Mondrian Stereo

Dylan Quenneville Daniel Scharstein
Middlebury College

Middlebury

This work is supported by NSF grant IIS-0917109

Motivation

- Stereo vision has come very far, but…

- Untextured regions still present difficulty
- Occlusions make this even harder

Modern Failure Cases – Plants

Left

GT (blurry)

3DMST (Li et al. 2017)

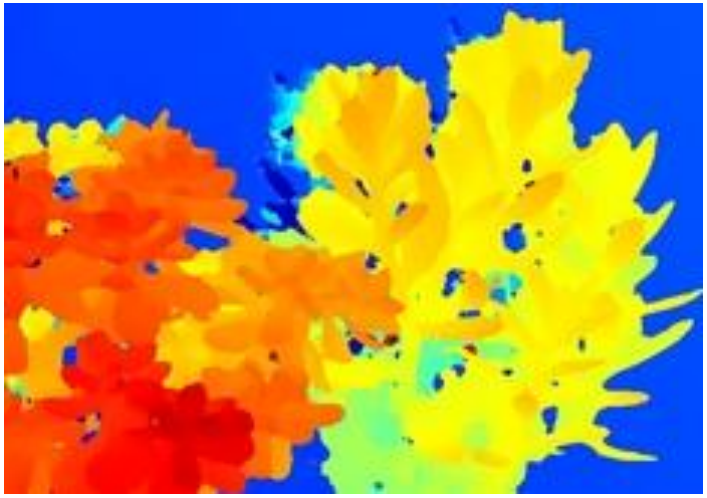

LW-CNN (Park and Lee 2016)

Modern Failure Cases – Staircase

Left

GT (blurry)

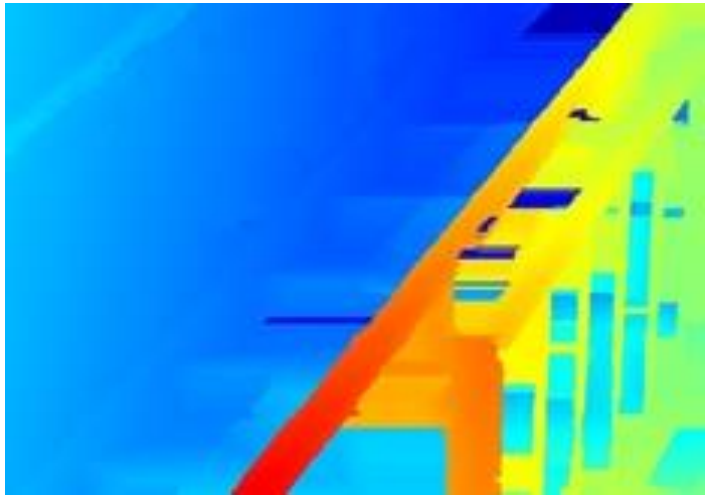

3DMST

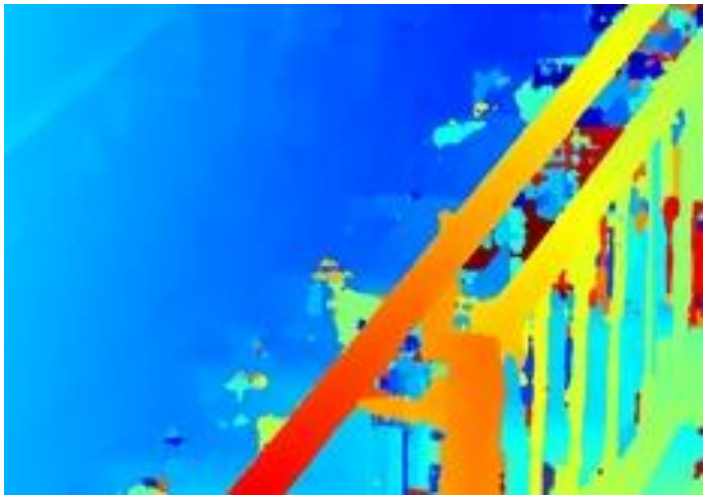

LW-CNN

Modern Failure Cases – What is Missing?

To tackle untextured areas we must consider…

- Edge disparities
- Segment color
- Occlusions

The Mondrian Problem

What if there is no texture at all?

