Filing Holes in Complex Surfaces Using Volumetric Diffusion

J. Davis, S. Marschner, M. Garr and M. Levoy

Presented by Owen Gray

Overview

- Challenges of hole filling
- Limitations of mesh-based approaches
- Constructing a distance function
  - VRIP
- Diffusion
- Results

Complications of Hole Filling

- Not always possible to fill holes from surface information alone
  - Holes may result in multiple connected components (red)
  - Accurate filling requires knowledge of surface topology

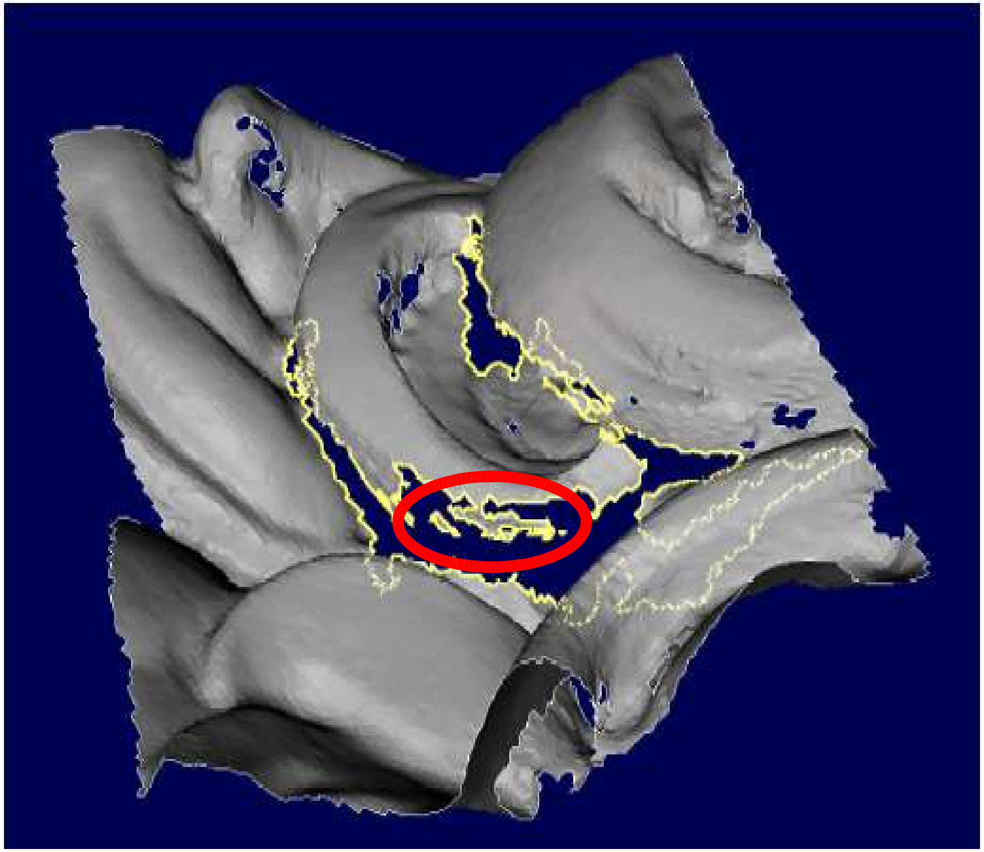

Complications of Hole Filling

- Scanner noise results in ragged hole boundaries
  - Can't correct for outliers at hole boundaries
  - Hole boundaries tend to be in regions where scanner noise is largest (oblique angle, poor surface characteristics)

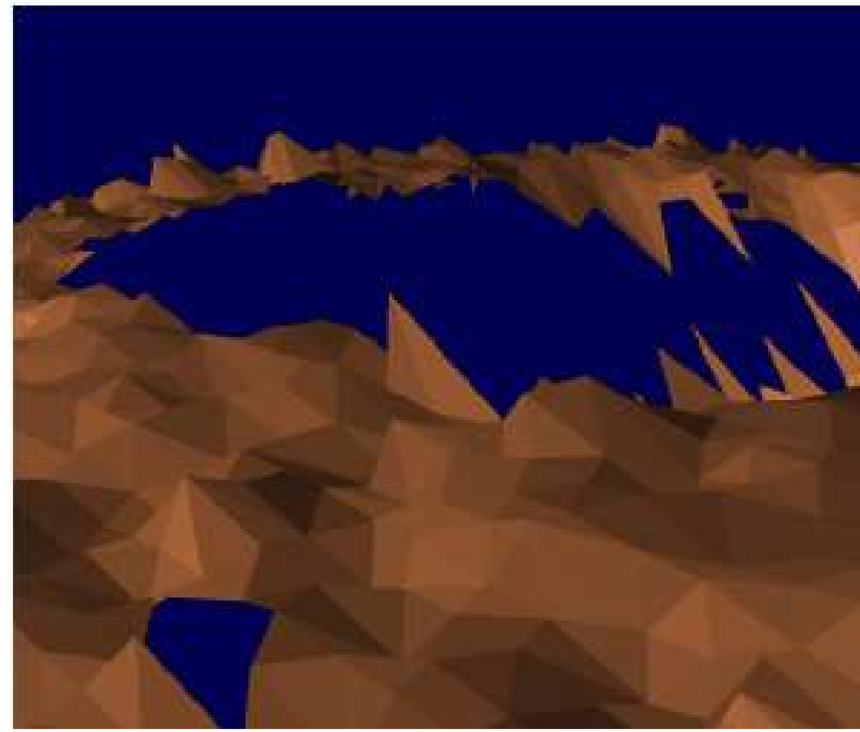

Limitations of Mesh-Based Approaches

- In the absence of data (holes) there is no single solution to the problem of mesh construction
  - Various heuristics exist for mesh construction
  - These tend to fail in the cases identified above
- Filling holes with simple disk topology may not be possible
  - Disconnected components
  - Complex hole geometry
- Construction of an arbitrary mesh can result in non-manifold surface
  - Defeats the purpose of filling holes (CAD etc.)

