Implicit Geometric Regularization for Learning Shapes

Sajal Tyagi 2019MT60761

COV887

Content

- Motivation and Problem Statement
- Proposed Methodology
- Evaluations and Limitations

Motivation and Problem Statement

3D Reconstruction

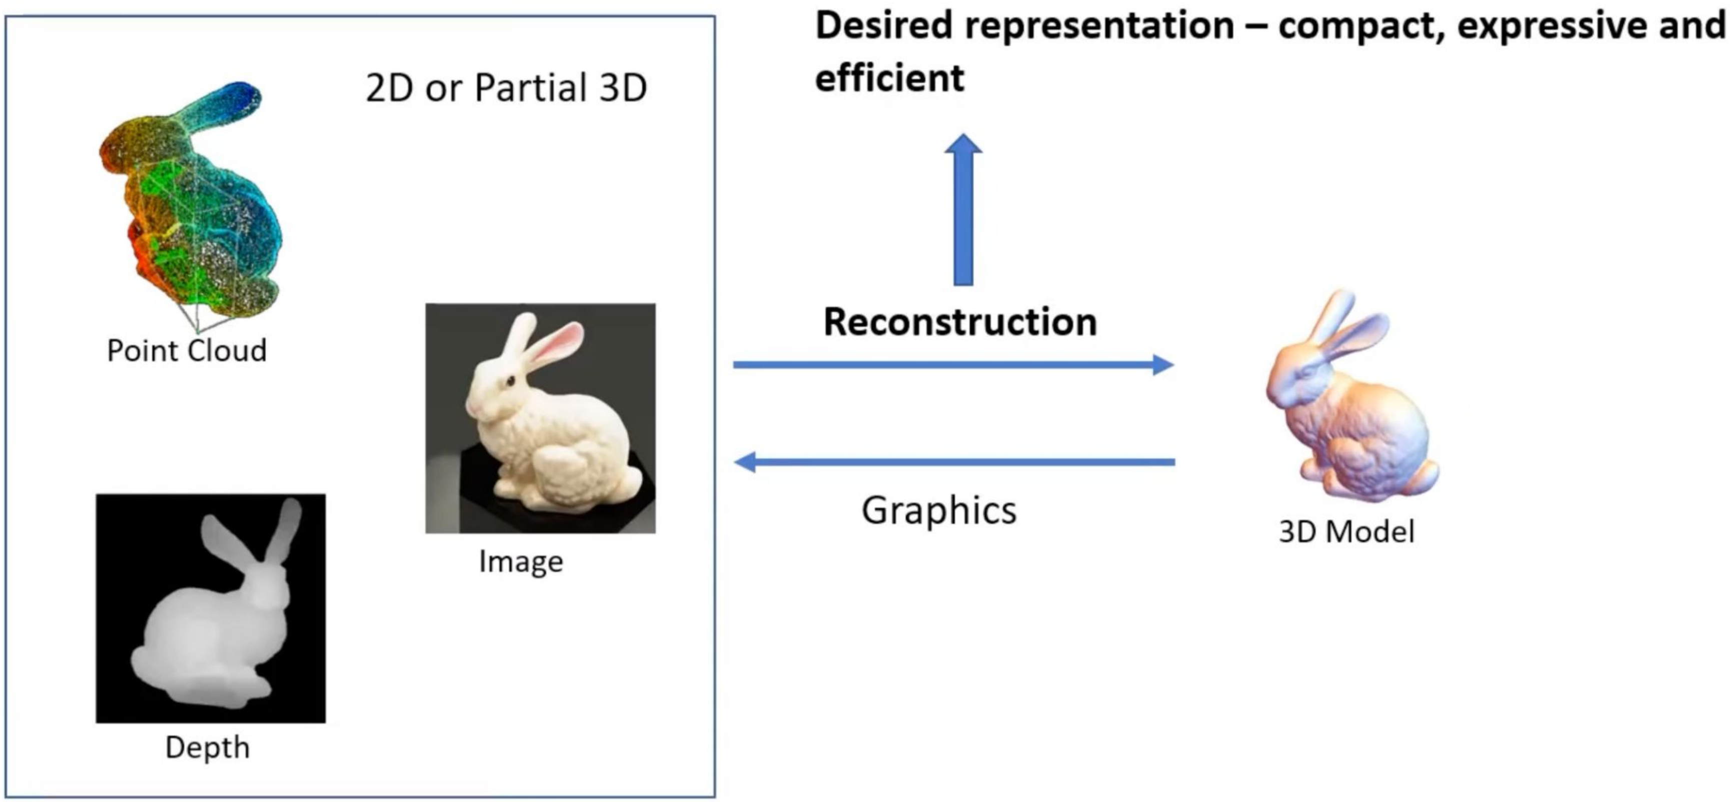

3D Representations

Voxel

Cubically growing compute and memory requirements

Point Cloud

Do not describe surface

Mesh

Limited to the typology of the template at that time.

