appendix C.1.

Theorem 4.7 (Intersection of two hypersurfaces and a diagonal plane). *Let $\tilde{\nu}_0(\mathbf{x}) = 0$ and $\tilde{\nu}_1(\mathbf{x}) = 0$ be two hypersurfaces passing two points \mathbf{x}_0 and \mathbf{x}_7 such that $\tilde{\nu}_0(\mathbf{x}_0) = \tilde{\nu}_1(\mathbf{x}_0) = 0$ and $\tilde{\nu}_0(\mathbf{x}_7) = \tilde{\nu}_1(\mathbf{x}_7) = 0$, $\mathbf{P}_i := \tilde{\nu}_0(\mathbf{x}_i)$ and $\mathbf{Q}_i := \tilde{\nu}_1(\mathbf{x}_i)$, $\mathbf{P}_\alpha = [\mathbf{P}_0; \mathbf{P}_1; \mathbf{P}_4; \mathbf{P}_5]$, $\mathbf{P}_\beta = [\mathbf{P}_2; \mathbf{P}_3; \mathbf{P}_6; \mathbf{P}_7]$, and $\mathbf{X} = [(1-x)^2; x(1-x); (1-x)x; x^2]$. Then, $x \in [0, 1]$ of the intersection point of the two hypersurfaces and a diagonal plane of $x = z$ is the solution of the following quartic equation:*

$$\mathbf{X}^\top (\mathbf{P}_\alpha \mathbf{Q}_\beta^\top - \mathbf{P}_\beta \mathbf{Q}_\alpha^\top) \mathbf{X} = 0, \quad \text{while} \quad y = \frac{\mathbf{X}^\top \mathbf{P}_\alpha}{\mathbf{X}^\top (\mathbf{P}_\alpha - \mathbf{P}_\beta)} \quad (\mathbf{P}_\alpha \neq \mathbf{P}_\beta). \quad (8)$$

Proof. Please refer to Theorem D.7 for the detailed proof, which uses linear algebra to rearrange two trilinear equations with two variables to get a solution. \square

Note that finding roots of a polynomial is the eigenvalue decomposition of a companion matrix [29], which can be parallelized (ref. Lemma D.8). Given that the companion matrix remains small for a quartic equation in $\mathbb{R}^{4 \times 4}$, the overall complexity remains $\mathcal{O}(|\mathcal{V}|)$. For the detailed complexity analysis on the proposed method, please refer to Appendix E.

4.3 Skeletonization, faces, and normals

Skeletonization selects the vertices such that: $\mathcal{V}^* = \{\mathbf{x} \mid |\tilde{\nu}(\mathbf{x})| \leq \epsilon, \mathbf{x} \in \mathcal{V}\}$, and the edges \mathcal{E}^* such that their two vertices are among \mathcal{V}^* . Recall that because a decision boundary is a subset of the tropical hypersurface (Definition 3.1 and Proposition 3.3), this step ensures accurate mesh extraction.

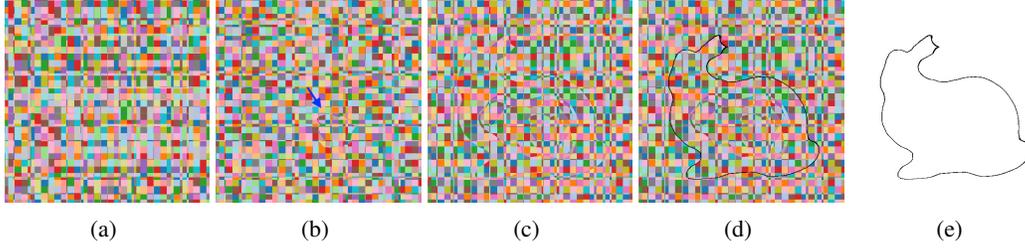


Figure 2: Trilinear regions in the xy -plane at $z = 0.04$, identified by the sign-vectors (Definition 3.4), are represented with random colors. (a) Grids described in Section 5 and Algorithm 4. (b) The neurons of the first layer representing folded hypersurfaces (blue arrow). (c) All neurons representing every nonlinear boundary. (d) Select all zero-set vertices and edges. (e) Skeletonized as in Section 4.3.

Faces are formed by edges lying on the same plane and sharing a common region. This is identified by evaluating the sign-vectors (*ref.* Definition 3.4 and Section 5). Note that faces are not inherently triangular; however, if needed, they can be triangulated via *triangularization* for further analysis.

Normals are conventionally described by the order of the vertices as each face has ambiguity with two opposite directions. Let $\mathbf{x}_{0,1,2}$ be the vertices forming a triangular face. The normal is defined using cross-product as $\mathbf{n} = (\mathbf{x}_1 - \mathbf{x}_0) \times (\mathbf{x}_2 - \mathbf{x}_0)$ where the normal is orthogonal to both vectors, with a direction given by the right-hand rule. Please refer to Appendix B.2 for the details.

5 Implementation

We employ the HashGrid [12] for τ , while our discussion on its generalizability is provided in Appendix C.2. We illustrate our method in Algorithm 1, 2, and Figure 2, inspired by Hanin and Rolnick [30] for its visualization of linear regions. We describe the implemental details of the algorithms in the following sections and the corresponding caption of Figure 6.

Initialization. While the original edge subdivision algorithm [4] starts with a unit cube of eight vertices and a dozen edges, we know that the unit cube is subdivided by orthogonal planes consisting of the grid. Let an input coordinate be $\mathbf{x} \in [0, 1]^3$, and there is a unit cube such that one of its corners is at the origin. Taking into account multi-resolution grids, we obtain the marks m_i , such that the grid planes are $x = m_i$, $y = m_i$, and $z = m_i$, following Algorithm 4 in Appendix. Notice that we carefully consider the offset $s/2$, which prevents the zero derivatives from aligning across all resolution levels (*ref.* Appendix A of Müller et al. [12]). Leveraging the obtained marks, we derive the associated grid vertices and edges, which serve as the initial sets for \mathcal{V} and \mathcal{E} , respectively. Please refer to Figure 2a where its multi-resolution grids are visualized using randomized colors.

Optimizing sign-vectors. The number of orthogonal planes may significantly surpass the number of neurons, making it impractical to allocate elements for the sign-vectors (Definition 3.4). To address this, leveraging the orthogonality of the grid planes, we store a plane index along with an indicator of whether the input is on the plane ($\gamma = 0$) or not ($\gamma = 1$).

Piecewise trilinear region. In Theorem 4.7, we need the outputs of corners P and Q in a common piecewise trilinear region; however, some corners may be outside of the region. For this, we replace all ReLU activations using the mask \mathbf{m} that: $\mathbf{m}_j^{(i)} = (\tilde{\nu}_j^{(i)}(\mathbf{x}_0) > \epsilon) \mid (\tilde{\nu}_j^{(i)}(\mathbf{x}_7) > \epsilon)$, where \mathbf{x}_0 and \mathbf{x}_7 are the vertices of an edge of interest, for the calculated cubic corners $\mathbf{x}_{0\dots7}$. This implies that ν exhibits linearity, except for instances where two vertices simultaneously deviate from linearity. For P and Q , we select the last two pairs of i and j such that $\gamma_{\phi(i,j)}(\mathbf{x}_0) = \gamma_{\phi(i,j)}(\mathbf{x}_7) = 0$, indicating two common hypersurfaces passing two points \mathbf{x}_0 and \mathbf{x}_7 .

6 Experiment

Objective. For a given trilinear neural network with the eikonal constraint, our goal is to get a boundary mesh of the zero-set of the SDF. (Note that NeuS [6] showed how to convert the density

Algorithm 1 Polyhedral Complex Derivation

Input: Sets of vertices \mathcal{V} and edges \mathcal{E} (from Section 5), # of layer L , hidden size H , trilinear networks $\tilde{\nu}$

for $i = 0$ **to** $L - 1$ **do**

for $j = 0$ **to** $H - 1$ **do**

$\mathcal{V}, \mathcal{E} = \text{CESub}(\mathcal{V}, \mathcal{E}, \tilde{\nu}, i, j)$ \triangleright Algorithm 2

end for

end for

$\mathcal{V}, \mathcal{E} = \text{CESub}(\mathcal{V}, \mathcal{E}, \tilde{\nu}, L, 0)$

$\mathcal{V}^*, \mathcal{E}^* = \text{Skeletonize}(\mathcal{V}, \mathcal{E}, \tilde{\nu}, \epsilon)$ \triangleright Section 4.3

$\mathcal{F}^* = \text{ExtractFaces}(\mathcal{V}^*, \mathcal{E}^*)$ \triangleright Section 4.3

return $\mathcal{V}^*, \mathcal{E}^*, \mathcal{F}^*$

Algorithm 2 Curved Edge Subdivision (CESub)

Input: Vertices \mathcal{V} , Edges \mathcal{E} , trilinear networks $\tilde{\nu}$, layer index i , neuron index j

$d = \tilde{\nu}(\mathcal{V}; i, j), \mathcal{V}_{\text{new}} = \emptyset$

for e **in** \mathcal{E} **do**

if $d_{e_0} \cdot d_{e_1} < 0$ **then**

$\mathbf{P}, \mathbf{Q} = \text{Corners}(e)$ \triangleright Section 5

$\mathbf{x}_{\text{new}} = \text{Intersect}(\mathbf{P}, \mathbf{Q}, e)$ \triangleright Theorem 4.7

$\mathcal{V}_{\text{new}} \leftarrow \mathcal{V}_{\text{new}} + \{\mathbf{x}_{\text{new}}\}$

$e \leftarrow [e_0, \text{Index}(\mathbf{x}_{\text{new}})]$

$\mathcal{E} \leftarrow \mathcal{E} + \{[e_1, \text{Index}(\mathbf{x}_{\text{new}})]\}$

end if

end for

$\mathcal{V} \leftarrow \mathcal{V} + \mathcal{V}_{\text{new}}$

$\mathcal{E} \leftarrow \mathcal{E} + \text{FindPolygons}(\mathcal{V}_{\text{new}})$ \triangleright Section 3.2

return \mathcal{V}, \mathcal{E}

networks of NeRFs to an SDF in the frameworks of volume rendering.) To evaluate this, we utilize marching cubes with excessive sampling to get a pseudo-ground truth mesh to remove Bayes error caused by underfitting. We explicitly specify the samplings in the results.

Hyperparameters. We used the number of layers L of 3 and hidden size H of 16 for the networks, and ϵ of 1e-4 for the sign-vectors (ref. Definition 3.4). The weight for the eikonal loss is 1e-2. For HashGrid, the resolution levels of 4, feature size of 2, base resolution N_{\min} of 2, and max resolution N_{\max} of 32 *by default*. N_{\min} and N_{\max} are doubled (x2) or quadrupled (x4) for *Medium* or *Large* settings in Figure 4. We use the official Python package of `tinycudnn`⁴ for the HashGrid module.

Chamfer distance and efficiency. The chamfer distance is a metric used to evaluate the similarity between two sets of sampled points from two meshes. Let \mathcal{S}_0 and \mathcal{S}_1 be the two sets of points. Specifically, the bidirectional chamfer distance (CD) is defined as follows:

$$\text{CD}(\mathcal{S}_0, \mathcal{S}_1) = \frac{1}{2} \left(\overrightarrow{\text{CD}}(\mathcal{S}_0, \mathcal{S}_1) + \overrightarrow{\text{CD}}(\mathcal{S}_1, \mathcal{S}_0) \right), \quad \overrightarrow{\text{CD}}(\mathcal{S}_0, \mathcal{S}_1) = \frac{1}{|\mathcal{S}_0|} \sum_{\mathbf{x} \in \mathcal{S}_0} \min_{\mathbf{y} \in \mathcal{S}_1} \|\mathbf{x} - \mathbf{y}\|_2^2. \quad (9)$$

We randomly select 100K points on the faces through ray-marching from the origin, directing toward a randomly chosen point on a unit sphere.

Table 1 and Figure 4 show the results for the Stanford 3D Scanning repository [31]. Varying the number of samples in marching cubes [9], we plot the baseline with respect to the number of generated vertices. Given that our method extracts vertices from the intersection points of hypersurfaces, it enables parsimonious representations using the lower number of vertices having better CD, summarized by chamfer efficiency (CE)⁵ in Table 1, and ours are located in the lower-left region compared with the baselines in Figure 4. In particular, we observe a strong CE in a simple object, *e.g.*, Drill Bit, which achieved a CE of 47.8 versus 19.2 (MC).