Signed Distance Function

- Voxels are assigned a boolean value
  - 1 (full) or 0 (empty)

    $v_i(\boldsymbol{x}) \in 0,1$
- Each voxel is assigned a weight representing the confidence in that voxel's membership in the interior or empty space set
  - Weights are determined using the VRIP algorithm, but other weighting schemes could be substituted

    $w_s(\boldsymbol{x}) \in [0,1]$

Signed Distance Function

- Algorithm is essentially a post-processing step, applicable to any input that allows the definition of a signed distance function **d**
  - **d** is defined only in a narrow band of arbitrary thickness at the known surface boundary
  - **d** ranges from -1 (inside) to 1 (outside) at the extremes of this band

$$d_i(\boldsymbol{x}) \in [-1, 1] \qquad d_s(\boldsymbol{x}) \in [-1, 1]$$

$$d_i = w_s d_s + (1 - w_s)\hat{d}_i$$

VRIP

- VRIP (Volumetric Range Image Processing Package) was used to generate the initial volumes for diffusion
  - Freely available implementation of the algorithm described by Curless and Levoy (A Volumetric Method for Building Complex Models From range Images, '96)
  - Applies line of sight constraints based on scanner geometry to generate volume and weight values for each voxel
  - Diffusion would also be applicable to volumes generated by other algorithms

Initialization and Iteration

- All voxels within range of the surface are assigned an initial distance from the surface
  - v_0=1 only if it has non zero weight and is in the initial source volume

$$(d_{0,} v_0) = (d_s, [w_s > 0])$$

  - The distance function is updated iteratively based on the values of it's neighbors

$$d_i = w_s d_s + (1 - w_s)\hat{d}_i$$

  - Results of subsequent iterations are combined using a variation of alpha blending

Convolution

- Convolution kernel **h** is a 7-voxel "Plus" operator
  - Result is similar to a Gaussian blur of the original boolean volume
    - Applying a true Gaussian approximation as the kernel may improve preservation of sharp edges
  - Effect is a low-pass filter, removing fine detail

$$(\hat{d}_i, v_i) = \boldsymbol{h} * (d_{i-1}, v_{i-1})$$

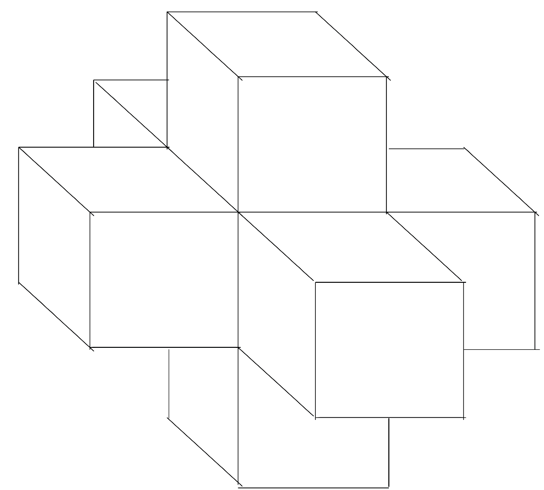

Volume Blending

- The original distance value of each voxel $\mathbf{d}_s$ is composed with the resulting distance value after each iteration
  - Composition is accomplished with the alpha-blending algorithm presented by Tom Porter and Tom Duff (SIGRAPH '84) where the voxel weight $\mathbf{w}_s$ is used as the alpha value

$$\alpha_{out} = \alpha_{foreground} + (1 - \alpha_{foreground}) * \alpha_{background}$$

$$C'_{out} = C'_{foreground} + (1 - \alpha_{foreground}) * C'_{background}$$

$$where$$

$$C'_{foreground} = C_{foreground} * \alpha_{foreground}$$

$$C'_{background} = C_{background} * \alpha_{background}$$

$$C'_{out} = C_{out} * \alpha_{out}$$

Termination

- Algorithm terminates when the change in the distance Δd_i is less than the known noise of the scanner
  - Not clear how this threshold is determined since scanner accuracy varies
    - Local geometry (angle of incidence of scanner beam)
    - Surface properties (specularity, reflectance, subsurface scattering)

Results

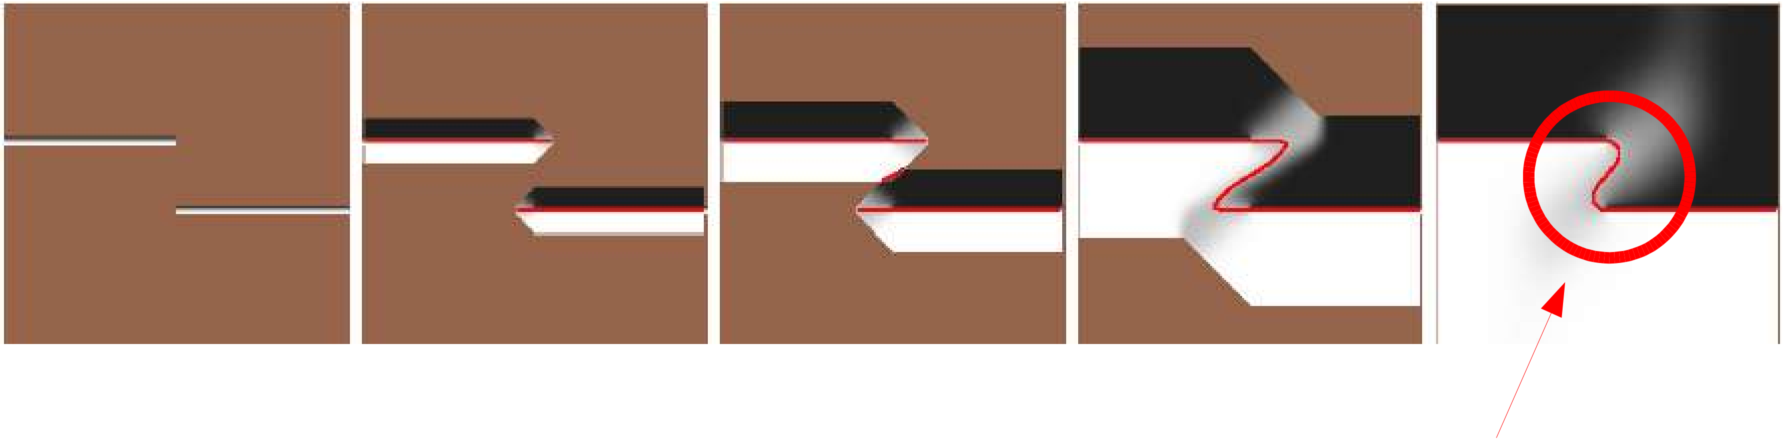

- Without applying line-of-sight constraints when generating the weighting function, blending results in marked splines at (unobserved) hole boundaries
  - Diffusion cannot reproduce sharp features

