Light-Field Retargeting with Focal Stack Seam Carving

Clemens Birklbauer[1] and Oliver Bimber[1]
Johannes Kepler University Linz

SIGGRAPH2011

JOHANNES KEPLER UNIVERSITY LINZ | JKU

Institute of Computer Graphics

Abstract

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 refocusing and perspective changes after capturing. Rescaling light fields to different resolutions and aspect ratios, 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.

Light-Field Retargeting
We present a first approach to light-field retargeting. It allows compressing or stretching light fields while retaining angular consistency. Therefore, we convert the light field to a focal stack, apply a seam-carving algorithm to the stack and convert it back to a light field. This approach does not rely on high-quality depth information like stereo-pair or 3D retargeting.

Related Work

Focal Stack to Light-Field Conversion
A method for converting focal stacks to light fields is presented in [1]. For computing the different perspective images of the resulting light field, the slices of a given focal stack are sheared, added and the result is deconvolved with a spatially invariant, but correspondingly sheared PSF. We apply this method in our algorithm.

Seam Carving
Seam carving algorithms remove or add pixel paths of least importance from or to an image for resizing. In [2] 2D seam carving is extended to 3D space-time volumes of videos. We apply the forward energy calculations that are described in [2] for computing seams that can be carved.

Stereo-Pair Retargeting
Stereo-pair retargeting approaches, such as [3], calculate corresponding seams in left and right stereo images. They introduce additional energy terms to preserve consistency in the resulting stereo-pair. In contrast to our approach, these methods rely on robust image correspondences.

3D Retargeting
In [4] a technique for 3D retargeting is presented. To avoid undesired deformations of a 3D model during resizing, it automatically detects vulnerable regions of the object's mesh. The vulnerability information constraints resizing, which is realized with a space-deformation technique. Converting light fields to an intermediate 3D mesh representation to apply 3D retargeting would require high-quality depth reconstruction.

References

1. Levin, A., and Durand, F. 2010. Linear view synthesis using a dimensionality gap light field prior. In CVPR, 1831–1838.
2. Rubinstein, M., Shamir, A., and Avidan, S. 2008. Improved seam carving for video retargeting. ACM Trans. Graph. 27 (August), 16:1–16:9.
3. Utsugi, K., Shibahara, T., Koike, T., Takahashi, K., and Naemura, T. 2010. Seam carving for stereo images. In 3DTVConference: The True Vision - Capture, Transmission and Display of 3D Video (3DTV-CON), 2010, 1 –4.
4. Kraevoy, V., Sheffer, A., Shamir, A., and Cohen-Or, D. 2008. Non-homogeneous resizing of complex models. ACM Trans. Graph. 27 (December), 111:1–111:9.

