Dual Contouring of Signed Distance Data

Xiana Carrera[1], Ningna Wang[1], Christopher Batty[2], Oded Stein[3,4], Silvia Sellán[1]

Problem

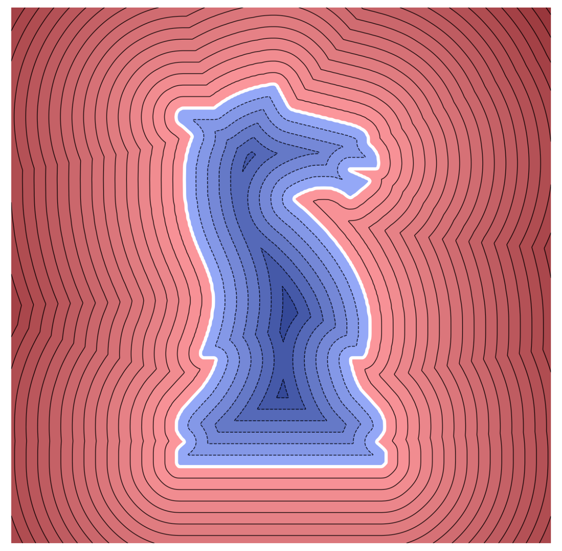

A **S**igned **D**istance **F**unction (SDF) stores, at each point in space, the distance to a surface (negative inside of it and positive outside).

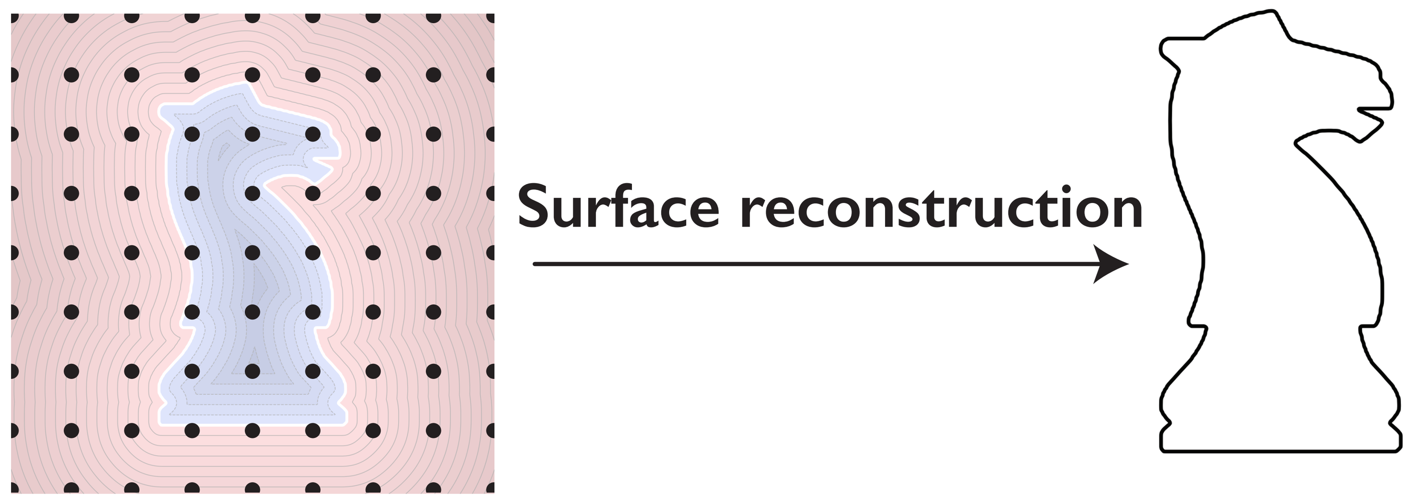

In some applications, we only have access to samples on a grid, not the full function...

... which complicates recovering the true surface represented by the SDF.

"Marching" (cubes/tetrahedra) methods cannot recover **sharp features**, unless these are aligned with their volumetric discretization.

Point clod methods smooth edges and merge gaps, unlike ours.

