

Learning a Generative Model of Human Faces Through Inverse Rendering

Skylar Sutherland

MIT, Cambridge, Massachusetts, United States

Bernhard Egger

Massachusetts Institute of Technology, Cambridge, Massachusetts, United States

Josh Tenenbaum

MIT, Cambridge, Massachusetts, United States

Abstract

Generative models in an inverse graphics framework are appealing models for visual perception. How might children acquire them? We present a computational procedure for learning generative models of human faces using developmentally plausible input. Our statistical model of shape and appearance initially uses the average face as a template with a simple Gaussian process model of deformations. We iteratively learn the statistical distribution of faces by performing analysis-by-synthesis on a small number of images and combine the results to construct an improved generative model. Our analysis-by-synthesis framework combines a convolutional neural network for fast inference with a Markov chain Monte Carlo process for detailed refinement. This learning strategy quickly captures the variation of natural faces and demonstrates an efficient way to learn the distribution of faces.