Our Approach

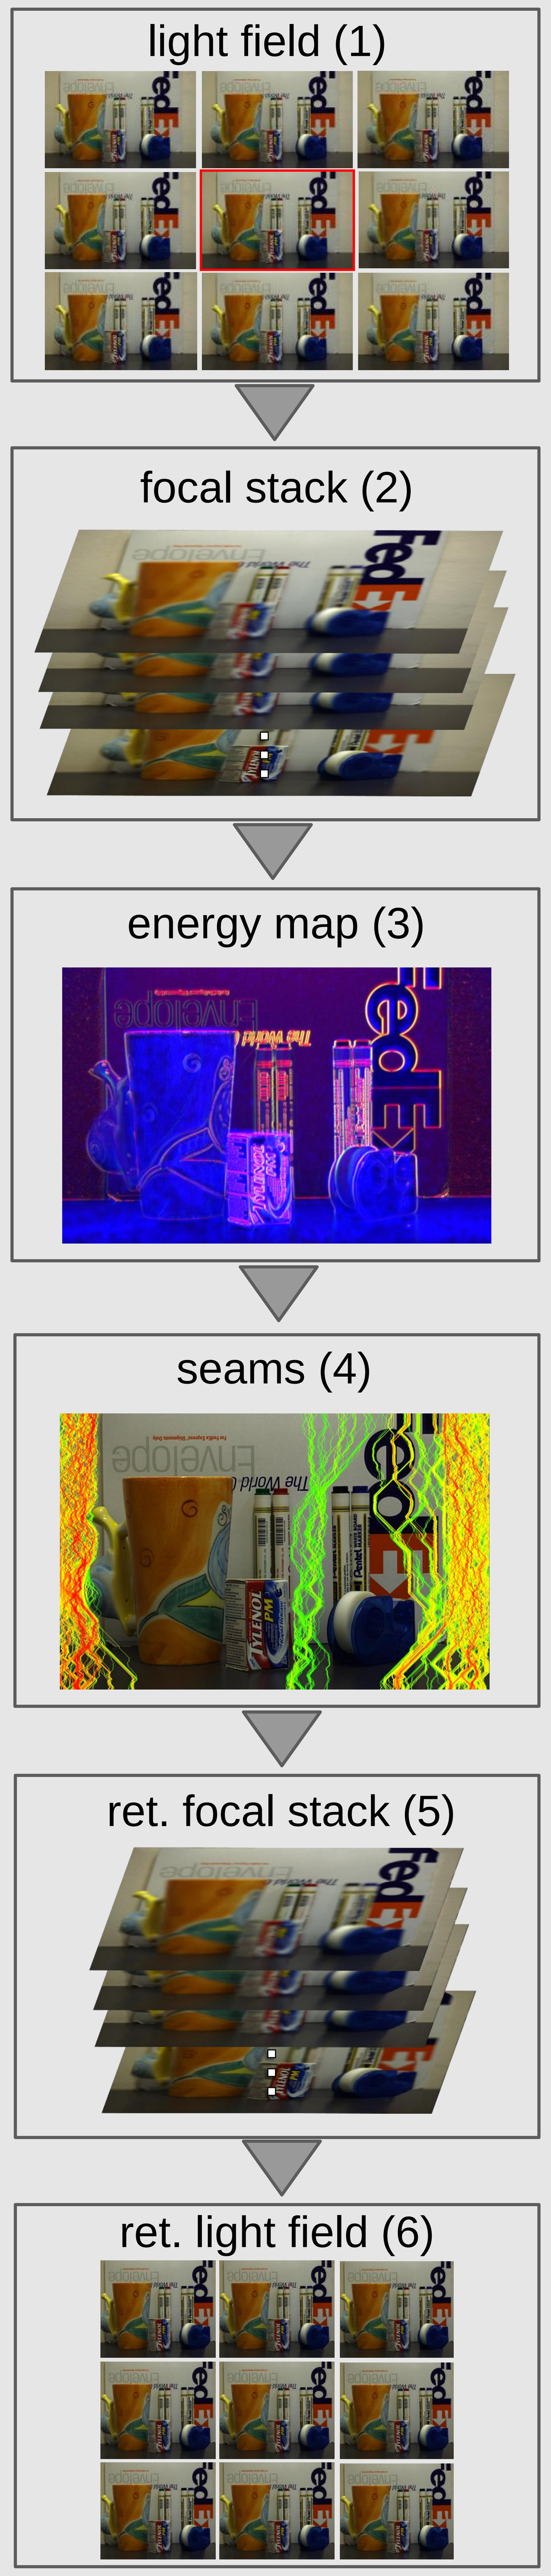

Figure 1: Processing pipeline.

Figure 2: Image rendered from the original (i.e., unmodified) light field.

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.

Then 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).

Additional energy terms are added (3, red and blue) to increase seam carving and reconstruction quality (see next section).

Since the central perspective matches the perspective of the focal stack, we can remove the same calculated seams in all of its slices without violating angular consistency.

The focal stack can then be converted back to a light field (5,6) 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 [1].

Figure 3: Image rendered from the retargeted light field.

Additional Energy Terms

Avoidance of Critical Seems
Removing pixels where seams are cutting through out-of-focus regions in slices of the focal stack, causes deconvolution errors especially around edges during the reconstruction of the light-field perspectives. Therefore, we avoid cutting seams through these critical regions.

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 (red in step 3 of our processing pipeline).

