CS635
Deep Visual Computing:

Deep Single Image Reconstruction and more

Daniel G. Aliaga

Spring 2025

Topics

- Single image -> 3D model
- Single image -> code fragment
- Single image -> better single image

Procedural Modeling of a Building from a Single Image

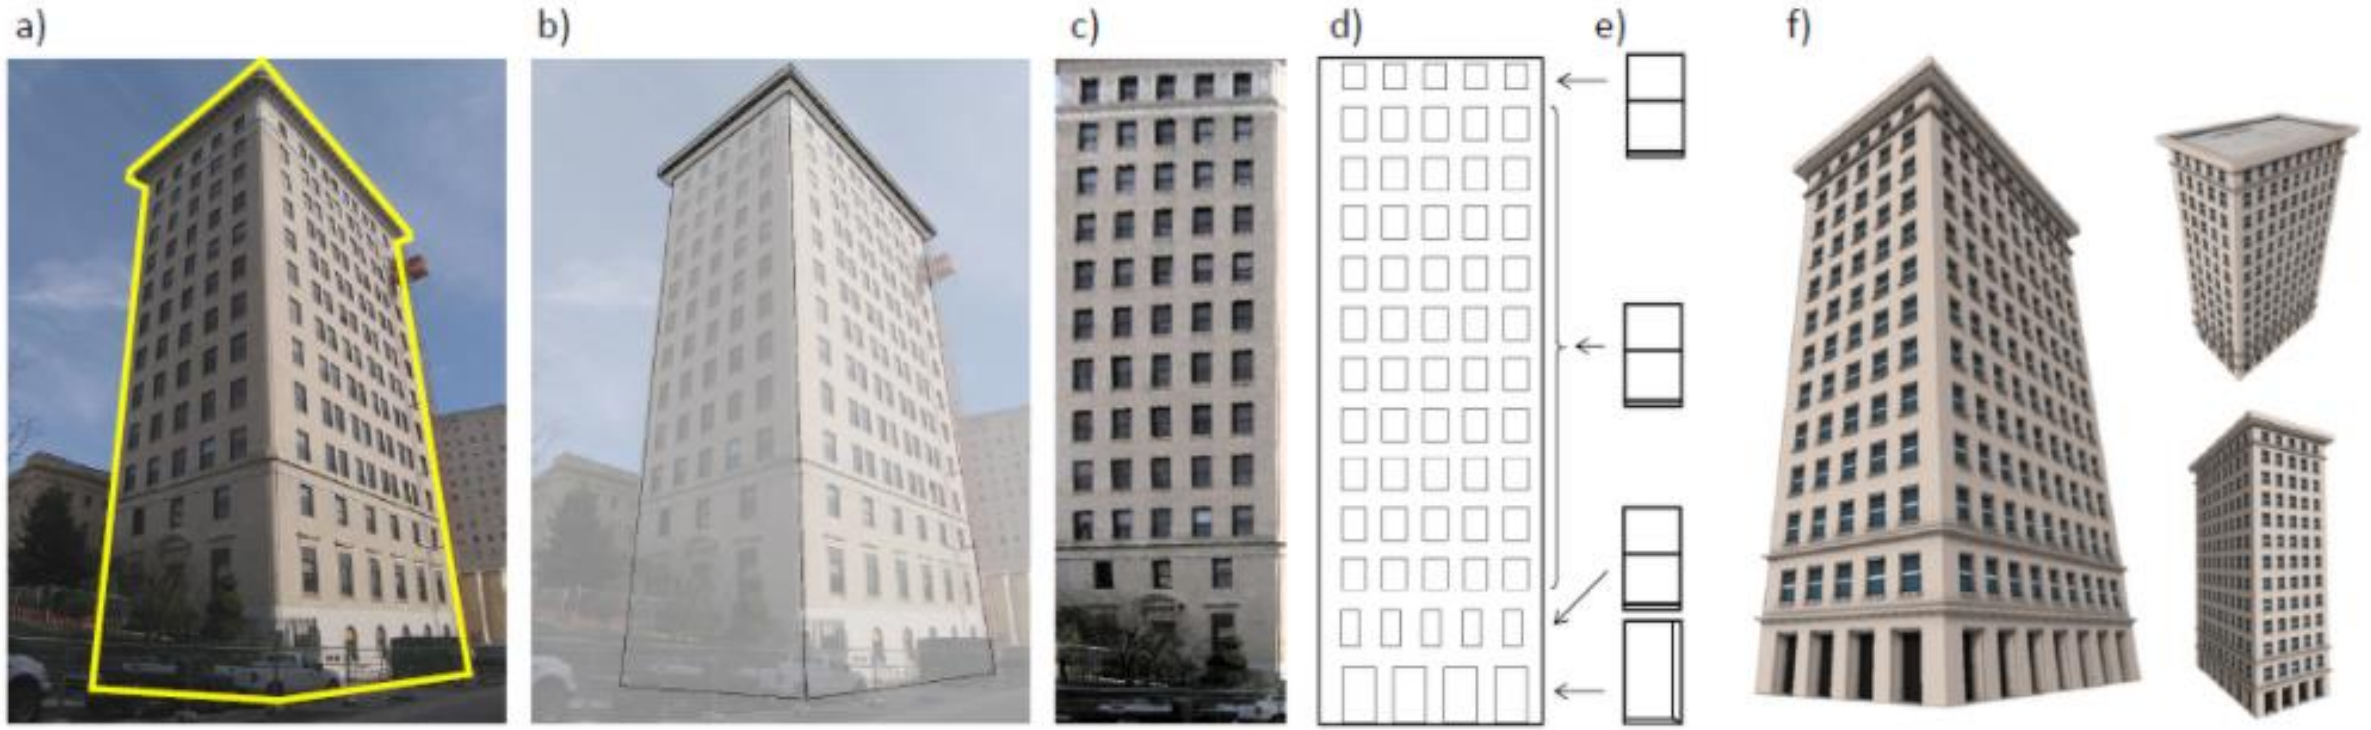

Deep Mesh Reconstruction from Single RGB Images via Topology Modification Networks

