Deep Parametric Shape Predictions using Distance Fields

Dmitriy Smirnov[1] Matthew Fisher[2] Vladimir G. Kim[2] Richard Zhang[2] Justin Solomon[1]

[1]Massachusetts Institute of Technology [2]Adobe Research

Summary

While CNNs are well-adapted to raster graphics, they are less suited for **parametric shapes**, which use sparse sets of parameters to express geometry. We formulate an **Eulerian version of Chamfer distance**, a common metric for geometric similarity, by **analytically computing a distance field to parametric primitives**. We apply our new framework to a variety of 2D and 3D vectorization tasks.

Distance Field Geometry Loss

We propose a framework for computing geometric loss functions using distance fields. We optimize these loss functions by analytically computing distances to parametric shapes during training.

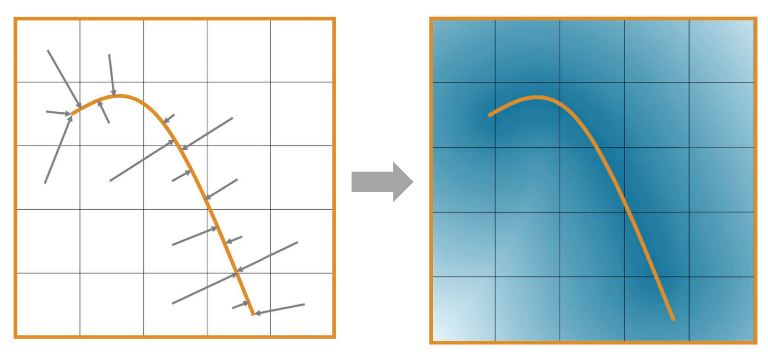

We define **general distance field loss** between two shapes *A* and *B* as

$$\mathcal{L}_\Psi[A,B] = \frac{1}{\mathrm{Vol}(S)}\int_{x\in S}\Psi_{A,B}(x)\,\mathrm{d}V(x)$$

for some measure of discrepancy Ψ.

Surface discrepancy measures overlap between two shapes:

$$\Psi^{\mathrm{surf}}_{A,B}(x) = \delta\{\ker \mathrm{d}_A^2\}(x)\mathrm{d}_B^2(x) + \delta\{\ker \mathrm{d}_B^2\}(x)\mathrm{d}_A^2(x)$$

Normal alignment discrepancy aligns the normals of the two shapes to each other:

$$\Psi^{\mathrm{align}}_{A,B}(x) = 1 - \langle \nabla \mathrm{d}_A(x), \nabla \mathrm{d}_B(x)\rangle^2.$$

We compute the loss function integral during training by sampling over a uniform grid *G*:

$$\mathcal{L}_\Psi[A,B] \approx \frac{1}{|G|}\sum_{x\in G}\Psi_{A,B}(x)$$

Comparison to Chamfer Distance

Chamfer distance is commonly used as a loss function in deep learning pipelines that produce geometry.

$$\mathrm{Ch}_{\mathrm{dir}}(X,Y) = \frac{1}{|X|}\sum_{x\in X}\min_{y\in Y}\|x-y\|_2^2$$

Because it requires sampling points from geometry, it is sensitive to misalignment of normal and nonuniform sampling. We propose loss functions that do not suffer from these artifacts, allowing us to improve deep predictions of parametric shapes.