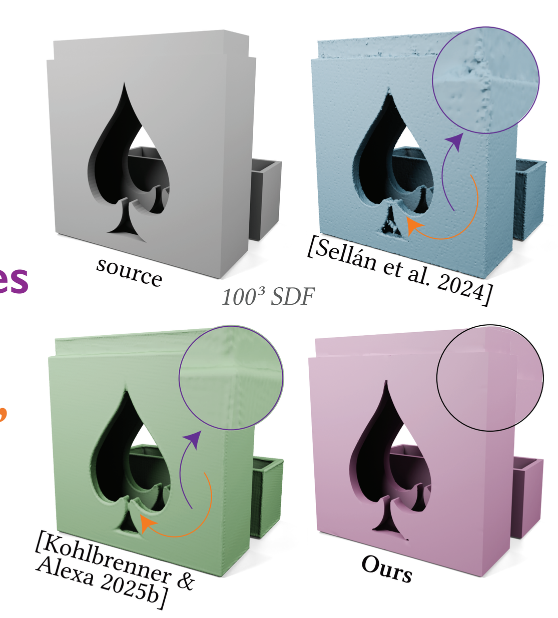

And Dual Contouring relies on **gradient information,** which we do not require.

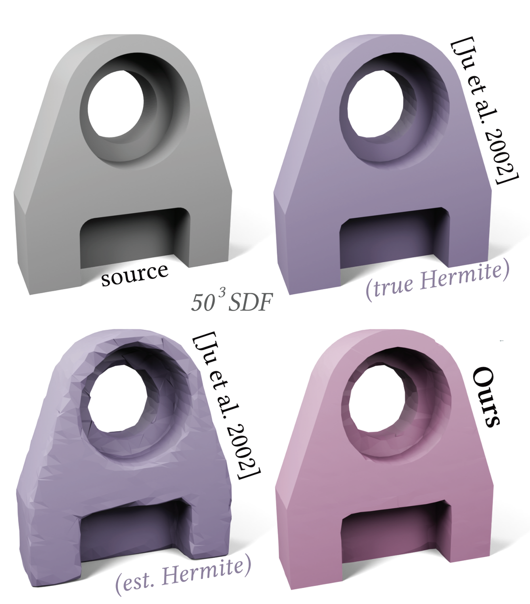

Method

Goal: reconstruct a surface from grid-sampled distances (no gradients, no learning, no additional queries).

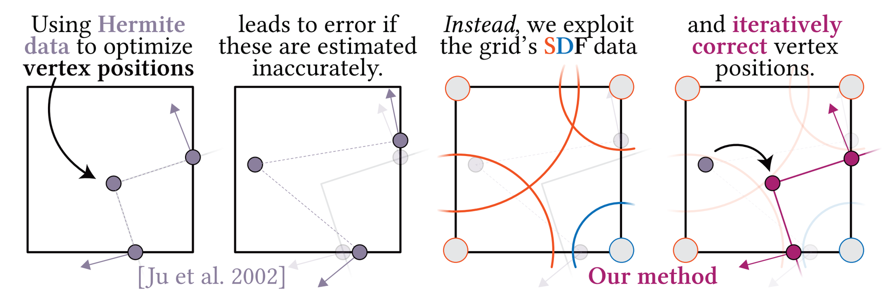

We place **one vertex per grid cell** and **optimize its position** so that the estimated surface is consistent with the distances encoded by the SDF.

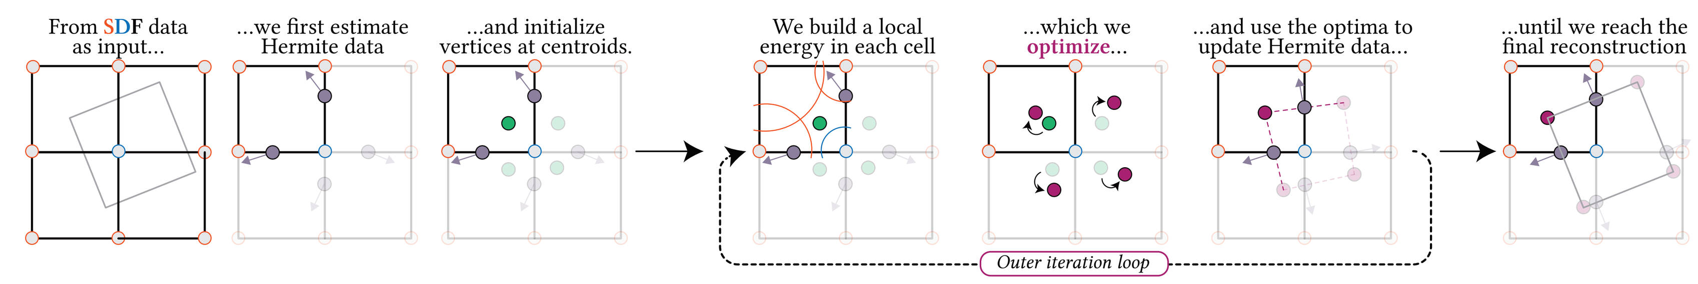

We alternate between locally optimizing vertex positions and globally updating geometric estimates (Hermite data: intersections and normals).