Mondrian

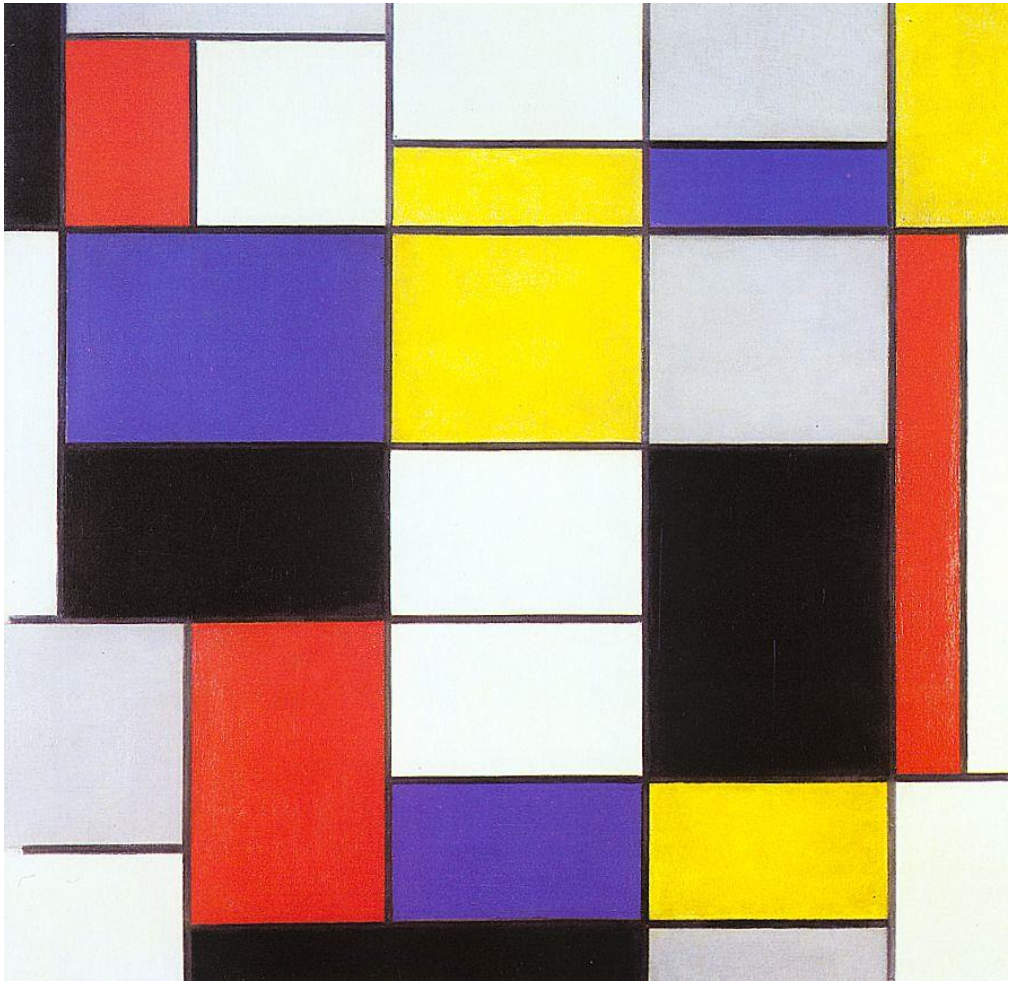

Piet Mondrian, *Composition A*, 1923

Existing Algorithms

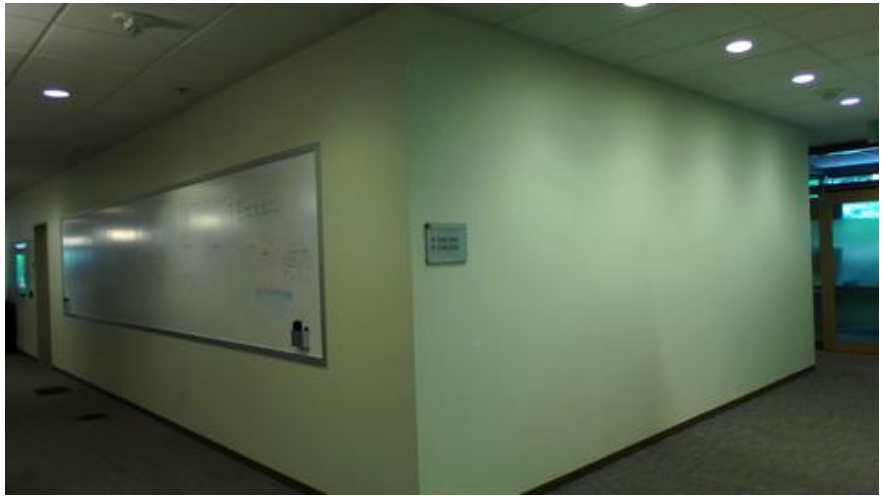

Left

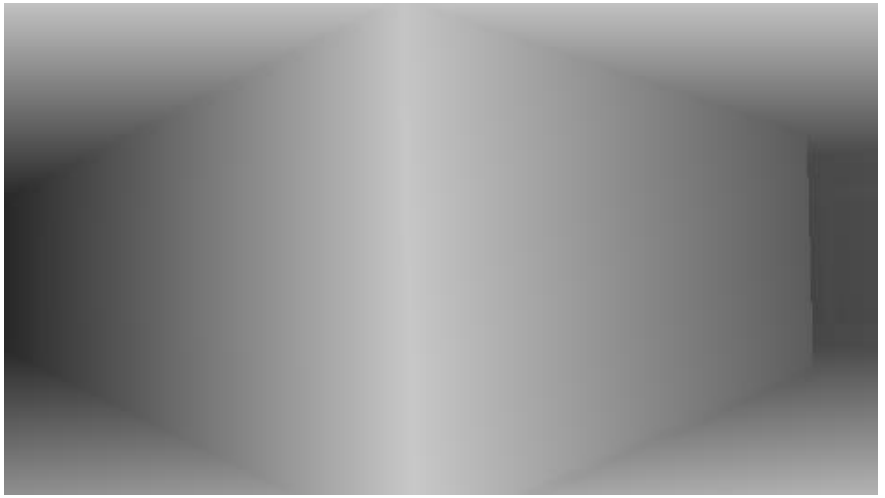

Artificial Ground Truth

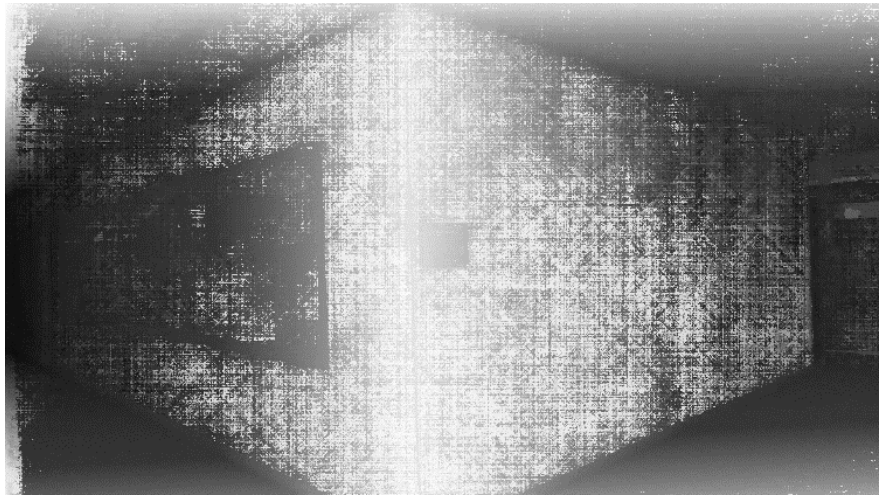

SGM (Hirschmüller 2005)

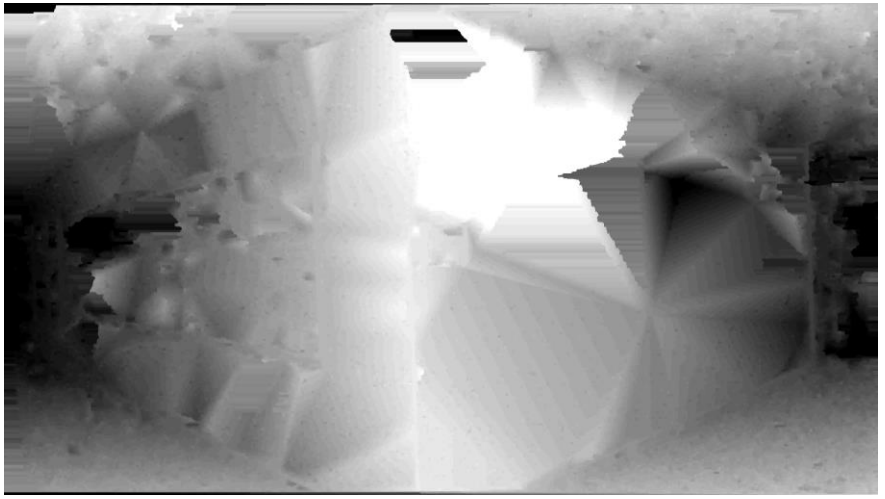

ELAS (Geiger et al. 2010)

Existing Algorithms

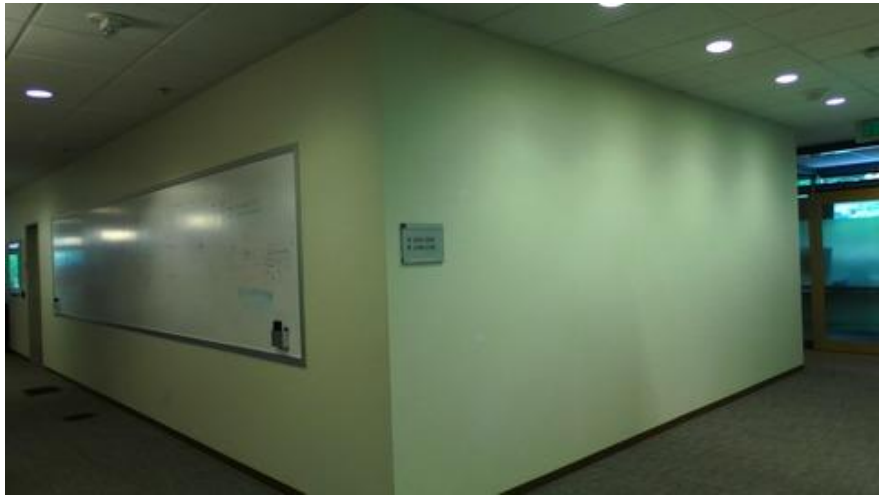

Right

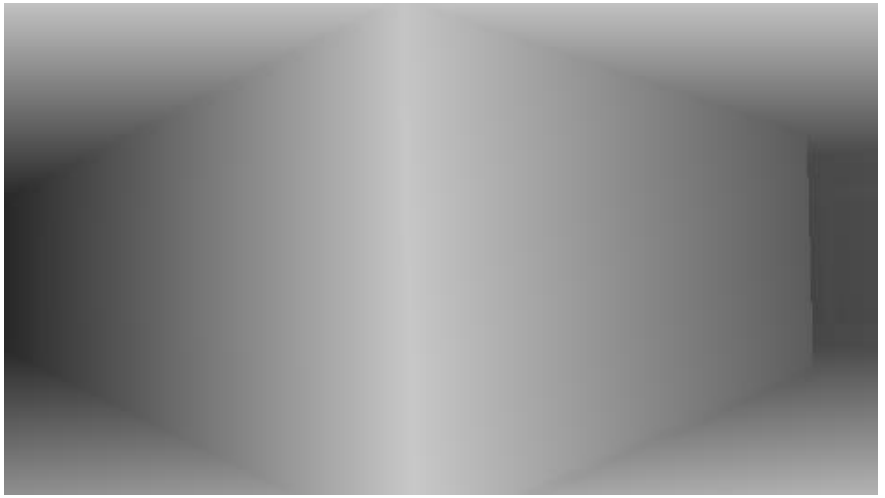

Artificial Ground Truth

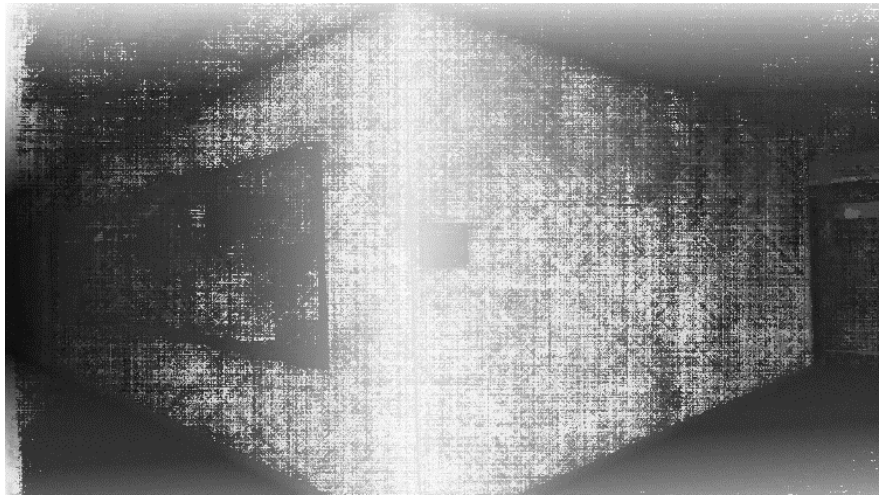

SGM (Hirschmüller 2005)

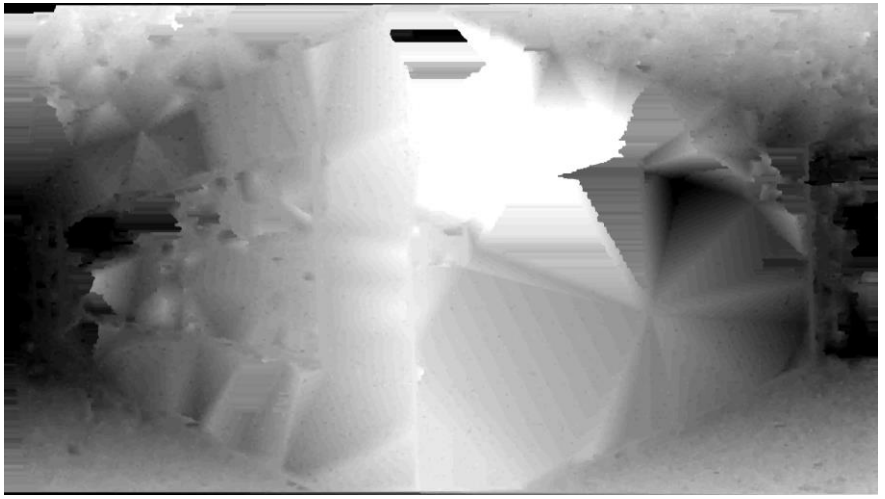

ELAS (Geiger et al. 2010)

Our Algorithm

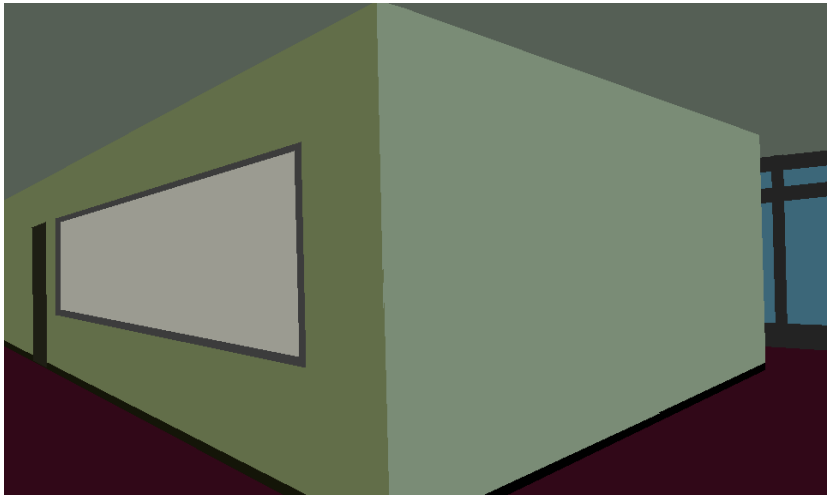

1.

