# A Heat Method for Generalized Signed Distance

Nicole Feng, Keenan Crane
Carnegie Mellon University

## Algorithm

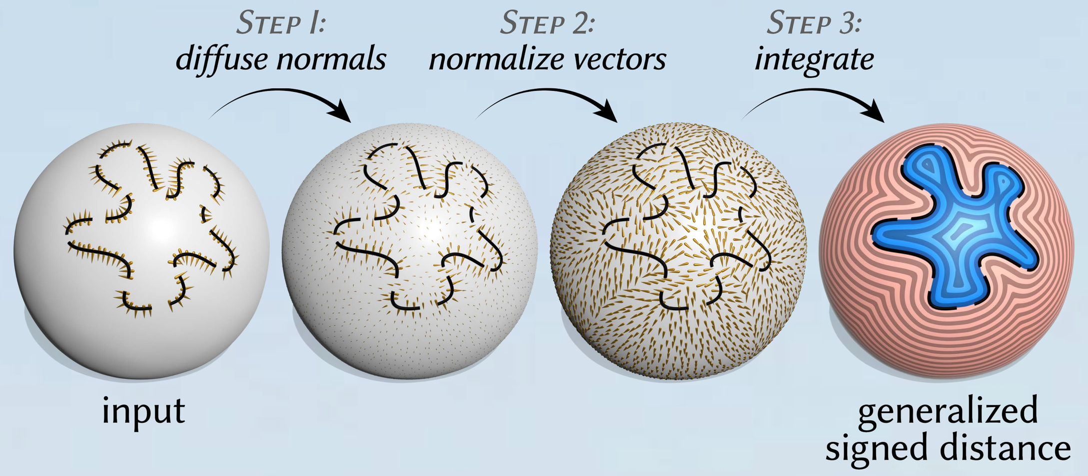


## Theory

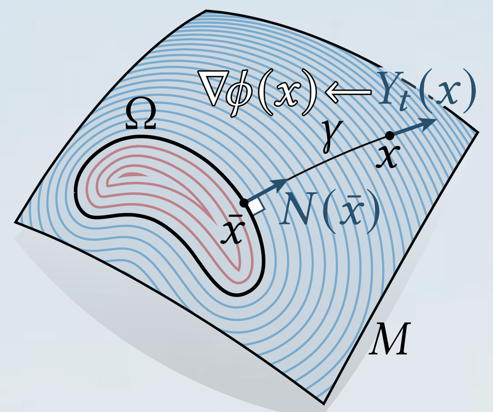


As diffusion time $t \to 0$, the diffused vector $X_t(x)$ at each point $x \in M$ aligns with the vector obtained via parallel transport along a minimal geodesic $\gamma$ of the normal $N(\bar{x})$ at the closest point $\bar{x} \in \Omega$. [Berline et al. 1992, Theorem 2.30].

Since parallel transport along geodesics preserves tangency, this vector will be tangent to $\gamma$ — and since traveling along $\gamma$ is the quickest way back to $\Omega$, it is parallel to the unsigned distance gradient. Since we transport oriented normals, we get the correct sign. We can hence normalize $X_t$ to obtain an approximation $Y_t := X_t/\|X_t\|$ of the signed distance gradient.

## Discretization

In Step 1, we use the edge-based *Crouzeix-Raviart* basis functions, which makes it straightforward to discretize curve sources, and enables a purely intrinsic formulation. We identify tangent vectors with complex numbers.
In Step 3, we use basis functions at vertices, so we obtain distance values at vertices.