Our method inherently holds an advantage over MC in capturing optimal vertices and faces, especially in cases where the target surfaces exhibit planar characteristics. When dealing with curved surfaces, the smaller grid size of the *Large* model in Figure 4 views a curved surface in a smaller interval, and it tends to approach the plane. From the plotting of *Small* to

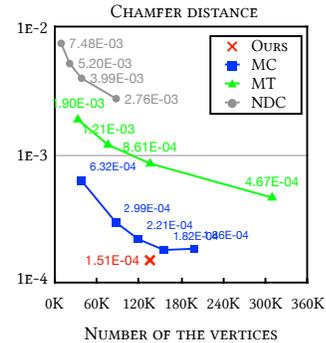


Figure 3: Chamfer distance for the bunny with the *Large* model comparing with MC, MT, and NDC.

⁴<https://github.com/NVlabs/tiny-cuda-nn>

⁵CE is defined as $100 / (\sqrt{|\mathcal{V}|} \times \text{CD})$, reflecting the trade-off between the number of vertices and the CD.

Table 1: Chamfer distance (CD) and chamfer efficiency (CE) for the Stanford 3D Scanning repository [31]. The CD ($\times 1e-6$) is evaluated using the marching cubes (MC) with 256^3 grid samples as the reference ground truth for all measurements. Please refer to Table 2 and Table 3 for the full results, and Table 4 for standard deviation and time spent in Appendix (Ours took 0.97 ± 0.21 while MC also took within 1 sec for the bunny.)

MC #	Bunny			Dragon			Happy Buddha			Armadillo	Drill	Lucy
	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow	CE \uparrow	CE \uparrow	CE \uparrow
32	1283	4106	19.0	1416	6330	11.2	1032	6951	13.9	11.2	7.5	24.9
64	5367	1371	13.6	6090	1936	8.5	4362	2465	9.3	8.3	19.2	16.0
128	21825	393	11.6	25236	542	7.3	18219	722	7.6	6.6	17.4	10.7
196	49569	141	14.3	57341	203	8.6	41351	276	8.8	7.5	15.3	11.7
Ours	4341	900	25.6	5104	1243	15.8	3710	1738	15.5	13.9	47.8	23.6

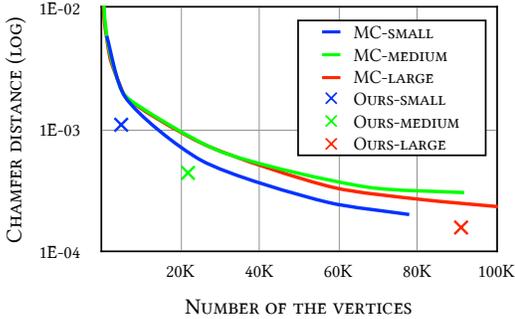


Figure 4: Chamfer distances for the Stanford dragon varying the number of vertices and the model sizes, showing consistent efficiency.

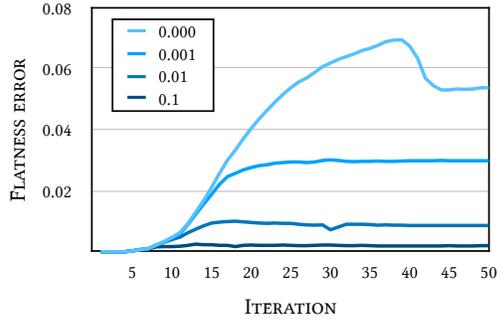


Figure 5: Effect of the weight of eikonal loss on the flatness error in Equation (11). For the plot, we conduct experiments using the unit sphere.

Large, our CDs (colored crosses) are decreasing to zero, consistently with better CEs, confirming this speculation. The complete results of the *Large* models and its visualization can be found in Tables 6 to 8 in Appendix F.1 and the right side of Figure 10 in Appendix, respectively. We also compare ours with Marching Tetrahedra (MT) [10, 16] and Neural Dual Contour (NDC) [32] in Figure 3, where our method achieved the lowest chamfer distance and the most efficiency method with respect to the number of vertices. For the rationale behind our selection, please refer to Appendix F.3.

Angular distance. The angular distance measures how much the normal vectors deviate from the ground truths. Analogous to the chamfer distance, we sample 100K points on the surface and calculate the normal vectors as described in Section 4.3. The angular distance is defined as follows:

$$AD(\mathcal{N}_0, \mathcal{N}_1) = \mathbb{E}_i \left[\frac{180}{\pi} \cdot \cos^{-1} (\langle \mathcal{N}_0^{(i)}, \mathcal{N}_1^{(i)} \rangle) \right] \quad (10)$$

In Table 5 in Appendix, the angular distances for the Stanford bunny are shown. As we expected, our approach efficiently estimates the faces using a parsimonious number of vertices compared to the sampling-based method. We confirm this trend is consistent across other objects.

Eikonal constraint for planarity. To assess the efficacy of the eikonal constraint, which enforces the planarity of hypersurfaces as outlined in Theorem 4.5, we quantify the flatness error associated with the planarity constraints presented in the proof of Theorem D.5. Let \mathbf{x}_0 and \mathbf{x}_7 be two vertices consisting of a diagonal edge, while $\mathbf{x}_{1..6}$ are its remaining cubic corners. We obtain $\tilde{\nu}(\mathbf{x}_{1..6})$ in a piecewise trilinear region as described in Section 5. This evaluation is conducted as follows:

$$\Delta_{\text{flat}} = \mathbb{E}_{\mathcal{E}} \left[\frac{1}{4} (\|\sum_{i \in \{1,2,4\}} \tilde{\nu}(\mathbf{x}_i)\|_1 + \|\sum_{i \in \{3,5,6\}} \tilde{\nu}(\mathbf{x}_i)\|_1) + \frac{1}{6} \left(\sum_{i=1}^3 \|\sum_{j \in \{i,7-i\}} \tilde{\nu}(\mathbf{x}_j)\|_1 \right) \right]. \quad (11)$$

In Figure 5, we empirically validate that the eikonal loss enforces planarity in the hypersurfaces within piecewise trilinear regions, as described in Theorem 4.5. The absence of the eikonal loss results in elevated planarity errors, leading to the construction of an inaccurate mesh.

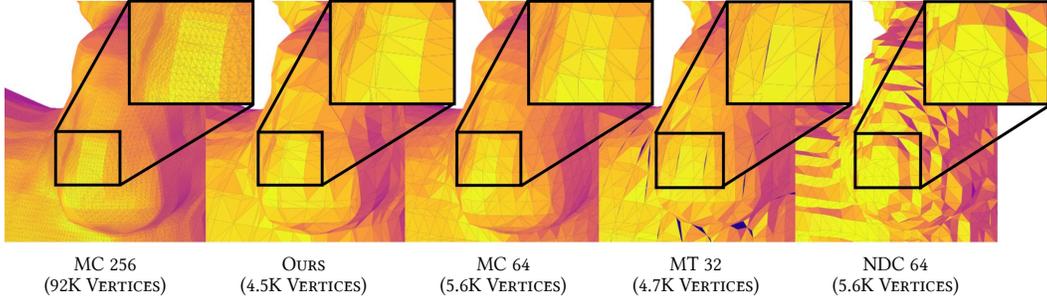


Figure 6: This figure provides a detailed visualization of Figure 10 in the Appendix and its competitors. MC 64, MT 32 (Marching Tetrahedra), and NDC (Neural Dual Contour [32]) suffer over-smoothing or inaccuracy of the actual surface in the Small networks within the nose, whereas our method reflects these details (see the boundaries of a nose) with consistent normals (in colors). MT efficiently demands SDF values by utilizing six tetrahedra within grids to get intermediate vertices. However, it is inefficient in terms of the number of vertices extracted. NDC faces a generalization issue for a zero-shot setting for the given networks, producing unsmooth surfaces.

Visualization. We provide the qualitative analyses in Figure 6 comparing with MT [10, 16] and NDC [32], along with MC. While the competitors exhibit over-smoothing or inaccuracies, particularly in the nose region, our method preserves surface details with consistent normals and outperforms in both accuracy and generalization. In addition, we present Figures 10 to 13 in Appendix G encompassing the comparison between the *Small* and *Large* models (Figure 10), visualizing the marching cubes varying the number of samples (Figure 11), detailed comparisons (Figure 12), and the gallery of the Stanford 3D Scanning meshes from the *Large* models (Figure 13).

Limitations. Although we implement the algorithm using pre-defined parallel tensor operations in PyTorch [33], achieving time and memory complexity of $\mathcal{O}(|\mathcal{V}|)$ as detailed in Appendix E, our method inherently requires forward passing the intermediate vertices during edge subdivisions iterating over neurons, which depends on previous steps (Algorithm 1). For the Stanford bunny using the *Small* model, the intermediate counts for vertices and edges exceed 11K and 19K, respectively, although they are eventually reduced to 4.3K vertices for face extraction. This overhead is negligible for the *Small* models, taking 0.97 seconds; however, for the *Large* models, the process takes 4.15 ± 0.61 seconds for $140.3 \pm 5.6K$ vertices with our optimized algorithm, compared to 1.31 and 11.88 seconds for marching cubes with 256 and 512 samples per axis, respectively. We note that parallelizing over multiple GPUs would be a reliable option to improve latency. For the discussions on satisfying the eikonal equation in practice and the generalization to other positional encodings using trilinear interpolation, please refer to Appendix C.1 and Appendix C.2, respectively.

7 Conclusions

In conclusion, our exploration of novel techniques in visualizing trilinear neural networks unveils its intricate structures and extends mesh extraction applications beyond CPWA functions. Focused on trilinear interpolating methods, our study yields novel theoretical insights and a practical methodology for precise mesh extraction. Notably, Theorem 4.5 shows the transformation of hypersurfaces to planes within the trilinear region under the eikonal constraint. Empirical validation using chamfer distance and efficiency, and angular distance affirms the correctness and parsimony of our methodology. The correlation analysis between eikonal loss and hypersurface planarity confirms the theoretical findings. Our proposed method for approximating intersecting points among three hypersurfaces, coupled with a vertex arrangement methodology, broadens the applications of our techniques. This work establishes a foundation for future research and practical applications in the evolving landscape of deep neural network analysis and toward real-time rendering mesh extraction.

Acknowledgments

I would like to express my sincere appreciation to my brilliant colleagues, Sangdoon Yun, Dongyoon Han, and Injae Kim, for their contributions to this work. Their constructive feedback and guidance have been instrumental in shaping the work. I also express my sincere thanks to the anonymous reviewers for their help in improving the manuscript. The NAVER Smart Machine Learning (NSML) platform [34] had been used for experiments.

References

- [1] Zheyang Zhang, Yongxing Wang, Peter K Jimack, and He Wang. Meshingnet: A new mesh generation method based on deep learning. In *International Conference on Computational Science*, pages 186–198. Springer, 2020.
- [2] Jiabao Lei and Kui Jia. Analytic marching: An analytic meshing solution from deep implicit surface networks. In *International Conference on Machine Learning*, pages 5789–5798. PMLR, 2020.
- [3] Jiabao Lei, Kui Jia, and Yi Ma. Learning and meshing from deep implicit surface networks using an efficient implementation of analytic marching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):10068–10086, 2021.
- [4] Arturs Berzins. Polyhedral complex extraction from relu networks using edge subdivision. In *International Conference on Machine Learning*. PMLR, 2023.
- [5] J Elisenda Grigsby and Kathryn Lindsey. On transversality of bent hyperplane arrangements and the topological expressiveness of relu neural networks. *SIAM Journal on Applied Algebra and Geometry*, 6(2):216–242, 2022.
- [6] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021.
- [7] Lior Yariv, Peter Hedman, Christian Reiser, Dor Verbin, Pratul P Srinivasan, Richard Szeliski, Jonathan T Barron, and Ben Mildenhall. Baked sdf: Meshing neural sdfs for real-time view synthesis. *arXiv preprint arXiv:2302.14859*, 2023.
- [8] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8456–8465, 2023.
- [9] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, 1987.
- [10] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. *Advances in Neural Information Processing Systems*, 34:6087–6101, 2021.
- [11] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [12] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4): 1–15, 2022.
- [13] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. TensorRF: Tensorial radiance fields. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022.
- [14] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020.

