Shading Annotations in the Wild

goo.gl/EFebP9

Balazs Kovacs, Sean Bell, Noah Snavely, Kavita Bala
Cornell University

1. Motivation

- Large-scale datasets fuel research progress
  - ImageNet, Places, SUN, NYUv2, MINC, ...
- **Missing:** large-scale dataset of shading annotations
- **Missing:** large-scale benchmark for intrinsic images shading component

2. Contributions

- **Shading annotations in the Wild (SAW)**
  - New large-scale dataset of shading annotations in real-world images
- New deep-learning based shading prediction
  - Smooth/non-smooth shading
- Benchmark for shading decomposition performance of intrinsic images

3. Data Collection

- We identify three shading annotation types:
  - Smooth/constant shading
  - Shadow boundaries
  - Depth/normal discontinuities
- How to collect shading annotations?
  - **Pilot study:** Ask people to compare shading at predetermined point pairs, similarly to [1].
    - Expected output: shading is <, >, =
    - People are not good at this task
  - **Idea 1:** Let people pick point pairs
    - We collect <, >, people still fail on =
    - Generate and filter shadow boundaries between these point pairs
  - **Idea 2:** Collect regions of constant shading
  - Automatically find depth/normal discontinuities
    - From depth maps of NYUv2 Depth [2]

4. Annotations

5. Crowdsourcing Pipeline

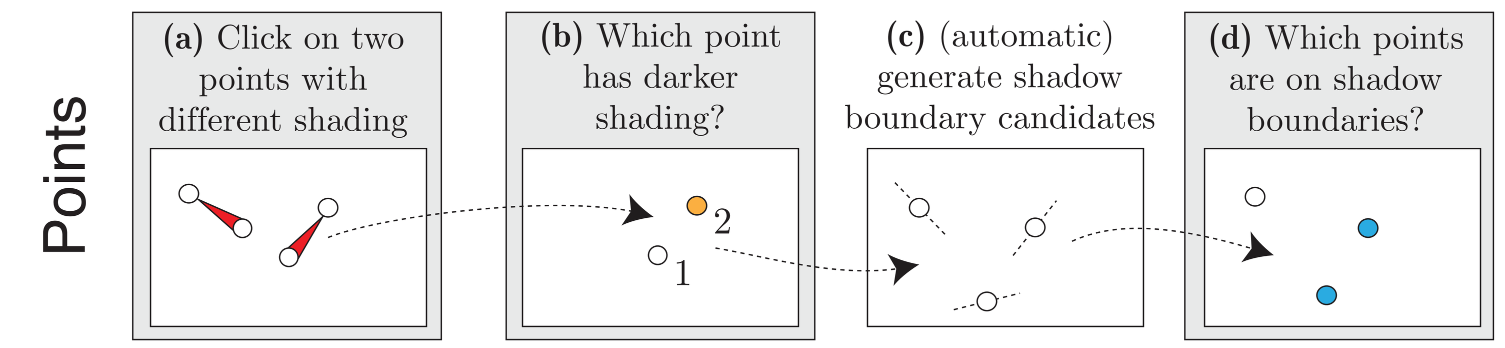

We use a two step pipeline to obtain <, > shading point comparisons. Then we ask people to filter shadow boundary candidates we automatically generated using these comparisons.