Results

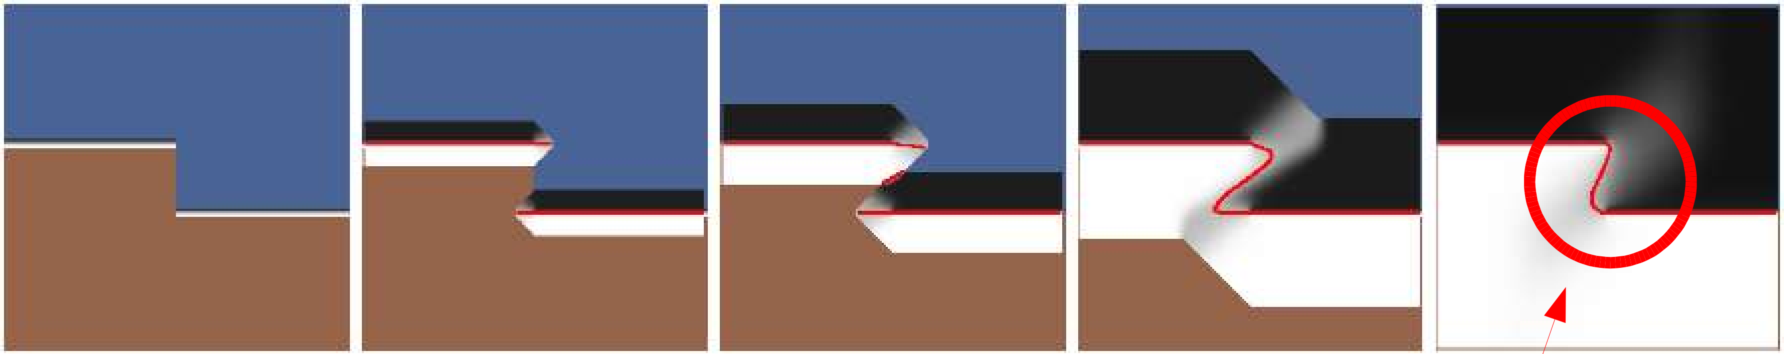

- Accounting for line-of-sight in the weights and boolean volume results in sharper features
  - Still produces chamfers and fillets at sharp boundaries
  - Geometry is plausible, but not accurate

Results

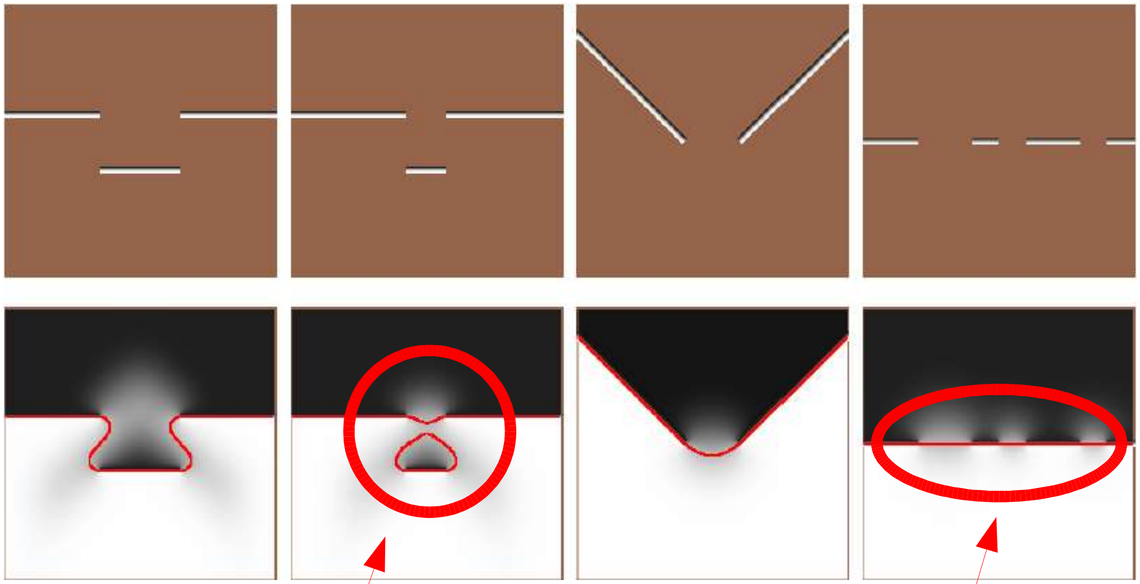

- The artifacts produced by diffusion can seal concave structures (alter topology) if line-of-sight constraints are not applied
- Performance on planar holes is good, producing s plausible fit with the original surface

Results

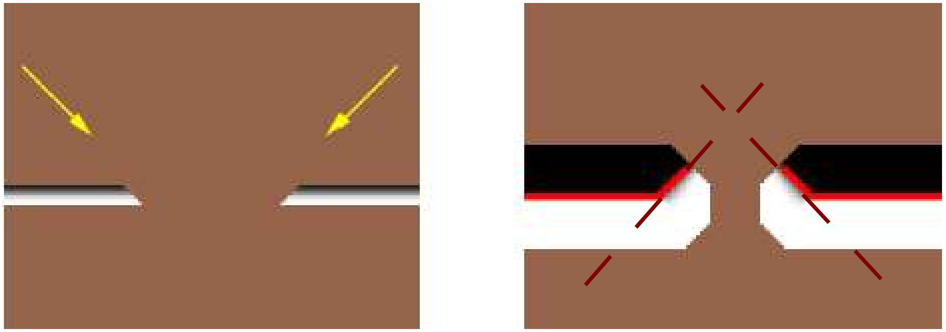

- Diffusion produces artificial contours for scans at acute angles
  - This cannot be corrected by line-of-sight constraints, but is in fact a byproduct of them