$$\mathbf{x}^{k,r+1} = \arg\min_{\mathbf{x}} \tilde{E}_d^{k,r}(\mathbf{x}) + w_H E_H^k(\mathbf{x}) + \mu\|\mathbf{x} - \mathbf{x}^{k,r}\|^2$$

The local optimization minimizes a **quadratic energy** combining a **linearized distance term**, the original **Dual Contouring energy**, and a **regularization term**.

$$E_H^{k,i}(\mathbf{x}) = \sum_{e_j \in c_i} \left((\mathbf{x} - \mathbf{h}_j^k) \cdot \mathbf{n}_j^k\right)^2$$

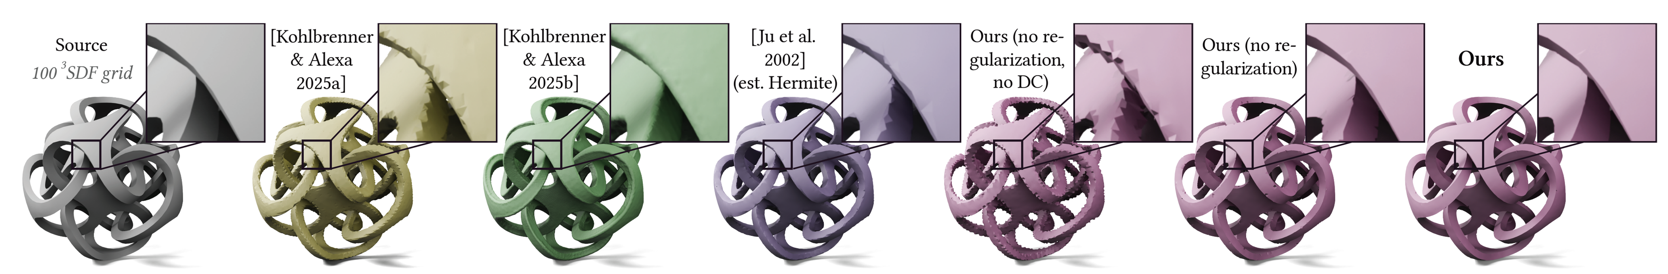

An SDF sample at a given point can be interpreted as a sphere that must be tangent to the true surface.

The linearized distance term measures how far the mesh is from satisying the sphere constraints.

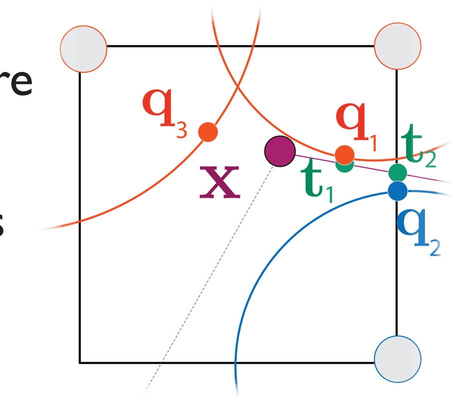

Removing tangential motion penalization

$$(|s_j| - d(\mathbf{u}_j, \mathcal{M}_i))^2 \approx \|\mathbf{t}_j - \mathbf{q}_j\|^2 \quad\longrightarrow\quad (|s_j| - d(\mathbf{u}_j, \mathcal{M}_i))^2 \approx ((\mathbf{t}_j - \mathbf{q}_j) \cdot \mathbf{d}_j)^2$$