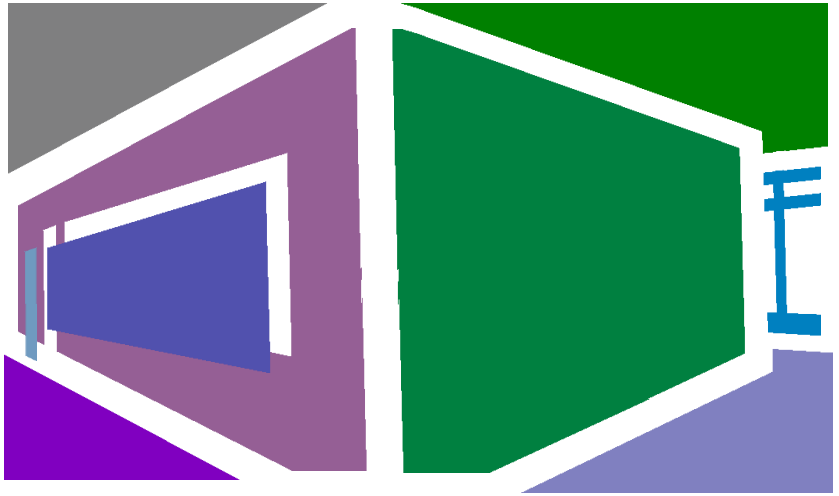

2.

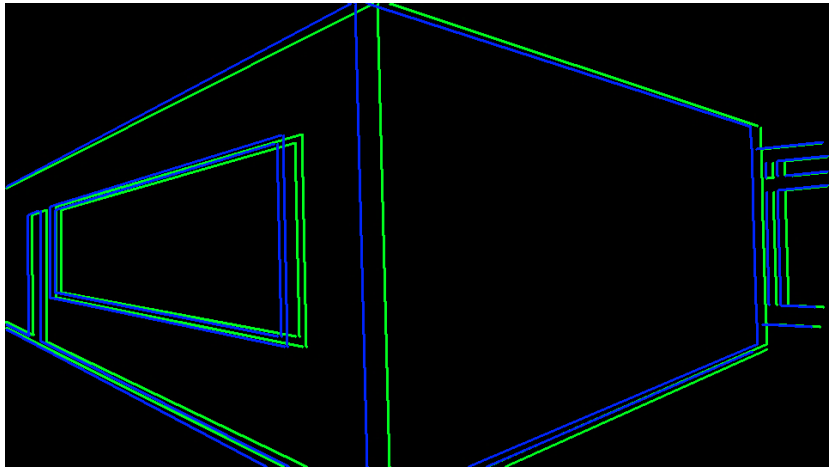

3.

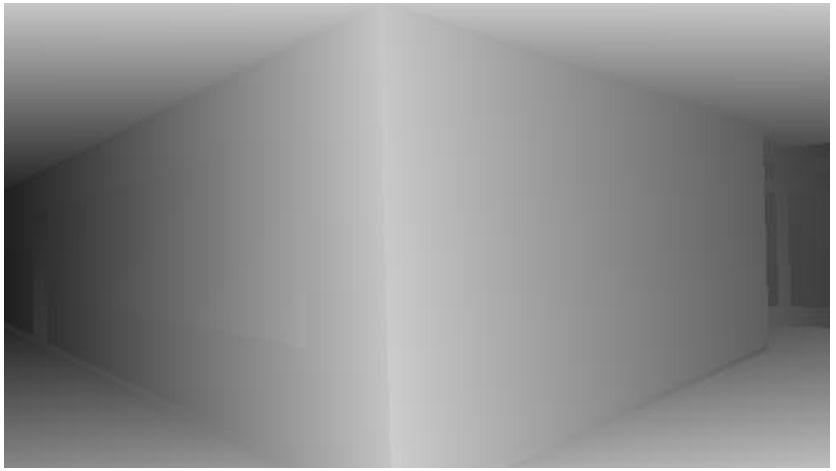

Determining Edge Ownership

- Edge is owned by both adjacent surfaces
  - Color edge

  - Surface crease

- Edge is owned by only one surface
  - Occlusion (depth) edge

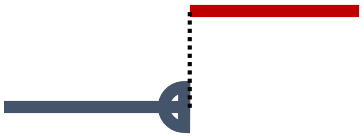

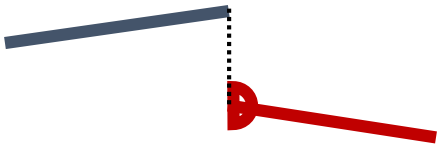

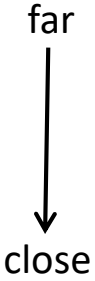

Resolving ambiguities

- **Assumption 1:**

  In the absence of other evidence, each surface is assumed to be as close as possible

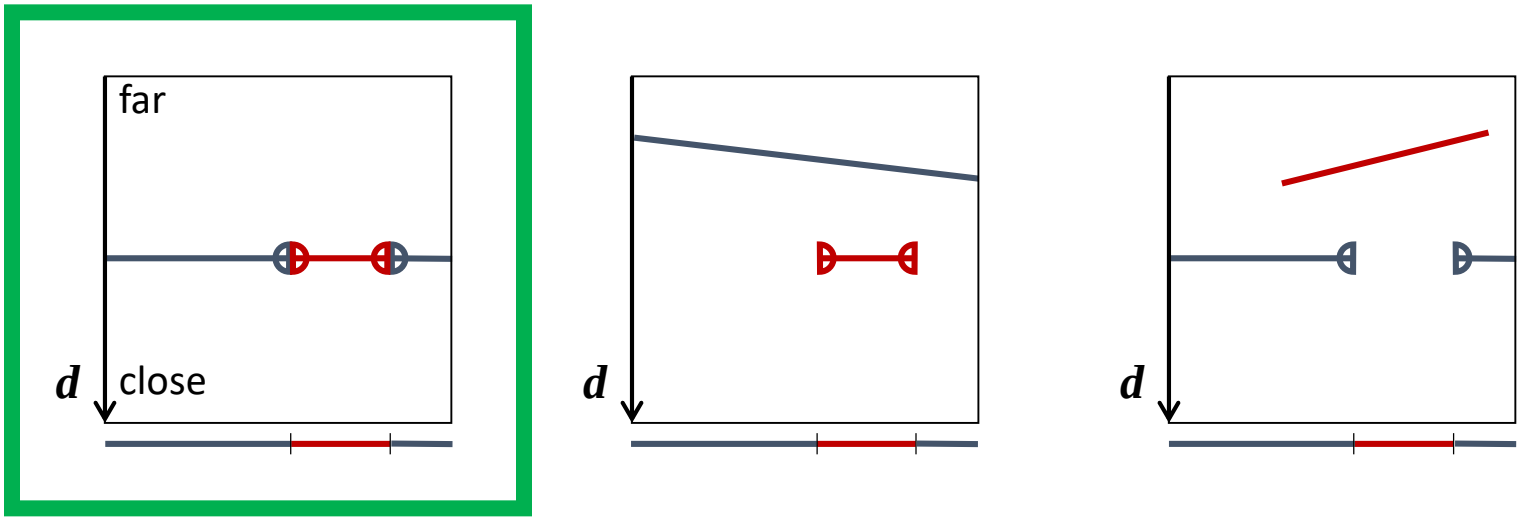

Resolving ambiguities

- **Assumption 1:**

  In the absence of other evidence, each surface is assumed to be as close as possible

- **Assumption 2:**

  Prefer assigning disconnected segments with same color to same surface if valid occluder present (edges must be owned by the other component)