<END>

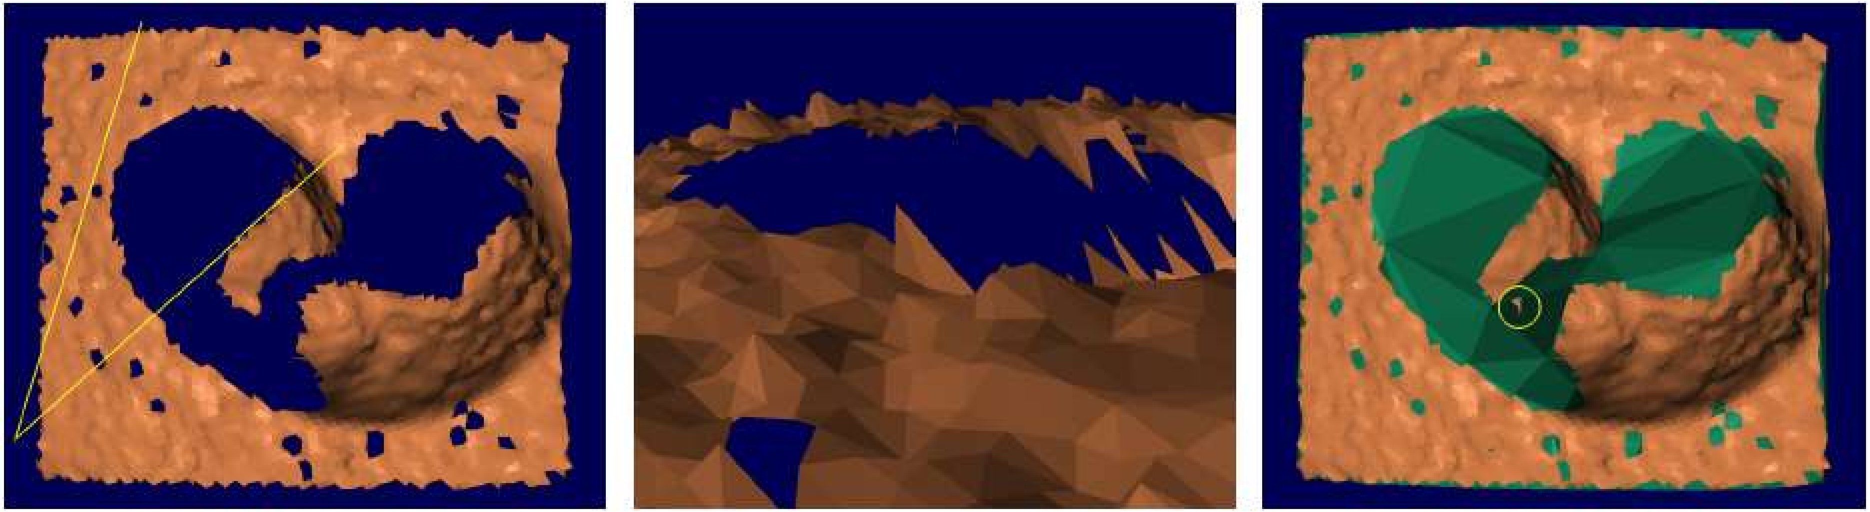

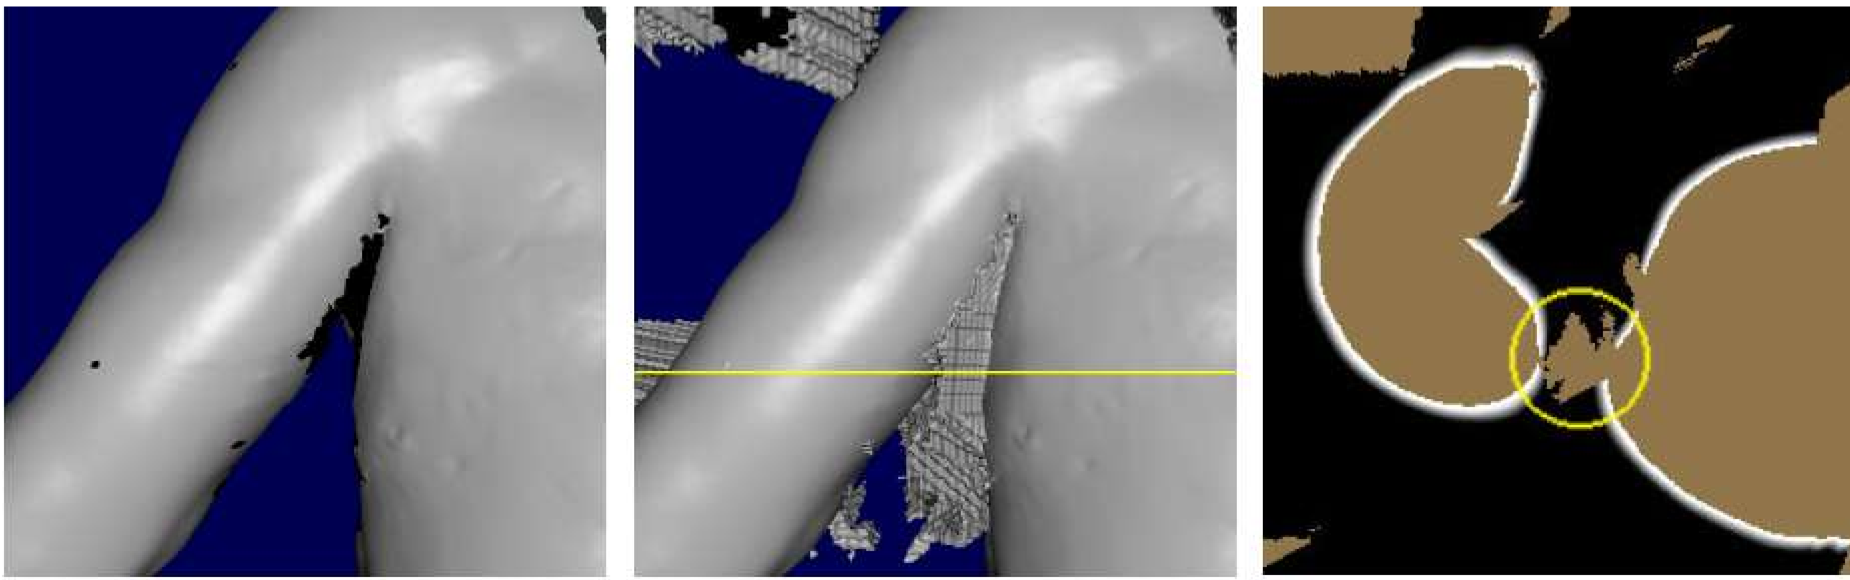

