

Light Field Layer Matting

Juliet Fiss¹, Brian Curless¹, Richard Szeliski²

¹University of Washington. ²Microsoft Research.




Figure 1: We model the input light field as foreground layer composited over a background light field.

In this paper, we use matting to separate foreground layers from light fields captured with a plenoptic camera. We represent the input 4D light field as a 4D background light field, plus a 2D spatially varying foreground color layer with alpha. Our method can be used to both pull a foreground matte and estimate an occluded background light field. Our method assumes that the foreground layer is thin and fronto-parallel, and is composed of a limited set of colors that are distinct from the background layer colors. Our method works well for thin, translucent, and blurred foreground occluders. Our representation can be used to render the light field from novel views, handling disocclusions while avoiding common artifacts.

Many photographs of natural scenes are composed of layers, where a foreground layer partially occludes background content. A common and important case of this type of composition occurs when the foreground layer is frontoparallel and spatially complex. A familiar example is the photograph taken through a window covered by a thin layer of translucent dust. In this paper, we propose a method for matting such layers from light field images.

We model an input light field L as a composite of two layers: a background light field K , plus a foreground layer with constant depth, spatially varying color F , and spatially varying alpha α . The foreground depth is user specified. Our method recovers the background light field, foreground color, and foreground alpha automatically.

Our method takes as user input two depth parameters: d_f , which specifies the depth at which the foreground layer is most in focus, and d_τ , a threshold depth that separates the foreground layer from the background layer. To select these parameters, the user sweeps a synthetic focal plane through the scene (e.g. with a depth slider) and makes selections by visual inspection. These parameters are necessary to signal user intent: what content should be considered part of the foreground, and what content part of the background? Our method uses these depth parameters as an alternative to trimaps or other forms of image annotation.

The foreground and background layers may also be used together to render novel views. In this paper, we compute our results on images from a Lytro camera, and compare our renderings to those from the Lytro perspective shift feature.



Figure 2: Removing dust from a light field captured through a window. Top left: original light field rendered with bird in focus. Notice the dust partially occluding the chest and face of the bird. Top right: background light field (our result) rendered with bird in focus. Bottom left: original light field rendered with window dust in focus. Bottom right: detail comparison from images in top row

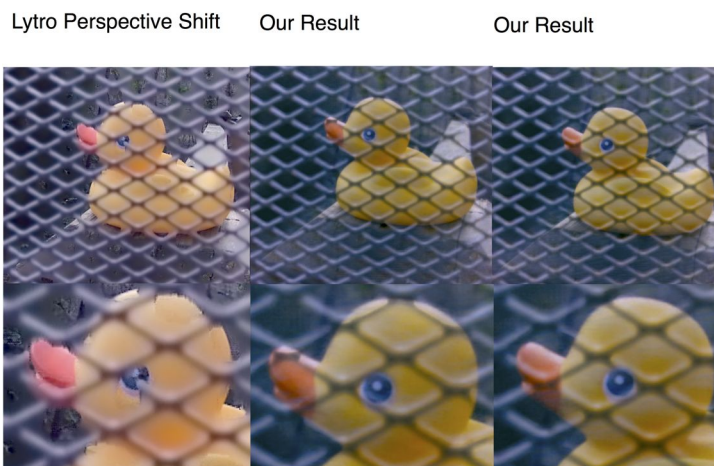


Figure 3: In this scene, a wire mesh occludes a rubber duck. We compare our novel view renderings to the Lytro perspective shift feature.