Signed Distance Function

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

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

$$f_\theta(x, y, z) \approx SDF(x, y, z)$$

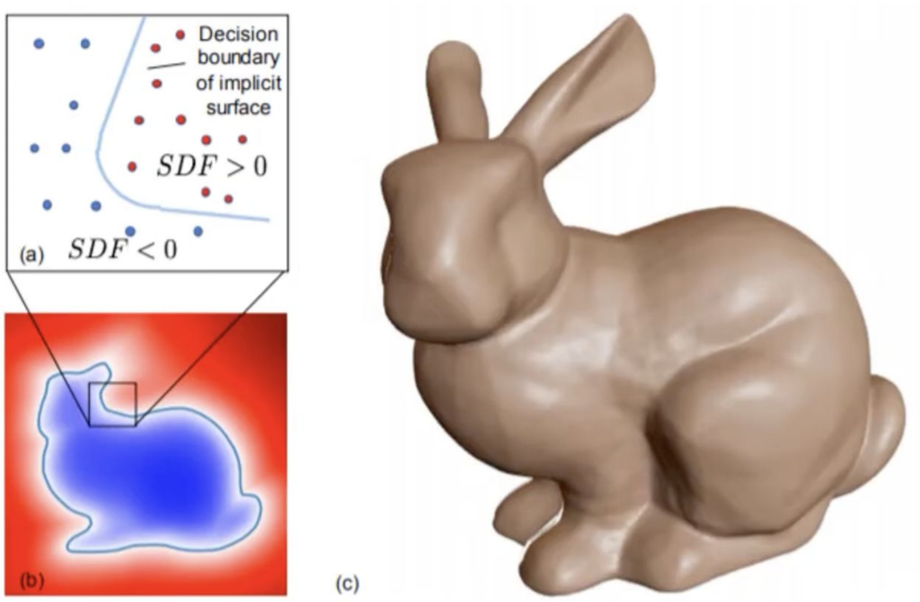

Problem Statement

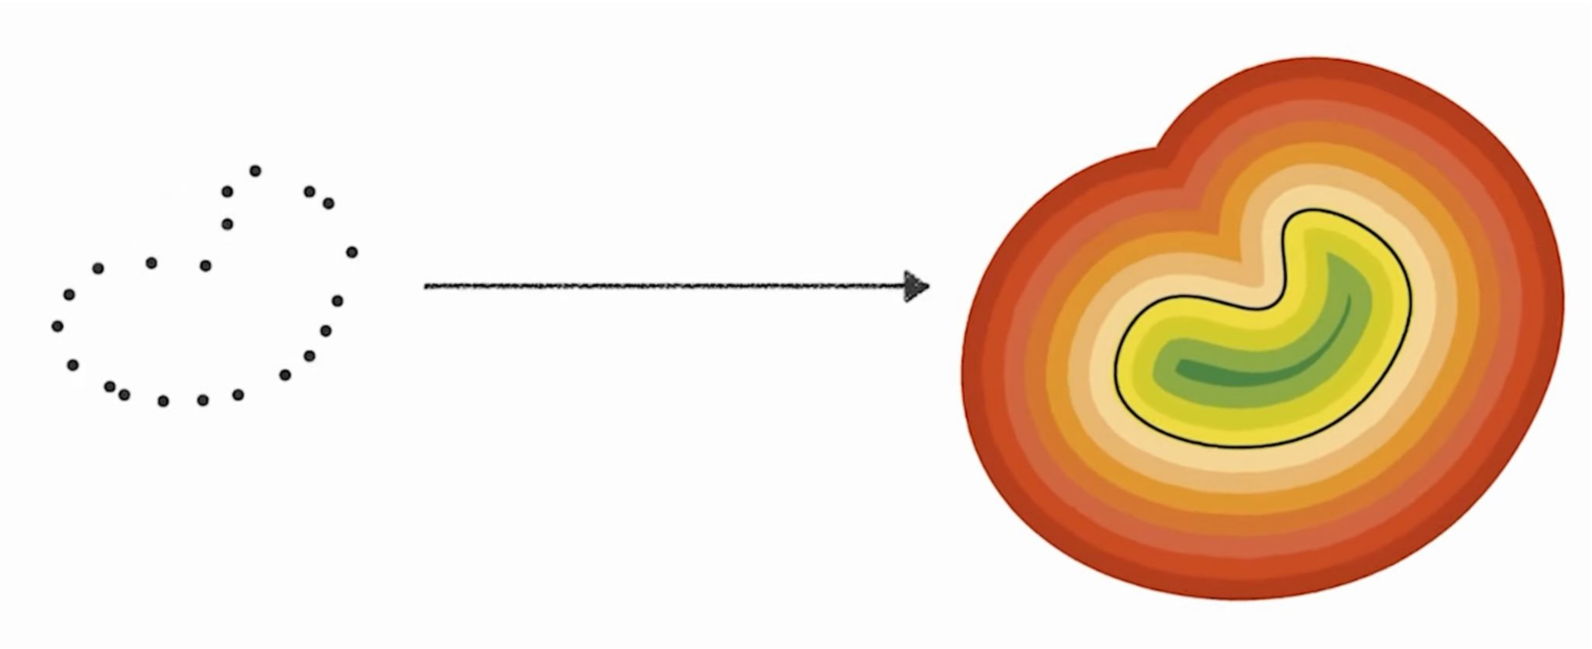

Raw Point Cloud
(with/without normal)

Signed Distance Function

Proposed Methodology

Method Overview

Signed Distance Function

$$f(x) = \begin{cases} -d(x, \partial\Omega) & x \in \Omega \\ d(x, \partial\Omega) & x \notin \Omega \end{cases}$$