- [15] Heinrich Bruns. *Das eikonol*, volume 21. S. Hirzel, 1895.
- [16] Akio Doi and Akio Koide. An efficient method of triangulating equi-valued surfaces by using tetrahedral cells. *IEICE TRANSACTIONS on Information and Systems*, 74(1):214–224, 1991.
- [17] Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Mesh optimization. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 19–26, 1993.
- [18] Thiago Serra, Christian Tjandraatmadja, and Srikumar Ramalingam. Bounding and counting linear regions of deep neural networks. In *International Conference on Machine Learning*, pages 4558–4566. PMLR, 2018.
- [19] Maithra Raghu, Ben Poole, Jon Kleinberg, Surya Ganguli, and Jascha Sohl-Dickstein. On the expressive power of deep neural networks. In *international conference on machine learning*, pages 2847–2854. PMLR, 2017.
- [20] Boris Hanin and David Rolnick. Complexity of linear regions in deep networks. In *International Conference on Machine Learning*, pages 2596–2604. PMLR, 2019.
- [21] Ahmed Imtiaz Humayun, Randall Balestriero, Guha Balakrishnan, and Richard G Baraniuk. Splinecam: Exact visualization and characterization of deep network geometry and decision boundaries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3789–3798, 2023.
- [22] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [23] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020.
- [24] Eric Katz. What is... tropical geometry. *Notices of the AMS*, 64(4), 2017.
- [25] Ilia Itenberg, Grigory Mikhalkin, and Eugenio I Shustin. *Tropical algebraic geometry*, volume 35. Springer Science & Business Media, 2009.
- [26] Diane Maclagan and Bernd Sturmfels. *Introduction to tropical geometry*, volume 161. American Mathematical Society, 2021.
- [27] Liwen Zhang, Gregory Naitzat, and Lek-Heng Lim. Tropical geometry of deep neural networks. In *International Conference on Machine Learning*, pages 5824–5832. PMLR, 2018.
- [28] Thomas Friedrich et al. *Elementary geometry*, volume 43. American Mathematical Soc., 2008.
- [29] Alan Edelman and Hiroshi Murakami. Polynomial roots from companion matrix eigenvalues. *Mathematics of Computation*, 64(210):763–776, 1995.
- [30] Boris Hanin and David Rolnick. Deep relu networks have surprisingly few activation patterns. *Advances in neural information processing systems*, 32, 2019.
- [31] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996.
- [32] Zhiqin Chen, Andrea Tagliasacchi, Thomas Funkhouser, and Hao Zhang. Neural dual contouring. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022.
- [33] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

- [34] Hanjoo Kim, Minkyu Kim, Dongjoo Seo, Jinwoong Kim, Heungseok Park, Soeun Park, Hyunwoo Jo, KyungHyun Kim, Youngil Yang, Youngkwan Kim, Nako Sung, and Jung-Woo Ha. NSML: Meet the MLAAS platform with a real-world case study. *arXiv preprint arXiv:1810.09957*, 2018.
- [35] 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*, pages 209–216, 1997.
- [36] Edoardo Remelli, Artem Lukoianov, Stephan Richter, Benoit Guillard, Timur Bagautdinov, Pierre Baque, and Pascal Fua. MeshSDF: Differentiable iso-surface extraction. *Advances in Neural Information Processing Systems*, 33:22468–22478, 2020.
- [37] 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 Trans. Graph.*, 42(4):37–1, 2023.

Appendix

Table of contents

A Tropical geometry and tropical algebra of neural networks	14
A.1 Tropical geometry	14
A.2 Tropical hypersurface	15
A.3 Tropical algebra of neural networks	15
B Implementation details	16
B.1 Zero-trailing binarization function	16
B.2 Sorting polygon vertices and HashGrid marks	16
C Further discussions	17
C.1 Satisfying the eikonal equation and approximation	17
C.2 Generalization to other positional encodings using trilinear interpolation	17
D Theoretical proofs	17
E Complexity analysis	23
F Supplementary experiments	24
F.1 Complete results	24
F.2 Mesh simplification using QEM	26
F.3 Comparison with alternate approaches	26
F.4 Error of the learned SDF	27
F.5 Effect of model sizes	27
G Visualizations	28

A Tropical geometry and tropical algebra of neural networks

In this section, we present an overview of tropical geometry (Appendix A.1) to provide a formal definition of the tropical hypersurface (Appendix A.2). This hypersurface is significant as it serves as the decision boundary for ReLU neural networks optimizing an SDF, resulting in the formation of a polyhedral complex. For this, we delve into the examination of the tropical algebra of neural networks (Appendix A.3). Furthermore, we discuss the edge subdivision algorithm [4] for extracting the tropical hypersurface (Section 3.2), and we extend this discussion to include considerations for trilinear interpolation in Section 4. Given that Appendix A.1 to A.3 offer insights into tropical geometry, illustrating how each neuron contributes to forming the polyhedral complex, you may choose to skip these subsections if you are already acquainted with these concepts.

A.1 Tropical geometry

Tropical geometry [25, 26] describes polynomials and their geometric characteristics. As a skeletonized version of algebraic geometry, tropical geometry represents piecewise linear meshes using *tropical semiring* and *tropical polynomial*.

The tropical semiring is an algebraic structure that extends real numbers and $-\infty$ with the two operations of *tropical sum* \oplus and *tropical product* \odot . Tropical sum and product are conveniently replaced with $\max(\cdot)$ and $+$, respectively. Notice that $-\infty$ is the tropical sum identity, zero is the tropical product identity, and these satisfy associativity, commutativity, and distributivity. For example, a classical polynomial $x^3 + 2xy + y^4$ would represent $\max(3x, x + y + 2, 4y)$. The

classical polynomial is rewritten for a tropical polynomial by naively replacing $+$ with \oplus and \cdot with \odot , respectively. After this, we rewrite by the definitions of tropical operations as $\max(3x, x+y+2, 4y)$. Roughly speaking, this is the tropicalization of the polynomial, denoted tropical polynomial or simply f . Note that this example (ours with \max notation) comes from the Wikipedia ⁶.

Additionally, we present a graphical diagram in Figure 7 illustrating how the tropical polynomial divides the 2D space for clearer understanding. Please see the caption for further details.

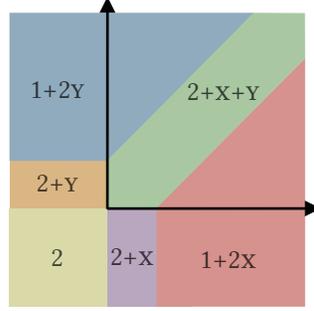


Figure 7: A classical polynomial $x^2 + y^2 + 2xy + 2x + 2y + 2$ turns into a tropical polynomial as $\max(1 + 2x, 1 + 2y, 2 + x + y, 2 + x, 2 + y, 2)$, taking into account the implicit coefficient of 1. We redraw Figure 1 of the tropical curve from Zhang et al. [27] as a region map. Each color indicates a region where each tropical monomial has maximum.

A.2 Tropical hypersurface

The set of points where a tropical polynomial f is non-differentiable, mainly due to tropical sum, $\max(\cdot)$, is called *tropical hypersurface*, denoted $V(f)$. Note that V is an analogy to the *vanishing set* of a polynomial.

Definition A.1 (Tropical hypersurface, restated). The tropical hypersurface of a tropical polynomial $f(\mathbf{x}) = c_1 \mathbf{x}^{\alpha_1} \oplus \dots \oplus c_r \mathbf{x}^{\alpha_r}$ is defined as:

$$V(f) = \{\mathbf{x} \in \mathbb{R}^d : c_i \mathbf{x}^{\alpha_i} = c_j \mathbf{x}^{\alpha_j} = f(\mathbf{x}) \text{ for some } \alpha_i \neq \alpha_j\}$$

where the tropical monomial $c_i \mathbf{x}^{\alpha_i}$ with d -variate \mathbf{x} is defined as $c_i \odot x_1^{\alpha_{i,1}} \odot \dots \odot x_d^{\alpha_{i,d}}$ with d -variate $\alpha_i = (a_{i,1}, \dots, a_{i,d})$ and *tropical product* $\odot := +$. In other words, this is the set of points where $f(\mathbf{x})$ equals two or more monomials in f .

A.3 Tropical algebra of neural networks

Definition A.2 (ReLU neural networks, restated). Assuming input and output dimensions are consistent, L -layer neural networks with ReLU activation are a composition of functions defined as:

$$\nu^{(L)} = \rho^{(L)} \circ \sigma^{(L-1)} \circ \rho^{(L-1)} \dots \sigma^{(1)} \circ \rho^{(1)} \quad (12)$$

where $\sigma(\mathbf{x}) = \max(\mathbf{x}, 0)$ and $\rho^{(i)}(\mathbf{x}) = \mathbf{W}^{(i)} \mathbf{x} + \mathbf{b}^{(i)}$.

Although the neural networks ν is generally non-convex, it is known that the difference of two tropical *signomials* can represent ν [27]. A tropical signomial takes the same form as a tropical polynomial (ref. Definition 3.1), but it additionally allows the exponentials α to be real values rather than integers. With tropical signomials F , G , and H ,

$$\nu^{(l)}(\mathbf{x}) = H^{(l)}(\mathbf{x}) - G^{(l)}(\mathbf{x}) \quad (13)$$

$$\sigma^{(l)} \circ \nu^{(l)}(\mathbf{x}) = F^{(l)}(\mathbf{x}) - G^{(l)}(\mathbf{x}) \quad (14)$$

⁶https://en.wikipedia.org/wiki/Tropical_geometry

where $F^{(l)}$, $G^{(l)}$, and $H^{(l)}$ are as follows:

$$F^{(l)}(\mathbf{x}) = \max(H^{(l)}(\mathbf{x}), G^{(l)}(\mathbf{x})) \quad (15)$$

$$G^{(l)}(\mathbf{x}) = \mathbf{W}_+^{(l)} G^{(l-1)}(\mathbf{x}) + \mathbf{W}_-^{(l)} F^{(l-1)}(\mathbf{x}) \quad (16)$$

$$H^{(l)}(\mathbf{x}) = \mathbf{W}_+^{(l)} F^{(l-1)}(\mathbf{x}) + \mathbf{W}_-^{(l)} G^{(l-1)}(\mathbf{x}) + \mathbf{b}^{(l)} \quad (17)$$

while $\mathbf{W}^{(l)} = \mathbf{W}_+^{(l)} - \mathbf{W}_-^{(l)}$ where $\mathbf{W}_+^{(l)}$ and $\mathbf{W}_-^{(l)} \in \mathbb{R}^+$.

Proposition A.3 (Decision boundary of neural networks, restated). *The decision boundary for L -layer neural networks with ReLU activation $\nu^{(L)}(\mathbf{x}) = 0$ is where two tropical monomials, or two arguments of max, equal (Definition 3.1), as follows:*

$$\mathcal{B} \subseteq \{\mathbf{x} : \nu_j^{(i)}(\mathbf{x}) = H_j^{(i)}(\mathbf{x}) - G_j^{(i)}(\mathbf{x}) = 0, i \in (1, \dots, L), j \in (1, \dots, H) \text{ if } i \neq L \text{ else } (1)\}$$

where the subscript j denotes the j -th neuron, assuming the hidden size of neural networks is H , while the L -th layer has a single output for rather discussions as an SDF. Remind that tropical operations satisfy the usual laws of arithmetic for recursive consideration.

This implies that piecewise hyperplanes corresponding to the preactivations $\{\nu_j^{(i)}\}$ are the candidates shaping the decision boundary. Each linear region has the same pattern of zero masking by $\text{ReLU}(\cdot)$, effectively making it linear as an equation of a plane $\mathbf{A}\mathbf{x} + \mathbf{b} = 0$ for some \mathbf{A} and \mathbf{b} . In turn, if the masking pattern changes as an input is continually changed, it moves to another linear region. From a geometric point of view, this is why the number of linear regions grows polynomially with the hidden size H and exponentially with the number of layers L [18, 27].

B Implementation details

B.1 Zero-trailing binarization function

B is zero-trailing binarization function, consisting of $D = 3$ bits. The zero-trailing means the number in bits is increasing left-aligned. For example, the number one is “100”, while the numbers two and four are “010” and “001”, respectively. The number three is “110”. From the left, each bit indicates x, y, and z axis in our notation. This convention seems to be widely used in the CUDA implementation in C language. Notice that we use the j subscript of B_j to select the axis. Naturally, $i = 0$ and $i = 7$ indicate the opposite two corner points.

B.2 Sorting polygon vertices and HashGrid marks

Algorithm 3 describes the algorithm in Section 4.3, which are sorting polygon vertices implicitly representing the direction of faces. If the viewing direction aligns opposite to the normal, the order of vertices follows a counter-clockwise arrangement. We determine the vertex order based on the normal derived from the Jacobian of the mean of vertices. We identify the counter-clockwise arrangement of vertices using dot product to compute relative angles $\theta \in [-\pi, \pi]$ and cross product for getting direction. Algorithm 4 describes the algorithm in Section 5, calculating the HashGrid marks from the hyperparameters of a HashGrid.