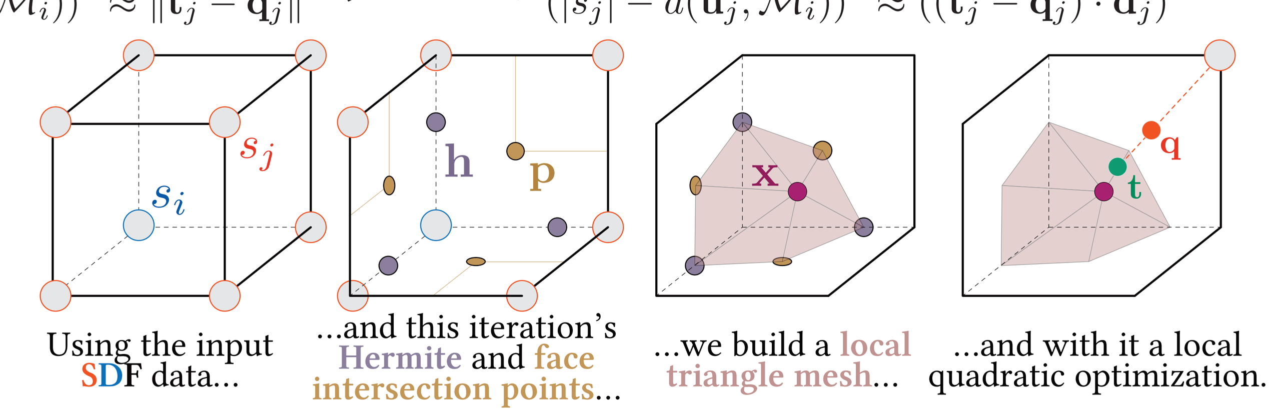

Results

We obtain **more faithful** reconstructions with a **better performance** than existing methods.