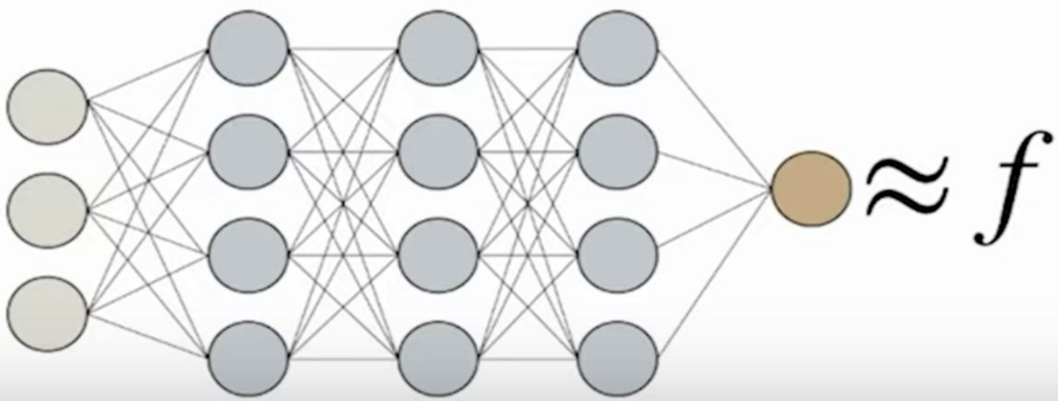

Method Overview

$$\text{loss}(\theta) = \sum_{i \in I} |f(x_i;\theta)|^2 + \lambda \mathbb{E}_x \left(\|\nabla_x f(x;\theta)\| - 1 \right)^2$$

vanish on surface samples

Eikonal on ambient space

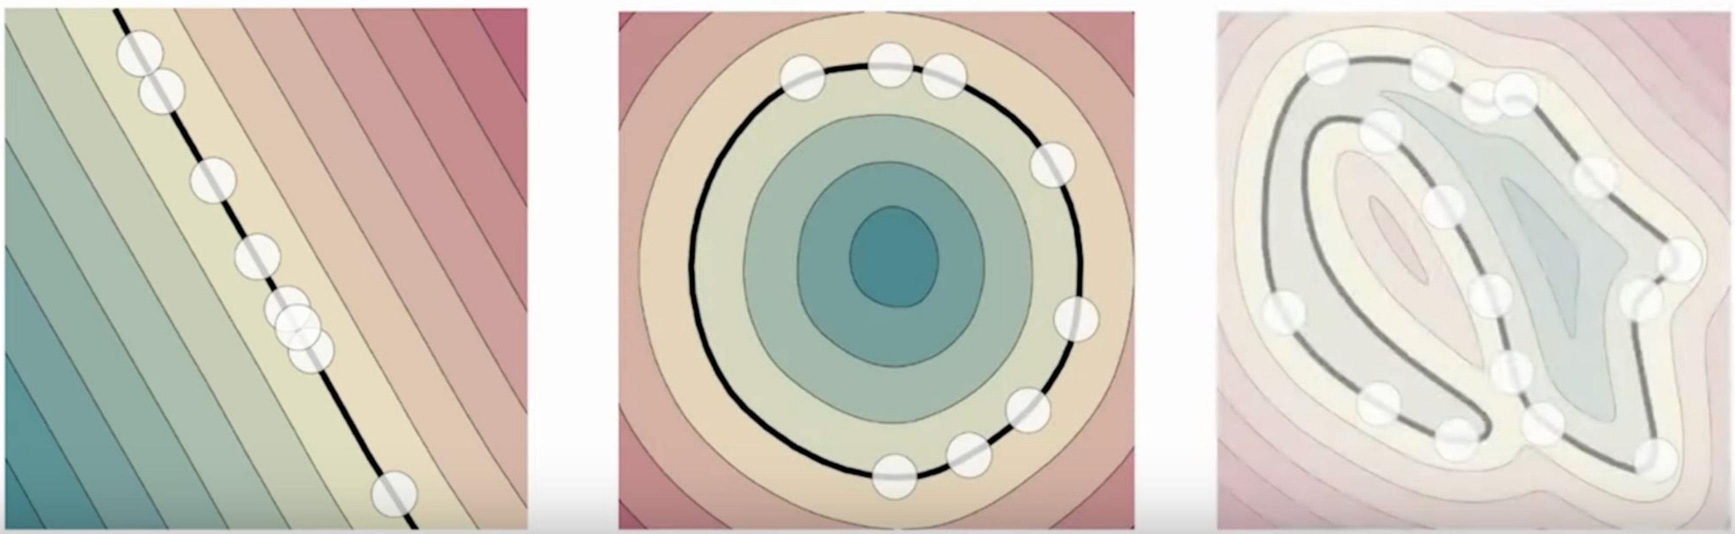

Note : For calculating the Expectation for Eikonal term the distribution used is average of a uniform distribution and a sum of Gaussians