Algorithm 3 Sorting Polygon Vertices

Input: vertices \mathbf{x}_i , normal \mathbf{n} , $N = |\{\mathbf{x}_i\}|$
 $\mu = \sum_i \mathbf{x}_i / N$
for $i = 0$ **to** $N - 1$ **do**
 $\mathbf{v}_i = (\mathbf{x}_i - \mu) / \|\mathbf{x}_i - \mu\|_2$
 $\mathbf{c}_i = \langle \mathbf{v}_i, \mathbf{v}_0 \rangle$, $\mathbf{d}_i = \langle \mathbf{v}_i \times \mathbf{v}_0, \mathbf{n} \rangle$
 $\theta_i = \mathbf{c}_i(2 \cdot \mathbf{1}_{\mathbf{d}_i \geq 0} - 1) + 2 \cdot \mathbf{1}_{\mathbf{d}_i < 0}$
end for
return $\text{ArgSort}(\{-\theta_i\})$

Algorithm 4 HashGrid Marks

Input: Resolution levels L , base resolution N_{\min} , max resolution N_{\max}
 $m = \{0, 1\}$
 $b = \log_2(N_{\max}/N_{\min}) / (L - 1)$
for $l = 0$ **to** $L - 1$ **do**
 $s = 1 / (N_{\min} \cdot \exp_2(l \cdot b) - 1)$
 $m \leftarrow m + \text{Max}(\text{Arange}(-s/2, 1, s), 0)$
end for
return $\text{Sort}(\text{Unique}(m))$

C Further discussions

C.1 Satisfying the eikonal equation and approximation

Theorem 4.7 handles where the eikonal equation does not make perfect hyperplanes. First, we observed the characteristics of trilinear hypersurfaces as illustrated in Figure 9 in Appendix. Among 64 edge scenarios within a trilinear region, a diagonal plane intersects with the hypersurface in nearly all cases. Moreover, we can take advantage of it since we can eliminate one variable by assuming the curved edge lies on the $x = z$ plane. In this light, Theorem 4.7 used linear algebra to rearrange two trilinear equations to get Equations (71) and (74). As we mentioned in the manuscript, the new vertices lie on at least two hypersurfaces, and the new edges exist on the same hypersurface, while the eikonal constraint minimizes any associated error. This error is tolerated by the hyperparameter $\epsilon=1e-4$ (Section 6) as specified in Definition 3.4.

Figure 5 also suggests empirical evidence. Empirically, the flatness error, derived from the proof of Theorem D.5, is effectively controlled by the learning rate of the eikonal loss. In the experiments, we used 0.01 to give our results, while many other neural implicit surface learning methods were used to exploit a higher learning rate of 0.1 (e.g., Yariv et al. [7]). Note that, in practice, the eikonal loss is only applied to near-surface using clamps (if the target position is obviously far from a surface, ignore it; it is wise since we are concerned with the surfaces, not wanting exact SDFs for all space) for effective learning, which is sufficient to our extraction method.

Moreover, the smaller grid size of the *Large* model (ref. Section 6) in Figure 4 views a curved surface in a smaller interval, and it tends to regard it as a plane. From the plotting of *Small* to *Large*, our CD (colored crosses) is decreasing to zero, consistently with better CE, confirming this speculation.

C.2 Generalization to other positional encodings using trilinear interpolation

We can say that *any encoding using trilinear interpolation (that) we defined can be applied in our method*. Here, the definition is implicitly referred to Definition 4.1 (Trilinear interpolation), where the cubic corners are represented by $\mathbf{H} \in \mathbb{R}^{D \times F}$. Since we do not rely on a specific mapping function $f : \mathbf{x} \rightarrow (\mathbf{H}_0, \mathbf{H}_1 \dots \mathbf{H}_7)$, we may apply our extraction method to any encoding using trilinear interpolation that we defined. Specifically, InstantNGP [12] used a hashing function to map a corner point to the arbitrary entry of hash tables, allowing gracious collisions with online adaptivity (refer to their Sec. 3), while TensorRF [13] can be interpreted that it used a mapping function to the summation of the multiplications of two vectors from the vector-matrix factorization. Therefore, if we have the full ranks for R_1, R_2, R_3 , in Eqn. 3 in Chen et al. [13], among the VM factorization, which can be equivalently represented by the HashGrid with a sufficient number of hash table sizes.

Then, the discussion should be directed toward their encoding characteristics, which impact our extraction method due to their online adaptivity (HashGrid) or low-rank factorization (TensorRF). When we train an SDF, the optimization to global minima is generally difficult to achieve. Instead, previous works on neural implicit surface learning methods resort to optimizing the SDF on near-surface since we are interested in modeling the zero-set surfaces, allowing some errors for the points that are far from the surfaces. Practically, this is done by defining a loss function using clamps and ignoring far-surface points. For this reason, the online adaptivity of HashGrid or low-rank factorization of TensorRF can exploit this nature: focused learning of the surfaces adaptively allocating their learnable parameters to the near-surface, which means there might be enough room for the regularization by the eikonal equation. The investigation of the efficiency of representing the implicit neural surfaces using various positional encodings is seemingly out of the scope of this work since our main interest lies in the theoretical exposition of the mesh extraction from it and setting up a feasible bridge between theory and practice.

D Theoretical proofs

Lemma D.1. (*Nested trilinear interpolation, restated*) *Let a nested cube be inside of a unit cube, where its positions of the eight corners deviate $-\mathbf{a}_j$ or $+\mathbf{b}_j$, such that $\mathbf{a}_j, \mathbf{b}_j \geq 0$, from \mathbf{w}_j for each dimension j , using the notations of Definition 4.1. The eight corners of the nested cube are the trilinear interpolations of the eight corners of the unit cube. Then, the trilinear interpolation with the unit cube and the nested cube for \mathbf{w} is identical.*

Proof. Let a left-aligned and zero-trailing binarization function be $\mathbf{B}(\cdot) \in \{0, 1\}^D$ and $D = 3$. The partial derivative of the trilinear interpolation with the nested cube $\hat{\mathbf{H}}$ with respect to \mathbf{H}_7 and \mathbf{H}_0 would be:

$$\frac{\partial \phi(\mathbf{w}; \hat{\mathbf{H}})}{\partial \mathbf{H}_7} = \sum_{i=0}^{2^D-1} \prod_{j=1}^D \frac{(1 - \mathbf{B}(i)_j) \mathbf{b}_j (\mathbf{w}_j - \mathbf{a}_j) + \mathbf{B}(i)_j \mathbf{a}_j (\mathbf{w}_j + \mathbf{b}_j)}{\mathbf{a}_j + \mathbf{b}_j} \quad (18)$$

$$= \sum_{i=0}^{2^{D-1}-1} \frac{\mathbf{b}_D (\mathbf{w}_D - \mathbf{a}_D) + \mathbf{a}_D (\mathbf{w}_D + \mathbf{b}_D)}{\mathbf{a}_D + \mathbf{b}_D} \prod_{j=1}^{D-1} \frac{(1 - \mathbf{B}(i)_j) \mathbf{b}_j (\mathbf{w}_j - \mathbf{a}_j) + \mathbf{B}(i)_j \mathbf{a}_j (\mathbf{w}_j + \mathbf{b}_j)}{\mathbf{a}_j + \mathbf{b}_j} \quad (19)$$

$$= \mathbf{w}_D \sum_{i=0}^{2^{D-1}-1} \prod_{j=1}^{D-1} \frac{(1 - \mathbf{B}(i)_j) \mathbf{b}_j (\mathbf{w}_j - \mathbf{a}_j) + \mathbf{B}(i)_j \mathbf{a}_j (\mathbf{w}_j + \mathbf{b}_j)}{\mathbf{a}_j + \mathbf{b}_j} \quad (20)$$

$$= \prod_{j=1}^D \mathbf{w}_j = \frac{\partial \phi(\mathbf{w}; \mathbf{H})}{\partial \mathbf{H}_7} \quad (21)$$

and

$$\frac{\partial \phi(\mathbf{w}; \hat{\mathbf{H}})}{\partial \mathbf{H}_0} = \sum_{i=0}^{2^D-1} \prod_{j=1}^D \frac{(1 - \mathbf{B}(i)_j) \mathbf{b}_j (1 - \mathbf{w}_j + \mathbf{a}_j) + \mathbf{B}(i)_j \mathbf{a}_j (1 - \mathbf{w}_j - \mathbf{b}_j)}{\mathbf{a}_j + \mathbf{b}_j} \quad (22)$$

$$= (1 - \mathbf{w}_D) \sum_{i=0}^{2^{D-1}-1} \prod_{j=1}^{D-1} \frac{(1 - \mathbf{B}(i)_j) \mathbf{b}_j (1 - \mathbf{w}_j + \mathbf{a}_j) + \mathbf{B}(i)_j \mathbf{a}_j (1 - \mathbf{w}_j - \mathbf{b}_j)}{\mathbf{a}_j + \mathbf{b}_j} \quad (23)$$

$$= \prod_{j=1}^D (1 - \mathbf{w}_j) = \frac{\partial \phi(\mathbf{w}; \mathbf{H})}{\partial \mathbf{H}_0}, \quad (24)$$

respectively. In a similar way, we can describe the other cases \mathbf{H}_1 to \mathbf{H}_6 of the nested trilinear weights since there are two cases for each dimension inside the product. \square

Proposition D.2 (Cubic Bézier curve). *Trilinear interpolation from Definition 4.1 with the weights on the diagonal line $\mathbf{t} = (t, t, t)^\top$ where $t \in [0, 1]$ is a cubic Bézier curve.*

Proof. We rewrite the trilinear interpolation from Definition 4.1 with the weight \mathbf{t} as follows:

$$\psi(\mathbf{t}; \mathbf{H}) = \sum_{i=0}^{2^3-1} \omega(i, \mathbf{t}) \cdot \mathbf{H}_i \quad (25)$$

$$= \sum_{i=0}^{2^3-1} \prod_{j=1}^3 \left[(1 - \mathbf{B}(i)_j) (1 - t) + \mathbf{B}(i)_j t \right] \cdot \mathbf{H}_i \quad (26)$$

$$= (1 - t)^3 \mathbf{H}_0 + (1 - t)^2 t (\mathbf{H}_1 + \mathbf{H}_2 + \mathbf{H}_4) + (1 - t) t^2 (\mathbf{H}_3 + \mathbf{H}_5 + \mathbf{H}_6) + t^3 \mathbf{H}_7 \quad (27)$$

$$= (1 - t)^3 \mathbf{H}_0 + 3(1 - t)^2 t \left(\frac{\mathbf{H}_1 + \mathbf{H}_2 + \mathbf{H}_4}{3} \right) + 3(1 - t) t^2 \left(\frac{\mathbf{H}_3 + \mathbf{H}_5 + \mathbf{H}_6}{3} \right) + t^3 \mathbf{H}_7 \quad (28)$$

$$= (1 - t)^3 \mathbf{P}_0 + 3(1 - t)^2 t \mathbf{P}_1 + 3(1 - t) t^2 \mathbf{P}_2 + t^3 \mathbf{P}_3 \quad (29)$$

Four points $\mathbf{P}_0 := \mathbf{H}_0$, $\mathbf{P}_1 := \frac{\mathbf{H}_1 + \mathbf{H}_2 + \mathbf{H}_4}{3}$, $\mathbf{P}_2 := \frac{\mathbf{H}_3 + \mathbf{H}_5 + \mathbf{H}_6}{3}$, and $\mathbf{P}_3 := \mathbf{H}_7$ define the cubic Bézier curve. \square

As t is changed from zero to one, the curve starts at P_0 going toward P_1 and turns in the direction of P_2 before arriving at P_3 . Usually, the curve does not pass P_1 nor P_2 ; but these give the directional guide for the curve.

Lemma D.3 (Hypersurfaces in a trilinear space). *In a piecewise trilinear region where \mathbf{x} lies, hypersurfaces of the piecewise trilinear networks are defined as follows, following the notation of Definition 4.1:*

$$0 = \tilde{\nu}(\mathbf{x}) = \sum_{i=0}^{2^D-1} \omega(i, \mathbf{w}) \cdot \tilde{\nu}(\mathbf{x}_i) \quad (30)$$

where $\mathbf{w} \in [0, 1]^D$ is the relative position in a grid cell and $\{\mathbf{x}_i\}$ and $\{\mathbf{H}_i\}$ are the corresponding corner positions and representations, respectively, in a trilinear space.