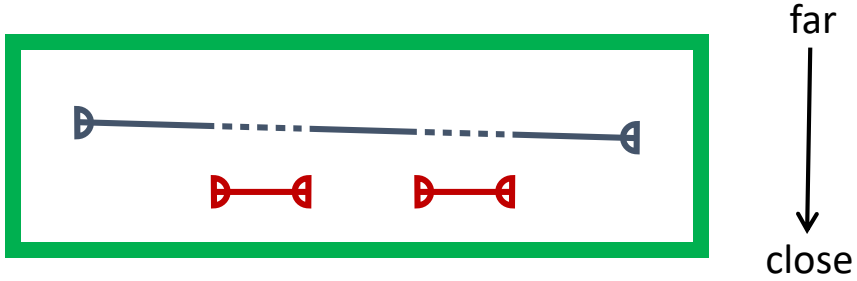

Results – Mondrian-esque

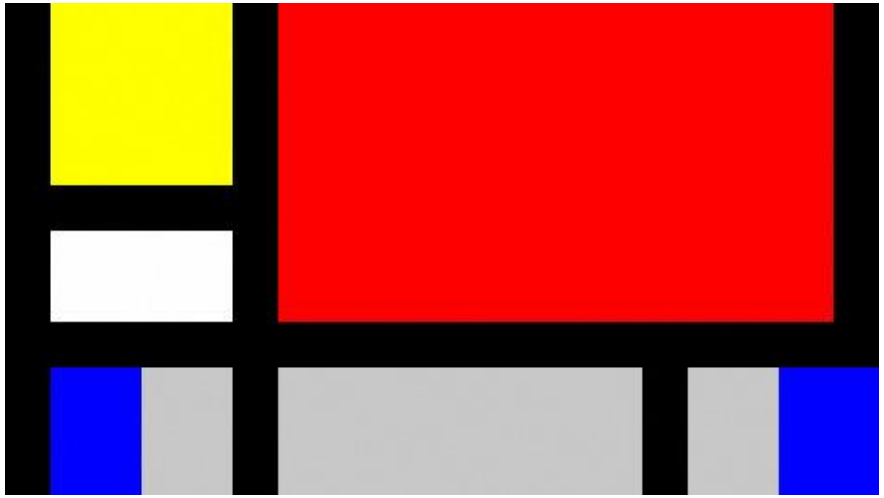

Left

GT

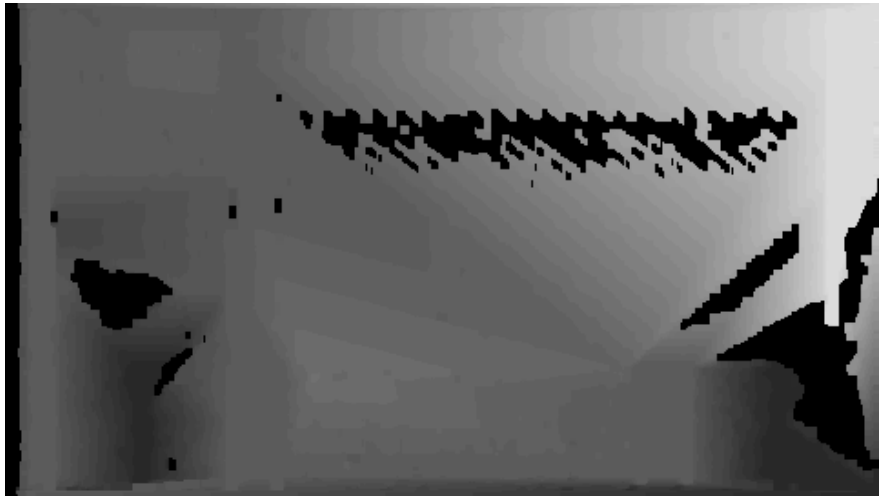

ELAS

Ours

Results – Mondrian-esque

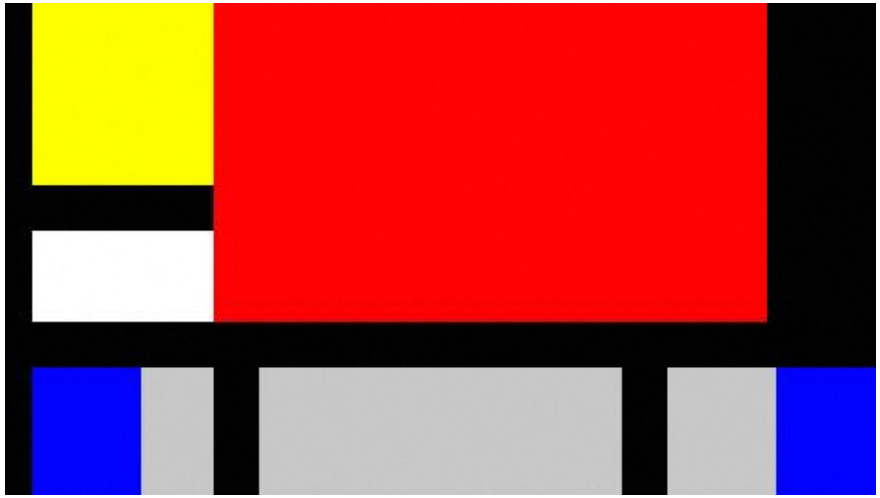

Right

GT

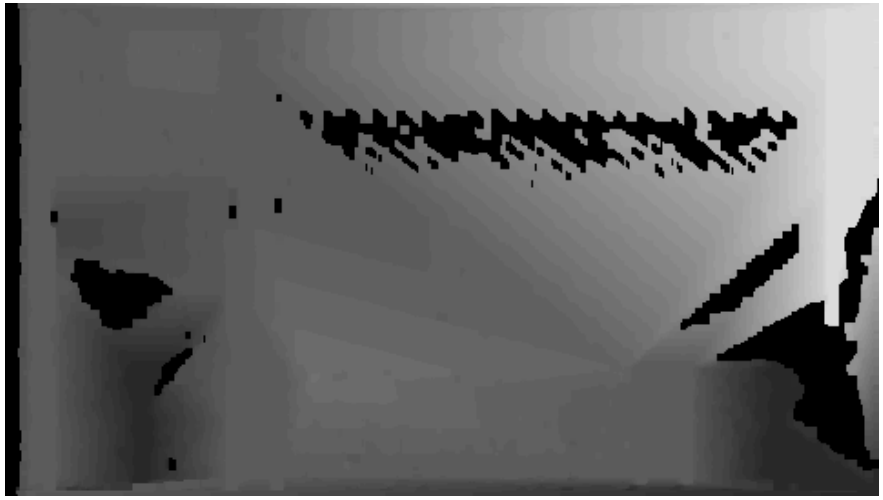

ELAS

Ours

Results – Corner

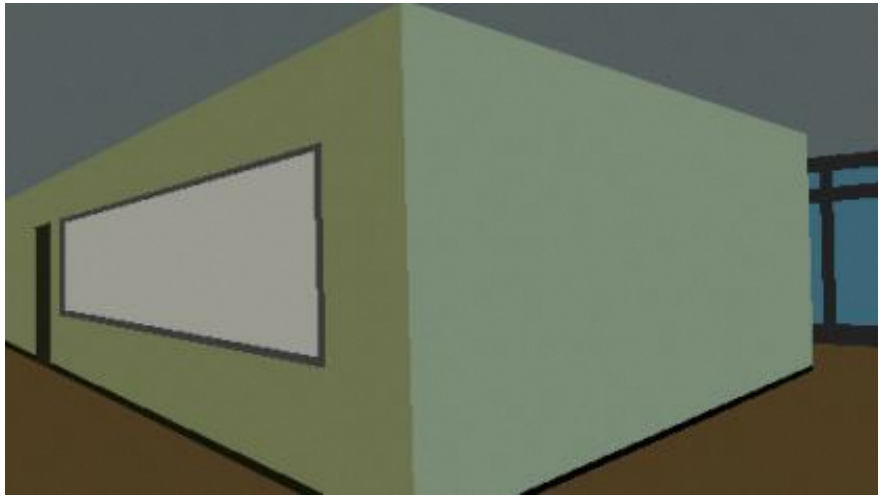

Left

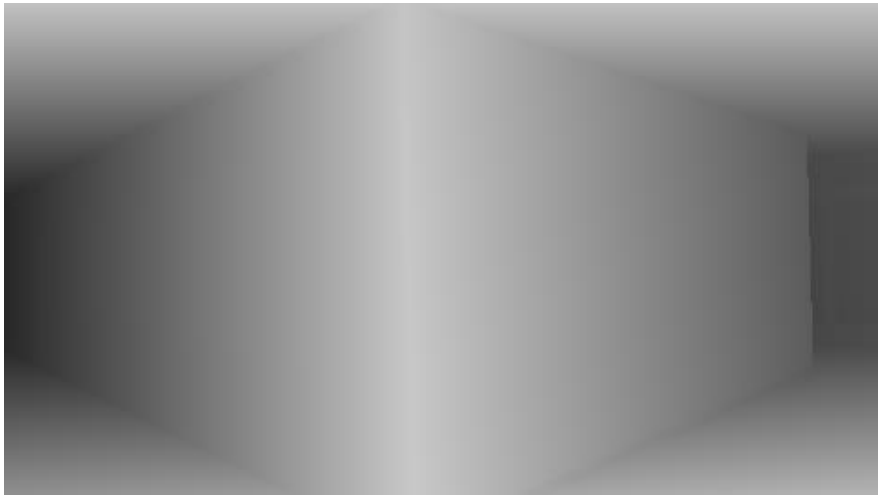

GT

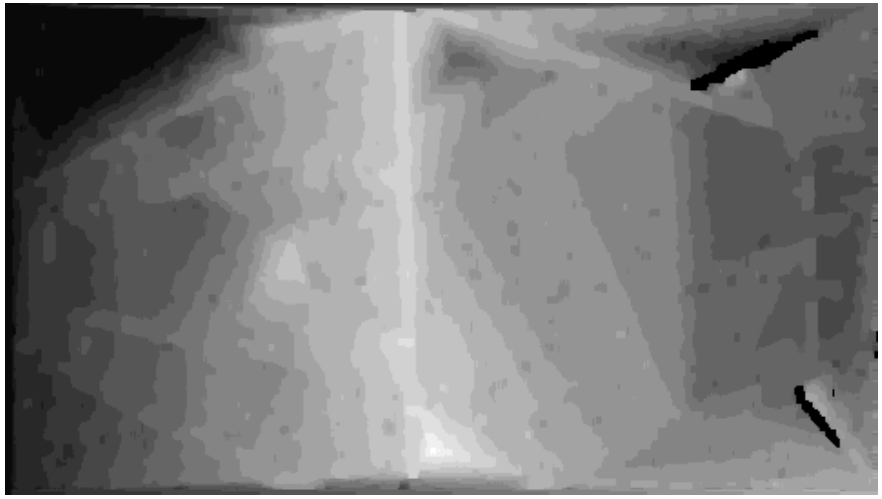

ELAS

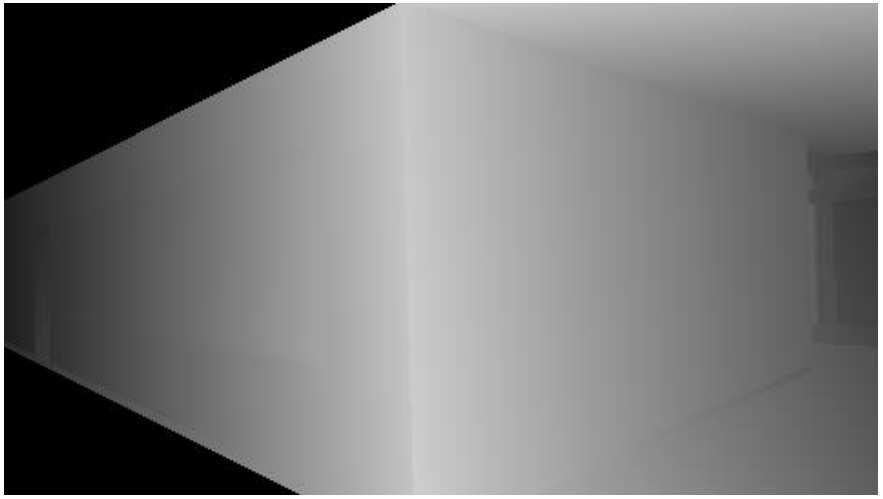

Ours (Pass 1)