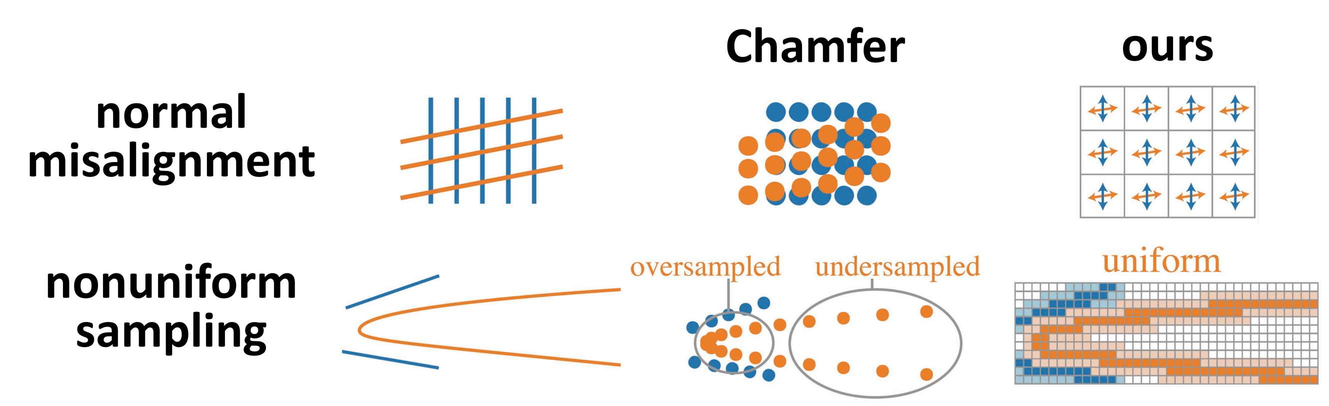

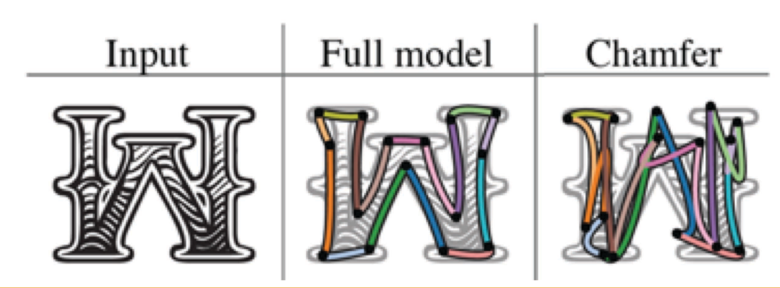

Pipeline Overview

Given an image or distance field, our network directly outputs parameters, which define Bézier curves in 2D or other primitives in 3D that vectorize the input. In 2D, we use templates that constrain the topology of the predicted curves. During training, we compute output distance fields and optimize a surface loss as well as an explicit normal alignment loss.