Proof. In a piecewise trilinear region, we can use the linearity of ν as follows:

$$\tilde{\nu}(\mathbf{x}) = (\nu \circ \tau)(\mathbf{x}) = \nu\left(\sum_{i=0}^{2^D-1} \omega(i, \mathbf{w}) \cdot \mathbf{H}_i\right) \quad (31)$$

$$= \sum_{i=0}^{2^D-1} \omega(i, \mathbf{w}) \cdot \nu(\mathbf{H}_i) \quad (32)$$

$$= \sum_{i=0}^{2^D-1} \omega(i, \mathbf{w}) \cdot \tilde{\nu}(\mathbf{x}_i) \quad (33)$$

which concludes the proof. \square

Lemma D.4 (Curved edge of two hypersurfaces, restated). *In a piecewise trilinear region, let an edge $(\mathbf{x}_0, \mathbf{x}_7)$ be the intersection of two hypersurfaces $\tilde{\nu}_i(\mathbf{x}) = \tilde{\nu}_j(\mathbf{x}) = 0$, while \mathbf{x}_0 and \mathbf{x}_7 are on the two hypersurfaces. Then, the edge is defined as:*

$$\{\mathbf{x} : \tau(\mathbf{x}) = (1-t)\tau(\mathbf{x}_0) + t\tau(\mathbf{x}_7) \quad \text{where } \mathbf{x} \in \mathbb{R}^D \quad \text{and } t \in [0, 1]\}. \quad (34)$$

Proof. By Definition 4.3,

$$\tilde{\nu}_i(\mathbf{x}_0) = \tilde{\nu}_i(\mathbf{x}_7) \Leftrightarrow \nu_i(\tau(\mathbf{x}_0)) = \nu_i(\tau(\mathbf{x}_7)), \quad (35)$$

while ν_i is linear in a piecewise trilinear region. Since it similarly goes for ν_j , the intersection line of two hyperplanes ν_i and ν_j is $(1-t)\tau(\mathbf{x}_0) + t\tau(\mathbf{x}_7)$ where $t \in \mathbb{R}$. In other words, the edge is on the hypersurface $\nu_i((1-t)\tau(\mathbf{x}_0) + t\tau(\mathbf{x}_7)) = 0$ where $t \in \mathbb{R}$. The edge is defined as follows:

$$\{\mathbf{x} : \tau(\mathbf{x}) = (1-t)\tau(\mathbf{x}_0) + t\tau(\mathbf{x}_7) \quad \text{where } \mathbf{x} \in \mathbb{R}^D \quad \text{and } t \in \mathbb{R}\} \quad (36)$$

which concludes the proof. \square

Notice that when $\tau(\mathbf{x})_1 + \tau(\mathbf{x})_2 + \tau(\mathbf{x})_4 = \tau(\mathbf{x})_3 + \tau(\mathbf{x})_5 + \tau(\mathbf{x})_6 = 0$, at least one connecting edge can be identified. This occurs along the diagonal line $\mathbf{x} = t \cdot \mathbf{1}$, where $t \in [0, 1]$. Along this line, $\tau(\mathbf{x})$ represents a cubic Bézier curve, as shown in Proposition D.2 with the control points of P_1 and P_2 set to zeros, which is the line.

Theorem D.5 (Hypersurface and eikonal constraint, restated). *A hypersurface $\tau(\mathbf{x}) = 0$ intersects two points $\tau(\mathbf{x}_0) = \tau(\mathbf{x}_7) = 0$ while $\tau(\mathbf{x}_{1\dots 6}) \neq 0$ for the remaining six points. These points form a cube, with \mathbf{x}_0 and \mathbf{x}_7 positioned on the diagonal of the cube. The hypersurface satisfies the eikonal constraint $\|\nabla\tau(\mathbf{x})\|_2^2 = 1$ for all $\mathbf{x} \in [0, 1]^3$. Then, the hypersurface of $\tau(\mathbf{x}) = 0$ is a plane.*

Proof. By the definitions of τ in Definition 4.3 and the eikonal constraint,

$$\|\nabla\tau(\mathbf{x})\|_2^2 = \left(\frac{\partial\tau(\mathbf{x})}{\partial x}\right)^2 + \left(\frac{\partial\tau(\mathbf{x})}{\partial y}\right)^2 + \left(\frac{\partial\tau(\mathbf{x})}{\partial z}\right)^2 = 1 \quad (37)$$

$$\frac{\partial\|\nabla\tau(\mathbf{x})\|_2^2}{\partial x} = 2\left(\frac{\partial\tau(\mathbf{x})}{\partial y}\right)\left(\frac{\partial^2\tau(\mathbf{x})}{\partial y\partial x}\right) + 2\left(\frac{\partial\tau(\mathbf{x})}{\partial z}\right)\left(\frac{\partial^2\tau(\mathbf{x})}{\partial z\partial x}\right) = 0 \quad (38)$$

$$\frac{\partial\|\nabla\tau(\mathbf{y})\|_2^2}{\partial y} = 2\left(\frac{\partial\tau(\mathbf{x})}{\partial x}\right)\left(\frac{\partial^2\tau(\mathbf{x})}{\partial x\partial y}\right) + 2\left(\frac{\partial\tau(\mathbf{x})}{\partial z}\right)\left(\frac{\partial^2\tau(\mathbf{x})}{\partial z\partial y}\right) = 0 \quad (39)$$

$$\frac{\partial\|\nabla\tau(\mathbf{z})\|_2^2}{\partial z} = 2\left(\frac{\partial\tau(\mathbf{x})}{\partial x}\right)\left(\frac{\partial^2\tau(\mathbf{x})}{\partial x\partial z}\right) + 2\left(\frac{\partial\tau(\mathbf{x})}{\partial y}\right)\left(\frac{\partial^2\tau(\mathbf{x})}{\partial y\partial z}\right) = 0 \quad (40)$$

where

$$\begin{aligned} \frac{\partial\tau(\mathbf{x})}{\partial y} &= (1-x)\left((1-z)(\tau(\mathbf{x}_2) - \tau(\mathbf{x}_0)) + z(\tau(\mathbf{x}_6) - \tau(\mathbf{x}_4))\right) \\ &\quad + x\left((1-z)(\tau(\mathbf{x}_3) - \tau(\mathbf{x}_1)) + z(\tau(\mathbf{x}_7) - \tau(\mathbf{x}_5))\right). \end{aligned} \quad (41)$$

However, since $\tau(\mathbf{x}_0) = \tau(\mathbf{x}_7) = 0$ while $\tau(\mathbf{x}_2) \neq 0$ and $\tau(\mathbf{x}_5) \neq 0$, there is no solution $\tau(\mathbf{x}_{1\dots 6})$ of $\frac{\partial\tau(\mathbf{x})}{\partial y} = 0$ for all \mathbf{x} . Likewise, the same goes for $\frac{\partial\tau(\mathbf{x})}{\partial x}$ and $\frac{\partial\tau(\mathbf{x})}{\partial z}$. Additionally, the three partial derivatives can be rewritten as follows:

$$bC + cB = aC + cA = aB + bA = 0 \quad (42)$$

where $a \neq 0$, $b \neq 0$, and $c \neq 0$. By substituting $B = -bC/c$ and $A = -aC/c$, we can show that $-abC/c - abC/c = 0 \Leftrightarrow C = 0$, which makes $A = B = C = 0$. Therefore, we rewrite A, B, and C as follows:

$$\frac{\partial^2\tau(\mathbf{x})}{\partial y\partial z} = (1-x)(-\tau(\mathbf{x}_2) - \tau(\mathbf{x}_4) + \tau(\mathbf{x}_6)) + x(+\tau(\mathbf{x}_1) - \tau(\mathbf{x}_3) - \tau(\mathbf{x}_5)) \quad (43)$$

$$\frac{\partial^2\tau(\mathbf{x})}{\partial z\partial x} = (1-y)(-\tau(\mathbf{x}_1) - \tau(\mathbf{x}_4) + \tau(\mathbf{x}_5)) + y(+\tau(\mathbf{x}_2) - \tau(\mathbf{x}_3) - \tau(\mathbf{x}_6)) \quad (44)$$

$$\frac{\partial^2\tau(\mathbf{x})}{\partial x\partial y} = (1-z)(-\tau(\mathbf{x}_1) - \tau(\mathbf{x}_2) + \tau(\mathbf{x}_3)) + z(+\tau(\mathbf{x}_4) - \tau(\mathbf{x}_5) - \tau(\mathbf{x}_6)), \quad (45)$$

respectively. The following conditions would satisfy the eikonal constraint:

$$-\tau(\mathbf{x}_2) - \tau(\mathbf{x}_4) + \tau(\mathbf{x}_6) = 0 \quad (46)$$

$$\tau(\mathbf{x}_1) - \tau(\mathbf{x}_3) - \tau(\mathbf{x}_5) = 0 \quad (47)$$

$$-\tau(\mathbf{x}_1) - \tau(\mathbf{x}_4) + \tau(\mathbf{x}_5) = 0 \quad (48)$$

$$+\tau(\mathbf{x}_2) - \tau(\mathbf{x}_3) - \tau(\mathbf{x}_6) = 0 \quad (49)$$

$$-\tau(\mathbf{x}_1) - \tau(\mathbf{x}_2) + \tau(\mathbf{x}_3) = 0 \quad (50)$$

$$+\tau(\mathbf{x}_4) - \tau(\mathbf{x}_5) - \tau(\mathbf{x}_6) = 0 \quad (51)$$

From Equation (48) and Equation (51),

$$\tau(\mathbf{x}_1) = -\tau(\mathbf{x}_6) \quad (52)$$

From Equation (47) and Equation (50),

$$\tau(\mathbf{x}_2) = -\tau(\mathbf{x}_5) \quad (53)$$

From Equation (47) and Equation (48),

$$\tau(\mathbf{x}_4) = -\tau(\mathbf{x}_3) \quad (54)$$

From Equation (50), Equation (51), and Equation (54),

$$\tau(\mathbf{x}_1) + \tau(\mathbf{x}_2) + \tau(\mathbf{x}_4) = 0 \quad (55)$$

$$\tau(\mathbf{x}_3) + \tau(\mathbf{x}_5) + \tau(\mathbf{x}_6) = 0 \quad (56)$$

Using Equation (52), Equation (53), Equation (54), Equation (55), and Equation (56),

$$\begin{aligned}\tau(\mathbf{x}) &= (x(1-y)(1-z) - (1-x)yz)\tau(\mathbf{x}_1) \\ &\quad + ((1-x)y(1-z) - x(1-y)z)\tau(\mathbf{x}_2) \\ &\quad + ((1-x)(1-y)z - xy(1-z))\tau(\mathbf{x}_4)\end{aligned}\tag{57}$$

$$= (x + \alpha)\tau(\mathbf{x}_1) + (y + \alpha)\tau(\mathbf{x}_2) + (z + \alpha)\tau(\mathbf{x}_4)\tag{58}$$

$$= x\tau(\mathbf{x}_1) + y\tau(\mathbf{x}_2) + z\tau(\mathbf{x}_4) + \alpha(\tau(\mathbf{x}_1) + \tau(\mathbf{x}_2) + \tau(\mathbf{x}_4))\tag{59}$$

$$= x\tau(\mathbf{x}_1) + y\tau(\mathbf{x}_2) + z\tau(\mathbf{x}_4) = 0\tag{60}$$

where $\alpha = 2xyz - xy - yz - zx$.

