Implicit geometric regularization for learning shapes

Saturday, December 26, 2020 1:47 PM

level sets of neural networks have been used to represent 3D shapes,

$$M = \{ x \in R^3 \mid f(x;\theta) = 0 \}, \quad (1).$$

where $f: R^3 \times R^m \to R$ is a MLP.

Most previous works using implicit neural representations computed f with 3D supervision: by comparing f to a known implicit representation of some shape.

In this work we work directly with raw data:

- given an input point cloud $\mathcal{X} = \{x_i\}_{i\in I} \subset R^3$
- with or without normal data $\mathcal{N} = \{n_i\}_{i\in I} \subset R^3$.
- goal is to compute θ such that $f(x;\theta)$ is approximately the signed distance function to a plausible surface $\mathcal{M}$ defined by the point data $\mathcal{X}$ and normals $\mathcal{N}$.

We show that SOTA implicit neural representations can be achieved without 3D supervision and/or a direct loss on the zero level set $\mathcal{M}$.
SGD optimization of a simple loss that fits an MLP f to a point cloud data $\mathcal{X}$, with or without normal data $\mathcal{N}$, while encouraging unit norm gradients $\nabla_x f$, consistently reaches good local minima, favoring smooth, yet high fidelity, zero level set surfaces $\mathcal{M}$ approximating the input data $\mathcal{X}$ and $\mathcal{N}$.

Given an input cloud $\mathcal{X} = \{x_i\}_{i\in I} \subset R^3$
with or without normal data $\mathcal{N} = \{n_i\}_{i\in I} \subset R^3$,
goal is to compute parameters θ of an MLP $f(x;\theta)$.
$f: R^3 \times R^m \to R$ approximates a SDF with plausible surface $\mathcal{M}$ defined by $\mathcal{X}, \mathcal{N}$.

Consider loss:

$$\ell(\theta) = \ell_{\mathcal{X}}(\theta) + \lambda \underbrace{\mathbb{E}_x \left(\| \nabla_x f(x;\theta) \| - 1 \right)^2}_{\text{Eikonal term.}}, \quad (2).$$

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

$$\ell_{\mathcal{X}}(\theta) = \frac{1}{|I|} \sum_{i\in I} \left(|f(x_i;\theta)| + \tau \| \nabla_x f(x_i;\theta) - n_i \| \right)$$

encourages f to vanish on $\mathcal{X}$, and if $\mathcal{N}$ exists ($\tau = 1$), that $\nabla_x f$ is close to the supplied normals $\mathcal{N}$.

Eikonal term encourages the gradients $\nabla_x f$ to be of unit-2 norm, the expectation is taken wrt some probability distribution $x \sim \mathcal{D}$ in R^3.

from presentation.
implicit shape for a sphere: $f(x) = \|x\|^2 - 1$.

SDF:

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

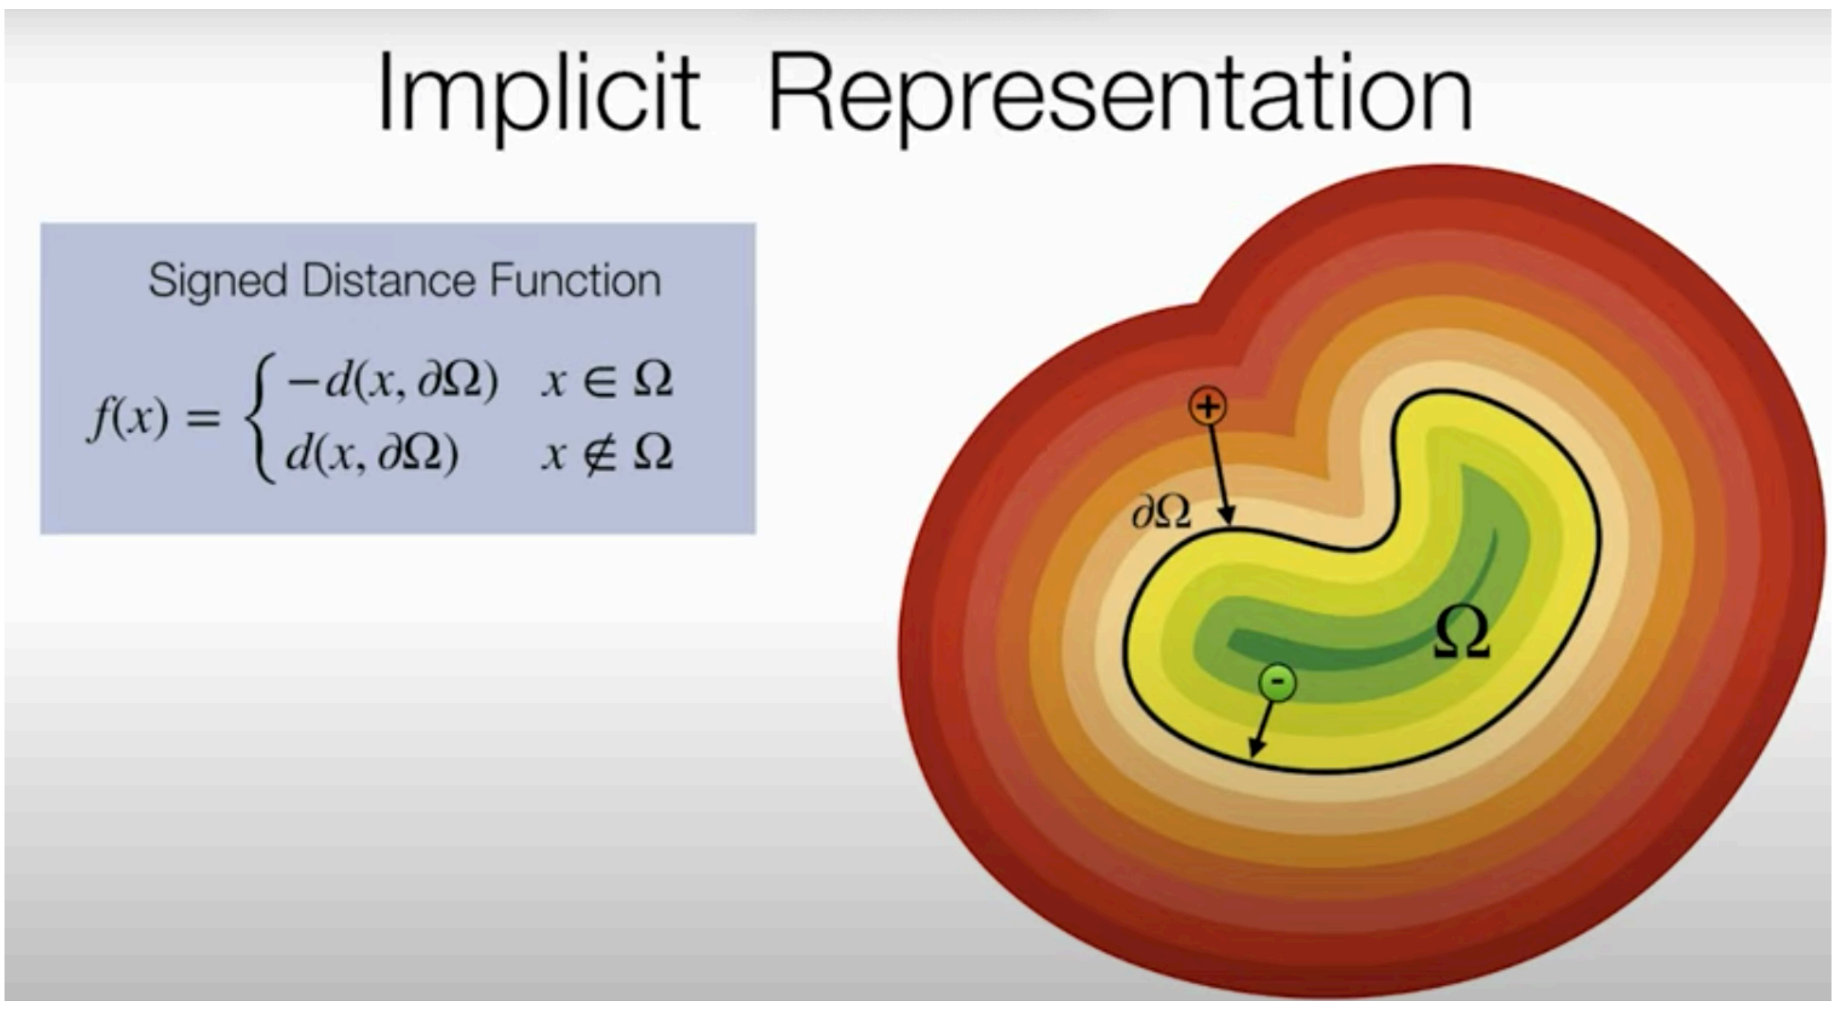

DeepSDF:

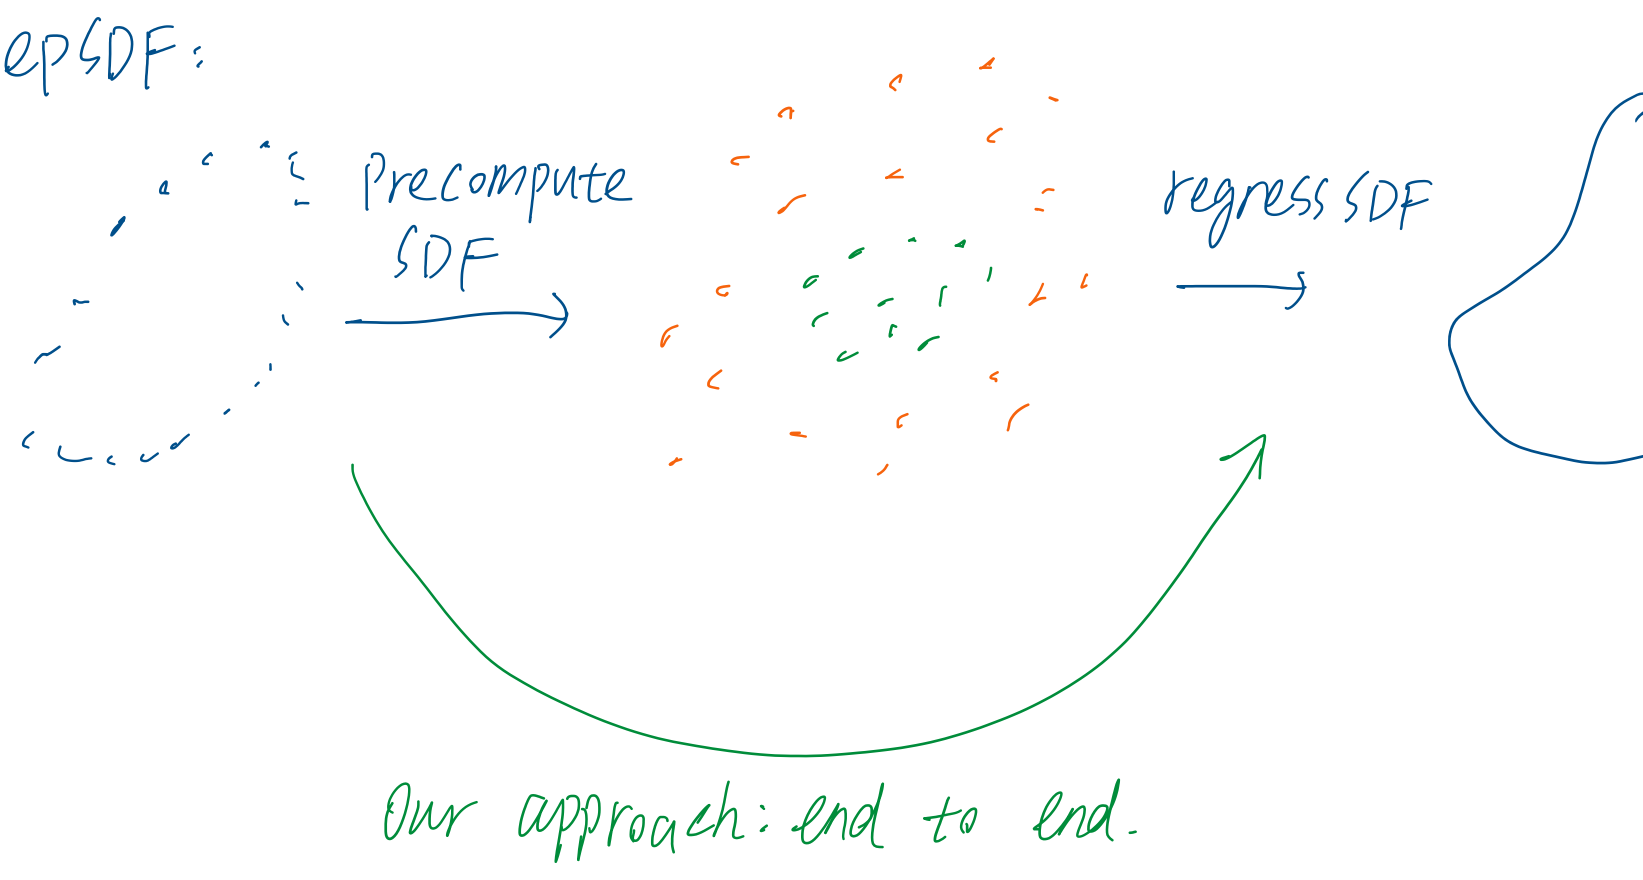

Eikonal PDE:

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

If we solve this PDE, we get the SDF.
from presentation ↑.

Analysis of the linear model and plane reproduction.

Consider a linear model

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

The loss takes the form

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