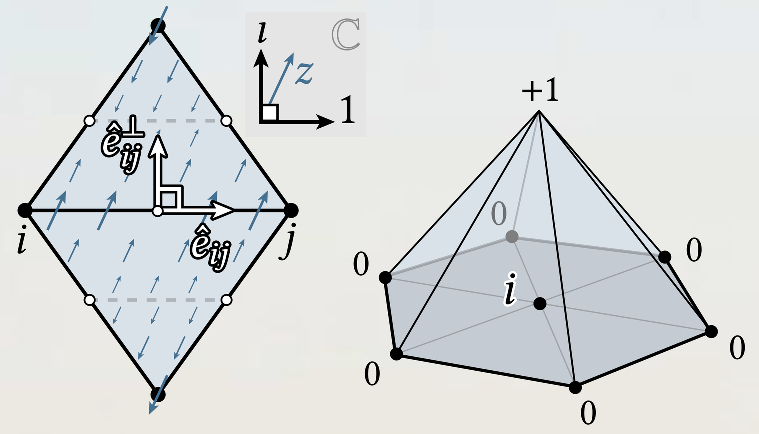

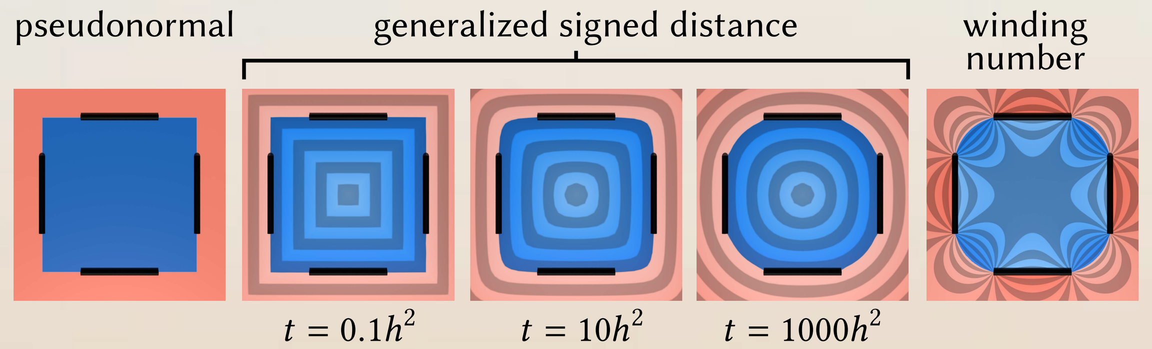


For different diffusion times, our method effectively interpolates between the "linear" prior used by the pseudonormal test (as $t \to 0$) and the "harmonic" prior used by generalized winding numbers (as $t \to \infty$).

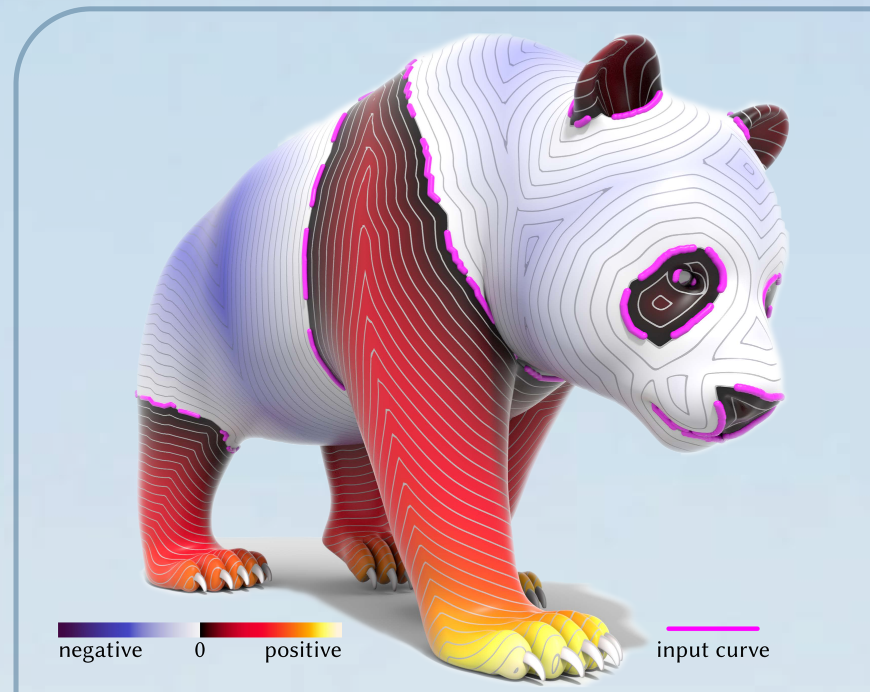


The **signed heat method** computes a well-behaved signed distance function (SDF), even for imperfect or incomplete input geometry.