Results – Corner

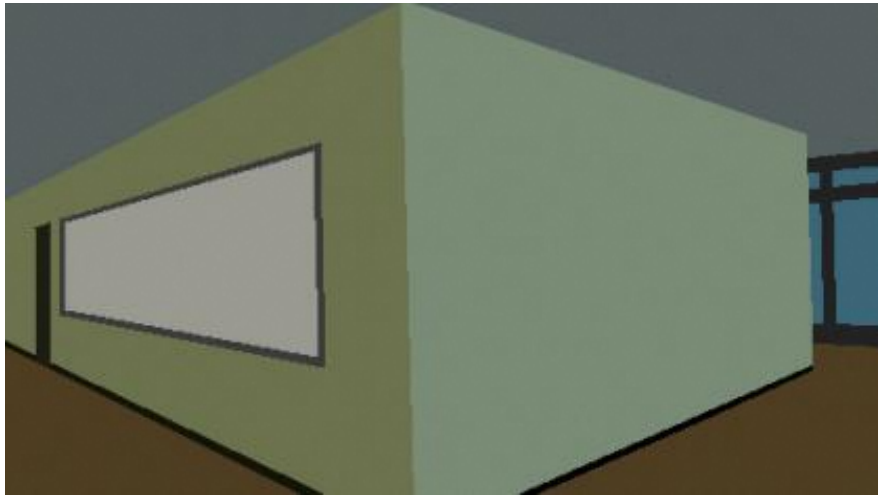

Left

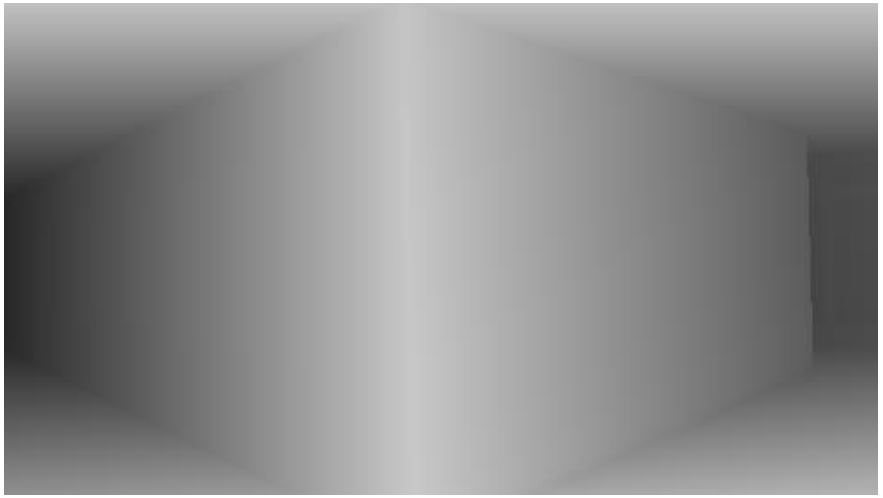

GT

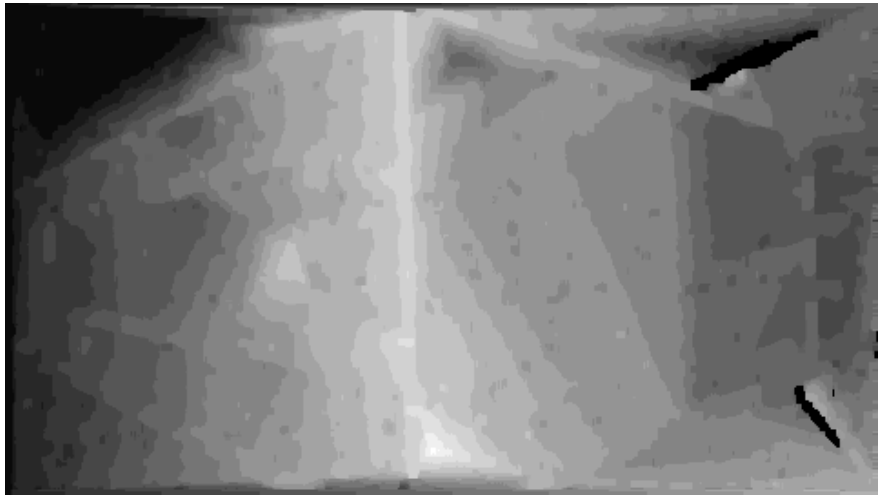

ELAS

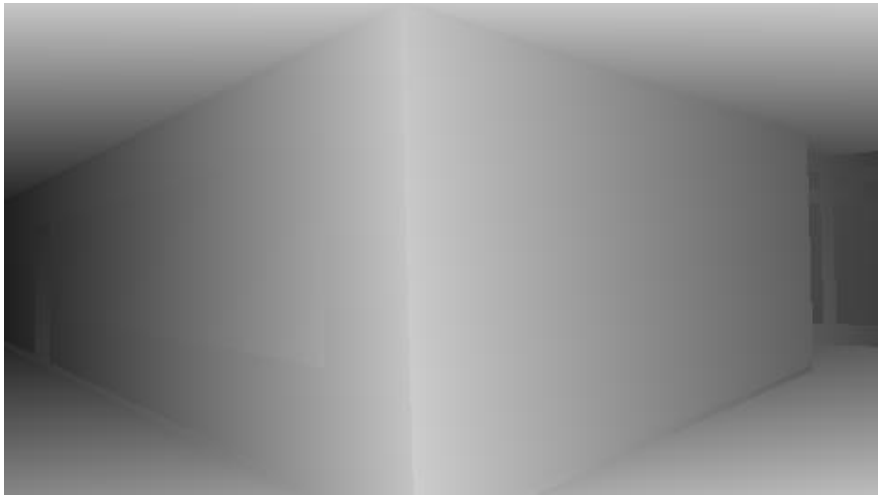

Ours (Pass 2)