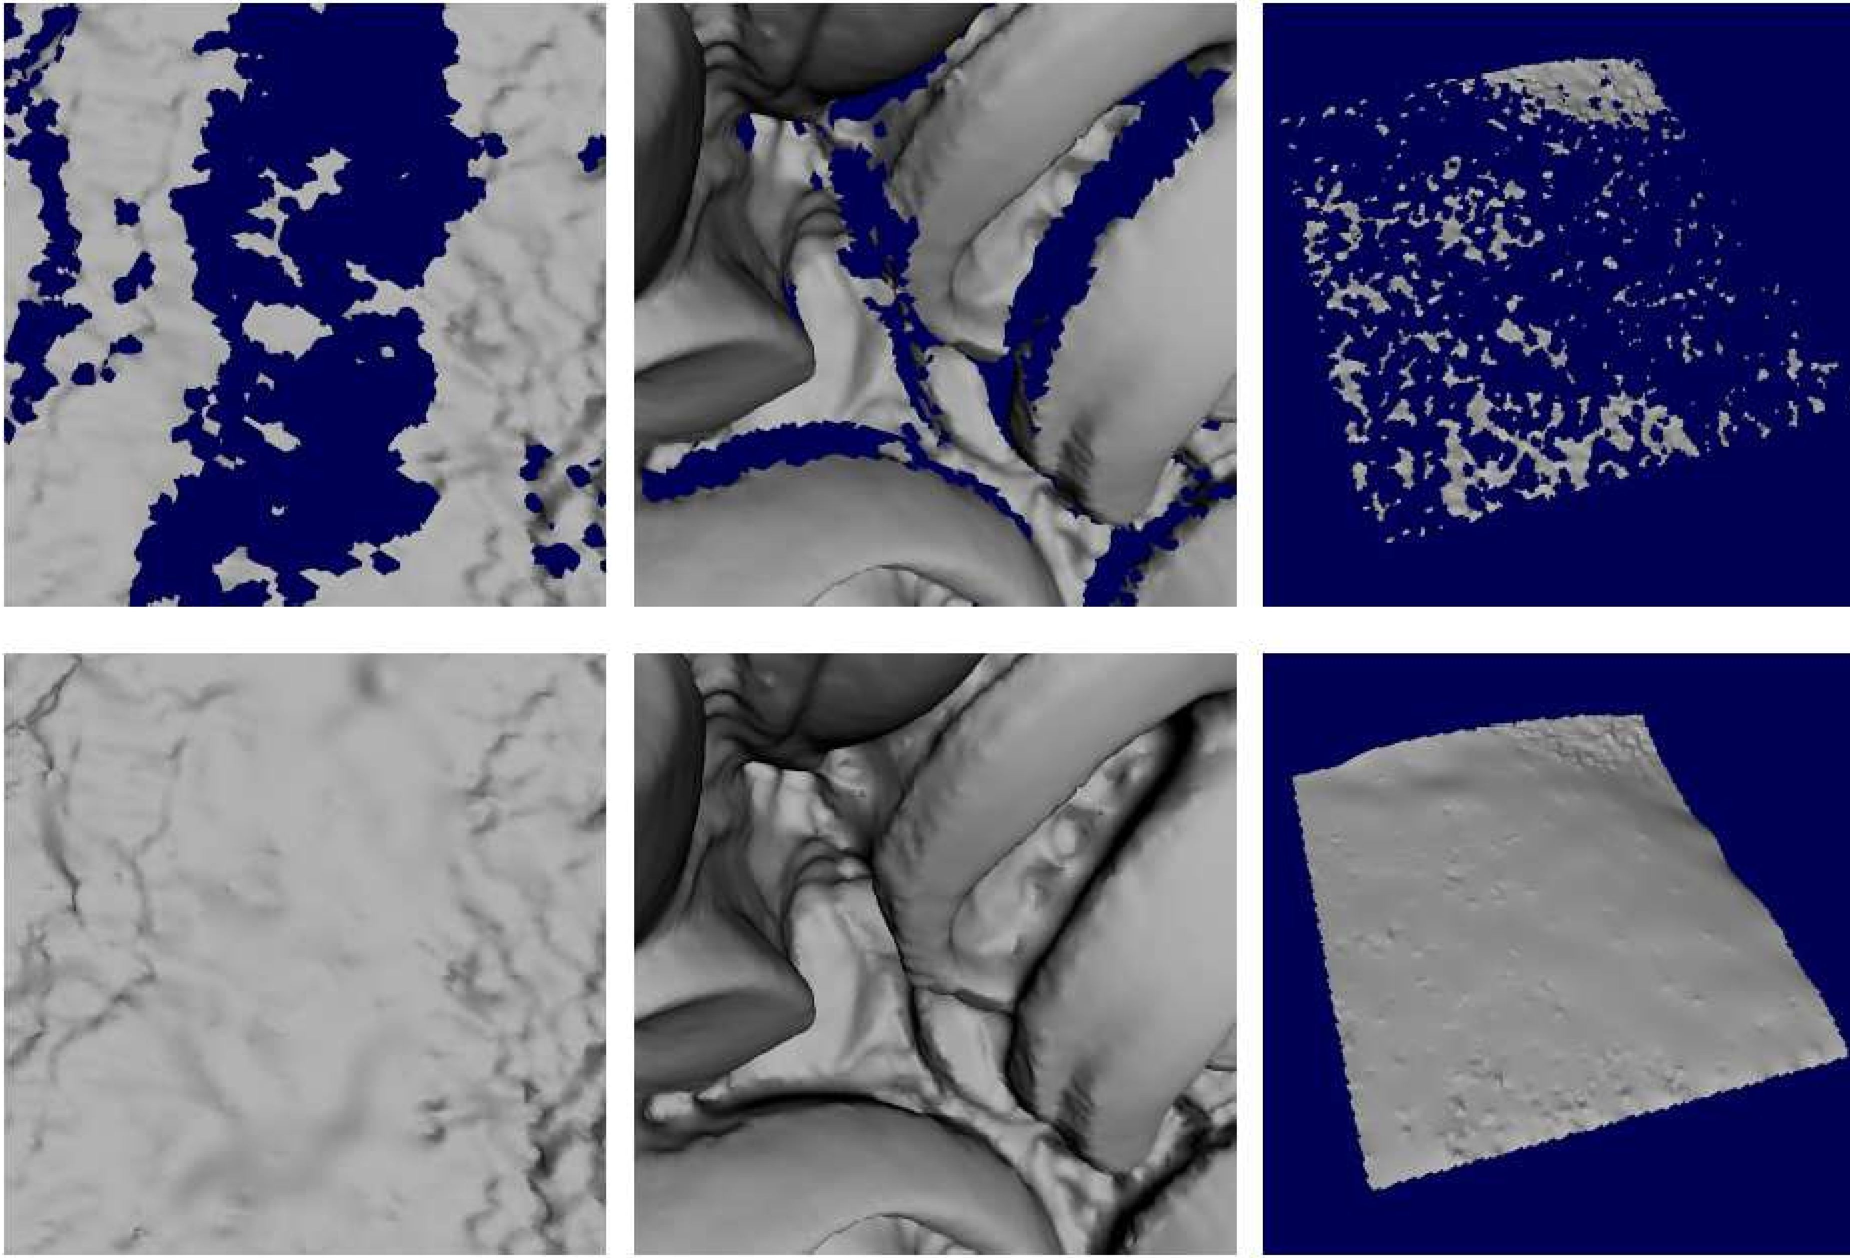

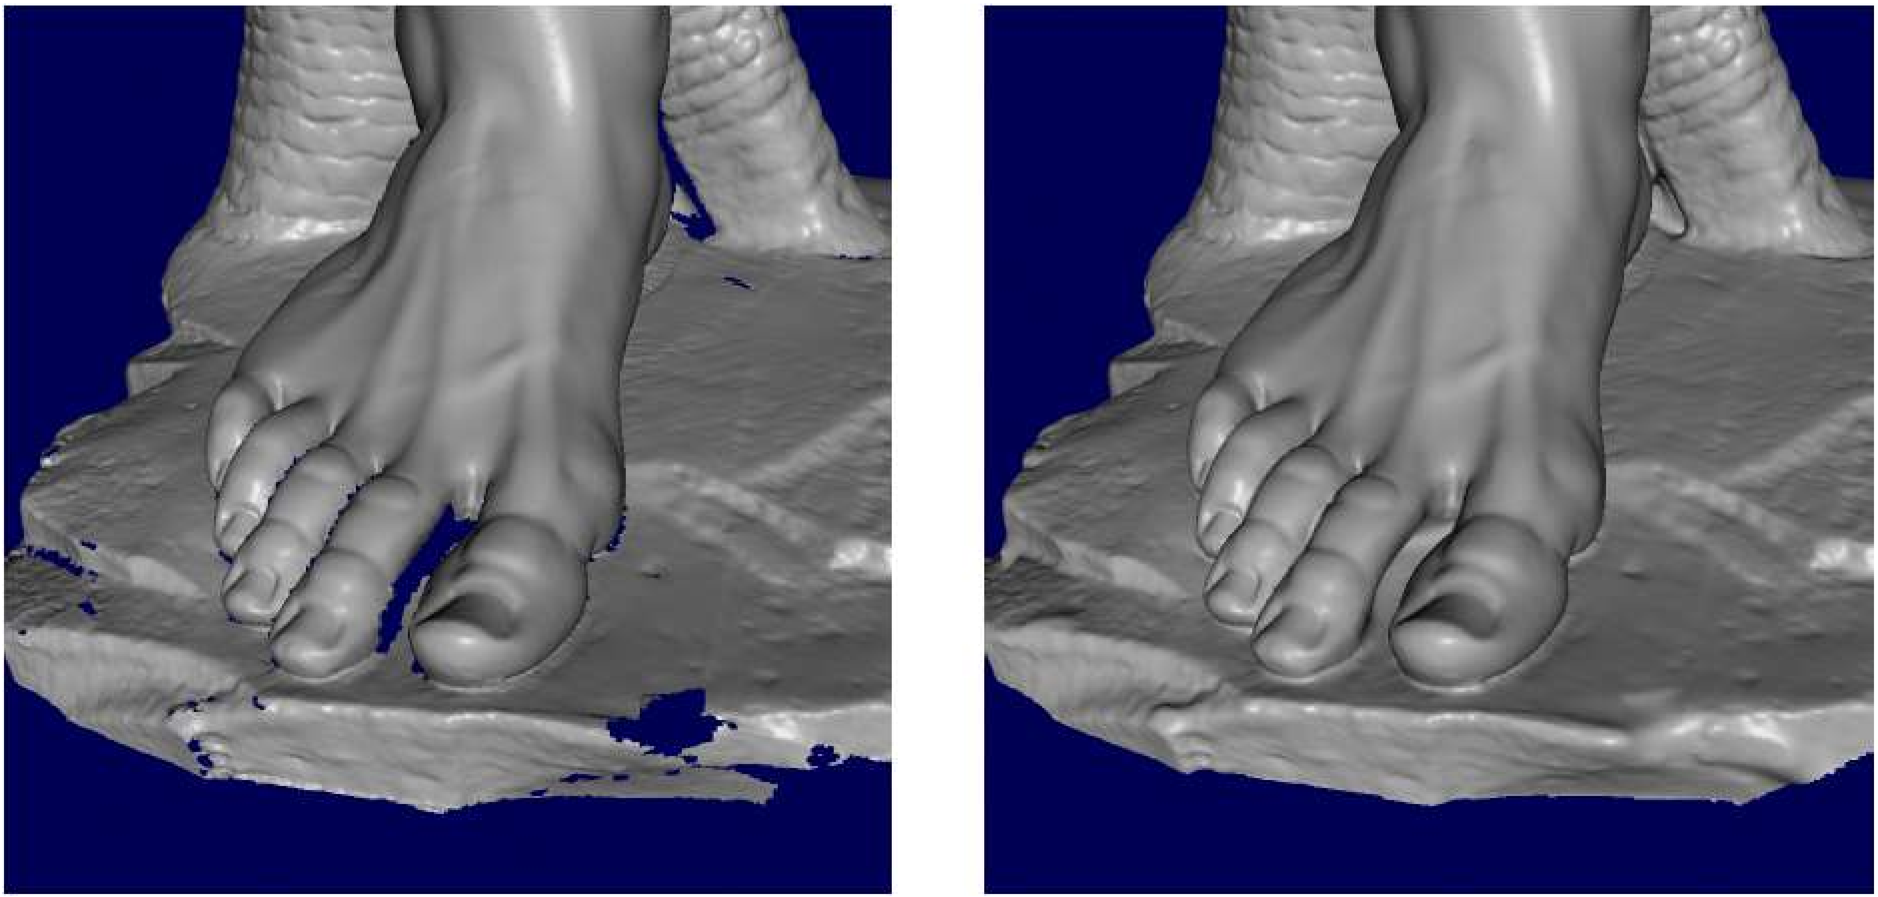