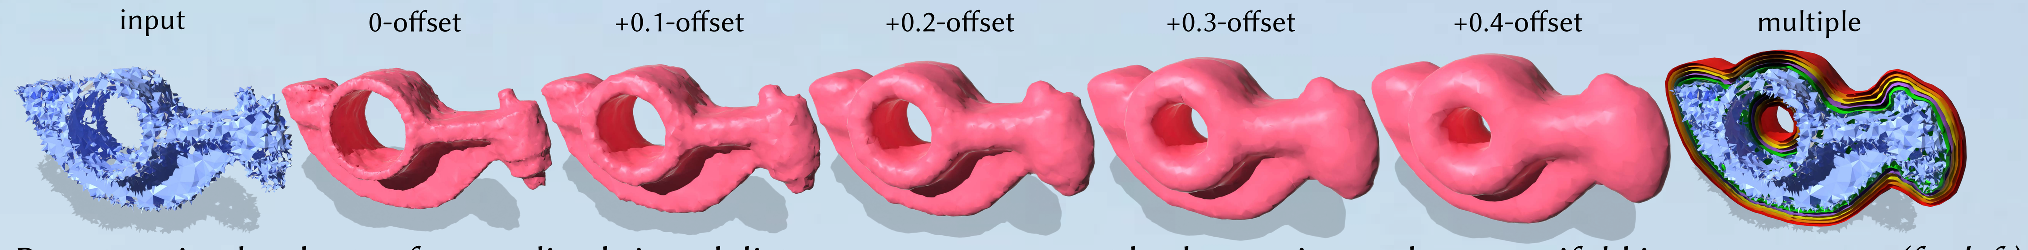


By extracting level sets of generalized signed distance, we can convert broken, noisy, and nonmanifold input geometry *(far left)* into closed, regular, manifold surfaces and evenly-spaced offset surfaces.

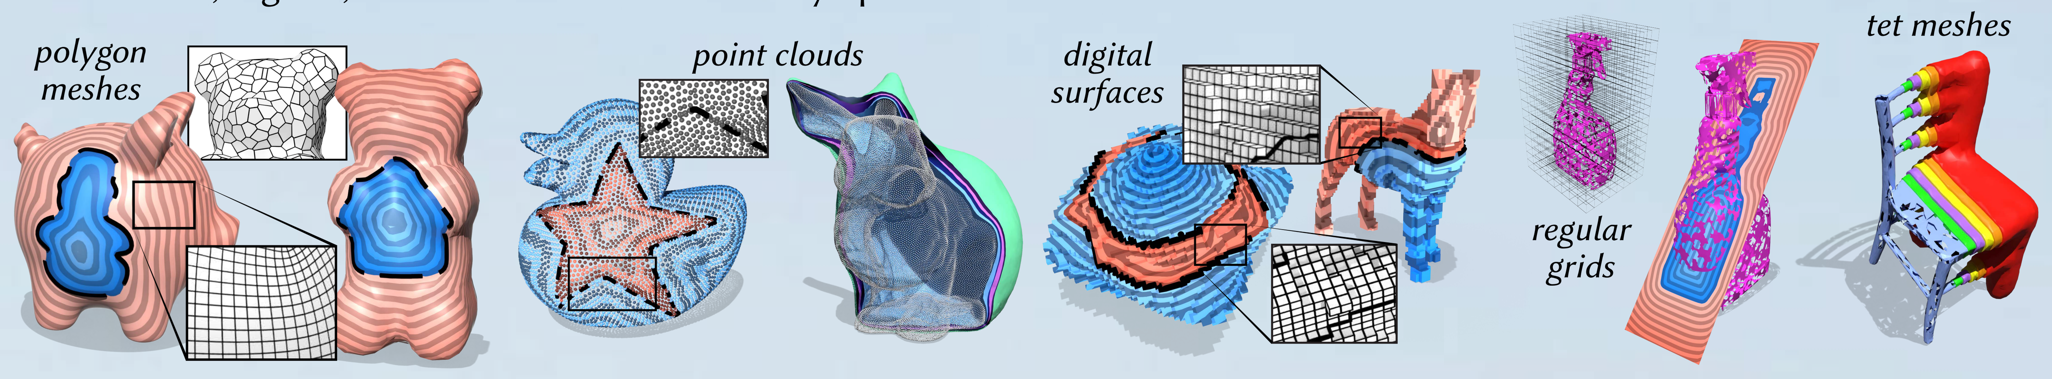


Our method can be applied in any dimension to any spatial discretization, including polygon meshes, point clouds, digital surfaces, tetrahedral meshes, and regular grids.