Results – Lattice Occluding

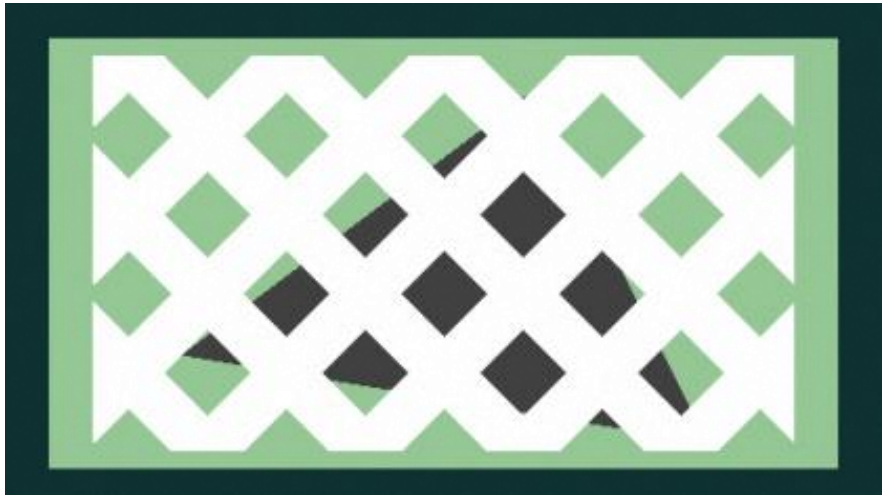

Left

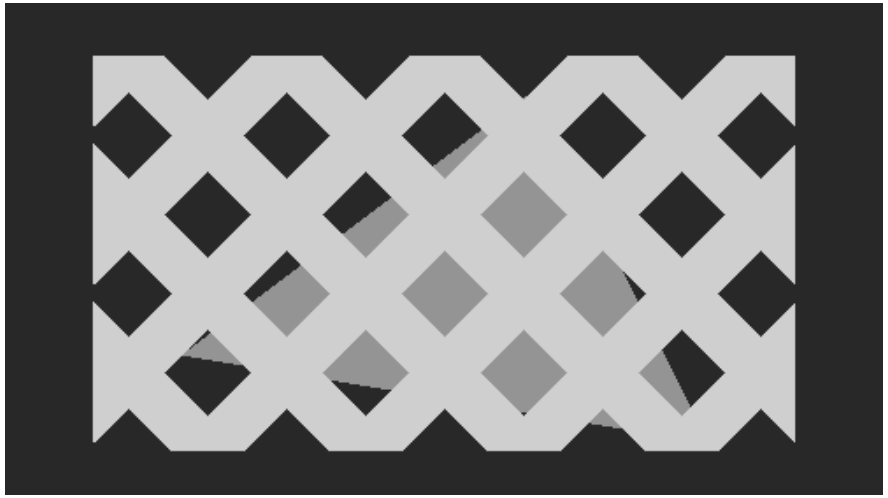

GT

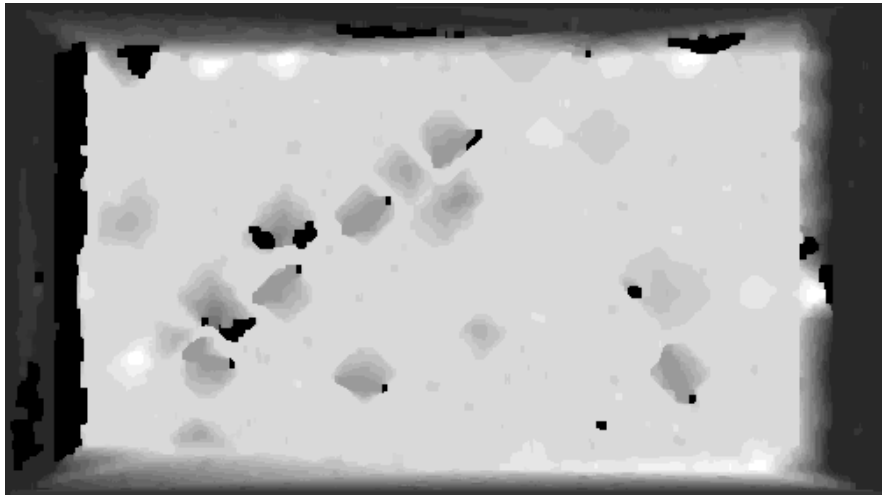

ELAS

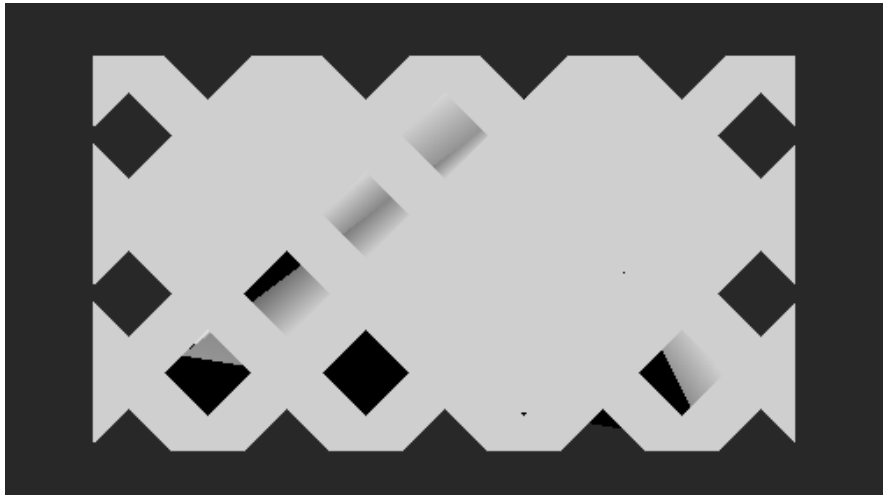

Ours (Pass 1)

Results – Lattice Occluding

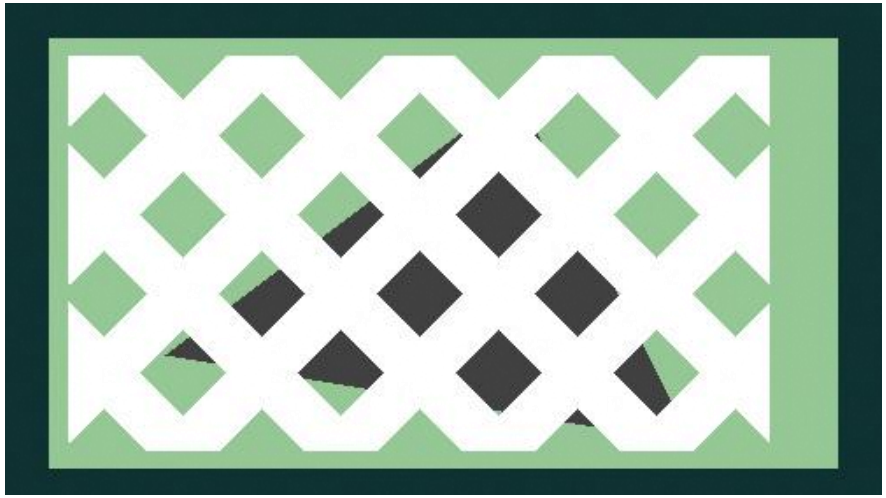

Right

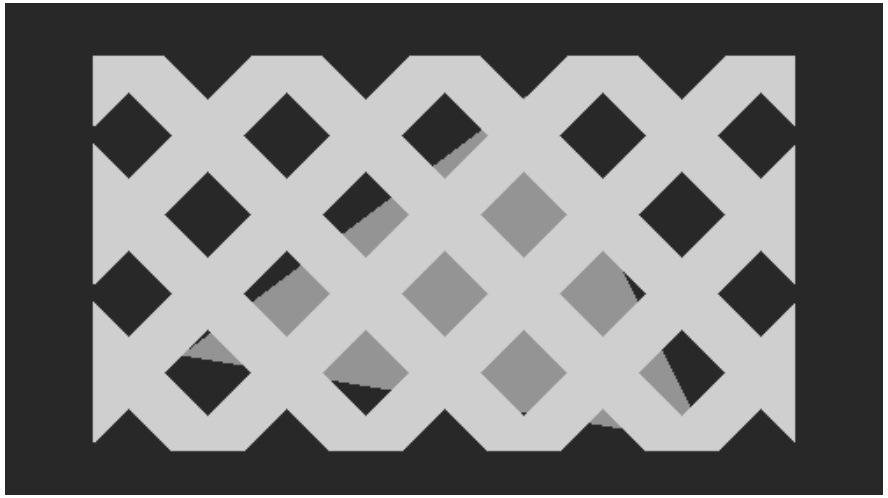

GT

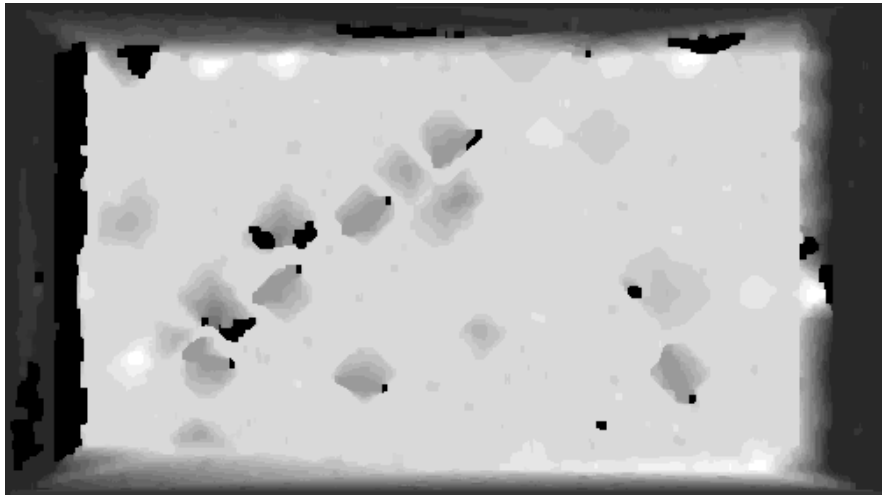

ELAS

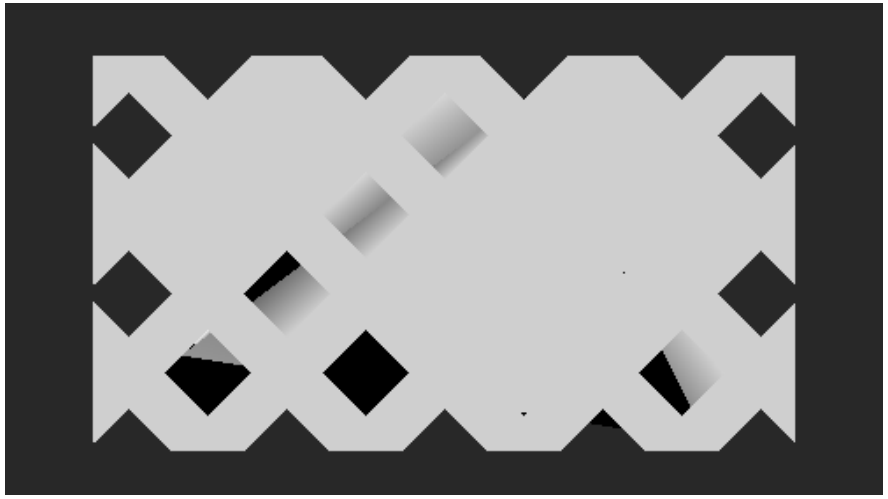

Ours (Pass 1)