Method Overview

$$\text{loss}(\theta) = \sum_{i \in I} \left(|f(x_i;\theta)|^2 + \tau \underbrace{\|\nabla_x f(x_i;\theta) - n_i\|^2}_{\text{fit surface normals}} \right) + \lambda \mathbb{E}_x \left(\|\nabla_x f(x;\theta)\| - 1 \right)^2$$

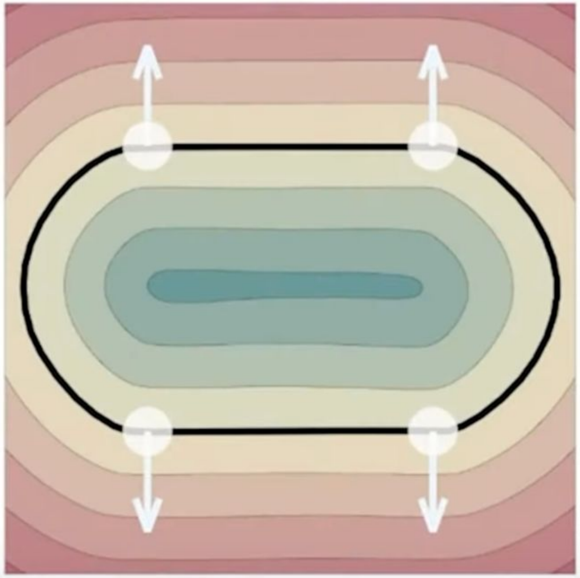

Eikonal PDE for SDFs

- A solution to the Eikonal equation will be a signed distance function and a global minimum of the loss in equation
- However the solutions to the Eikonal for a given point boundary data is not unique i.e. there are an infinite number of signed distance functions vanishing on arbitrary discrete sets of points