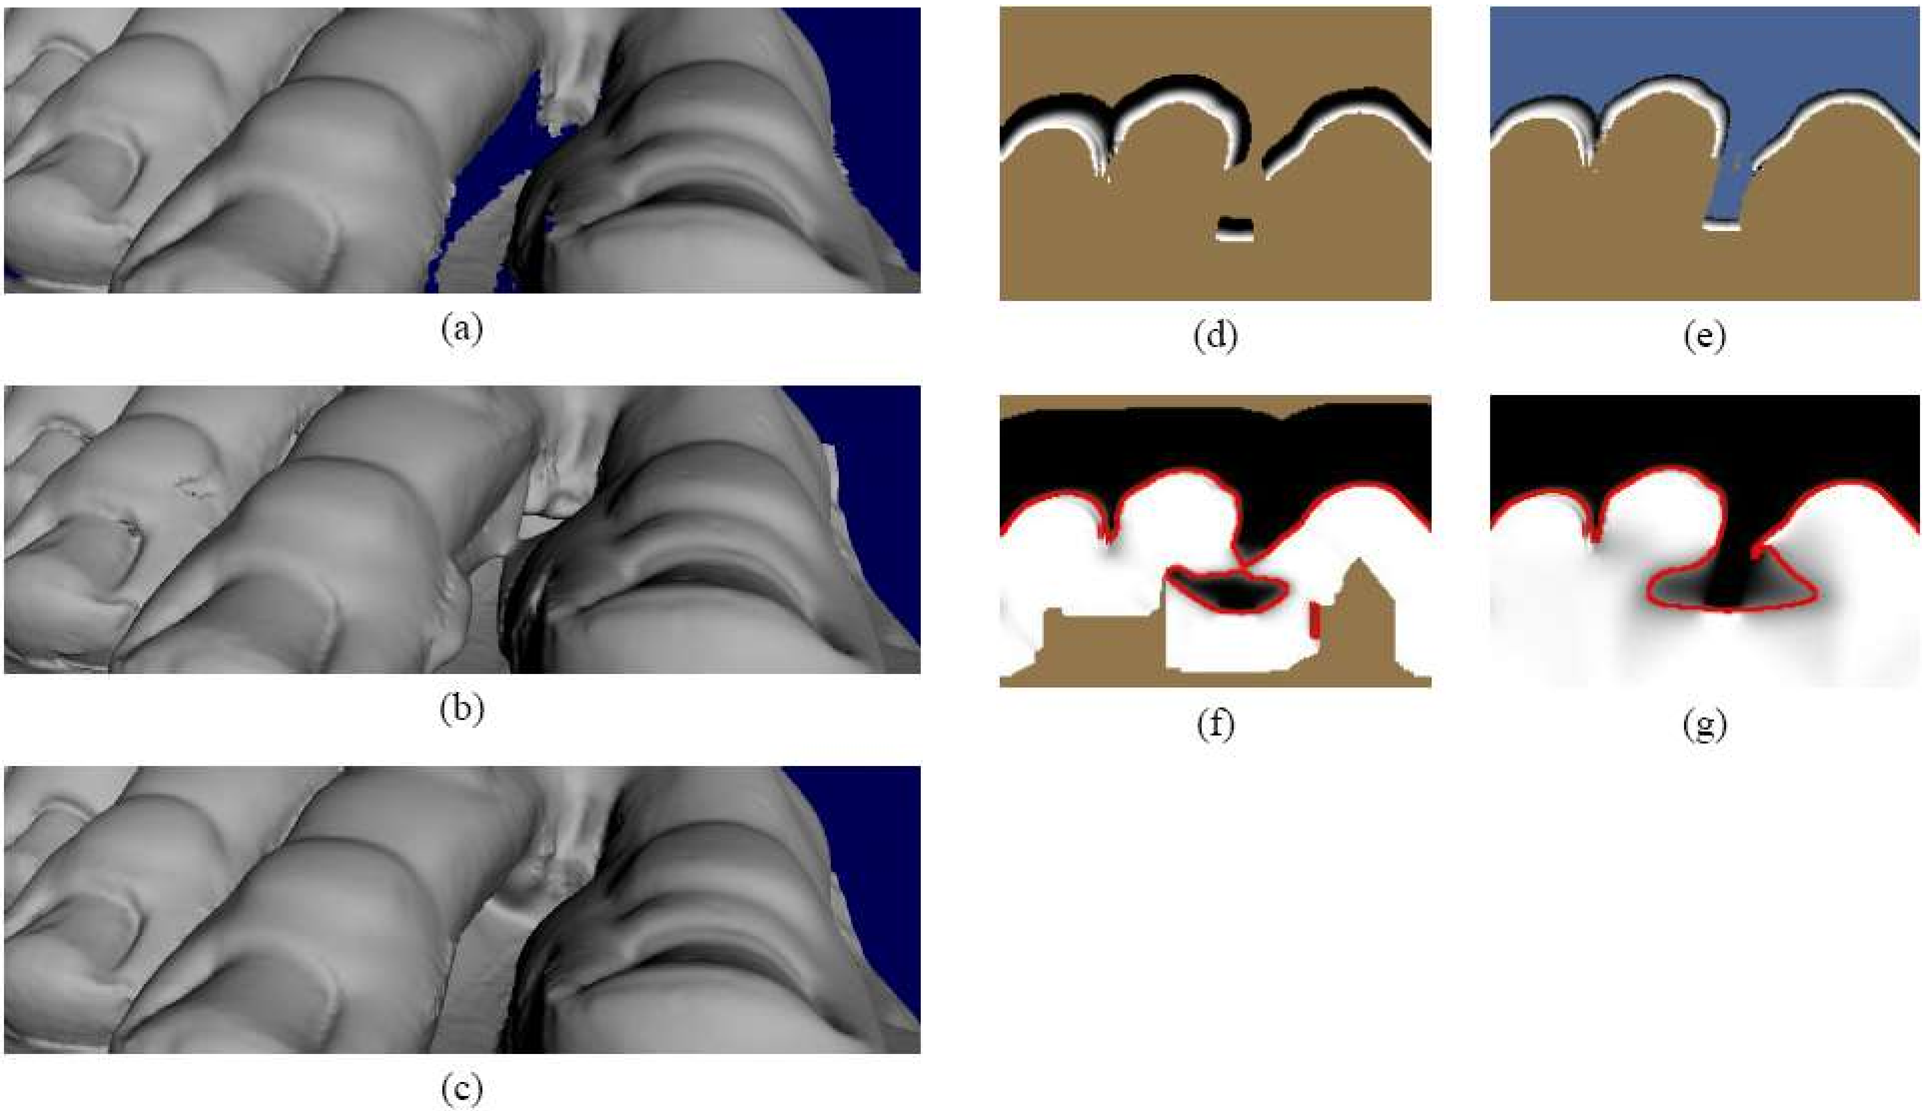
(a)
(d)
(e)
(b)
(f)
(g)
(c)

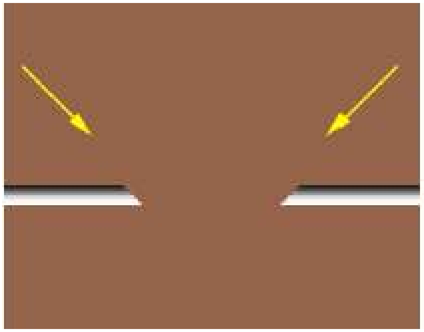
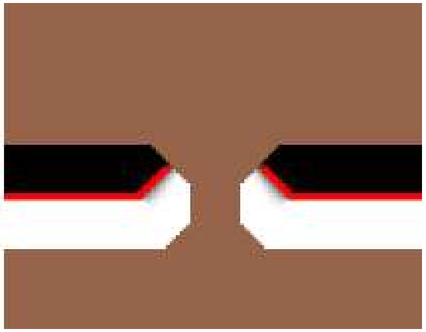

