1-Lipschitz Neural Distance Fields

Guillaume Coiffier
Louis Béthune

Signed Distance Function

Let Ω be a compact subset of $\mathbb{R}^n$ and $\partial\Omega$ be its boundary. The signed distance function S_Ω is defined over all $\mathbb{R}^n$ as:

$$S_\Omega(x) = y(x) \min_{p \in \partial\Omega} ||x - p||$$

where:

$$y(x) = \begin{cases} -1 \text{ if } x \in \Omega \\ 1 \text{ if } x \notin \Omega \end{cases}$$

S_Ω is a 1-Lipschitz function by construction.

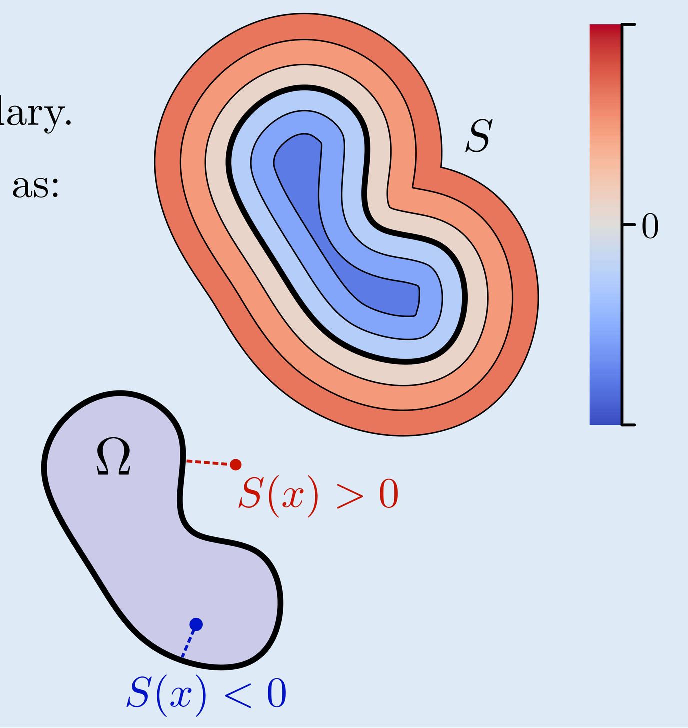

The 1-Lipschitz property garantees robustness to geometrical queries like projection. When dealing with *approximate* signed distance function, it is critical to always *underestimate* the true distance