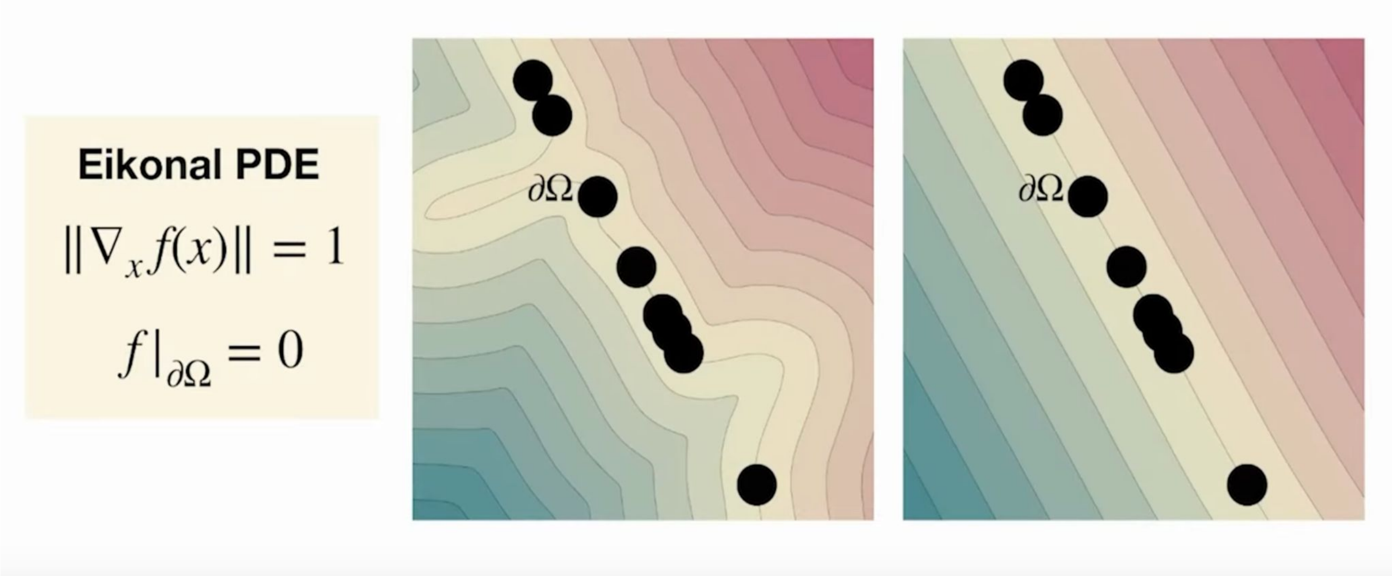

Eikonal PDE for SDFs

Theorem : If Ω is a subset of ℝ^{n} with piecewise smooth boundary, then every signed distance function Satisfies the Eikonal

Outline Of Proof :

For any $x \notin \partial\Omega$ we have $d(x, \partial\Omega) = d(x, y)$ for some $y \in \partial\Omega$ as $\partial\Omega$ is smooth.

The vector $x - y$ will be orthogonal to the tangent hyperplane at y ($T_y\partial\Omega$).

Now, if we take any point $\hat{x}$ that lies on the line $\alpha x + (1-\alpha)y$, we can easily see that $d(\hat{x}, \partial\Omega) = d(\hat{x}, y)$.

So, if we move the point by Δx units along the $x - y$, the value of $d(x, \partial\Omega)$ will increase by Δx units.

This implies the directional derivative $D_v f(x) = \nabla f(x).v = \pm 1$ (Let v be a unit vector along $x - y$), so $\|\nabla f(x)\| = 1$.

Implicit Geometric Regularization

- Optimizing the proposed loss function using stochastic gradient descent results in solutions that are close to a signed distance function with a smooth and surprisingly plausible zero level sets

- For the set of points available, both figures represent optimal solutions for the loss function

- The proposed model shows plane reproduction property

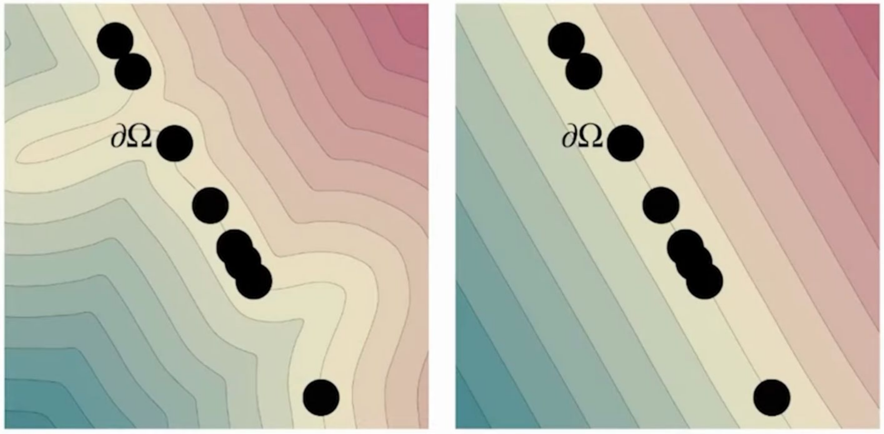

Plane Reproduction Property

Theorem:

Assume the data points span a $(d-1)$-dim hyperplane in $\mathbb{R}^d$ that contains the origin, then GD of the **linear model** with random initialization converges w.p 1 to the correct signed distance function.