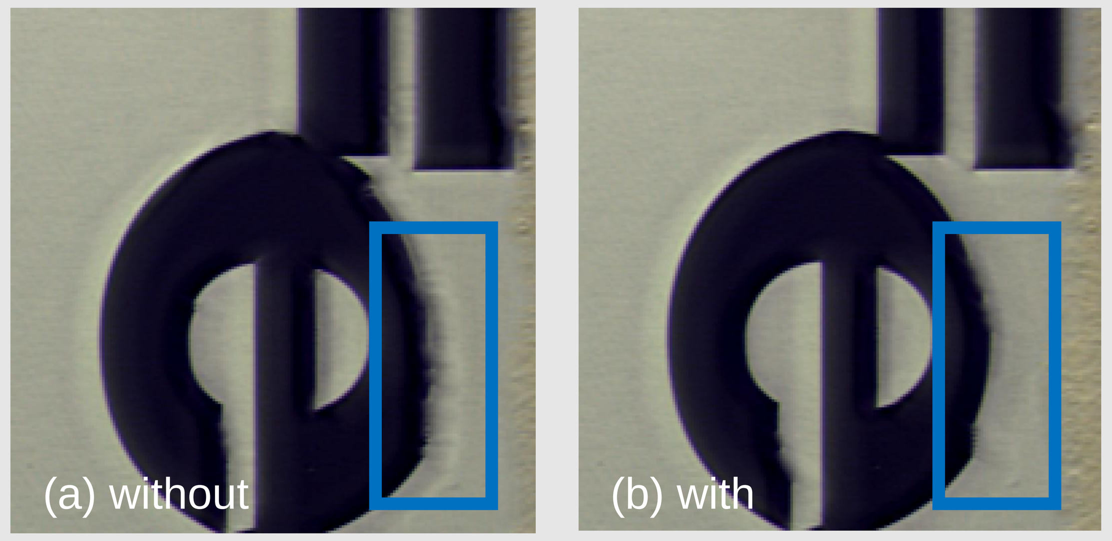

Figure 4: Deconvolution errors around edges (a) are reduced when avoiding critical seams (b).

Avoidance of Foreground Objects
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 (blue in step 3 of our processing pipeline).

Remaining Artifacts

We can not completely avoid deconvolution errors caused by critical seams. Therefore, remaining artifacts are attenuated by convolving the focal stack slices in regions with new gradients that are created by seam carving.

The new gradients are detected by comparing two gradient maps for each focal slice. First a gradient map is calculated before seam carving. This gradient maps is then seam-carved in the same way as its focal slice. After seam-carving a second gradient map is calculated from the retargeted focal slice and compared to the first map.

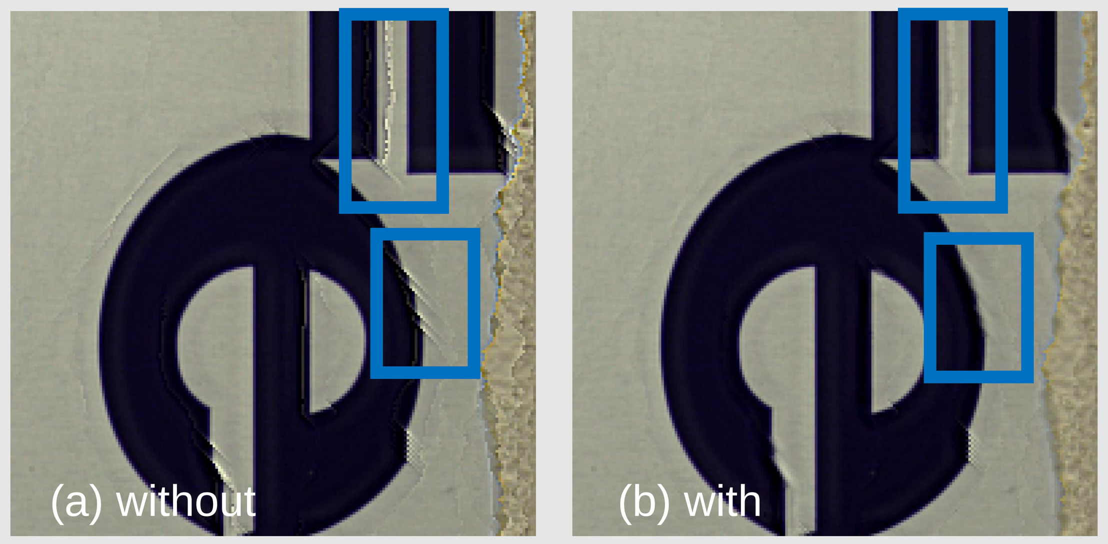

Figure 5: Deconvolution errors (a) are attenuated when convolving regions with new gradients (b).

Results

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.