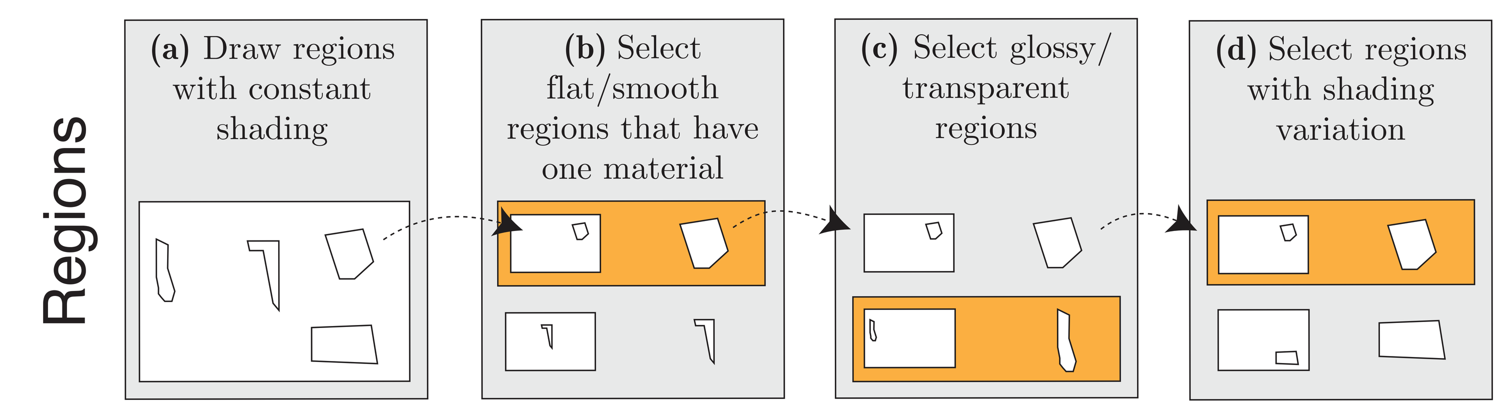

We asked workers to draw polygons around regions of constant shading. Then we pass these regions through three filtering tasks to ensure high quality annotations.

