Report on "Implicit Geometric Regularization for Learning Shapes" by Amos Gropp[2020] et al.

LALIT MEENA(2019CS50439)

This paper is presented by Sajal in COV877 course by Prof. Rahul Narain, CSE, IIT DELHI.

1 INTRODUCTION:

The paper introduces a groundbreaking approach to computing high-fidelity implicit neural representations of 3D shapes directly from raw data, such as point clouds with or without normal information. Previous methods either relied on precomputed implicit representations or explicit loss functions defined over neural level sets. The authors propose a novel paradigm leveraging implicit geometric regularization, resulting in state-of-the-art implicit neural representations.

2 METHOD

The authors' method aims to determine the parameters of a multilayer perceptron (MLP) that approximates a signed distance function to a given surface defined by input point cloud data. They utilize a loss function that encourages the MLP to vanish on the input point cloud while maintaining unit norm gradients. The Eikonal term in the loss enforces unit 2-norm gradients, aligning with the desired surface properties. The paper includes a theoretical analysis, particularly in the linear case, to demonstrate the effectiveness of their approach.

They considered a loss of the form

$$\ell(\theta) = \ell_{\mathcal{X}}(\theta) + \lambda \mathbb{E}_{\boldsymbol{x}} \left(\|\nabla_{\boldsymbol{x}} f(\boldsymbol{x}; \theta)\| - 1 \right)^2,$$

where $\lambda > 0$ is a parameter, $\|\cdot\| = \|\cdot\|_2$ is the euclidean 2-norm, and

$$\ell_{\mathcal{X}}(\theta) = \frac{1}{|I|} \sum_{i \in I} \left(|f(\boldsymbol{x}_i; \theta)| + \tau \|\nabla_{\boldsymbol{x}} f(\boldsymbol{x}_i; \theta) - \boldsymbol{n}_i\| \right)$$

encourages f to vanish on $\mathcal{X}$ and, if normal data exists (i.e., $\tau = 1$), that $\nabla_x f$ is close to the supplied normals $\mathcal{N}$. The second term in the equation is called the Eikonal term and encourages the gradients $\nabla_x f$ to be of unit 2-norm.

Here is the definition of signed distance function(SDF) used -

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

A solution to the Eikonal equation will be a signed distance function and a global minimum of the loss function. Here is Eikonal PDE involved-

$$\|\nabla_x f(x)\| = 1$$
$$f|_{\partial\Omega} = 0$$

3 IMPLICIT GEOMETRIC REGULARIZATION

The optimization process raises questions about the feasibility of critical solutions found by the algorithm and why they lead to signed distance functions. The authors address these questions, discussing how a quadratic penalty in the loss function with finite weight does not necessarily guarantee feasible critical solutions. They also explore the uniqueness of solutions when utilizing point boundary data as the proposed model shows plane reproduction property.

Author's address: Lalit Meena(2019CS50439).

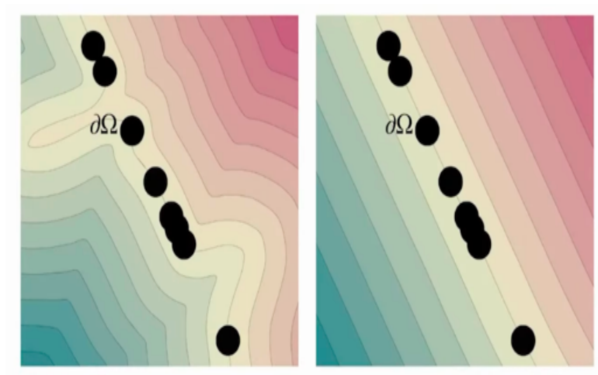

Fig. 1. The solutions to the Eikonal for a given point boundary data is not unique i.e. there are an infinite number of signed distance functions vanishing on arbitrary discrete sets of points

4 EVALUATIONS AND EXPERIMENTS

4.1 Surface Reconstruction

The paper presents experiments conducted on the task of surface reconstruction, where the goal is to approximate a surface from a given input point cloud. The authors compared their method to a recent deep learning chart-based surface reconstruction technique and evaluated performance using metrics such as **Chamfer distance and Hausdorff distance**. Results demonstrated that their method consistently outperforms the existing technique on most models in the dataset, offering improved fidelity and detail.

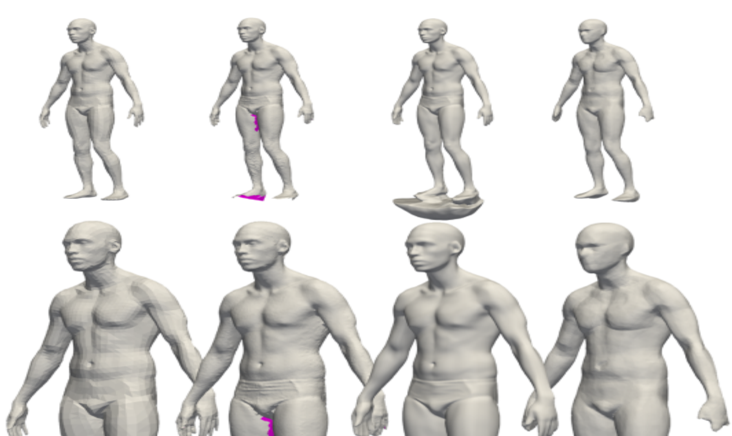

Fig. 2. A test result on D-Faust. Left to right: registrations, scans, this paper's results, SAL

4.2 Learning Shape Space

In this experiment, the authors tested their method on the task of learning shape space from raw scans using the D-Faust dataset, containing high-resolution raw scans of humans in various poses.

The results (Fig.2) showcased that their method can produce plausible approximations, even when trained on a limited number of shapes. They also explored shape space interpolation, providing insights into the learned shape representations.

5 DISCUSSIONS

There were interesting things discussed after the paper’s presentation by Sajal-

(1) The solutions to the Eikonal for a given point boundary data is not unique like Figure 1. It is avoided by the optimization algorithm that chooses to reconstruct a straight line in this case. This is a consequence of the plane reproduction property.
(2) There is nothing mentioned in paper about color properties and related work.
(3) There were a few clarification related questions asked for the Eikonal PDEs and also one in which the direction of SDF (which to choose +/-) will be chosen by model would be done in case only raw point cloud data.
(4) Some of limitations discussed are that due to using a Feed-Forward Neural Network as base architecture the model faces a problem when representing high-frequency features
(5) Also a few outlier points in data can lead to significant deviations in the predicted SDF.

6 CONCLUSION

The paper introduces a groundbreaking approach to learning implicit neural representations of shapes directly from raw data using implicit geometric regularization. The authors provide both theoretical analysis and practical demonstrations through experiments. Their method consistently achieves state-of-the-art results, offering higher accuracy and more detailed implicit neural representations compared to existing techniques.

7 REFERENCES

(1) Gropp, A. (2020). Implicit Geometric Regularization for Learning Shapes. *arXiv preprint arXiv:2002.10099*.
https://arxiv.org/abs/2002.10099
(2) Tayagi, S. (2023, September 19). Presentation on the paper "Implicit Geometric Regularization for Learning Shapes" by Amos Gropp. Google Slides.
https://docs.google.com/presentation/d/1E2cRodoBmYGFYEp3U0R2hqH1_drxFz0INr3iPvrz9a4/edit#slide=id.g243240f2022_0_17