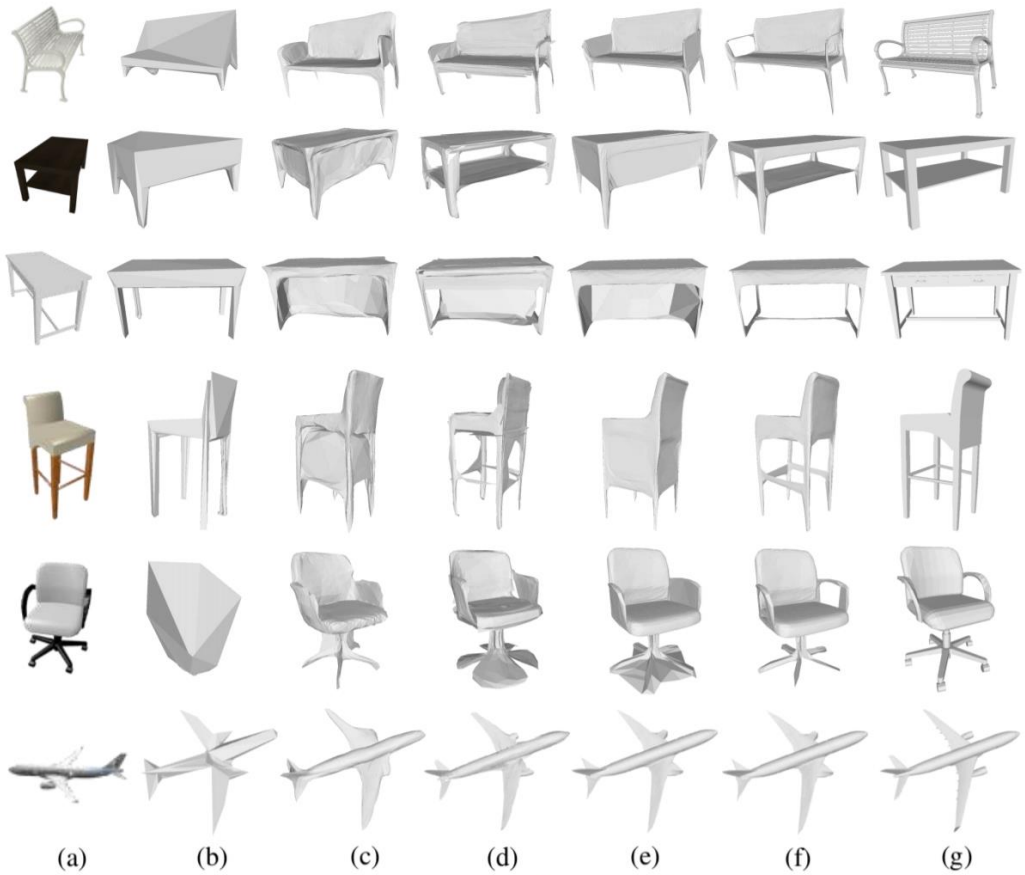

Figure 4. Qualitative results. (a) Input image; (b) N3MR; (c) Pixel2Mesh; (d) AtlasNet-25; (e) Baseline; (f) Ours; (g) Ground truth.