$$\pi(x) = x - f(x)\frac{\nabla f(x)}{||\nabla f(x)||}$$

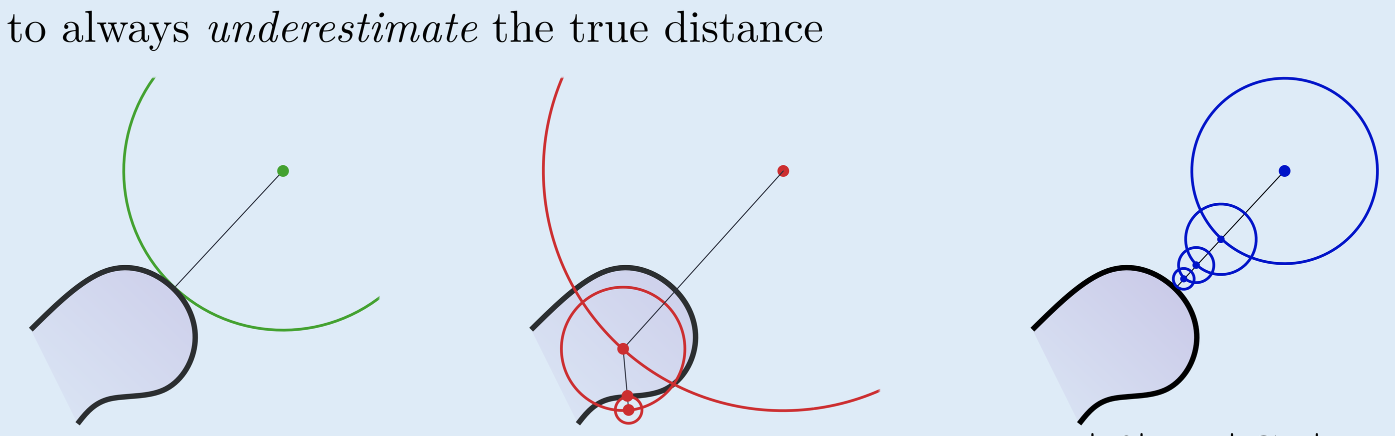

Learning a SDF: a binary classification problem

We propose to learn a signed distance function by optimizing the hinge-Kantorovitch-Rubinstein loss function over a 1-Lipschitz neural architecture:

$$\mathcal{L}_{hKR} = \mathcal{L}_{KR} + \lambda \mathcal{L}^m_{hinge} \qquad \lambda, m > 0$$

$$\mathcal{L}_{KR}(f_\theta, y) = \int_{\mathbb{R}^n} -y(x) f_\theta(x)\, dx$$

$$\mathcal{L}^m_{hinge}(f_\theta, y) = \int_{\mathbb{R}^n} \max\left(0, m - y(x) f_\theta(x)\right)\, dx$$

Achieving Robustness in Classification Using Optimal Transport with Hinge Regularization, Serrurier et al. (2021)

The SDF is the minimizer of the hKR loss

Let $m > 0$ and f^* be a minimizer of $\mathcal{L}_{KR}(f, y)$ under constraint that $\mathcal{L}^m_{hinge}(f, y) = 0$, where the minimum is taken over all possible 1-Lipschitz functions.

Then:

$$\forall x \in \mathbb{R}^n, \; \begin{cases} |S_\Omega(x)| > m & \implies \; f^*(x) = S_\Omega(x) \\ |S_\Omega(x)| \leqslant m & \implies \; |f^*(x) - S_\Omega(x)| \leqslant 2m \end{cases}.$$

Small m values implies high geometric fidelity but less stable training.

1-Lipschitz Neural Networks

Optimizing the hKR loss requires a neural architecture that is 1-Lipschitz by construction. This can be made by composing the following layer:

$$x_{k+1} = x_k - 2WD^{-1}\sigma(W^T x_k + b)$$

where:

$$D = \operatorname{diag}\left(\textstyle\sum_j |(W^T W)_{ij} \exp(q_j - q_i)|\right)$$

$$\sigma(x) = \max(0, x)$$

Parameters of the layers are W (matrix), b and q (vectors).

A Unified Algebraic Perspective on Lipschitz Neural Networks, Araujo et al. *(2023)*