This is an equation of a plane where its normal vector is $(\tau(\mathbf{x}_1), \tau(\mathbf{x}_2), \tau(\mathbf{x}_4))$. Again, to satisfy the eikonal constraint,

$$\tau(\mathbf{x}_1)^2 + \tau(\mathbf{x}_2)^2 + \tau(\mathbf{x}_4)^2 = 1\tag{61}$$

$$\tau(\mathbf{x}_3)^2 + \tau(\mathbf{x}_5)^2 + \tau(\mathbf{x}_6)^2 = 1,\tag{62}$$

which concludes the proof. \square

Corollary D.6 (Affine-transformed hypersurface and eikonal constraint, restated). *Let $\nu_0(\mathbf{x})$ and $\nu_1(\mathbf{x})$ be two affine transformations defined by $\mathbf{W}_i^\top \mathbf{x} + \mathbf{b}_i$, $i \in \{0, 1\}$. A hypersurface $\nu_0(\tau(\mathbf{x})) = 0$ passing two points $\nu_0(\tau(\mathbf{x}_0)) = \nu_0(\tau(\mathbf{x}_7)) = 0$ while $\nu_0(\tau(\mathbf{x}_{1\dots 6})) \neq 0$ satisfies the eikonal constraint $\|\nabla \nu_1(\tau(\mathbf{x}))\|_2^2 = 1$ for all $\mathbf{x} \in [0, 1]^3$. Then, the hypersurface is a plane.*

Proof. Similarly to Theorem D.5, we investigate the Jacobian of the eikonal constraint to be zero as follows:

$$\|\nabla \nu_1(\tau(\mathbf{x}))\|_2 = \left(\frac{\partial \nu_1}{\partial \tau} \frac{\partial \tau(\mathbf{x})}{\partial x}\right)^2 + \left(\frac{\partial \nu_1}{\partial \tau} \frac{\partial \tau(\mathbf{x})}{\partial y}\right)^2 + \left(\frac{\partial \nu_1}{\partial \tau} \frac{\partial \tau(\mathbf{x})}{\partial z}\right)^2 = 1\tag{63}$$

$$\frac{\partial \|\nabla \nu_1(\tau(\mathbf{x}))\|_2}{\partial x} = 2\left(\frac{\partial \nu_1}{\partial \tau} \frac{\partial \tau(\mathbf{x})}{\partial y}\right)\left(\frac{\partial \nu_1}{\partial \tau} \frac{\partial^2 \tau(\mathbf{x})}{\partial y \partial x}\right) + 2\left(\frac{\partial \nu_1}{\partial \tau} \frac{\partial \tau(\mathbf{x})}{\partial z}\right)\left(\frac{\partial \nu_1}{\partial \tau} \frac{\partial^2 \tau(\mathbf{x})}{\partial z \partial x}\right) = 0.\tag{64}$$

Since $\partial \nu_1 / \partial \tau$ is a constant, using Lemma D.3 and replacing ν_1 and τ with $\nu_1 \circ \nu_0^{-1}$ and $\nu_0 \circ \tau$, respectively, give us the same constraints for $\nu_0(\tau(\mathbf{x}_{1\dots 6}))$ regardless of ν_1 . \square

Theorem D.7 (Intersection of two hypersurfaces and a diagonal plane, restated). *Let $\tilde{\nu}_0(\mathbf{x}) = 0$ and $\tilde{\nu}_1(\mathbf{x}) = 0$ be two hypersurfaces passing two points \mathbf{x}_0 and \mathbf{x}_7 such that $\tilde{\nu}_0(\mathbf{x}_0) = \tilde{\nu}_1(\mathbf{x}_0) = 0$ and $\tilde{\nu}_0(\mathbf{x}_7) = \tilde{\nu}_1(\mathbf{x}_7) = 0$, $\mathbf{P}_i := \tilde{\nu}_0(\mathbf{x}_i)$ and $\mathbf{Q}_i := \tilde{\nu}_1(\mathbf{x}_i)$,*

$$\mathbf{P}_\alpha = [\mathbf{P}_0; \mathbf{P}_1; \mathbf{P}_4; \mathbf{P}_5], \quad \mathbf{P}_\beta = [\mathbf{P}_2; \mathbf{P}_3; \mathbf{P}_6; \mathbf{P}_7],\tag{65}$$

and

$$\mathbf{X} = \begin{bmatrix} (1-x)^2 \\ x(1-x) \\ (1-x)x \\ x^2 \end{bmatrix}.\tag{66}$$

Then, $x \in [0, 1]$ of the intersection point of the two hypersurfaces and a diagonal plane of $x = z$ is the solution of the following quartic equation:

$$\mathbf{X}^\top (\mathbf{P}_\alpha \mathbf{Q}_\beta^\top - \mathbf{P}_\beta \mathbf{Q}_\alpha^\top) \mathbf{X} = 0\tag{67}$$

while

$$y = \frac{\mathbf{X}^\top \mathbf{P}_\alpha}{\mathbf{X}^\top (\mathbf{P}_\alpha - \mathbf{P}_\beta)} \quad (\mathbf{P}_\alpha \neq \mathbf{P}_\beta).\tag{68}$$

Proof. We rearrange $\tilde{\nu}_0(\mathbf{x})$ using $x = z$ as follows:

$$\tilde{\nu}_0(\mathbf{x}) = (1-y) \cdot \mathbf{X}^\top [\mathbf{P}_0; \mathbf{P}_1; \mathbf{P}_4; \mathbf{P}_5] + y \cdot \mathbf{X}^\top [\mathbf{P}_2; \mathbf{P}_3; \mathbf{P}_6; \mathbf{P}_7]\tag{69}$$

$$= (1-y)\mathbf{X}^\top \mathbf{P}_\alpha + y\mathbf{X}^\top \mathbf{P}_\beta = 0\tag{70}$$

$$\Leftrightarrow y = \frac{\mathbf{X}^\top \mathbf{P}_\alpha}{\mathbf{X}^\top (\mathbf{P}_\alpha - \mathbf{P}_\beta)} \quad (\mathbf{P}_\alpha \neq \mathbf{P}_\beta).\tag{71}$$

In a similar way, we rewrite $\tilde{\nu}_1(x)$ as follows:

$$\mathbf{X}^\top(\mathbf{P}_\alpha - \mathbf{P}_\beta)\tilde{\nu}_1(x) = \mathbf{X}^\top\mathbf{P}_\alpha \cdot \mathbf{X}^\top\mathbf{Q}_\beta - \mathbf{X}^\top\mathbf{P}_\beta \cdot \mathbf{X}^\top\mathbf{Q}_\alpha \quad (72)$$

$$= \mathbf{X}^\top\mathbf{P}_\alpha\mathbf{Q}_\beta^\top\mathbf{X} - \mathbf{X}^\top\mathbf{P}_\beta\mathbf{Q}_\alpha^\top\mathbf{X} \quad (73)$$

$$= \mathbf{X}^\top(\mathbf{P}_\alpha\mathbf{Q}_\beta^\top - \mathbf{P}_\beta\mathbf{Q}_\alpha^\top)\mathbf{X}. \quad (74)$$

Notice that this is a *quartic* equation of x . To find the coefficients of the general form of quadratic equation, we consider a mapping \mathbf{C} from the interpolation coefficients τ_i to polynomial coefficients c_i :

$$\begin{bmatrix} c_0 \\ c_1 \\ c_2 \end{bmatrix} = \mathbf{C} \begin{bmatrix} \tau_0 \\ \tau_1 \\ \tau_2 \\ \tau_3 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ -2 & 1 & 1 & 0 \\ 1 & -1 & -1 & 1 \end{bmatrix} \begin{bmatrix} \tau_0 \\ \tau_1 \\ \tau_2 \\ \tau_3 \end{bmatrix} \quad (75)$$

where $\tau_0(1-x)^2 + \tau_1x(1-x) + \tau_2(1-x)x + \tau_3x^2 = c_0 + c_1x + c_2x^2$. Then, the coefficients of the quartic equation of Equation (74) are as follows:

$$\begin{bmatrix} c_{00} & c_{01} & c_{02} \\ c_{10} & c_{11} & c_{12} \\ c_{20} & c_{21} & c_{22} \end{bmatrix} = \mathbf{C}(\mathbf{P}_\alpha\mathbf{Q}_\beta^\top - \mathbf{P}_\beta\mathbf{Q}_\alpha^\top)\mathbf{C}^\top \quad (76)$$

where the anti-diagonal sum of the result is the quartic coefficients, *i.e.*,

$$c_{00} + (c_{10} + c_{01})x + (c_{20} + c_{11} + c_{02})x^2 + (c_{21} + c_{12})x^3 + c_{22}x^4. \quad (77)$$

The solutions of the quartic equation such that $x \in [0, 1]$ are the valid intersection points, and we choose one of them x if there are more than one solution. Note that we can compute a batch of quartic equations using the parallel computing of the eigenvalue decomposition with Lemma D.8.

We determine y by Equation (71) using x , and $z = x$. □

We reproduce the method to find polynomial roots from companion matrix eigenvalues [29].

Lemma D.8 (Polynomial roots from companion matrix eigenvalues, reproduced). *The companion matrix of the monic polynomial*⁷

$$p(x) = c_0 + c_1x + \cdots + c_{n-1}x^{n-1} + x^n \quad (78)$$

is defined as:

$$C(p) = \begin{bmatrix} 0 & 0 & \cdots & 0 & -c_0 \\ 1 & 0 & \cdots & 0 & -c_1 \\ 0 & 1 & \cdots & 0 & -c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & -c_{n-1} \end{bmatrix}. \quad (79)$$

The roots of the characteristic polynomial $p(x)$ are the eigenvalues of $C(p)$.

⁷The nonzero coefficient of the highest degree is one. If you are dealing with the equation, you can make it by dividing the nonzero coefficient of the highest degree for both sides.

E Complexity analysis

We provide a detailed exposition of the complexity analysis in Section 3.2 to ensure the comprehensiveness of our work. Notice that our complexity analysis closely aligns with the approach outlined in Berzins [4], as in their Appendix A.

Algorithm 2 operates on the vertices and edges, \mathcal{V} and \mathcal{E} , respectively. Therefore, our initial focus is on the complexity associated with these inputs. Let $|\mathcal{V}|$, $|\mathcal{E}|$, and $|\hat{\mathcal{E}}|$ ($\hat{\mathcal{E}} \subset \mathcal{E}$) be the number of vertices, edges, and splitting edges consisting of polygons (Section 3.2), at the iteration of i and j .

1. During the initialization step (refer to Section 5), we determine $|\mathcal{V}|$ and $|\mathcal{E}|$ to construct a grid in \mathbb{R}^D . In the case of having M HashGrid marks, the quantities are defined as M^D and $D(M-1)M^{D-1}$, respectively. In case of the large M , we consider vertex and edge pruning for efficient computations inspired by a pruning strategy proposed in Berzins [4]. If, for every future neuron, denoted as $\phi(i, j) > \phi(i^*, j^*)$, the two vertices of an edge exhibit identical signs, the edge can be safely pruned as it will not undergo splitting. We make the assumption of that $M^D < |\mathcal{V}|$ to avoid an excessively dense grid, with a complexity of $\mathcal{O}(|\mathcal{V}|)$.
2. Evaluate $\tilde{v}(\mathcal{V}; i, j)$ with a complexity of $\mathcal{O}(|\mathcal{V}|)$, assuming negligible inference costs as constants. The intermediate results from the inference can be used for the sign-vectors (Definition 3.4).
3. Identify splitting edges $\hat{\mathcal{E}}$ by comparing two signs per edge, with a complexity of $\mathcal{O}(|\mathcal{E}|)$.
4. Obtain P and Q for Theorem 4.7 with a complexity of $\mathcal{O}(|\hat{\mathcal{E}}|)$ in a similar way.
5. Inference for Theorem 4.7 involves matrix-vector multiplications and the eigenvalue decomposition of a companion matrix for each edge, with a complexity of $\mathcal{O}(|\hat{\mathcal{E}}|)$. Considering the maximum size of 4 for a quartic equation, this operation is regarded as linear.
6. Find intersecting polygons by pairing two sign-vectors with two common zeros in a shared region of the new vertices resulting from splitting, with a complexity of $\mathcal{O}(|\hat{\mathcal{E}}|)$.

In the study by Berzins [4], a linear relationship between $|\mathcal{V}|$ and $|\mathcal{E}|$ for a mesh is empirically observed (see Figure 8a in Berzins [4]). It is worth to note that the number of splitting edges can be effectively upper-bounded as $|\hat{\mathcal{E}}| < |\mathcal{E}|D/\phi(i, j)$ using Theorem 5 in [20]. Consequently, all the steps outlined in Algorithm 2 demonstrate a linear dependency on the number of vertices, denoted as $\mathcal{O}(|\mathcal{V}|)$.

To obtain faces, we sort the vertices within each region based on the Jacobian with respect to the mean of the vertices (see Algorithm 3). Considering triangularization, where each face has three vertices, the complexity remains $\mathcal{O}(|\mathcal{V}|)$, assuming the negligible cost of calculating the Jacobian, especially when performed in parallel.

F Supplementary experiments

F.1 Complete results

Table 2 and Table 3 are extended from the Stanford 3D Scanning dataset in Table 1. Table 4 extended from the Stanford bunny of Table 1, including standard deviations, and pseudo-ground truth results, obtained through marching cubes with 256 samples. Notice that we used the more excessive 512 samples for the *Medium* and *Large* models in Figure 4. We found that 256 samples were not enough for those. Table 5 shows the angular distance results. Tables 2 to 5 use the *Small* model, while Table 8 uses the *Large* model (ref. Section 6). We conducted all experiments with an NVIDIA V100 32GB.

Table 2: Chamfer distance (CD) and chamfer efficiency (CE) for the *Part A* of the Stanford 3D Scanning repository [31] using the *Small* model as the default setting (ref. Section 6). The chamfer distance ($\times 1e-6$) is evaluated using the marching cubes (MC) with 256^3 grid samples as the reference ground truth for all measurements. The chamfer efficiency is defined as $100/(\sqrt{|\mathcal{V}|} \times CD)$, providing a concise representation of the trade-off between the number of vertices and the chamfer distance. The results are averaged over three random seeds.