Raw

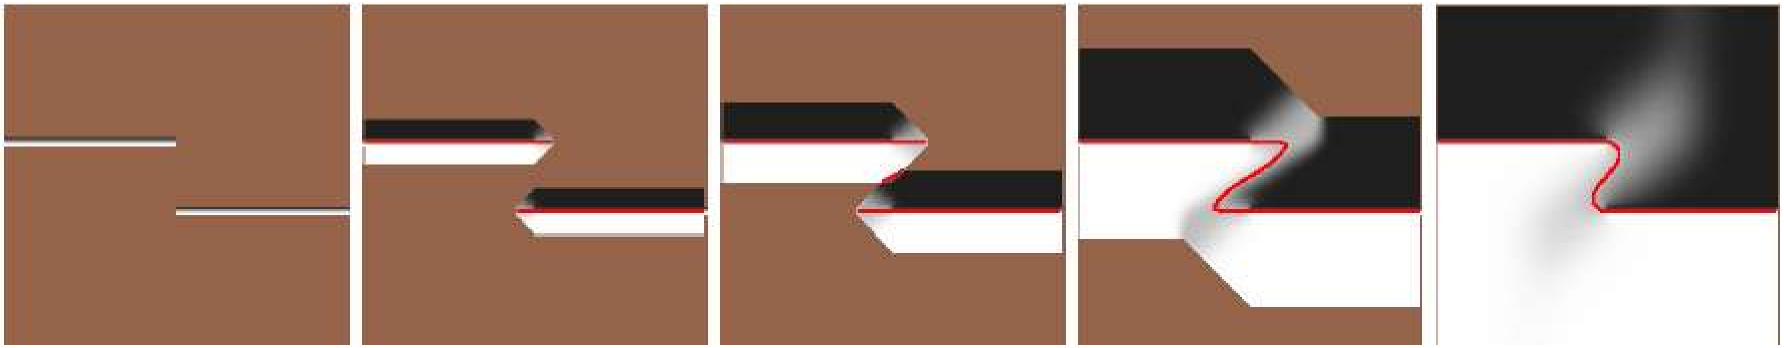

With Line of site
constraints

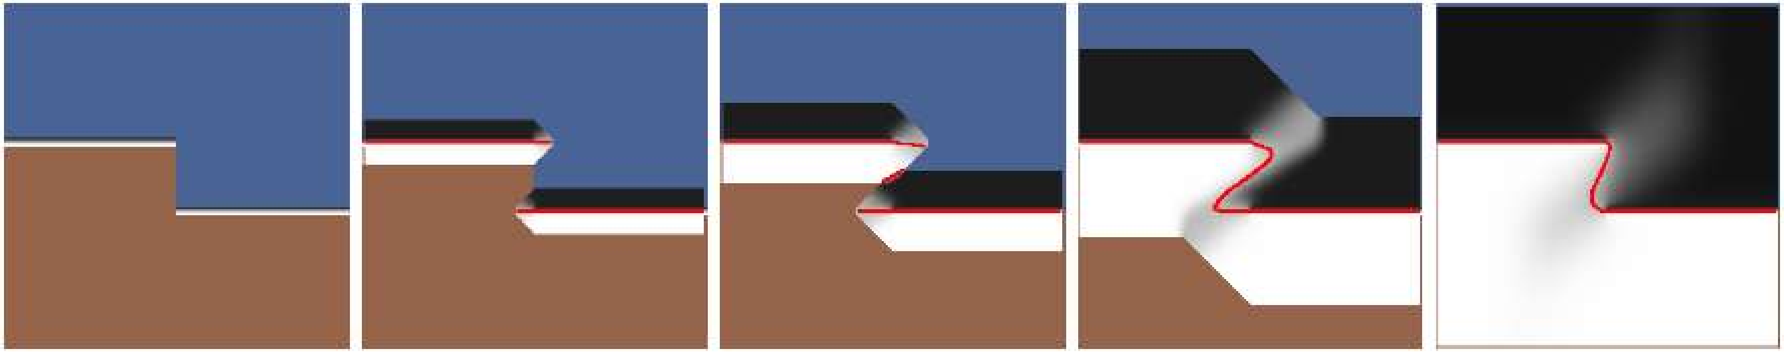

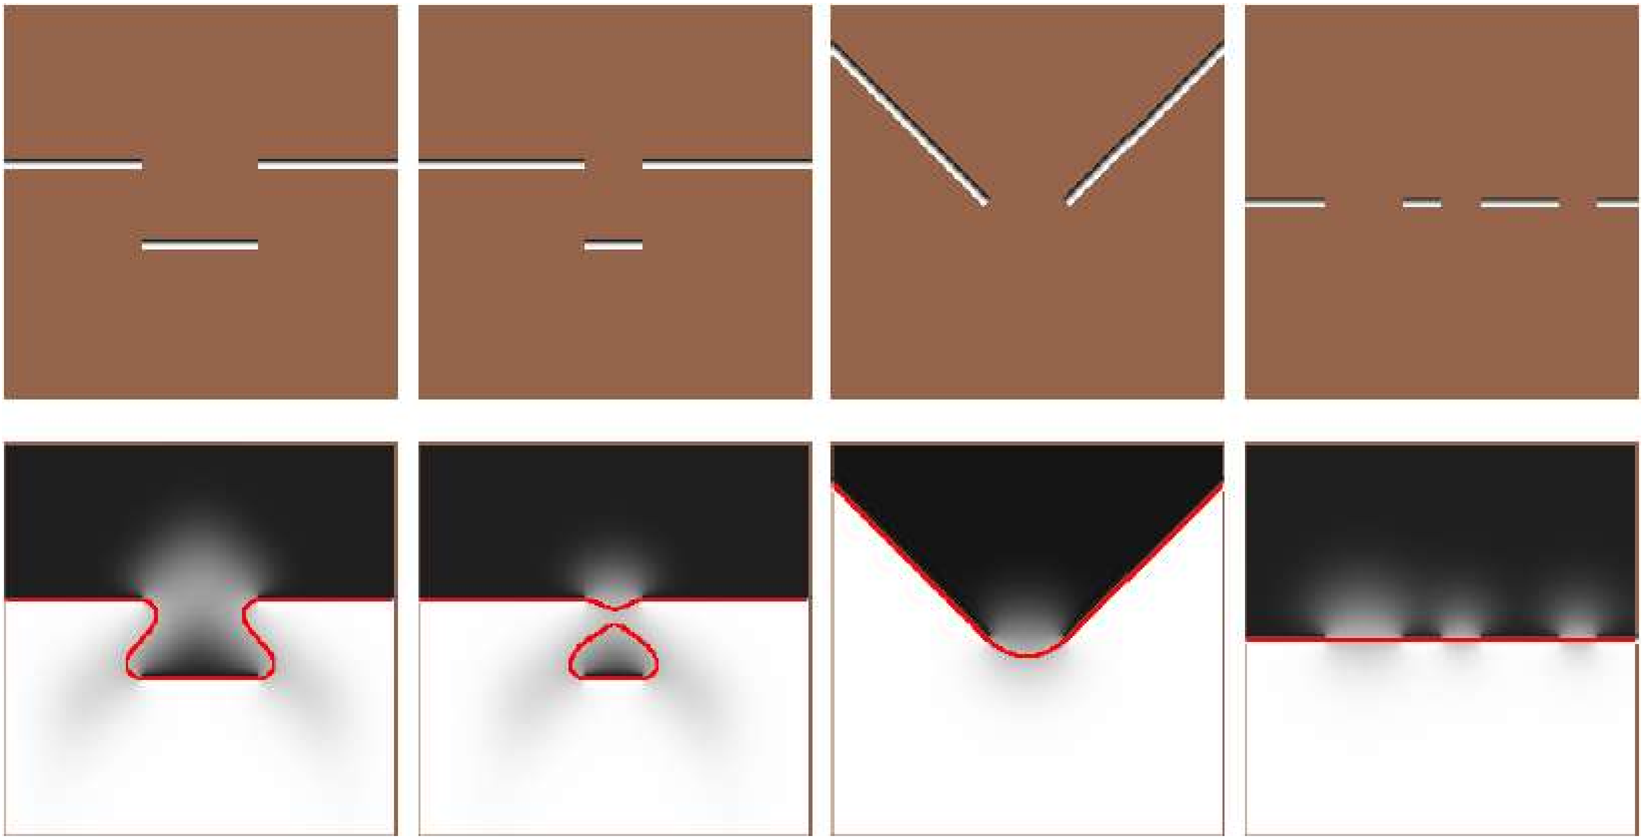

$$d_i(\boldsymbol{x}) \in [-1, 1]$$

$$d_s(\boldsymbol{x}) \in [-1, 1]$$

$$w_s(\boldsymbol{x}) \in [0, 1]$$

$$(d_{0,} v_0) = (d_s, [w_s > 0])$$

$$(\hat{d}_i, v_i) = h * (d_{i-1}, v_{i-1})$$

$$d_i = w_s d_s + (1 - w_s) \hat{d}_i$$