6. Pixel Labels

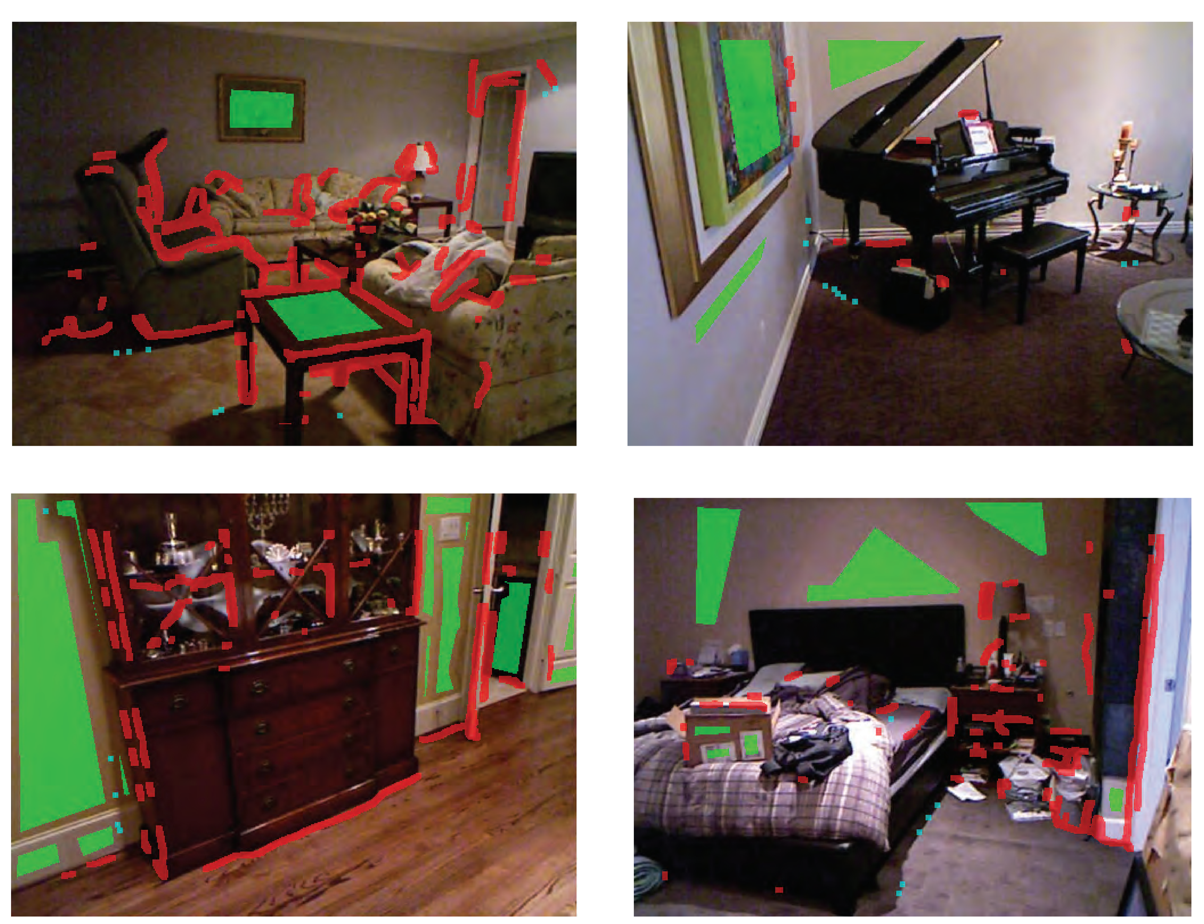

- Final pixel labels from mturk annotations and depth/normal discontinuities
- Green: smooth shading (mturk)
- Cyan: shadow boundary (semi-automatic)
- Red: depth/normal discontinuity (automatic)
- Use two classes for training:
  - Smooth shading: green
  - Non-smooth shading: cyan + red

7. Learning

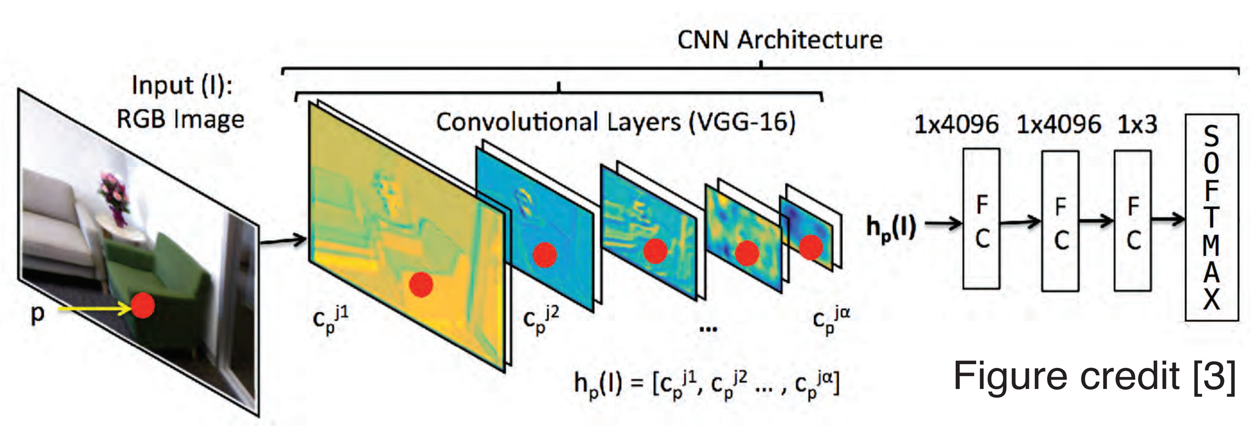

Figure credit [3]

- Fine-tune PixelNet [3] to predict smooth/non-smooth shading for each pixel
- Balance classes with 2 : 1 : 1 ratio

Smooth shading heatmaps

8. Shading Prior

Original image | Shading layer (Retinex) | Shading layer (Retinex with prior)

- Use smooth shading predictions as a prior in Retinex
- Promising initial results
- More research is needed to seamlessly incorporate prior

[1] Sean Bell, Kavita Bala, Noah Snavely. "Intrinsic Images in the Wild", SIGGRAPH 2014.
[2] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor segmentation and support inference from rgbd images. ECCV 2012.
[3] A. Bansal, B. Russell, and A. Gupta. Marr Revisited: 2D-3D model alignment via surface normal prediction. CVPR 2016.