Results – Lattice Occluding

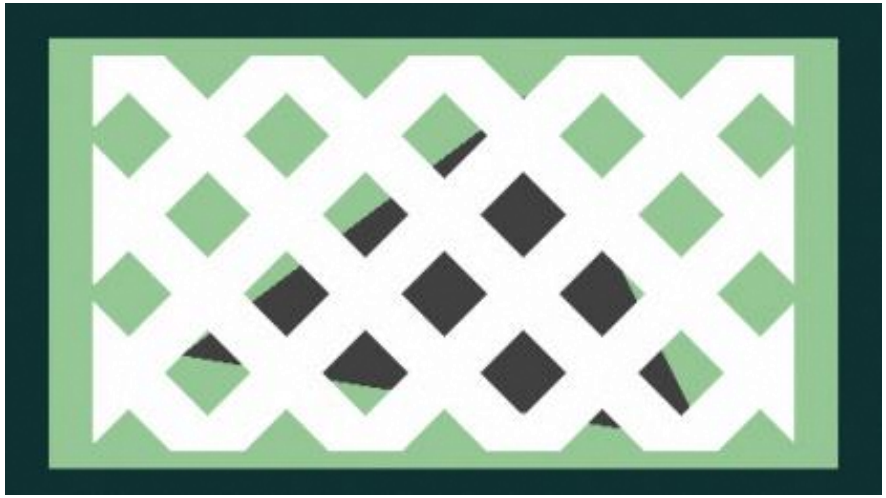

Left

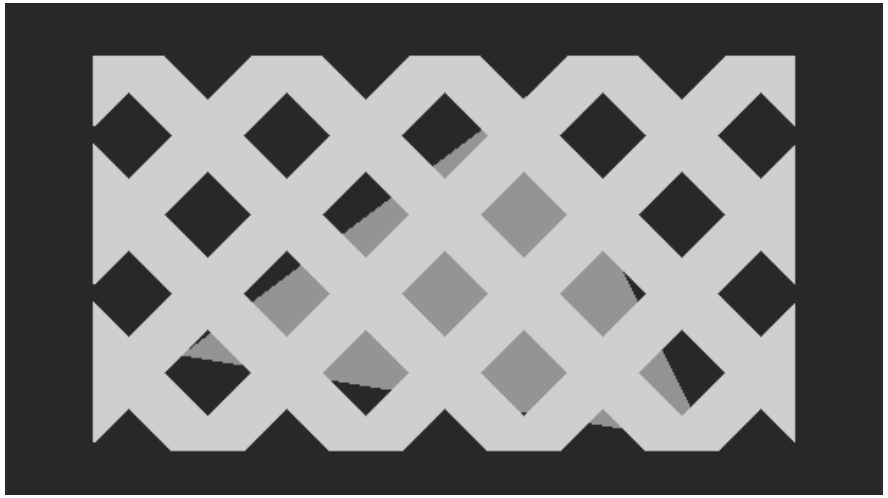

GT

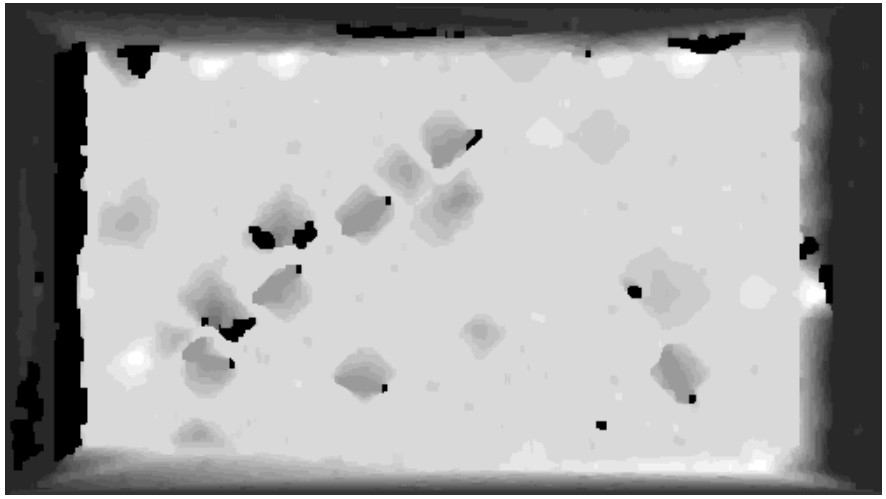

ELAS

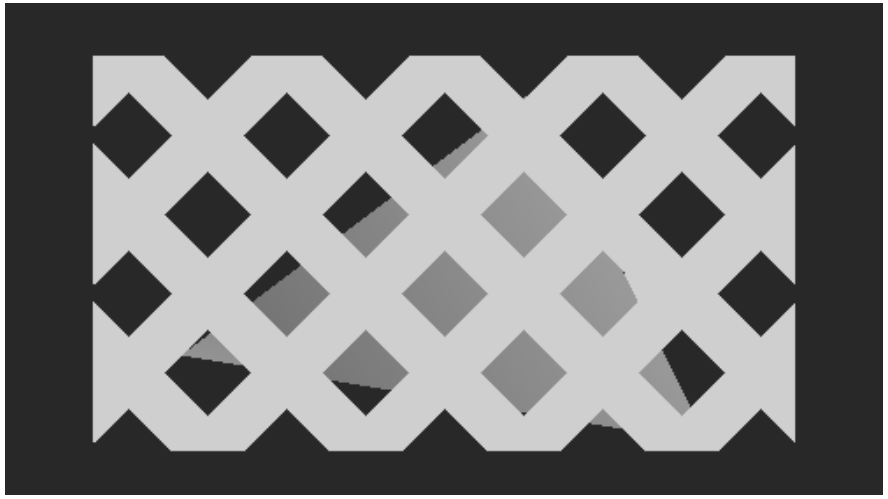

Ours (Pass 2)

Results – Texture Squares

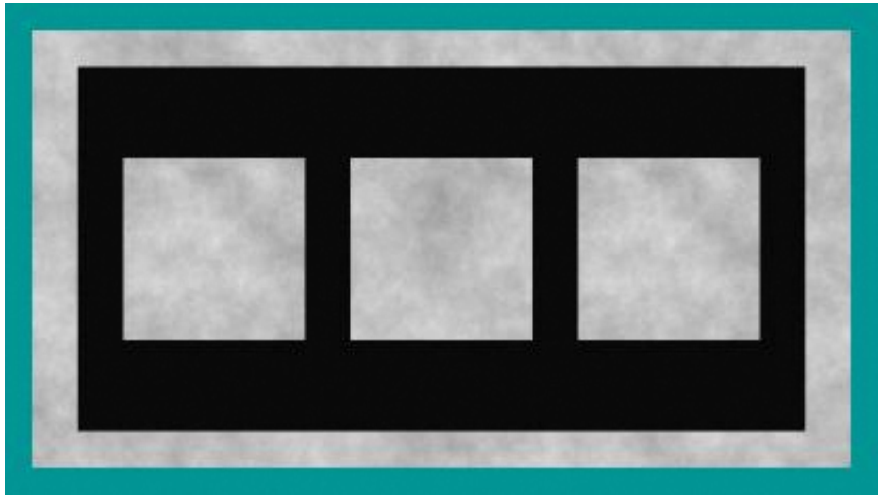

Left

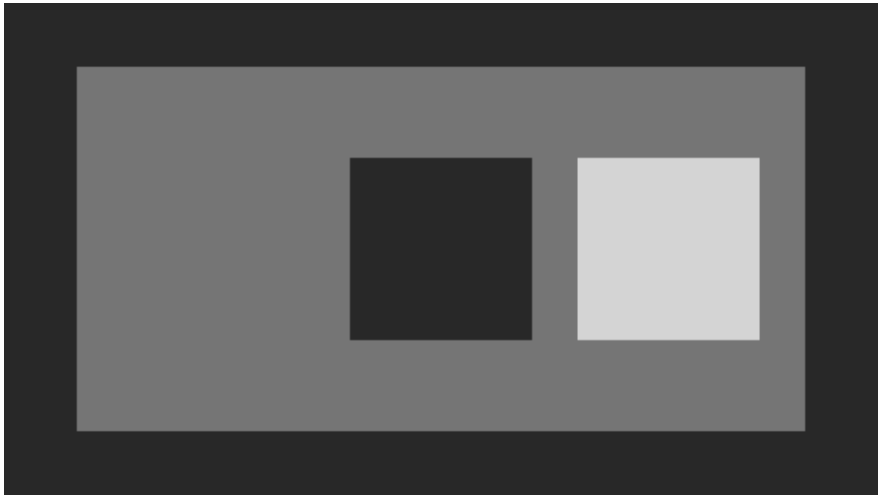

GT

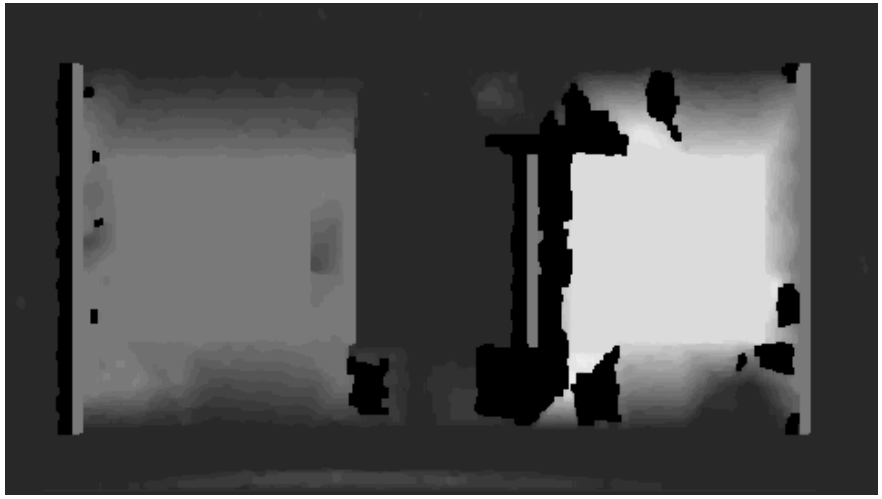

ELAS

Ours (Pass 1)