CSGNET

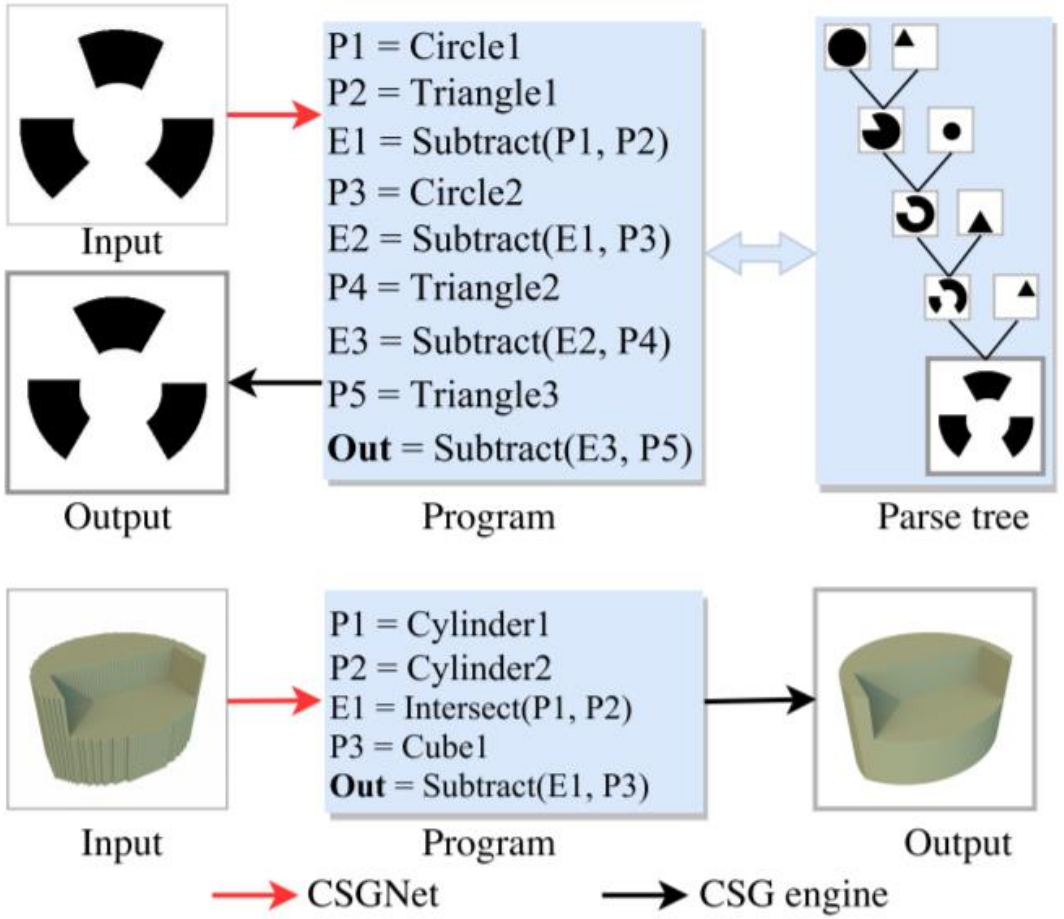

Figure 1. **Our shape parser produces a compact program that generates an input 2D or 3D shape**. On top is an input image of 2D shape, its program and the underlying parse tree where primitives are combined with boolean operations. On the bottom is an input voxelized 3D shape, the induced program, and the resulting shape from its execution.

Learning to Describe Scenes with Programs

(a) Input image

(sphere, small, metal, green, x=2, y=2, z=0) …

(b) Object parsing & group detection

```
for(i<3)
 for(j<3-i)
  sphere(pos=(1+i,1+j,0),
         color=6-j)
for(i<4)
  cylinder(pos=(3,3,i),
           color=7)
cylinder(pos=(4,2,0),
         color=3)
```