MC #	Bunny			Dragon			Happy Buddha		
	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow
32	1283	4106	19.0	1416	6330	11.2	1032	6951	13.9
64	5367	1371	13.6	6090	1936	8.5	4362	2465	9.3
128	21825	393	11.6	25236	542	7.3	18219	722	7.6
192	49569	141	14.3	57341	203	8.6	41351	276	8.8
Ours	4341	900	25.6	5104	1243	15.8	3710	1738	15.5

Table 3: Chamfer distance (CD) and chamfer efficiency (CE) for the *Part B* of the Stanford 3D Scanning dataset [31].

MC #	Armadillo			Drill Bit			Lucy		
	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow
32	1438	6212	11.2	144	91959	7.5	594	6753	24.9
64	6151	1960	8.3	873	5949	19.2	2701	2310	16.0
128	25164	605	6.6	3718	1542	17.4	11255	830	10.7
192	57306	231	7.5	8382	782	15.3	25589	335	11.7
Ours	4783	1499	13.9	694	3013	47.8	2406	1763	23.6

Table 4: Chamfer distance (CD) and chamfer efficiency (CE) for the Stanford bunny are evaluated using the marching cubes with 256^3 grid samples (\star) as the reference ground truth. The standard deviations are provided after the \pm notation, based on results obtained from three random seeds, which are used in training. Similar outcomes are observed for the remaining cases.

Method	Sample	Vertices \downarrow	CD ($\times 1e-6$) \downarrow	CE \uparrow	Time \downarrow
Marching Cubes	32	1283 \pm 89	4106 \pm 202	19.0	0.00 \pm 0.00
	48	2967 \pm 177	2178 \pm 94	15.5	0.01 \pm 0.06
	56	4057 \pm 245	1539 \pm 100	16.0	0.01 \pm 0.00
	64	5367 \pm 342	1371 \pm 93	13.6	0.02 \pm 0.00
	128	21825 \pm 1336	393 \pm 19	11.6	0.10 \pm 0.00
	192	49569 \pm 3043	141 \pm 4	14.3	0.44 \pm 0.10
Marching Cubes \star	224	67601 \pm 4146	114 \pm 6	12.9	0.60 \pm 0.03
	256	88492 \pm 5404	-	-	0.88 \pm 0.05
Ours	-	4341 \pm 253	900 \pm 171	25.6	0.97 \pm 0.21

Table 5: Angular distance (AD) for the Stanford bunny. Similarly to the chamfer distance, the angular distance ($^{\circ}$) is also evaluated using the marching cubes with 256^3 grid samples (\star) as the reference ground truth for all measurements. The standard deviations are provided after the \pm notation, based on results obtained from three random seeds, which is used in training.

Method	Sample	Vertices \downarrow	AD ($^{\circ}$) \downarrow
Marching Cubes	32	1283 \pm 89	6.97 \pm 0.45
	48	2967 \pm 177	5.07 \pm 0.15
	56	4057 \pm 245	4.23 \pm 0.25
	64	5367 \pm 342	4.10 \pm 0.17
	128	21825 \pm 1336	2.47 \pm 0.12
	192	49569 \pm 3043	1.63 \pm 0.06
Marching Cubes \star	224	67601 \pm 4146	1.53 \pm 0.06
	256	88492 \pm 253	-
Ours	-	4490 \pm 91	3.57 \pm 0.38

Table 6: Chamfer distance (CD) and chamfer efficiency (CE) for the *Part A* of the Stanford 3D Scanning repository [31] using the *Large* model having 4x-resolution of HashGrid compared to the default setting (*ref.* Section 6). The chamfer distance ($\times 1e-6$) is evaluated using the marching cubes (MC) with 512^3 grid samples as the reference ground truth for all measurements. The results are averaged over three random seeds.

MC #	Bunny			Dragon			Happy Buddha		
	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow
128	38395	632	4.12	25986	759	5.07	18621	1110	4.84
192	87245	299	3.83	58855	335	5.08	42793	533	4.38
224	119000	221	3.80	80470	268	4.64	58357	395	4.34
256	155484	182	3.53	105235	226	4.21	76183	322	4.07
Ours	135691	151	4.87	91147	157	6.99	67562	229	6.46

Table 7: Chamfer distance (CD) and chamfer efficiency (CE) for the *Part B* of the Stanford 3D Scanning dataset [31] using the *Large* model as the default setting (*ref.* Section 6).

MC #	Armadillo			Drill Bit			Lucy		
	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow	$ \mathcal{V} \downarrow$	CD \downarrow	CE \uparrow
128	25309	887	4.45	3590	1342	20.8	12224	1304	6.28
192	57751	399	4.34	8092	625	19.8	27863	687	5.23
224	78703	315	4.04	11000	588	15.5	38337	524	4.97
256	103085	267	3.63	14339	496	14.1	49838	435	4.61
Ours	90567	186	5.94	10275	342	28.5	43631	358	6.40

Table 8: Chamfer distance (CD) and chamfer efficiency (CE) for the Stanford dragon using the *Large* model having 4x-resolution of HashGrid compared to the default setting (*ref.* Section 6) are evaluated using the marching cubes with 512^3 grid samples (\star) as the reference ground truth.

Method	Sample	Vertices \downarrow	CD ($\times 1e-6$) \downarrow	CE \uparrow	Time \downarrow
Marching Cubes	128	25986 \pm 121	759 \pm 64	5.07	0.29 \pm 0.01
	192	58855 \pm 181	335 \pm 21	5.08	0.60 \pm 0.06
	256	105235 \pm 452	226 \pm 17	4.21	1.16 \pm 0.04
Marching Cubes \star	512	424311 \pm 1262	-	-	9.19 \pm 0.39
Ours	-	91147 \pm 203	157 \pm 18	6.99	14.9 \pm 1.66

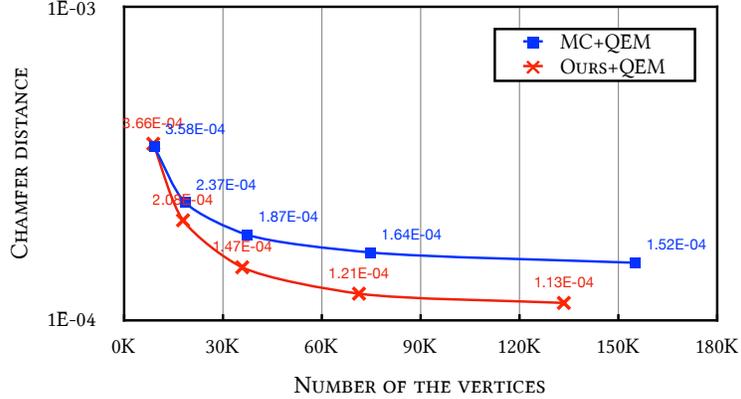


Figure 8: Chamfer distance for the Bunny with the *Large* model applying QEM (a mesh simplification). Both benefit from QEM while our method efficiently retains better chamfer distances.

F.2 Mesh simplification using QEM

In Figure 8, Quadric Error Metrics (QEM) [35] is a popular mesh simplification algorithm, it could potentially improve both MC and our method for efficiency. Overall, as shown, the mesh simplification favors our method. Notice that QEM entails a higher computational cost, at least $O(n \log n)$ where n is the number of vertices. Additionally, these methods can cause shape deformation or holes, sensitive to hyperparameters. Therefore, QEM may be applied with caution for further mesh simplification, acknowledging the potential for unpredictable side effects. We used MeshLab’s Quadric Edge Collapse Decimation with the default settings to halve the number of vertices for each step.

F.3 Comparison with alternate approaches

Marching Cubes (MC) remains the *de facto* sampling-based method. In contrast, our method is a white-box method grounded in the theoretical understanding of trilinear interpolation and polyhedral complex derivation. Here is our rationale for this evaluation:

1. Analytical approaches for CPWA functions are limited by their inability to effectively manage spectral bias and their exponential computational cost, which struggles with exponentially growing linear regions (infeasible to run our benchmarks.) For example, the time complexity of Analytic Marching [2] is $\mathcal{O}((n/2)^{2(L-1)}|\mathcal{V}_{\mathcal{P}}|n^4L^2)$, which grows exponentially with the number of layers L and the width n of the networks, where $|\mathcal{V}_{\mathcal{P}}|$ represents the number of vertices per face (Section 4.1 from Lei and Jia [2]). In contrast, our method has a linear time complexity with respect to the number of vertices, as discussed in Appendix E, allowing it to avoid visiting every linear region exponentially growing.
2. We also exclude optimization-based methods as they typically rely on MC or its variants initialization in their pipelines (our method can also be integrated into them), such as MeshSDF [36], Deep Marching Tetrahedra [10], and FlexiCubes [37]. Additionally, these methods result in prolonged computational costs.

We provided both rigorous quantitative Figure 3 and qualitative Figure 6 analyses on Marching Tetrahedra (MT) [9, 10] and Neural Dual Contour (NDC) [32], which serve as alternative representative approaches. MT is a popular variant of MC that utilizes six tetrahedra within each grid to identify intermediate vertices. NDC is a data-driven approach (a pretrained model) that uses Dual Contour. In short, MT efficiently requires SDF values but is less efficient regarding the number of vertices extracted. (MT is better than MC in the same resolution, producing more vertices, though.) NDC faces a generalization issue for a zero-shot setting, producing unsmooth surfaces. We used the code from Deep Marching Tetrahedra [10] for MT and the official code of NDC. And PyMCubes for our modernized MC library.

F.4 Error of the learned SDF

In Equation (11) and Figure 5, we confirmed our hypothesis that the eikonal loss induces piecewise linearity in trilinear networks. To further support this, we also present the error analysis of the learned SDF. Although the Chamfer distance virtually evaluates the SDF values based on sampling, concerns may arise regarding noise in the pseudo-ground truth. For this reason, one might wish to verify whether the predicted SDFs converge toward near-zero.

We randomly sampled 100K points on the surface of pretrained networks (the Stanford bunny *Large* model) and calculated the average of squares for predicted SDF (model’s outputs) on the surface. As we expected, the predicted SDFs on the reconstructed surface converge to zero as the weight of the eikonal loss increases, and the number is close to zero $\approx 3e-8$ as shown in Table 9. Note that, for a *Large* model, fine-grained trilinear regions from the dense grid marks (from the more dense multi-resolution grids) allow relatively small errors of $97 \times 1e-9$ without the eikonal loss, although the eikonal loss of $1e-2$ further decreases to $31 \times 1e-9$. Note that the stronger eikonal loss of $1e-1$ proved ineffective due to over-regularization.

Table 9: Error of the learned SDF. *Weight* stands for the weight of the eikonal loss and *SE* stands for the squared error of predicted SDFs.

Method	Weight	SE (1e-9) ↓
MC64	1e-2	2308
MC256	1e-2	66
	0	97
Ours	1e-4	91
	1e-3	68
	1e-2	31

F.5 Effect of model sizes

In Table 10, we explore the ablation study to determine whether model size impacts model fitting, which in turn affects the GT mesh (we used MC256 for the pseudo-ground truth of the *Small* models). With a smaller model size, the learned mesh becomes simpler due to underfitting. As our evaluation focuses on how accurately the extraction methods capture the underlying zero level-set decision boundary of the networks, the change in CD remains small.

Table 10: Ablation study varying the number of *layers* and the *width* of networks for the Stanford bunny *Small* models. And, *Note* indicates the architecture of the decoding networks.

Layers	Width	Vertices	CD (1e-6) ↓	Time (sec) ↓	Note
2	4	4565	738	0.11	
2	8	4863	768	0.13	
2	16	4566	734	0.12	
2	32	4577	754	0.22	
2	64	4611	739	0.50	<i>InstantNGP</i> [12]
3	4	4547	666	0.12	
3	8	4661	750	0.19	
3	16	4628	724	0.65	<i>Ours</i>
3	32	4567	754	0.57	
3	64	5588	759	2.86	

Notice that the model with the number of layers of 3 and the width of 64 has the longest runtime of 2.80, primarily because the number of intermediate vertices exceeds 80K. The time complexity is more sensitive to the number of vertices than to the number of layers or the width, which are only indirect factors. Note that the model capacity of the decoding networks is comparable to InstantNGP [12], where the representational efficiency is driven by the HashGrid.