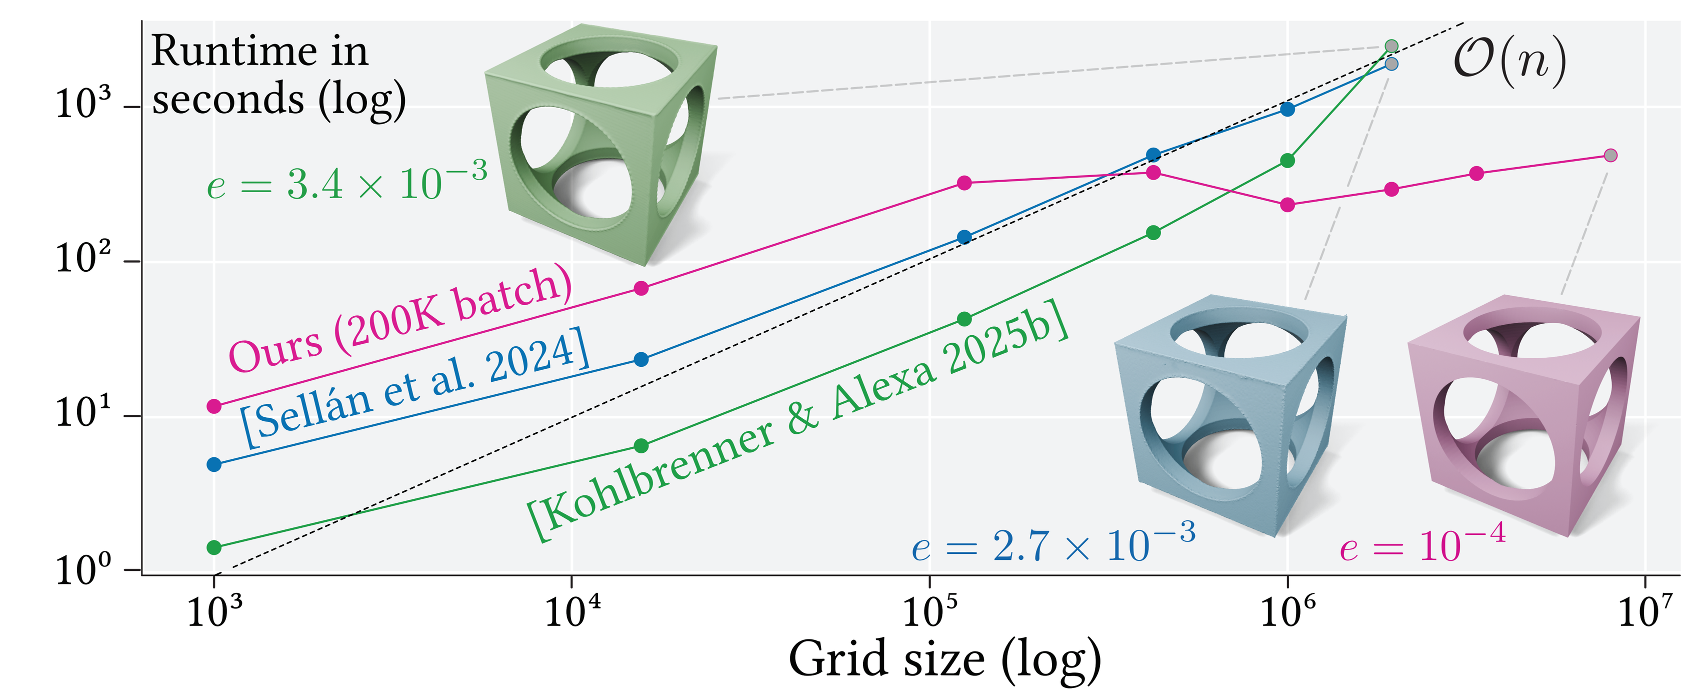

Our method demonstrates improved accuracy in applications such as **motion geometry reconstruction**...