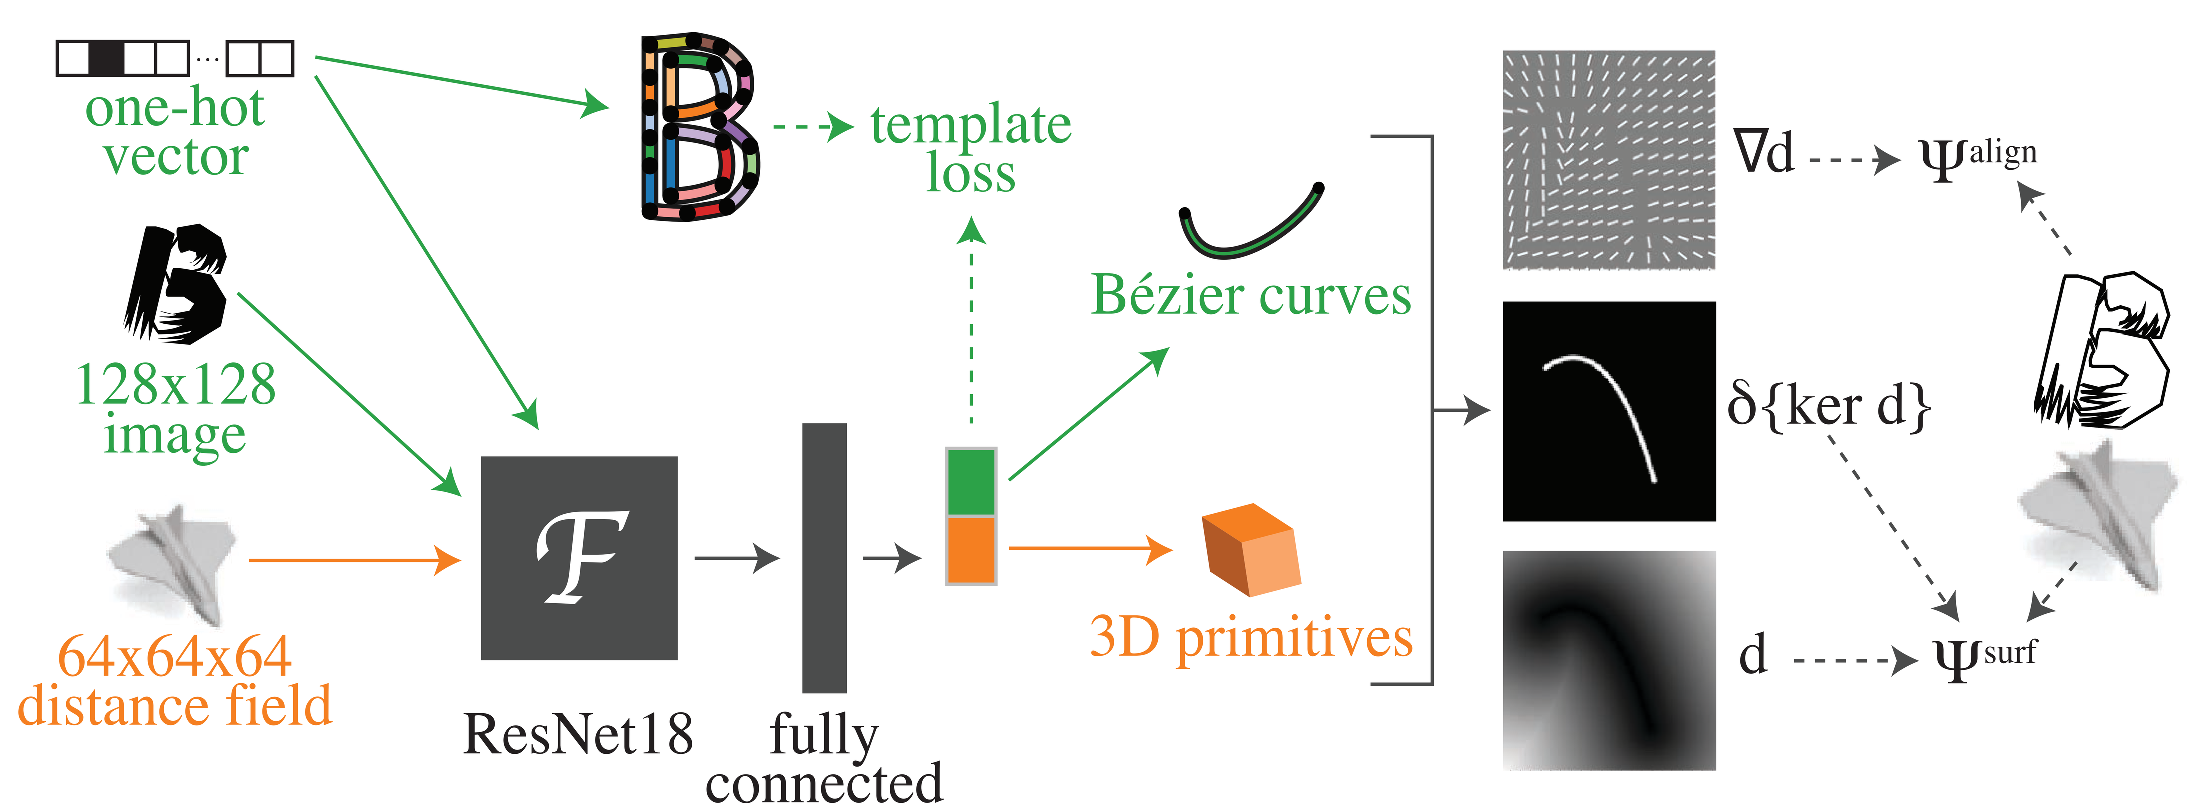

2D: Font Vectorization

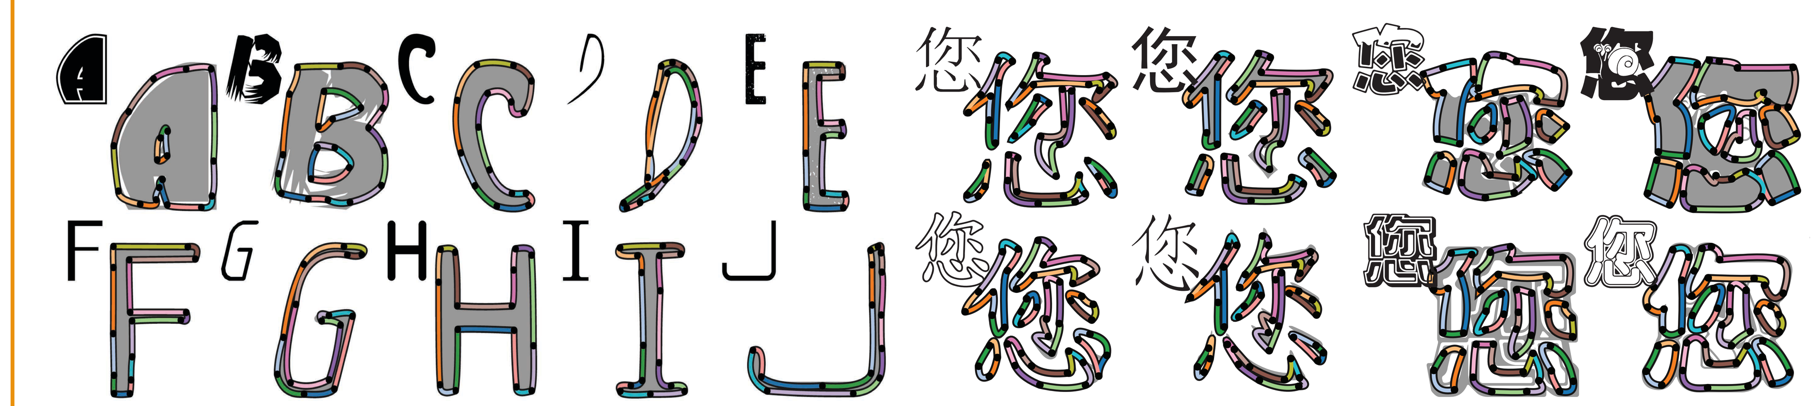

vectorization of font glyphs with Bézier curves

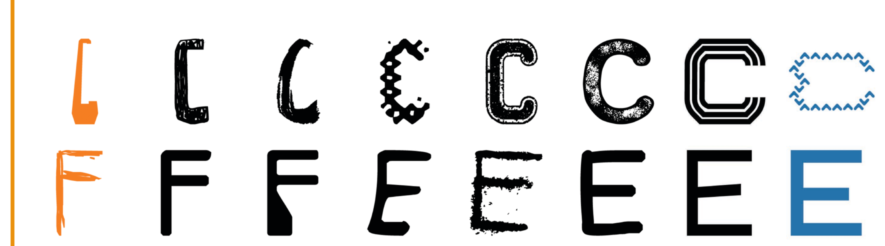

interpolation between glyphs using learned vectorizations

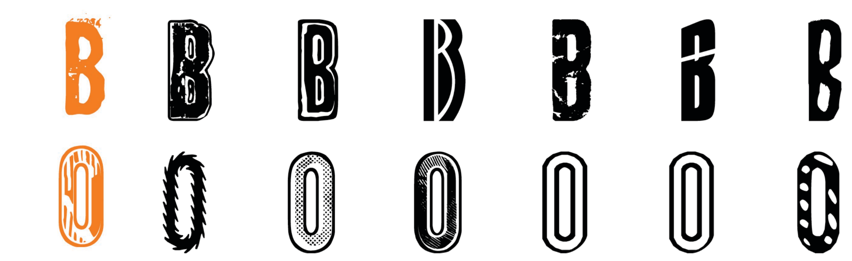

