Light-Field Retargeting with Focal Stack Seam Carving

Clemens Birklbauer and Oliver Bimber
Johannes Kepler University Linz *

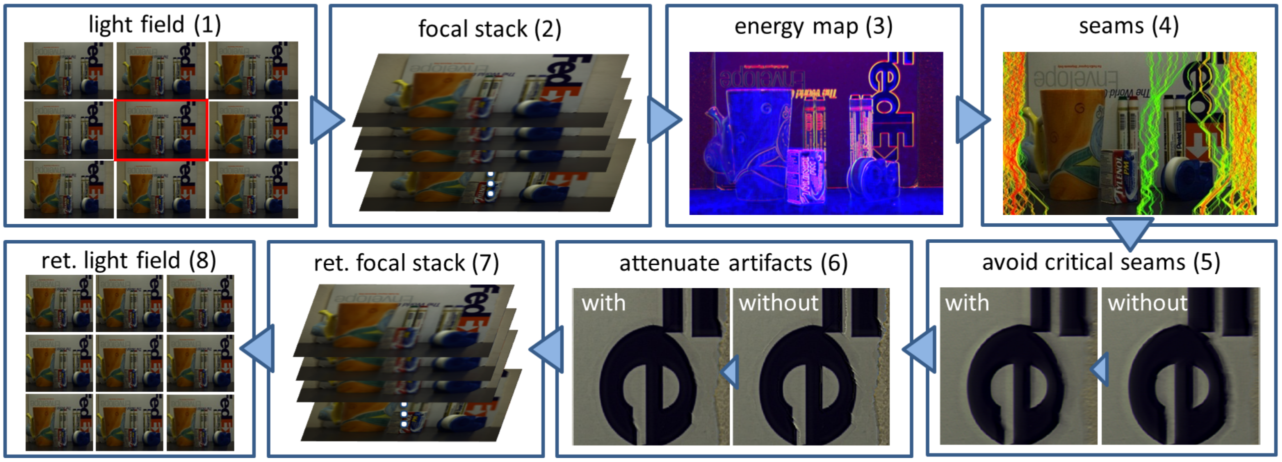

Figure 1: *Steps of our light-field retargeting process with an example of 24% width compression.*

1 Introduction and Motivation

With increasing sensor resolutions of digital cameras, light-field imaging is becoming more and more relevant, and might even replace classical 2D imaging in photography sooner or later. It enables, for instance, digital refocussing and perspective changes after capturing. Rescaling light fields to different resolutions and aspect rations, however, is challenging. As for regular image and video content, a linear scaling alters the aspect ratio of recorded objects in an unnatural way. In contrast, image and video retargeting utilizes a nonlinear and content-based scaling. Applying image retargeting to individual video frames independently does not retain temporal consistency. Similarly, applying image retargeting naively to the spatial domain of light fields will not retain angular consistency. We present a first approach to light-field retargeting. It allows compressing or stretching light-fields while retaining angular consistency.

2 Our Approach

While stereo-pair retargeting or 3D retargeting could potentially be extended to light-field retargeting, both would rely on high-quality depth reconstruction which is difficult to achieve for complex scenes. We follow an entirely different approach that does not depend on precise depth information (cf. figure 1):
First, we convert a light field (1) into a focal stack (2) by rendering its focal slices from the central perspective (1, red border) with correspondingly adjusted focal length. The focal stack can be converted back to a light field (7,8) by shifting and averaging its slices according to the desired perspective and by deconvolving the resulting image with a perspective-dependent point-spread function, as explained in [Levin and Durand 2010]. Before we do this for all perspectives of the light field we apply seam carving to the central perspective. The seams (4) are mainly derived from the L1-norm of the gradients with forward energy calculation (3, green). Since this central perspective matches the perspective of the focal stack, we remove the calculated seams in all of its slices without violating angular consistency. Removing pixels where these seams are cutting through out-of-focus regions in slices of the focal stack, however, causes deconvolution (ringing) errors around edges during the reconstruction of the remaining light-field perspectives. Therefore, we avoid cutting seams through these critical regions (5). They are determined by integrating the RGB differences between all focal stack slices and the deconvolved central perspective. This integral is used as an additional energy term for seam carving (3, red). It does not completely avoid critical seams. Remaining artifacts are attenuated by convolving the focal stack slices in regions with new gradients that are created by seam carving (6). We assume that background objects are visually less salient. Therefore, we penalize seams through foreground objects to better preserve them. This is achieved by reconstructing a coarse depth map from the focal stack. The depth is used as a third energy term for seam carving (3, blue).

3 Results and Limitations

To our knowledge, we presented the first approach to light-field retargeting. It enables stretching and compressing light fields while preserving angular consistency without the need for reconstructing precise depth information (see supplementary video for results). However, it has several limitations:
The conversion of focal stacks to light fields, as explained in [Levin and Durand 2010], leads to imperfect reconstructions at occlusion boundaries. It also requires shifting the focal slices according to the constructed perspective images. For the inverse process, a similar shift of the perspective images is required to construct the focal slices. Therefore, the spatial resolution of resulting light fields is cropped by the maximum shift that is applied during these two conversion processes.
In future, we want to investigating more advanced techniques for filtering critical seams.

References

LEVIN, A., AND DURAND, F. 2010. Linear view synthesis using a dimensionality gap light field prior. In *CVPR*, 1831–1838.

*{firstname.lastname}@jku.at