Shadows From Shading

Lahav Lipson **Columbia University** **Niles Christensen**

Curriculum Learning

As the network improves, we increase the complexity of the scene

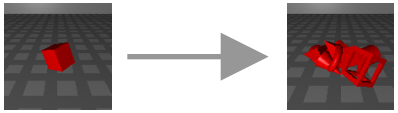

Our Idea

Shadow prediction can be used to approximate shadow borders in shadow-removal systems. Neural Networks trained to predict shadows may implicity learn 3D shape.

ShadowNet

ShadowNet is a convolutional autoencoder that predicts the likelihood of a shadow at each pixel given only the shading as input. For hard shadows, we set a likelihood threshold above which there is a shadow and below there is none. For soft shadows, we scale the shadow intensity by the likelihood.

P>10% **P>50%** **P>75%** **Soft**

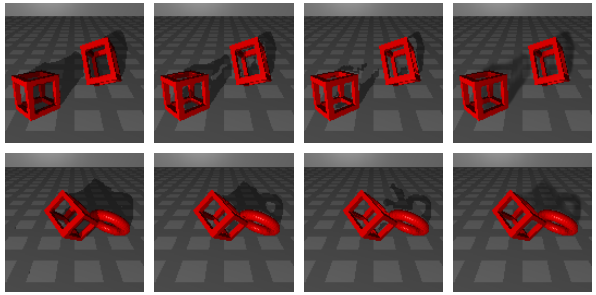

All other images on this poster have P > 50%

How well does ShadowNet work?

Left: Input Image **Center:** Predicted **Right**: Ground Truth

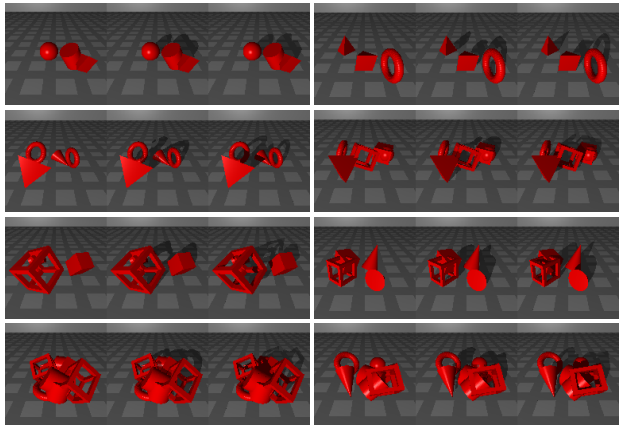

Scene Reconstruction

How much of the scene is encoded in the shadows? We trained a second network to estimate scene parameters (shape type and transformations) from the latent vectors of ShadowNet.

Left: Input Image **Right**: Reconstructed Scene

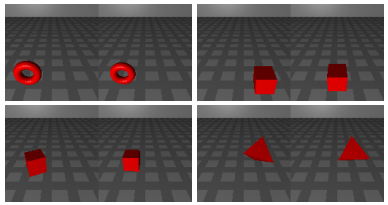

The second network could accurately estimate all shape parameters, except for orientation.

Random Lights

Our system is robust to changes in lighting

Left: Input **Center:** Predicted **Right:** Ground Truth

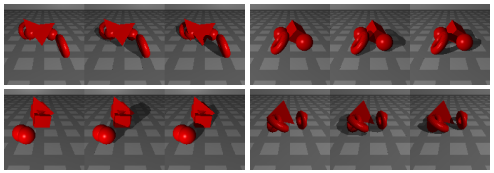

Can our system generalize?

We tested our network on shapes not seen during training

Left: Input **Center:** Predicted **Right**: Ground Truth

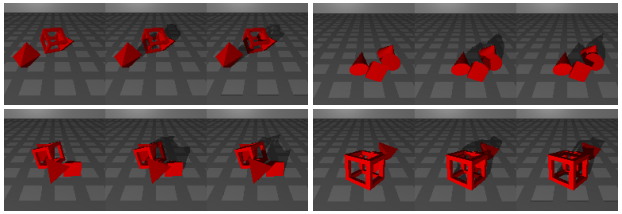

Stereo Vision

Stereo vision gives subjectively better results

Top: Input **Middle:** Predicted **Bottom:** Ground Truth

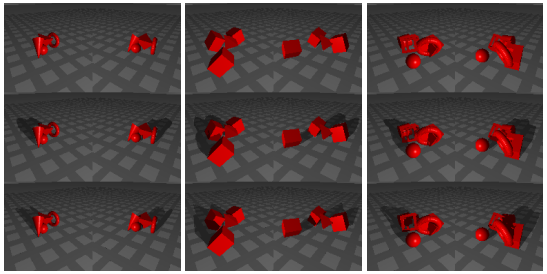

Shape Difficulty

It is harder to estimate shadows for some shapes than others. By evaluating on scenes containing only one type of shape (while still using the curriculum paradigm), we can see which shapes are harder.