$$f(x) = w^T x$$

$$\text{loss}(\theta) = \sum_{i \in I} (w^T x_i)^2 + \lambda(\|w\|^2 - 1)^2$$

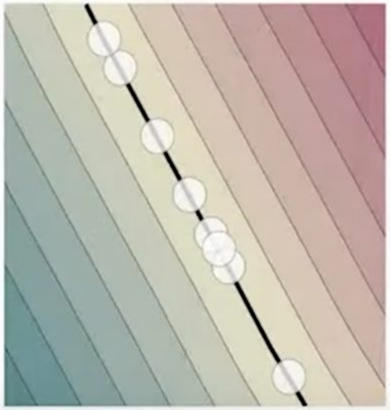

Neural Architecture

$$\mathrm{loss}(\theta) = \sum_{i\in I} \left(|f(x_i;\theta)|^2 + \tau \underbrace{\|\nabla_x f(x_i;\theta) - n_i\|^2}_{\text{fit surface normals}} \right) + \lambda \mathbb{E}_x \left(\|\nabla_x f(x;\theta)\| - 1 \right)^2$$

$$\boldsymbol{y}^{\ell+1} = \sigma(\boldsymbol{W}\boldsymbol{y}^{\ell} + \boldsymbol{b})$$

$$\nabla_{\boldsymbol{x}} \boldsymbol{y}^{\ell+1} = \mathrm{diag}\left(\sigma'\left(\boldsymbol{W}\boldsymbol{y}^{\ell} + \boldsymbol{b}\right)\right) \boldsymbol{W} \nabla_{\boldsymbol{x}} \boldsymbol{y}^{\ell}$$

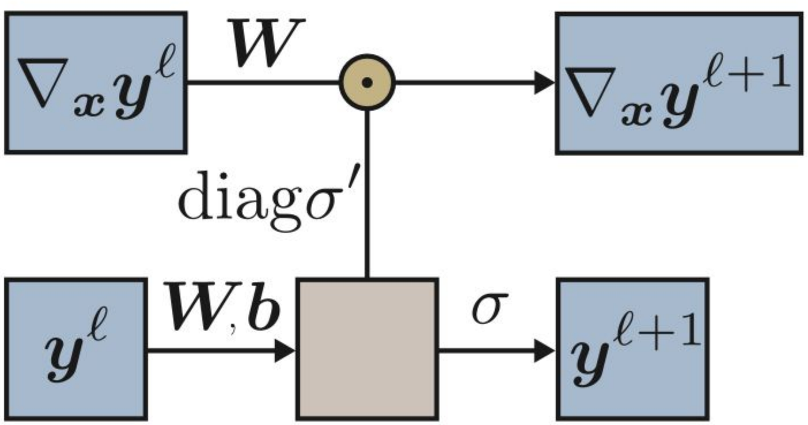

Figure Illustrating Layer-Wise Computation Network

Number of Layers : 8
Number of Hidden Units/layer : 512

*Contains a Single Skip Connection from Input to Middle Layer

Evaluations and Limitations

Model Evaluations

The authors focused on 2 major evaluations :

1. Signed distance function approximation :

   Goal : Testing the ability of our trained model f to reproduce a signed distance function of known manifold surfaces

   Evaluation Metric : Average of Abs. Relative Error at randomly sampled points

2. Fidelity and level of details :

   Goal : Testing the faithfulness or fidelity of the learning method on raw point cloud scans of humans

   Evaluation Metric : Two-sided Chamfer distance and Two-sided Hausdorff distance between the Reconstructed Surface and Ground Truth Surface; as well as One-sided Chamfer distance and One-sided Hausdorff distance from input point clouds to Reconstructed Surface

Evaluation Metrics

Let $\mathcal{X}_1, \mathcal{X}_2 \subset \mathbb{R}^3$ be 2 points sets :

$$\mathrm{d}_{\vec{\mathrm{C}}}\left(\mathcal{X}_1, \mathcal{X}_2\right) = \frac{1}{\left|\mathcal{X}_1\right|} \sum_{\boldsymbol{x}_1 \in \mathcal{X}_1} \min_{\boldsymbol{x}_2 \in \mathcal{X}_2} \left\|\boldsymbol{x}_1 - \boldsymbol{x}_2\right\|$$