9. Evaluation

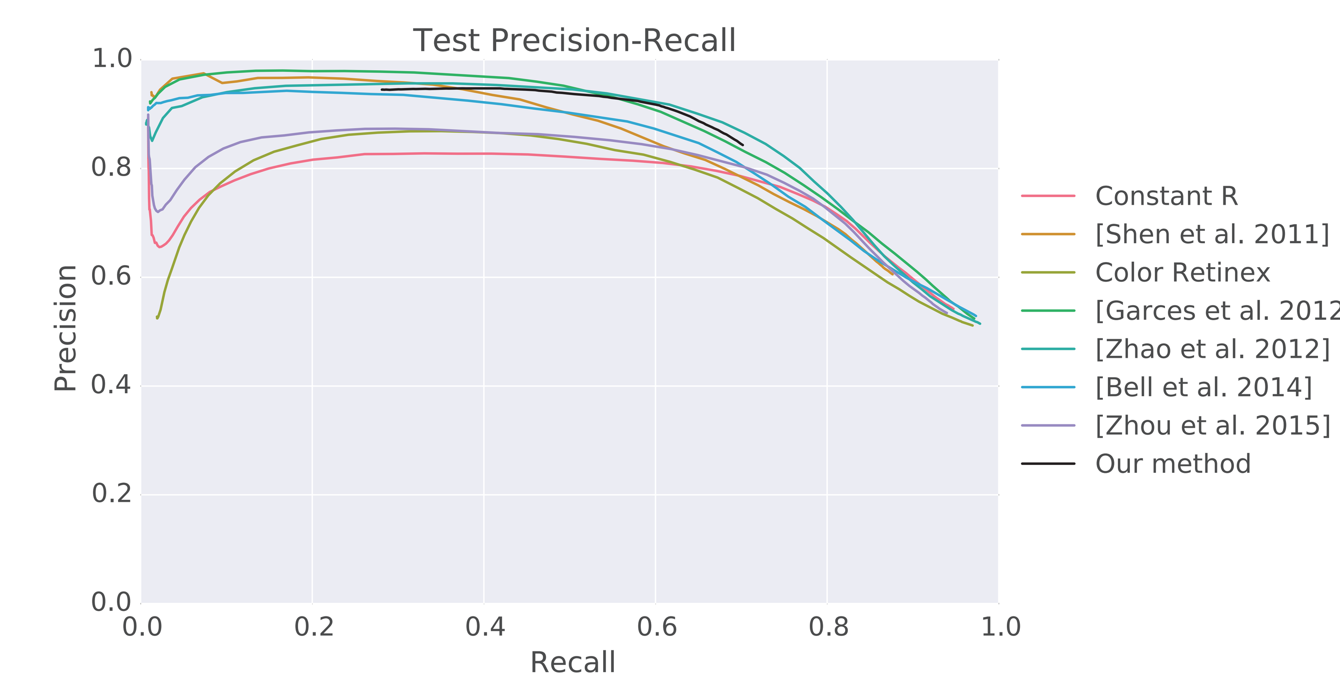

- To compare our smooth/non-smooth predictions to existing methods (which predict a full shading layer):
  - Threshold the gradient of shading
  - Compare the resulting 2-class labels
- Our method achieves competitive results
- **Future:**
  - New shading benchmark for intrinsic images that combines reflectance and shading
  - Improved fully convolutional training

This work was supported by the National Science Foundation (grants IIS-1617861, IIS-1011919, IIS-1161645, IIS-1149393), and by a Google Faculty Research Award.