



Paul Upchurch *Cornell University*

Recognizing and Editing Photorealistic Materials using Crowdsourcing and Deep Neural Networks

**** GE-104, 15:00-16:00, 23 April 2018 ****

Abstract: Colors in a photograph are the result of light interacting with materials so understanding materials is of central importance for creating machines which can reason like a human about real-world scenes given only a photo or a video. In this talk I describe two recent projects which can recognize or edit materials in an image. The Materials in Context project assigns material labels (metal, plastic, wood, etc) to every pixel of an image with the highest recognition accuracy to date. We also released a large-scale material database used by research and industrial groups with 3000x more labels and 2x as many categories as previous databases. In a follow-up work we developed a new crowdsourcing method which reduces error by 39% when gathering material labels for hard-to-label images. The second project, Deep Feature Interpolation, makes believable edits to a high-resolution photo. For example a person can be made to look older, have facial hair, open their eyes, etc. Our method produces dramatically higher-quality results compared to the previous state-of-the-art. These works have the potential to improve training data collection for autonomous vehicles, provide material labels for robotic action planning, and synthesize training variations for face recognition.

Biography: Paul Upchurch works on recognizing and editing materials in images. His work lies at the intersection of Computer Vision, Computer Graphics and Machine Learning. Paul is currently a PhD candidate at Cornell University under the supervision of Kavita Bala, Noah Snavely and Kilian Weinberger. Previously, he worked on real-time visualization software at the NASA Jet Propulsion Laboratory and videogames for the Sony PlayStation. He holds a B.S. in Computer Science from Caltech.