Algorithmic pipeline

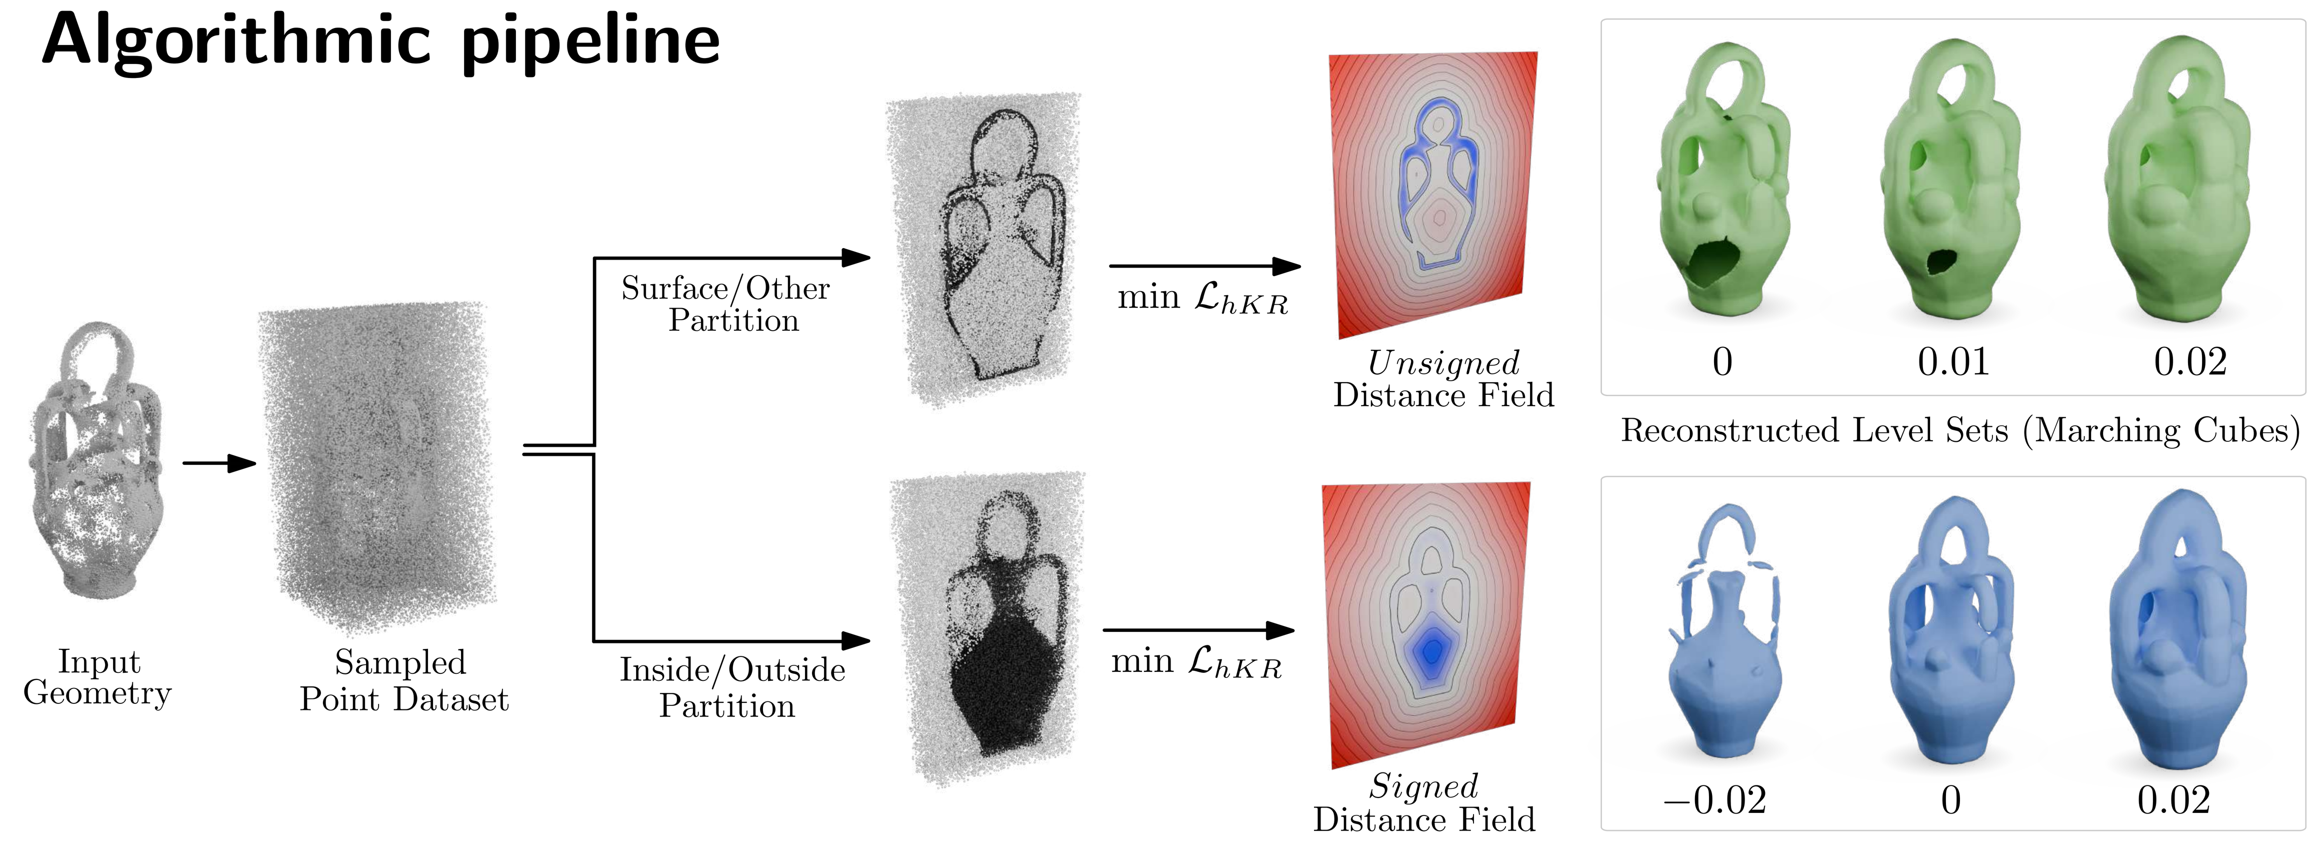

Inside/outside partitioning

The *generalized winding number* is the sum of solid angles over a surface. It allows the computation of y even for noisy, incomplete or faulty input data.

$$w_{\partial\Omega}(x) = \frac{1}{4\pi}\int_{\partial\Omega} d\Theta(x)$$

$w \in \{0; 1\}$ for closed geometries