G Visualizations

For Figure 6 in the main paper (closer looks at bunny’s nose), we want to leave a note about the small facets along the left-side of the nose line in the highlighted region. We assume these narrow facets are embedded in the facets of MC256 (pseudo-ground truth), while MC64 fails to capture the left and right edge of the nose and represents this with more facets cutting these two sides of edges – not aligned with nose lines and rather assign more facets in here.

In Figure 9, we explain the rationale behind adopting the diagonal plane assumption in Theorem 4.7 and Section 4.2. We illustrated 64 edge scenarios within a trilinear region, omitting isomorphic cases.

In Figure 10, we showcase the analytical generation of a normal map from piecewise trilinear networks, utilizing HashGrid [12] and ReLU neural networks. Our approach dynamically assigns vertices based on the learned decision boundary. Notably, the configuration of hash grids influences vertex selection, and each grid can be subdivided into multiple polyhedra for accurate representation of smooth curves.

In Figure 11, we present comparative visualizations of normal maps for marching cubes, demonstrating variations in the number of samplings alongside our method, with the respective vertex counts indicated in parentheses. This analysis highlights the trade-off between detail preservation and computational efficiency, showcasing the optimality of our analytically extracted mesh from decision boundary apex points.

In Figure 12, we compare marching cubes (MC) with 256 and 64 samplings to our method. MC 64 shows limitations in representing pointy areas, while our approach excels with analytical extraction from decision boundary apex points. Adaptive sampling leads to some vertices being in close proximity (10.1% within $1e-3$), prompting potential mesh optimization for improved efficiency.

In Figure 13, the gallery of the six Stanford 3D Scanning meshes from the *Large* models is shown.

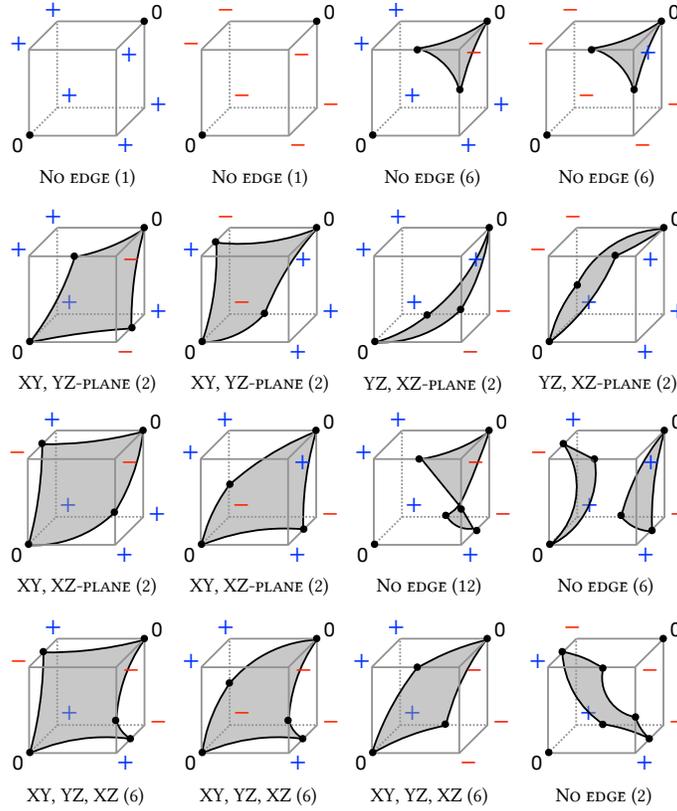


Figure 9: We depict the 64 scenarios of the edge within a trilinear region, excluding isomorph cases specifying the count within parentheses. *No edge* indicates where two diagonal zero corners are not connected. The labels *XY*, *YZ*, and *XZ* refer to $x = y$, $y = z$, and $z = x$ diagonal planes, respectively, form the shared edge, connecting the two diagonal zero corners, with the hypersurface in the schematic figure. In nearly all cases, a diagonal plane intersects with the hypersurface, except, at most, four cases, depending on the curvature determined by corner values.

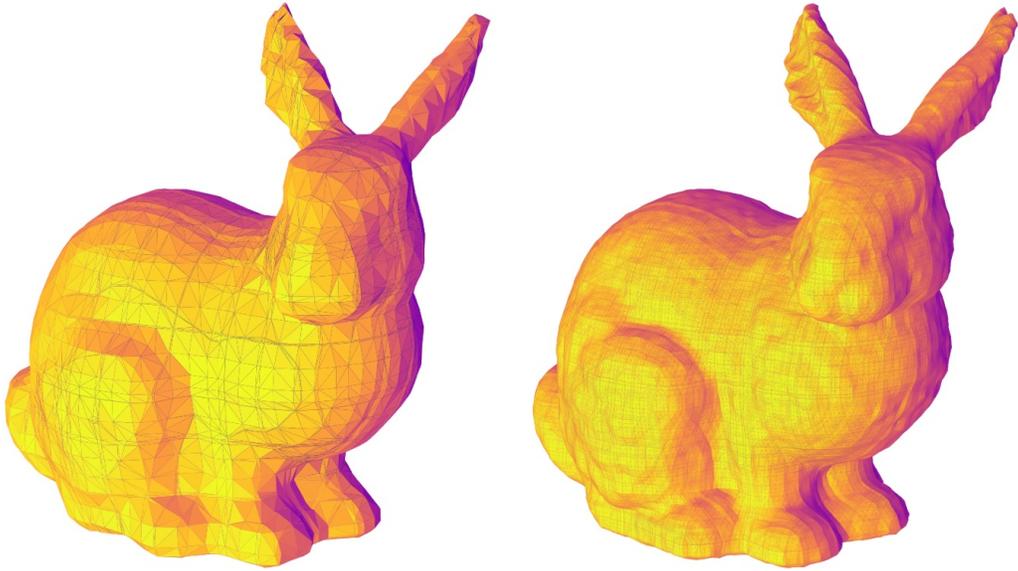


Figure 10: The normal map is analytically generated from piecewise trilinear networks, comprising both HashGrid [12] and ReLU neural networks, that have been trained to learn an SDF with the eikonal loss for the Stanford bunny [31]. Our approach dynamically assigns vertices based on the learned decision boundary. One notable observation is that the hash grids influence the selection of vertices. Additionally, each grid can be subdivided into multiple polyhedra to accurately represent smooth curves. The *Small* (left) and *Large* (right) models (*ref.* Section 6) are used to learn the SDF.

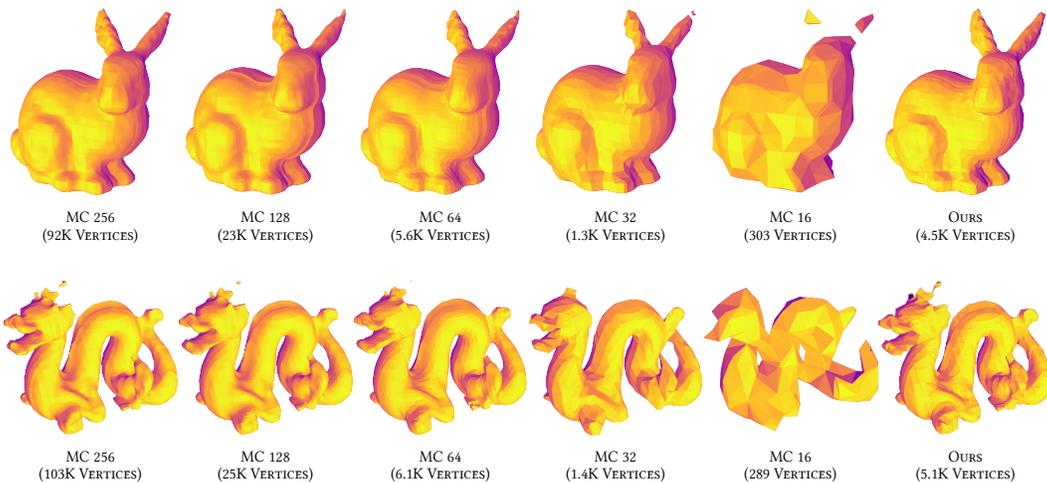


Figure 11: Visualizing normal maps for marching cubes varying the number of samplings, and ours, with the number of vertices denoted in parentheses. Low sampling has a risk of omitting sharp object details, while dense sampling results in redundant vertices and faces. As our mesh is analytically extracted from the decision boundary apex points of networks, it represents an optimal choice in balancing detail preservation and efficiency.

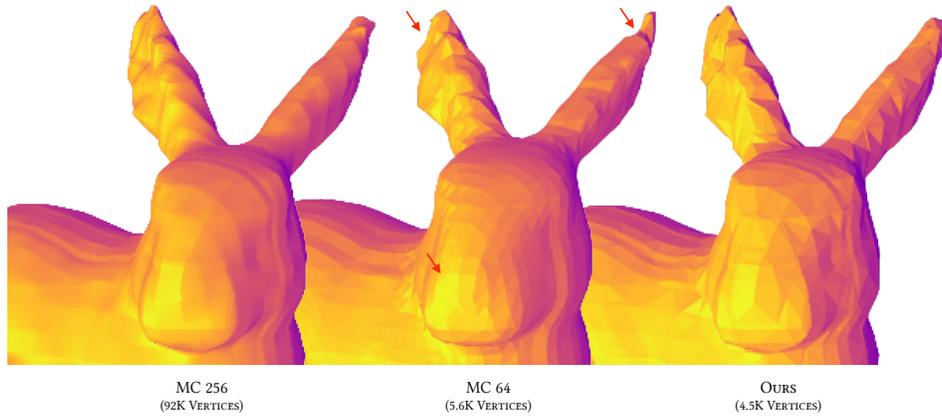


Figure 12: In a *face-to-face* comparison, we examine marching cubes with 256 and 64 samplings, denoted as MC 256 and MC 64, respectively, and our method. Notably, MC 64 struggles to accurately represent pointy areas of meshes due to its coarse sampling. In contrast, our approach has an advantage in this aspect, as we analytically extract the mesh from the decision boundary apex points of networks. However, it is worth noting that our adaptive sampling introduces a disadvantage where some vertices are in close proximity; specifically, we observed that 10.1% of vertices are within $1e-3$ proximity to others. We believe that optimizing the mesh can significantly improve the parsimony of the vertex count compared with marching cubes.

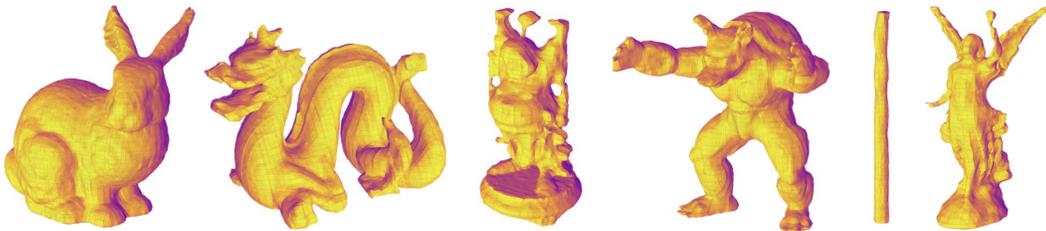


Figure 13: The extract mesh visualizations of the six from the Stanford 3D Scanning repository using the *Large* models. From left, Bunny, Dragon, Happy Buddha, Armadillo, Dril Bit, and Lucy.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Our main claims made in the abstract and introduction are accurately reflected in theoretical and experimental expositions. Theorem 4.5 and Figure 5 provide the theoretical and experimental supports of the planarity constraint of the eikonal equation, Theorem 4.7 provides a plausible approximation for the curved edge subdivision, and Table 1 shows the consistent chamfer efficiency across multiple meshes from the Stanford 3D Scanning repository, while Figure 4 shows that the higher resolution of HashGrid is favorable to the planarity constraint, achieving lower chamfer distances.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Please refer to Section 6.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Every Lemma and Theorem has its proof in the Appendix, noting the assumptions that we used, while they are properly referenced by numbering.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Section 6 provides the hyperparameters we used, Algorithms are provided in the main and the Appendix, while reproducible code can be found in the supplemental material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The reproducible code can be found in the supplemental material. The Stanford 3D Scanning repository can be freely download via <http://graphics.stanford.edu/data/3Dscanrep/>.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Section 6 provides the hyperparameters we used, and the reproducible code can be found in the supplemental material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Table 4 and Tables 5 and 8 in Appendix F.1 provides the standard deviations using three random seeds where learned SDFs may slightly differ. Notice that our algorithm is deterministic for given SDF networks.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Appendix F.1 provides the information on the computer resources.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We conform with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: There is no societal impact of the work performed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This work poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The Stanford 3D Scanning repository (MIT License) is used with a proper citation [31], along with HashGrid [12] for our trilinear module.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Not involved in crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Not involved in crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.