glyph nearest neighbor retrieval using learned vectorization

mixing of style and structure

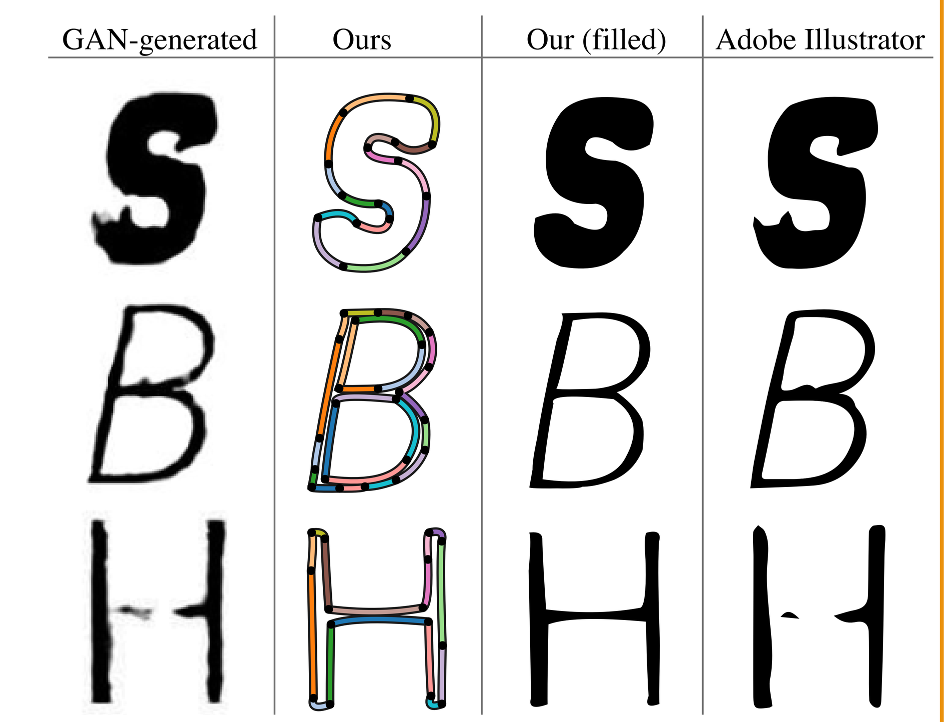

repair and vectorization of [1]

3D: Shape Abstraction

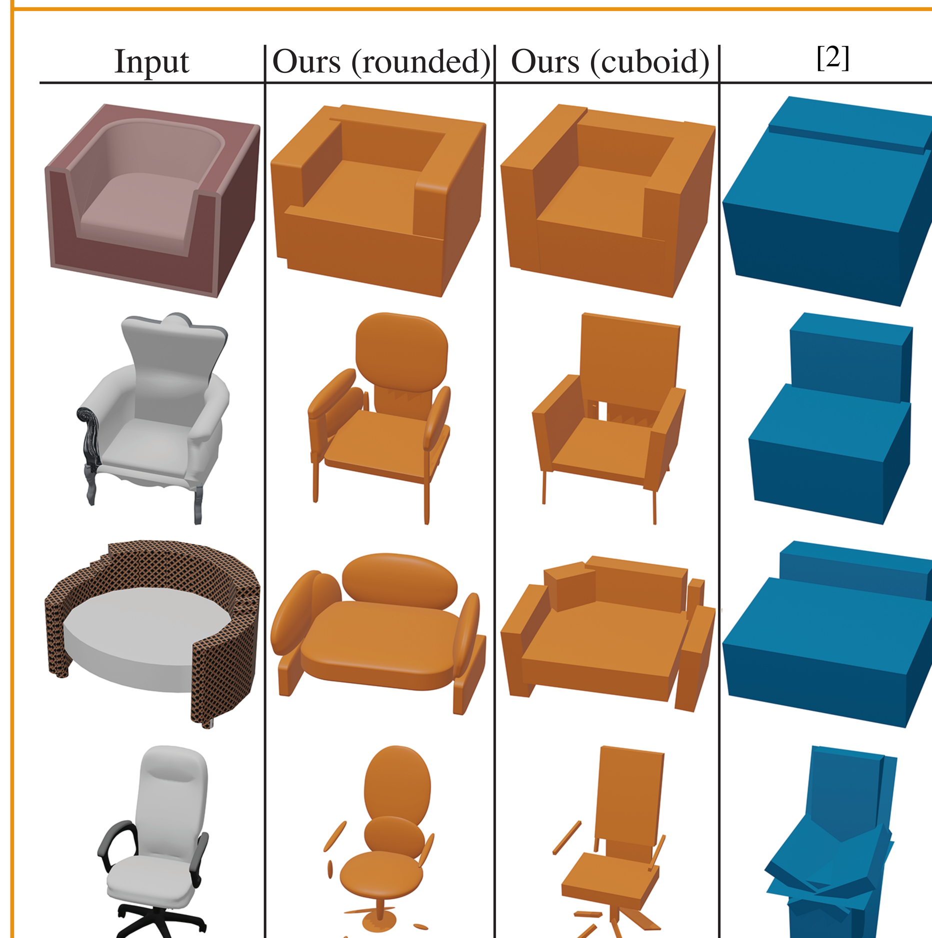

cuboid abstraction

CSG decomposition

[1] Samaneh Azadi, Matthew Fisher, Vladimir Kim, Zhaowen Wang, Eli Shechtman, and Trevor Darrell. "Multi-Content GAN for Few-Shot Font Style Transfer." CVPR 2018.

[2] Shubham Tulsiani, Hao Su, Leonidas J. Guibas, Alexei A. Efros, and Jitendra Malik. "Learning Shape Abstractions by Assembling Volumetric Primitives." CVPR 2017.

Project Webpage

https://people.csail.mit.edu/smirnov/deep-parametric-shapes/