$w \in [0; 1]$ for imperfect geometries

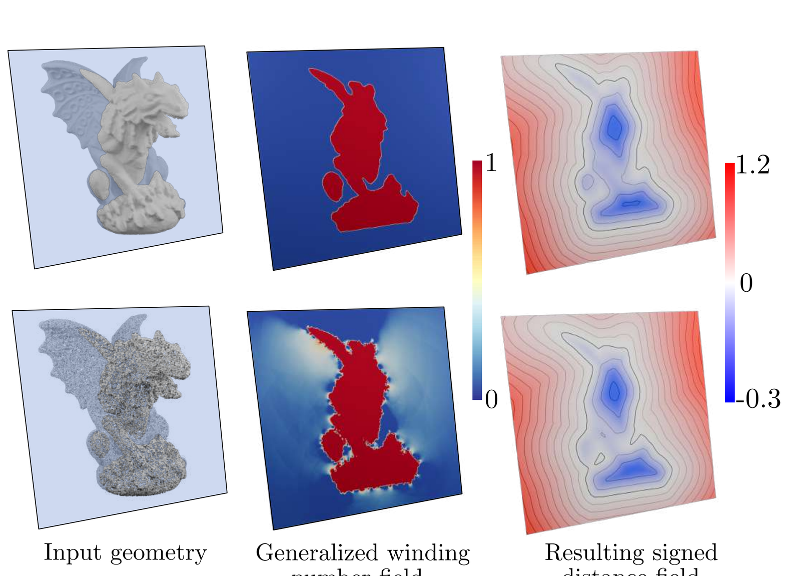

Fast Winding Numbers for Soups and Clouds, Barill et al. (2018)

Unsigned distance field

The hyperparameter m in the hinge loss plays the role of a small 'width' around the isosurface.

This allows the reconstruction of unsigned distance fields for open curves and surfaces.