(c) Program

Deep Learning for Single Image Super-Resolution: A Brief Review

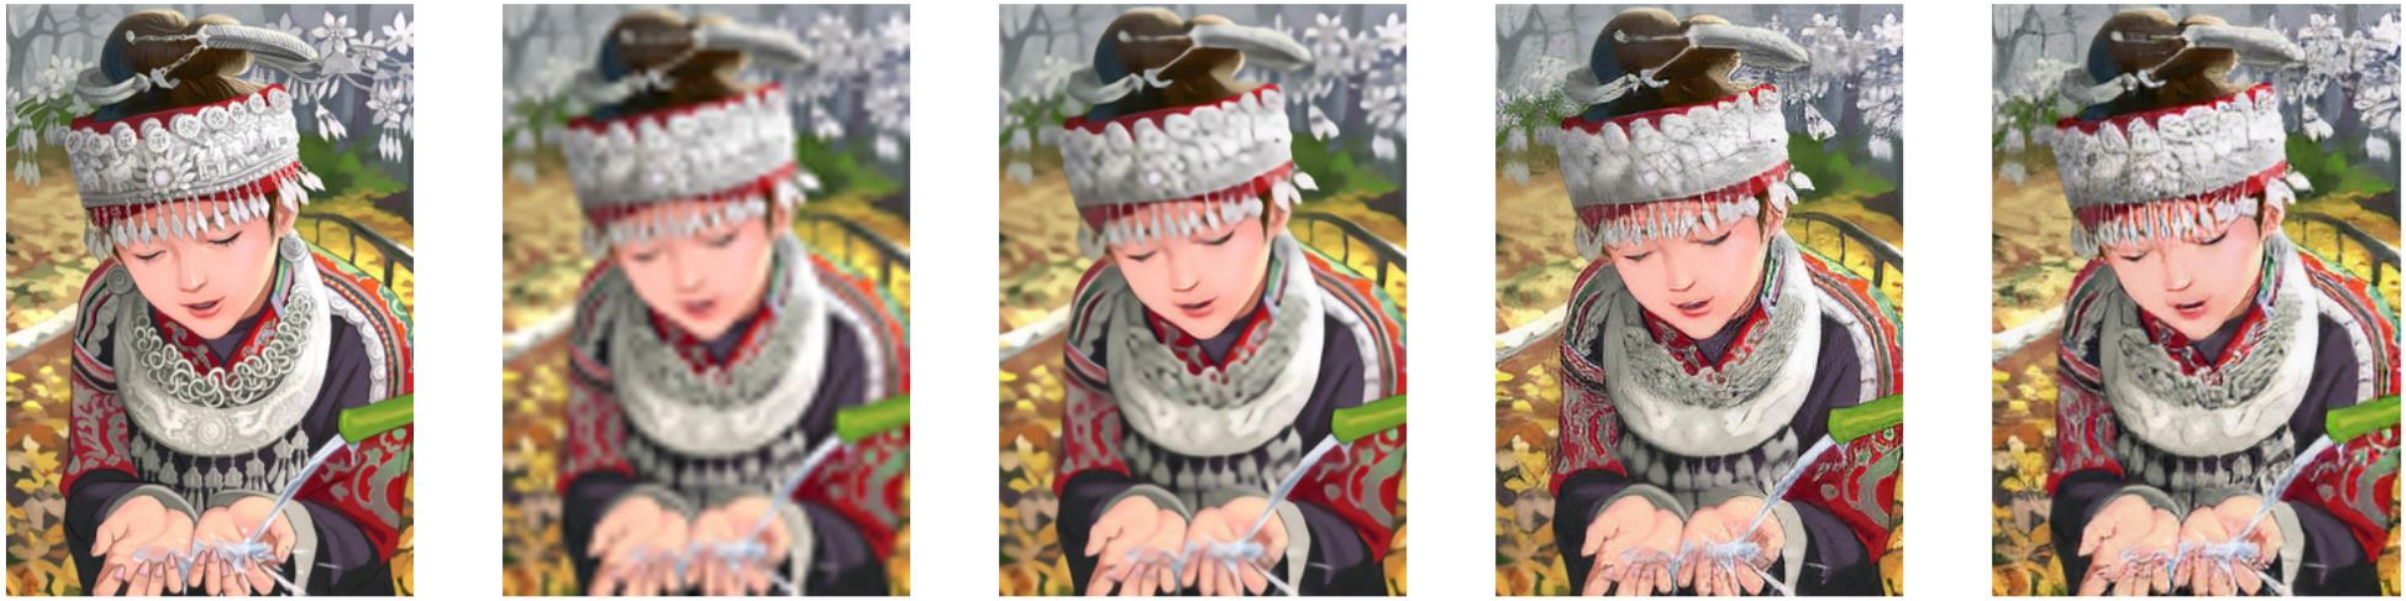

(a) HR (b) bicubic(21.59dB/0.6423) (c) SRResNet(23.53dB/0.7832) (d) SRGAN(21.15dB/0.6868) (e) SRCX(20.88dB/0.6002)

Figure 12: Visual comparisons between the MSE, MSE + GAN and MAE +GAN + Contextual Loss (The authors of [68] and [116] released their results.) We can see that the perceptual loss leads to a lower PSNR/SSIM but a better visual quality.

Fast Separation of Direct and Global Images Using High Frequency Illumination

- http://www1.cs.columbia.edu/CAVE/publications/pdfs/Krishnan_TOG06.pdf

HDR From Single Image

- https://cave.cs.columbia.edu/Statics/publications/pdfs/Nayar_CVPR00_2.pdf

Learning Single Camera Depth Estimation using Dual-Pixels

- https://arxiv.org/abs/1904.05822

https://www.youtube.com/watch?v=zem03fZWLrQ

Starting at beginning…

Deep Depth Completion of a Single RGB-D Image

- https://deepcompletion.cs.princeton.edu/

Depth Anything

- https://depth-anything.github.io/

Learning View Priors for Single-view 3D Reconstruction

- http://hiroharu-kato.com/projects_en/view_prior_learning.html
- https://arxiv.org/abs/1811.10719

Transformable Bottleneck Networks

- https://kyleolsz.github.io/TB-Networks/

(includes video)

- Also:

https://www.youtube.com/watch?v=zem03fZWLrQ

Starting at 11:10…