Results – Texture Squares

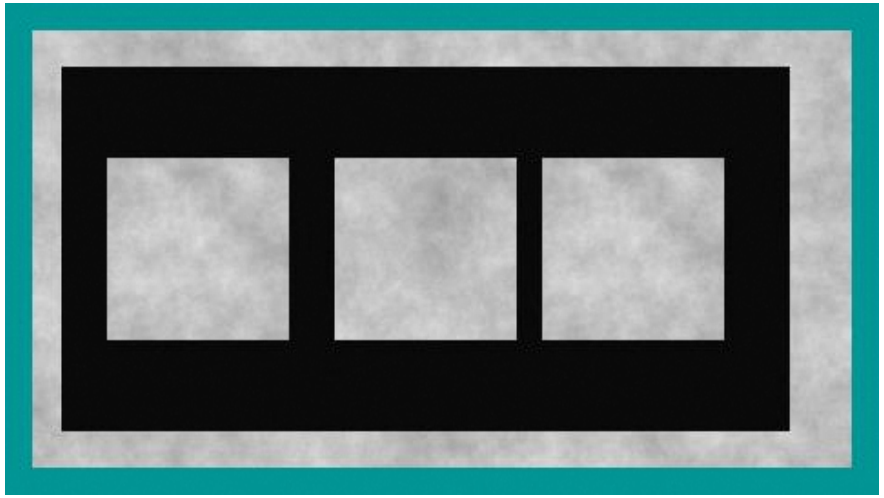

Right

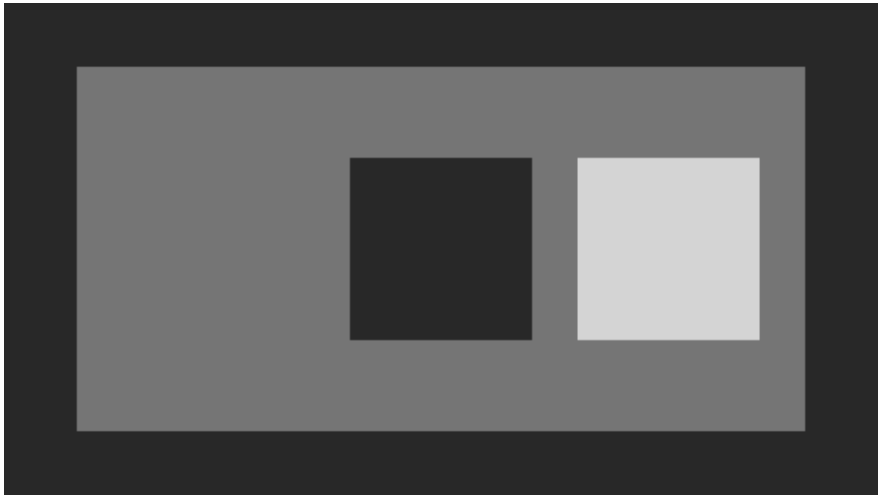

GT

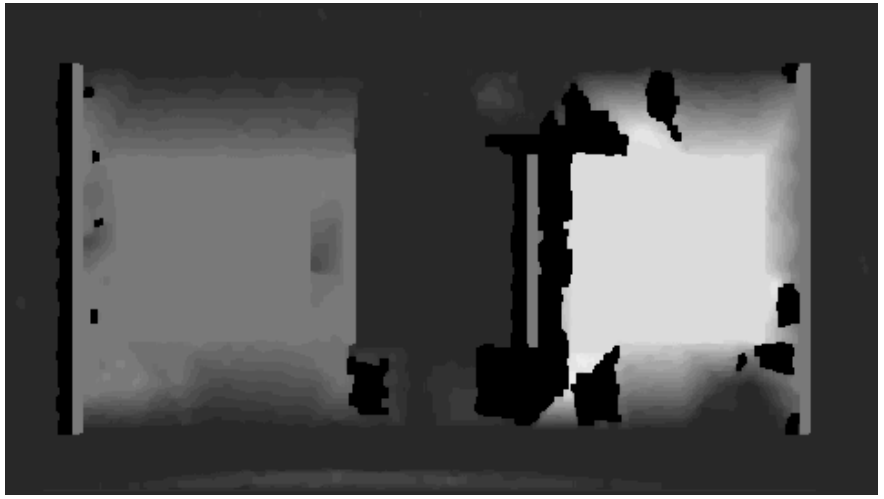

ELAS

Ours (Pass 1)

Results – Texture Squares

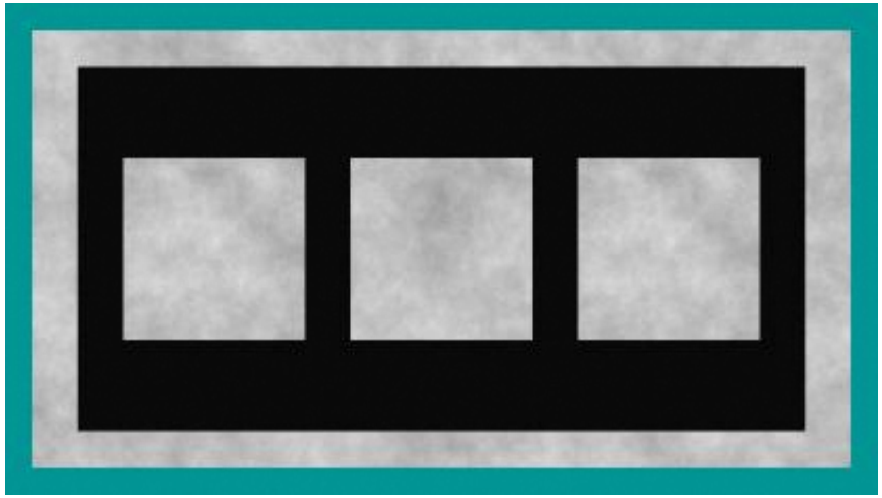

Left

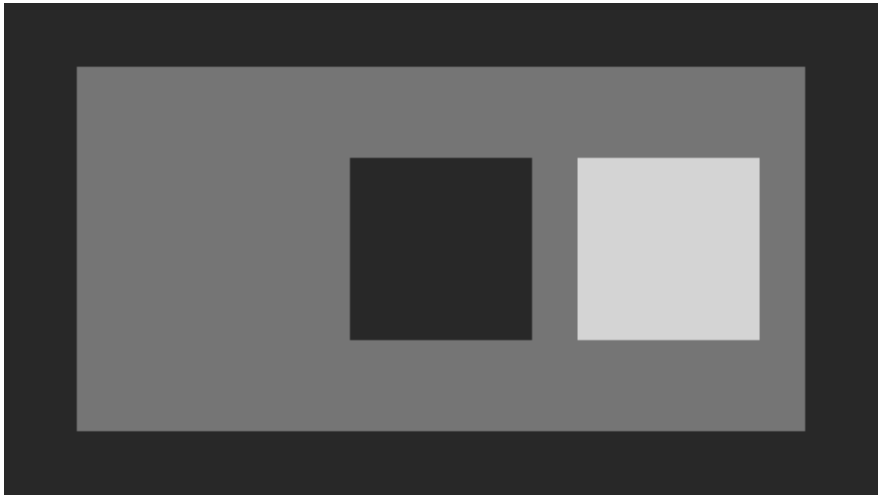

GT

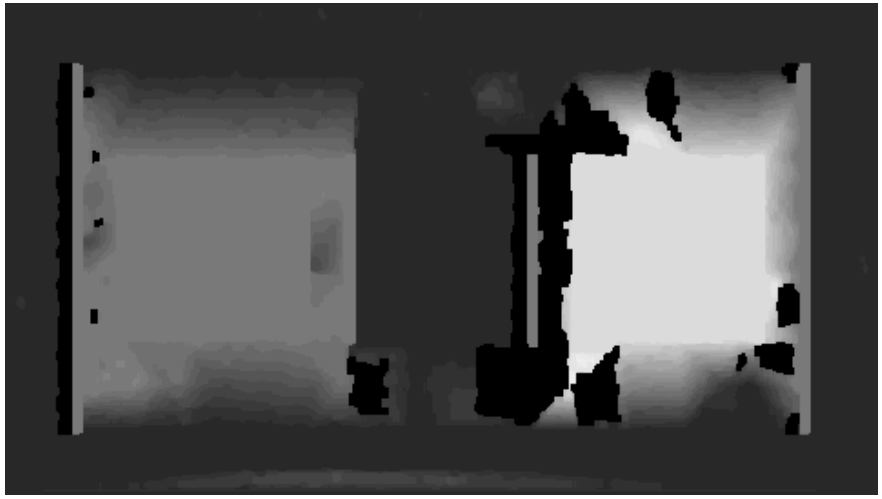

ELAS

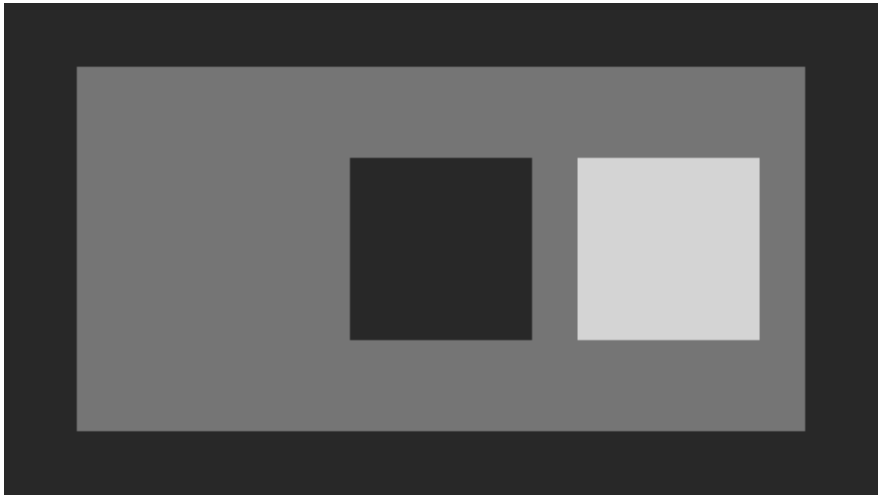

Ours (using texture)

Mondrian Stereo

✓Can we write an algorithm to solve these "pathological cases"?

- And if so, what can we learn that will help us improve real stereo algorithms?

Towards “real” Mondrian Stereo

- Avoid complex (yet insufficient) monolithic energy minimization
- Use sequence of simpler minimization problems
- First pass: “closest surfaces”
- Identify untextured segments
- “Push back” to farther surface with color and occlusion reasoning

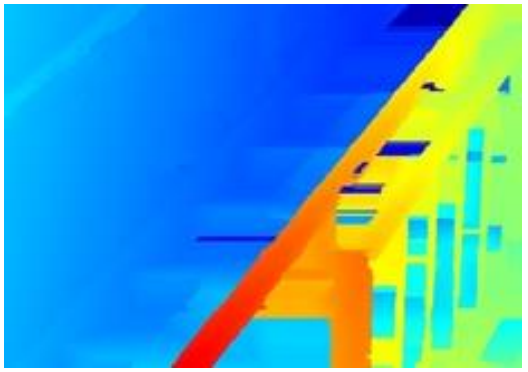

Conclusion

Stereo is not a solved problem!