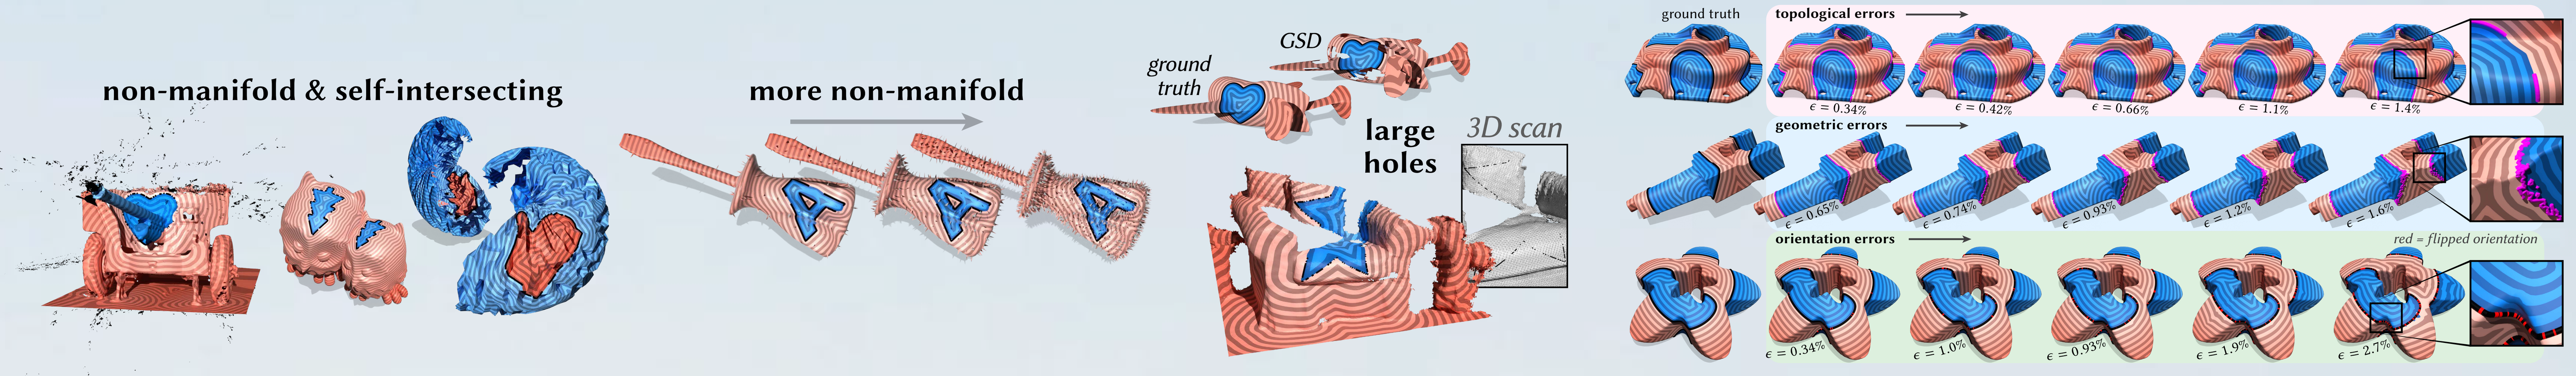


Our algorithm works out-of-the-box on non-manifold, non-orientable, and self-intersecting domains; and fails gracefully in the presence of significant topological, geometric, or orientation errors in the source geometry. Errors ε in geodesic distance are displayed relative to the exact polyhedral SDF of a finely sampled version of the original curve.