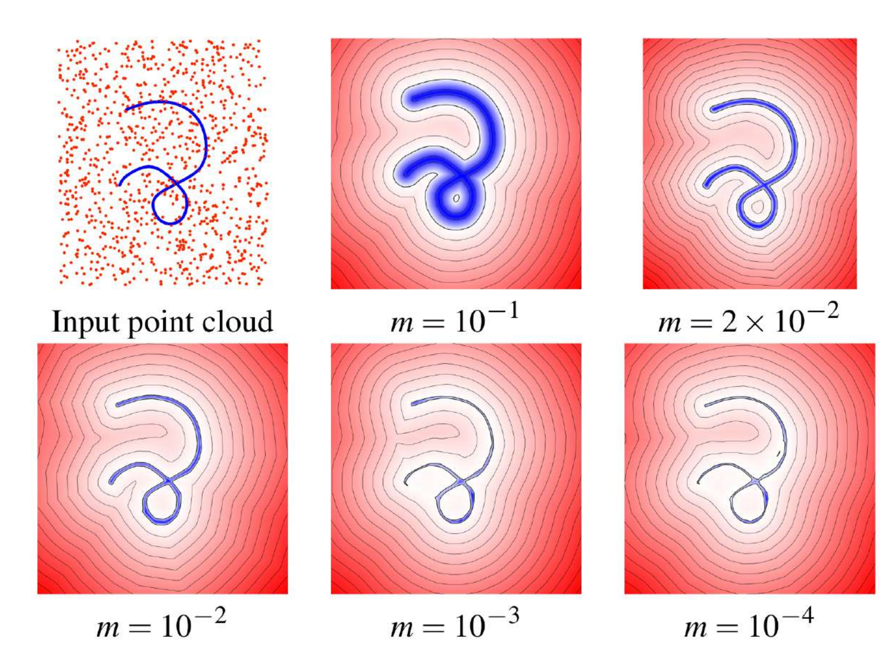

Isosurface reconstruction