: One Sided Chamfer Distance

$$\mathrm{d}_{\vec{\mathrm{H}}}\left(\mathcal{X}_1, \mathcal{X}_2\right) = \max_{\boldsymbol{x}_1 \in \mathcal{X}_1} \min_{\boldsymbol{x}_2 \in \mathcal{X}_2} \left\|\boldsymbol{x}_1 - \boldsymbol{x}_2\right\|$$

: One Sided Hausdorff Distance

$$\mathrm{d}_{\mathrm{C}}\left(\mathcal{X}_1, \mathcal{X}_2\right) = \frac{1}{2}\left(\mathrm{d}_{\vec{\mathrm{C}}}\left(\mathcal{X}_1, \mathcal{X}_2\right) + \mathrm{d}_{\vec{\mathrm{C}}}\left(\mathcal{X}_2, \mathcal{X}_1\right)\right)$$

: Two Sided Chamfer Distance

$$\mathrm{d}_{\mathrm{H}}\left(\mathcal{X}_1, \mathcal{X}_2\right) = \max\left\{\mathrm{d}_{\vec{\mathrm{H}}}\left(\mathcal{X}_1, \mathcal{X}_2\right), \mathrm{d}_{\vec{\mathrm{H}}}\left(\mathcal{X}_2, \mathcal{X}_1\right)\right\}$$

: Two Sided Hausdorff Distance

Limitations

- Due to using a Feed-Forward Neural Network as base architecture the model faces a problem when representing high-frequency features

- A few outlier points in data can lead to significant deviations in the predicted SDF

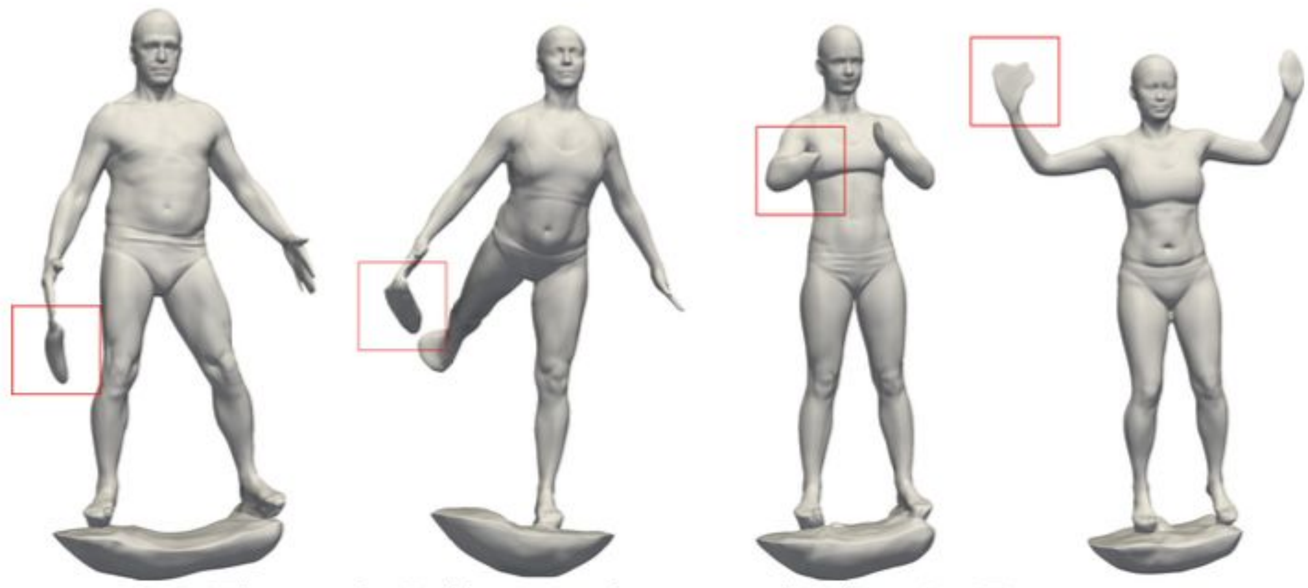

Thank you

Please Feel Free to Ask any Questions