### mix/match signed & unsigned distance

### morphological operations

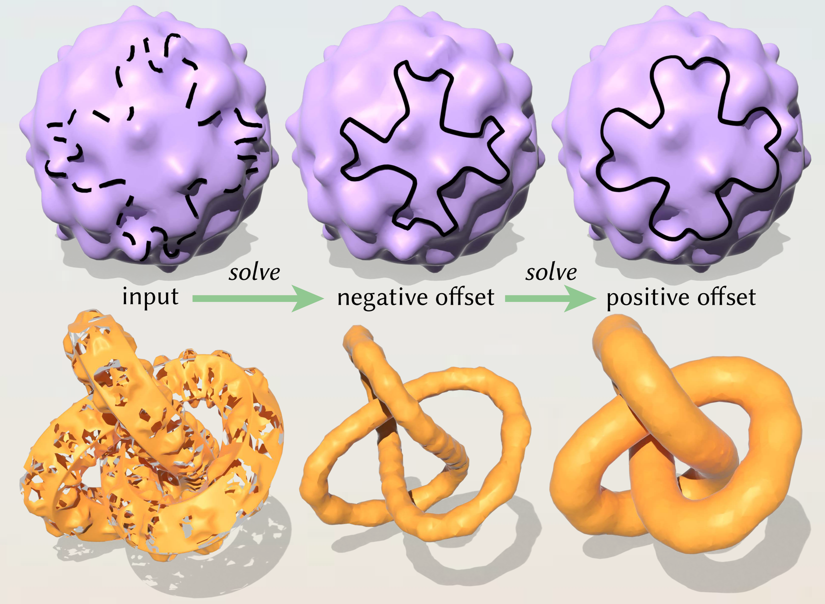


### robust inside/outside & distance

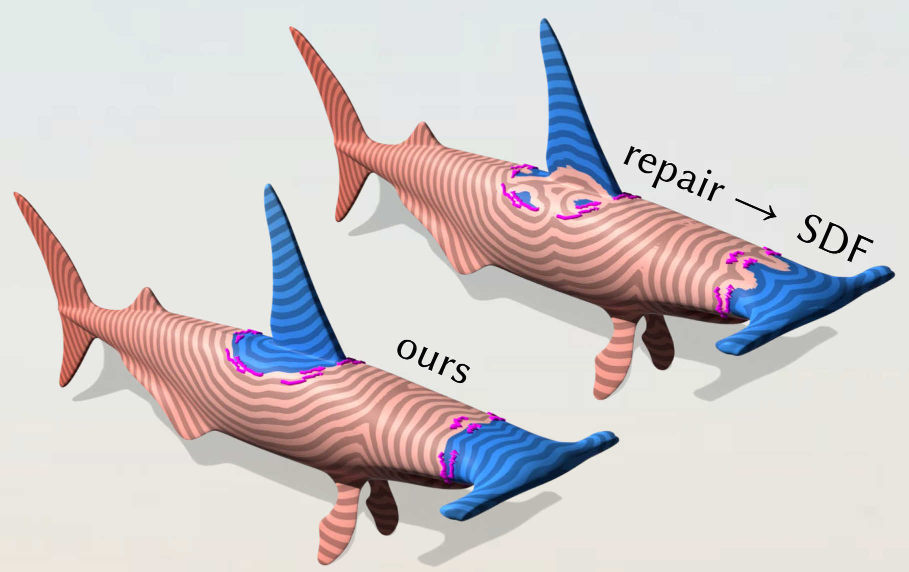


### correct boundary conditions

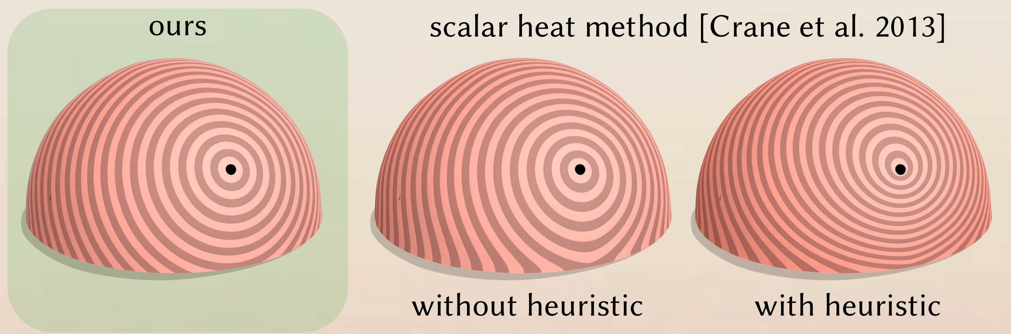


### surface illustration

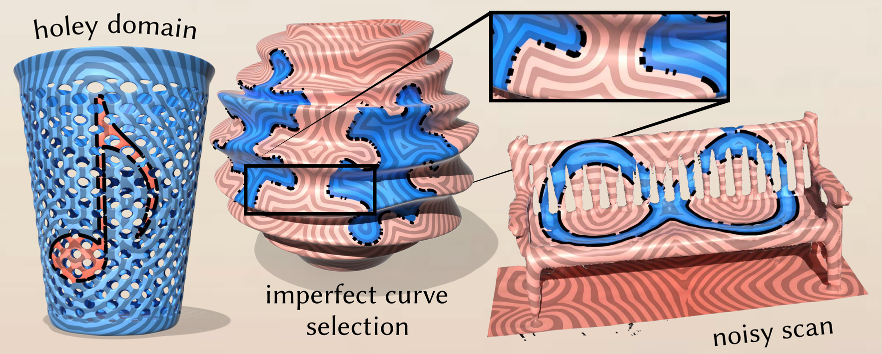


### non-orientable domains

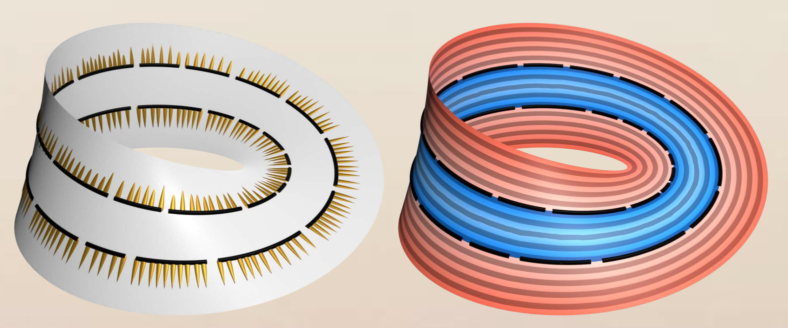

## Nonbounding loops and discontinuous distance

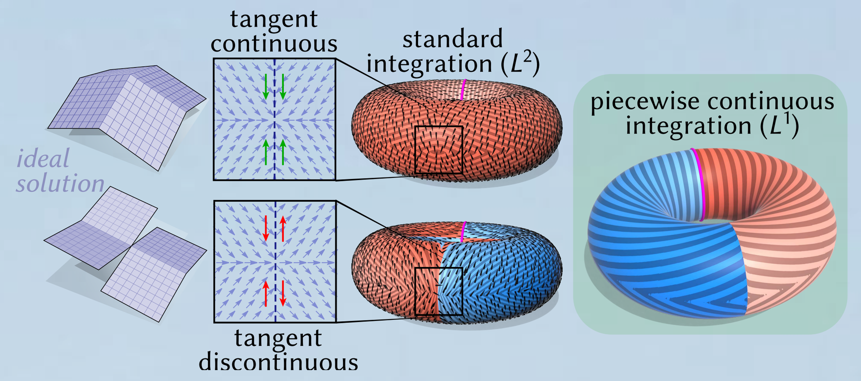