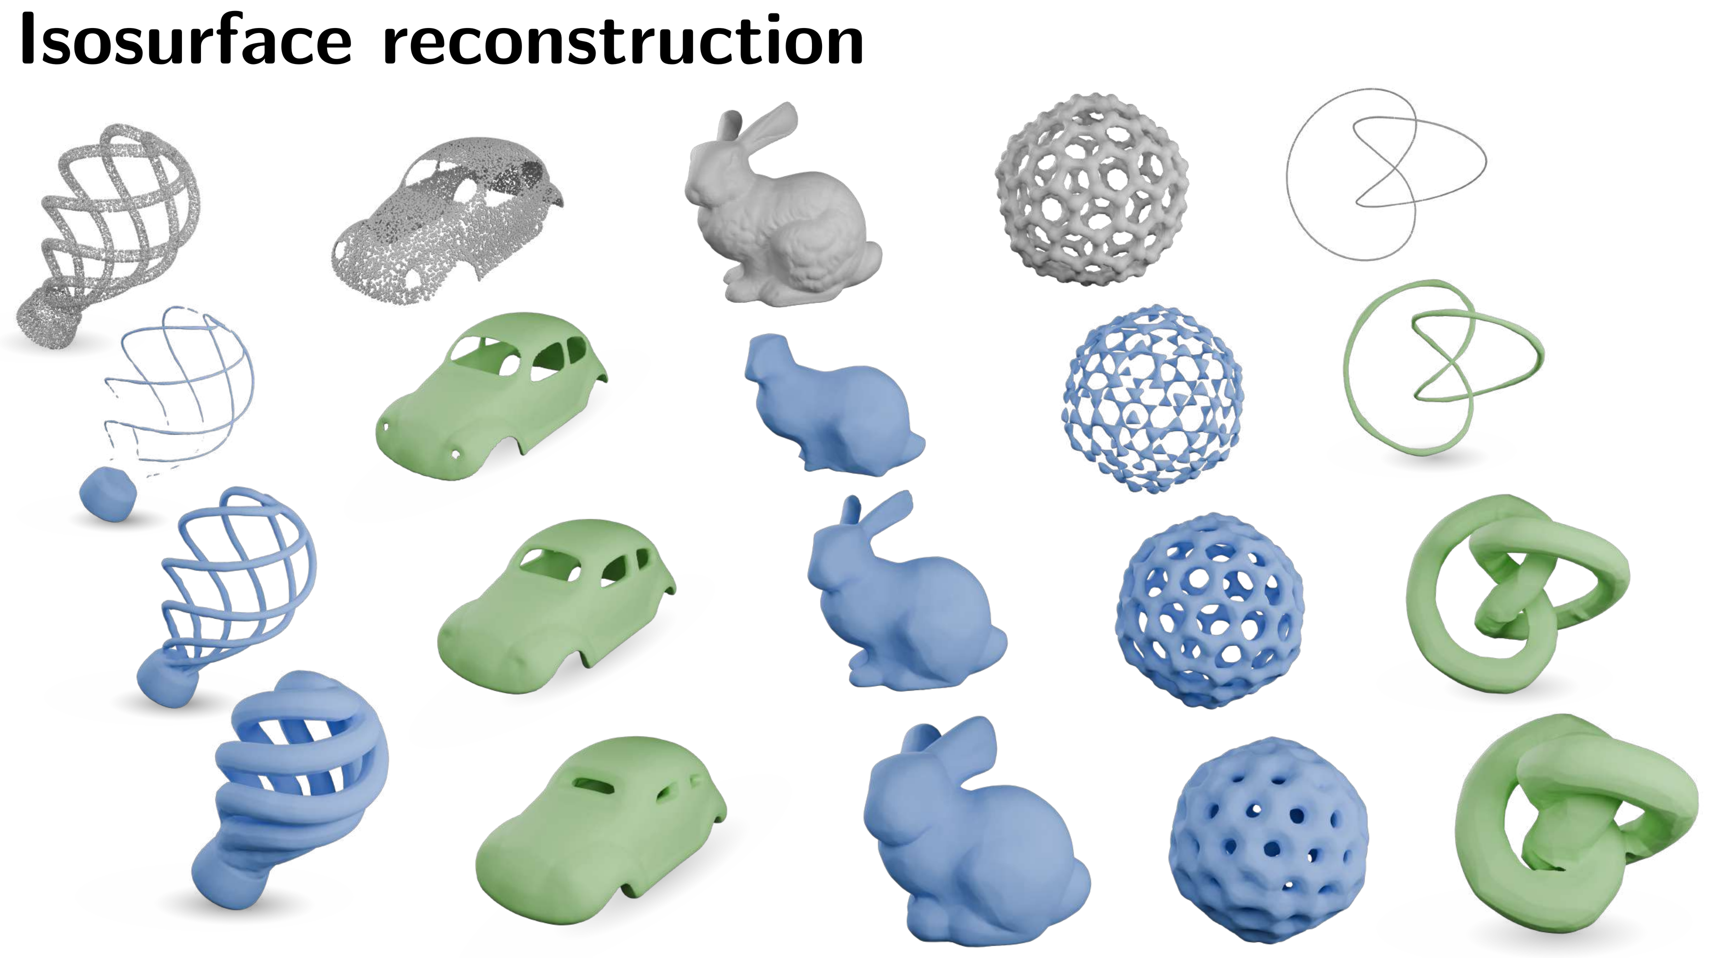

Robust Geometric Queries

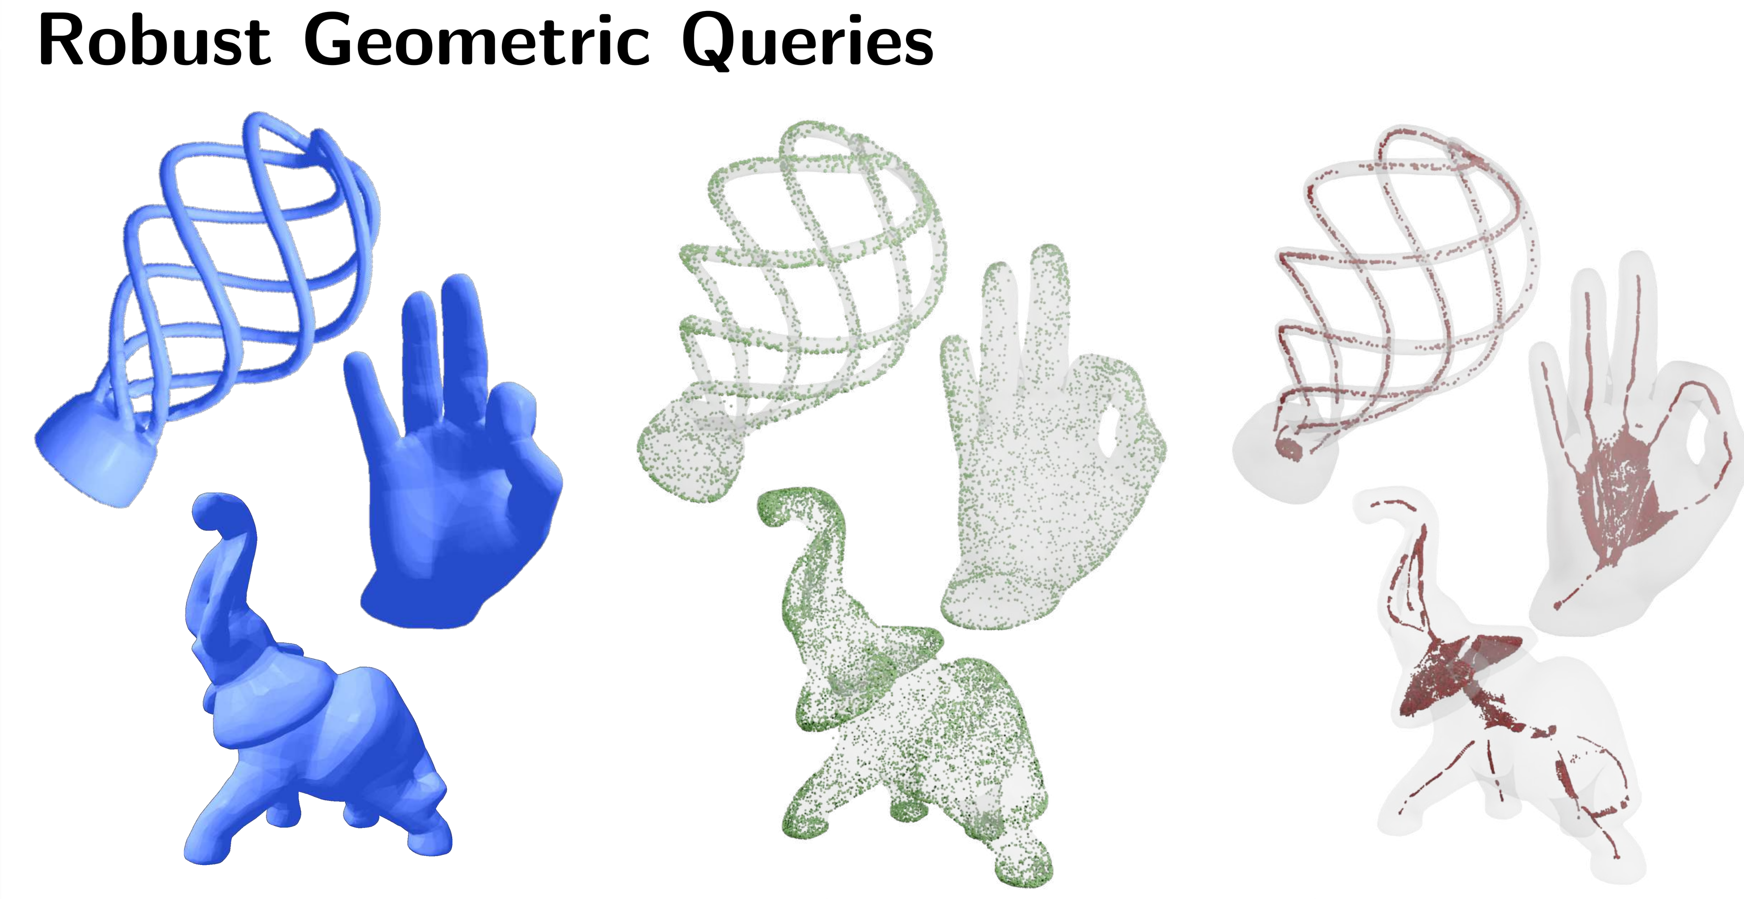

UCLouvain
IRIT
Symposium on Geometry Processing 2024
Best Paper Award