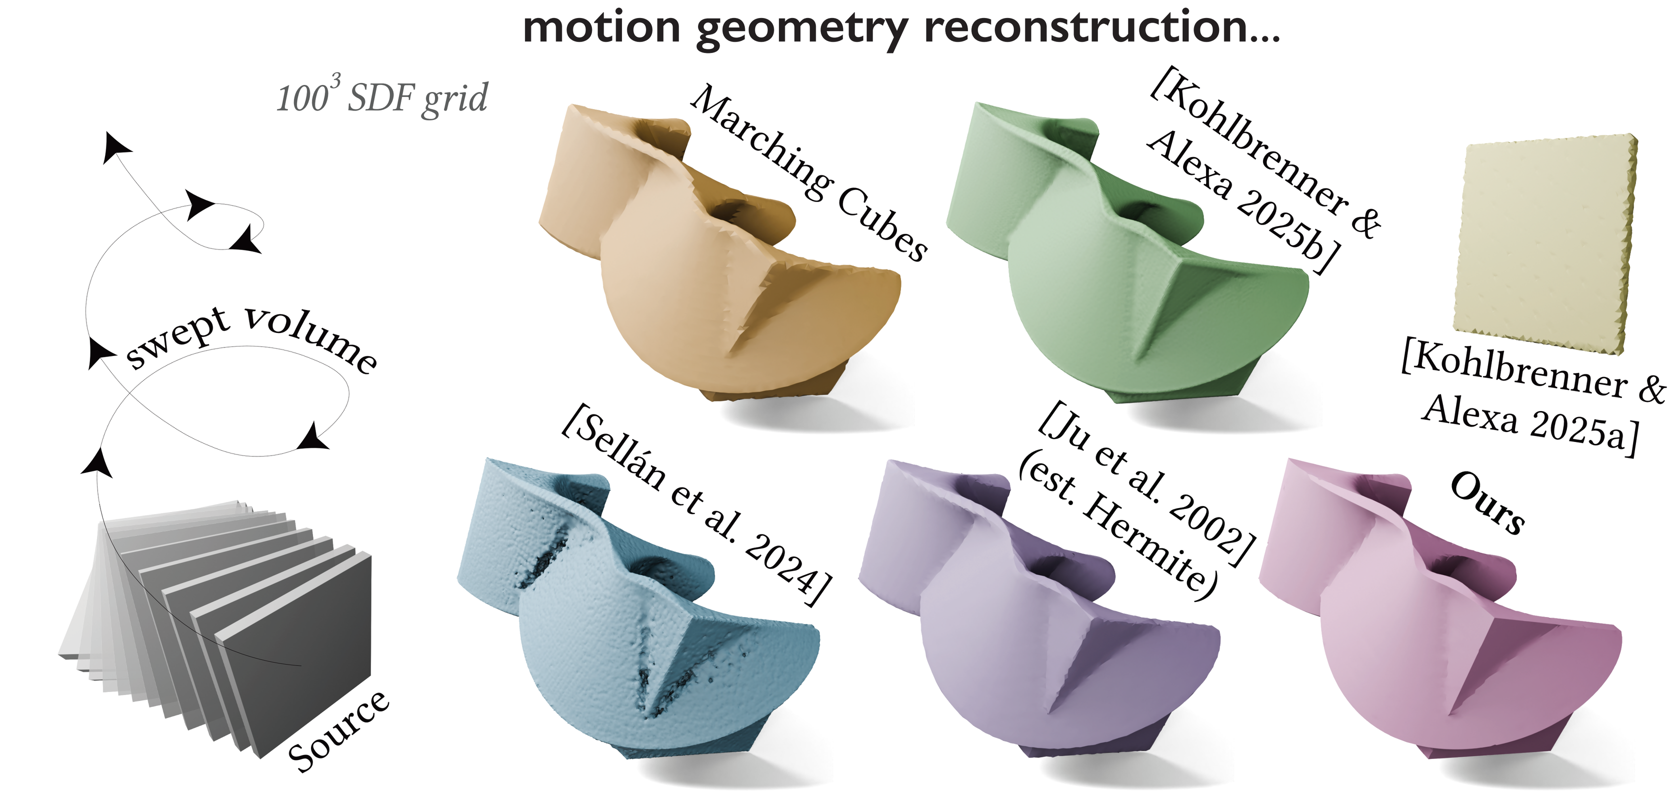

... and is successful even when restricting the amount of information available to a **narrow band** arond the input.

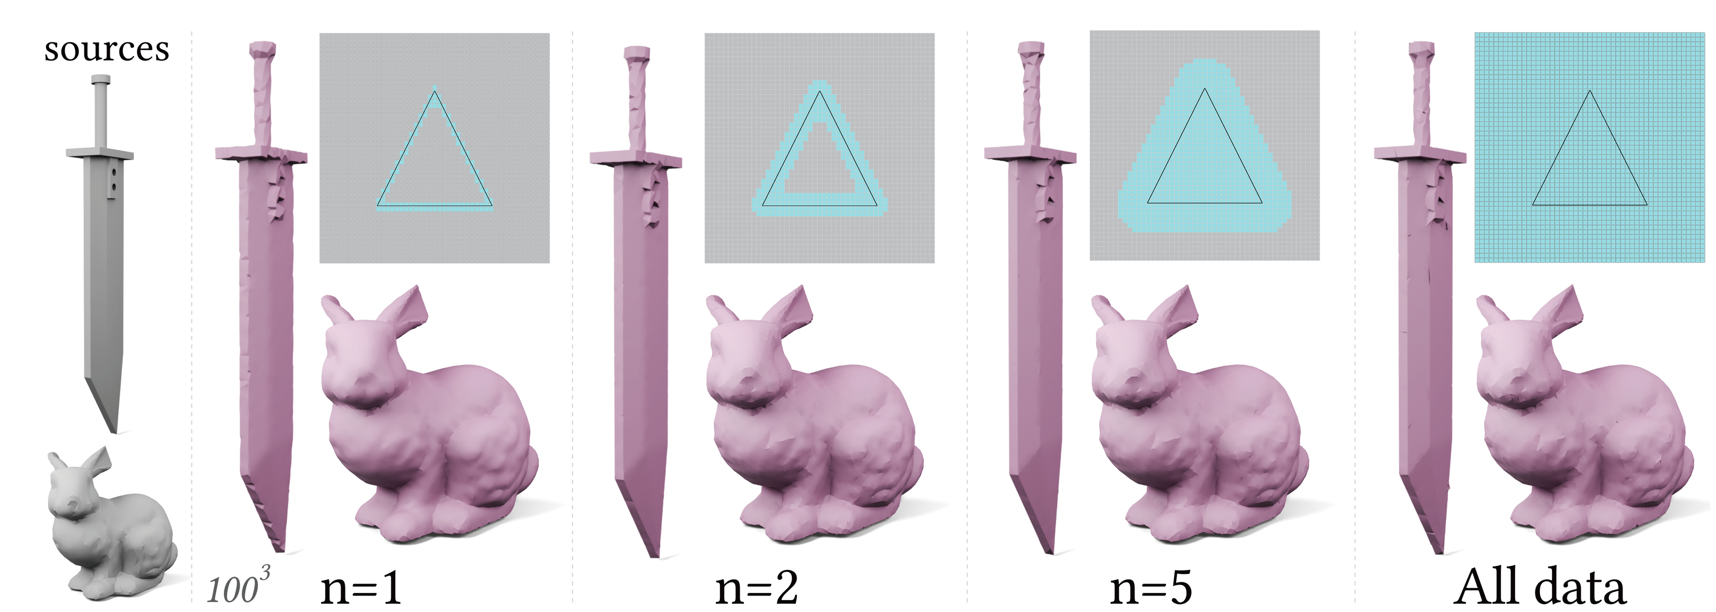

geometry and the city

[1]Columbia University
[2]University of Waterloo
[3]Technion
[4]University of Southern California

Conditionally accepted to SIGGRAPH 2026