Unlike the gradient of an unsigned distance function *(top)*, the unit vector field we compute for signed distance might not be integrated by any continuous function *(bottom)*. In such scenarios, we can instead compute a piecewise continuous SDF *(right)*.

## Faithful surface reconstruction

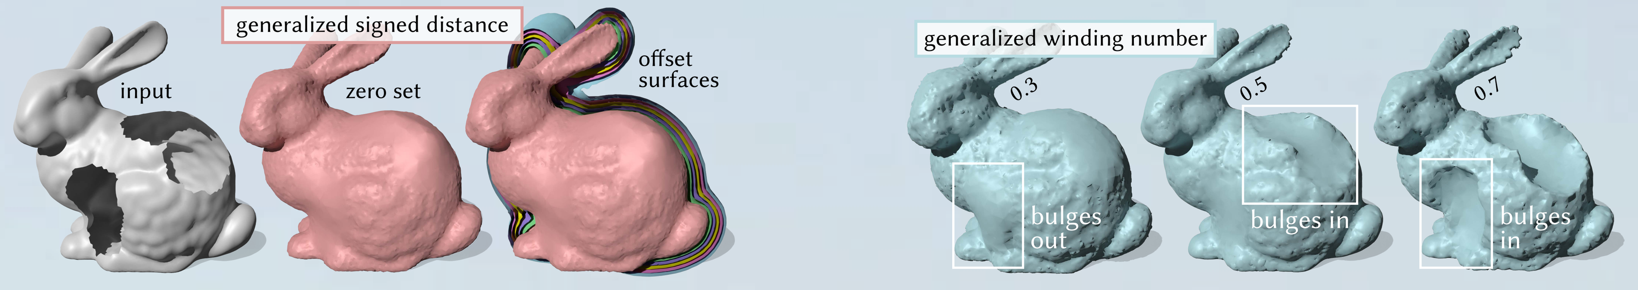


GWN/PSR completes surfaces with saddle-shaped harmonic patches that exhibit poor normal continuity with the observed geometry (across many contour values). Our method directly incorporates normal information, providing more plausible reconstruction even for large holes.

## Optional extensions

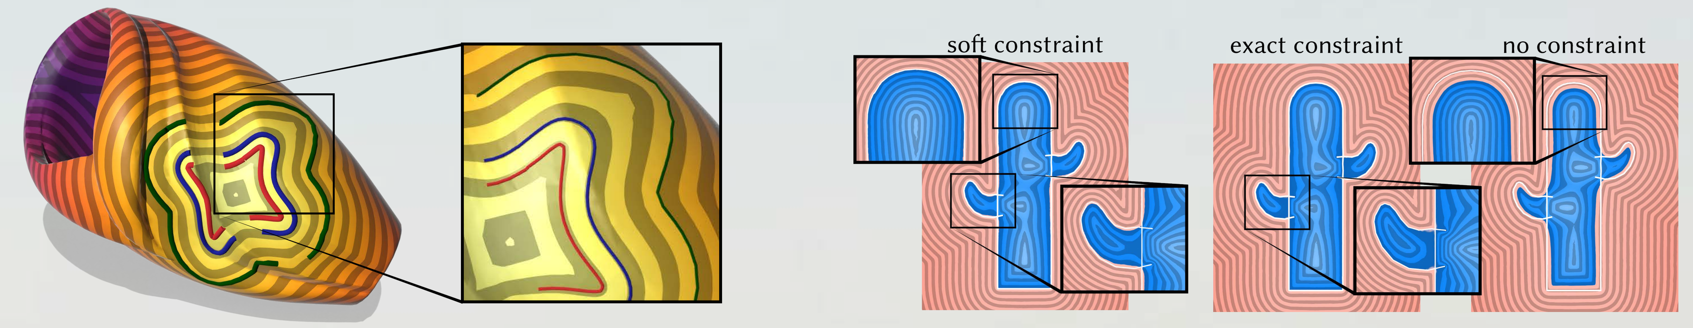


*simultaneous fitting of multiple level sets* *soft level set constraints*

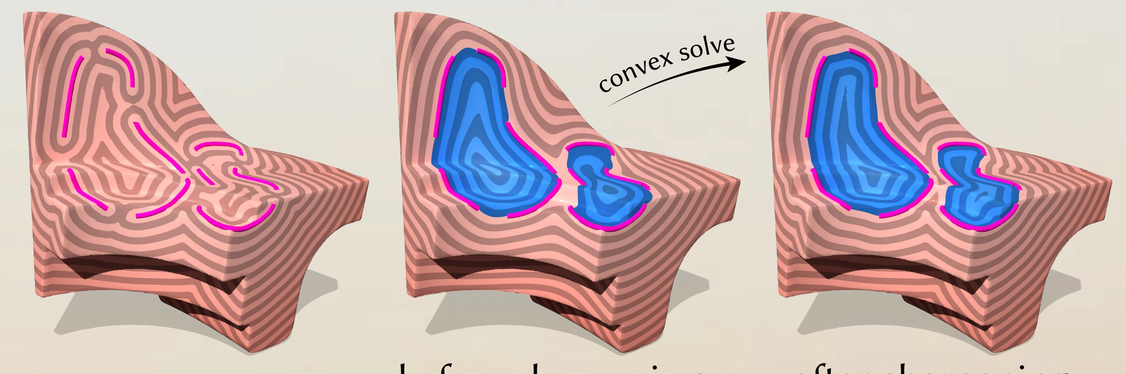


*Left:* Methods based on convex optimization yield more accurate distances, but compute only unsigned distance. *Right:* Using our method as a warm start, we can "sharpen" distance while preserving the inside-outside classification. Here we start with a large diffusion time ($t = 100h^2$) to visually emphasize the effect.