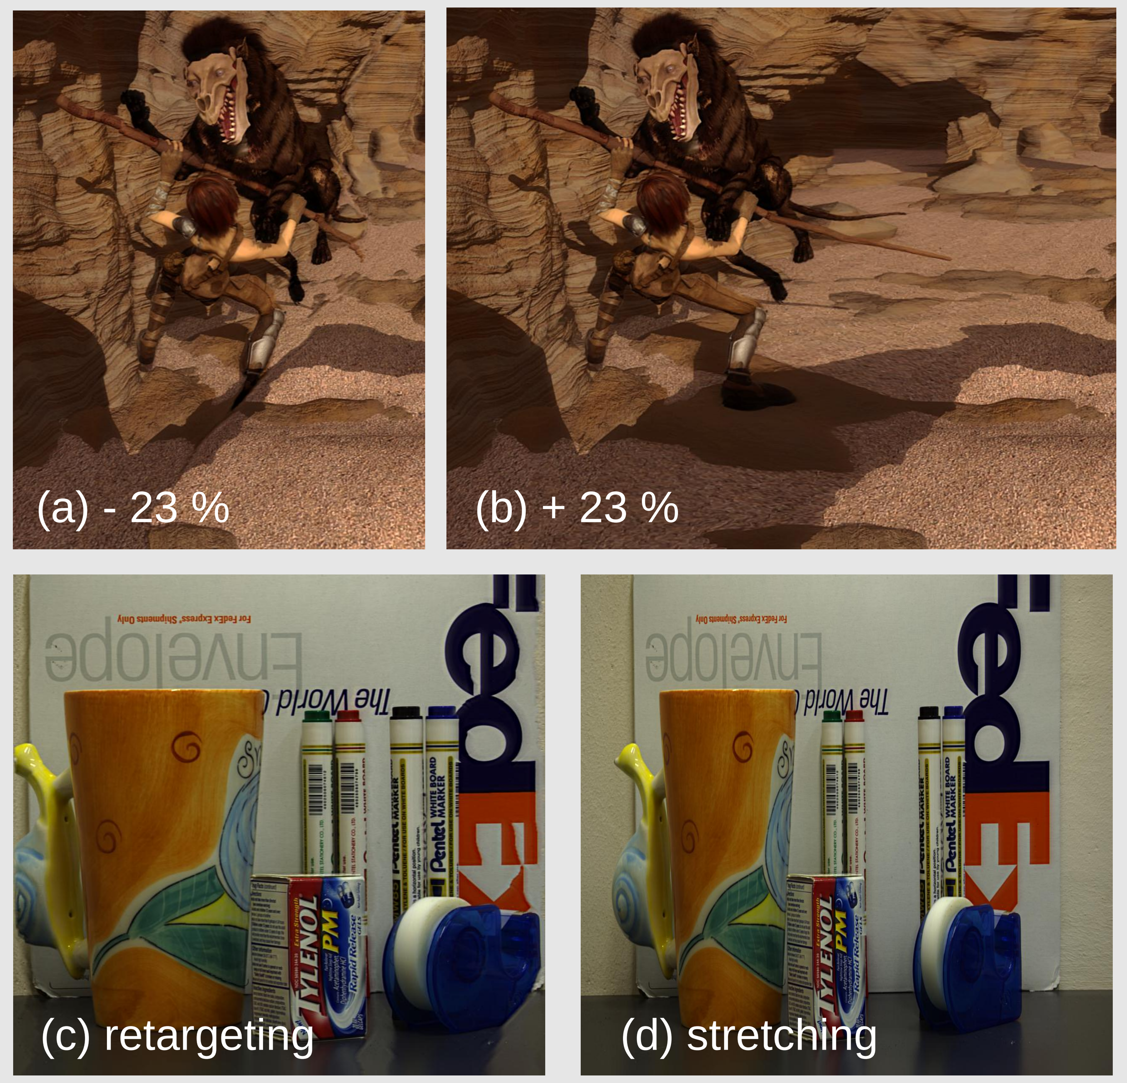

Figure 6: Retargeting results for image reduction (a) and enlargement (b); comparison between retargeting (c) and linear stretching (d).

Limitations

The conversion of focal stacks to light fields, as explained in [1], 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.

Further the deconvolution errors due to critical seams are still visible in some cases. The errors get especially visible in perspectives near the center, because the seams (and therefore the errors created by them) add up at one place as there is nearly no shift between the individual slices.

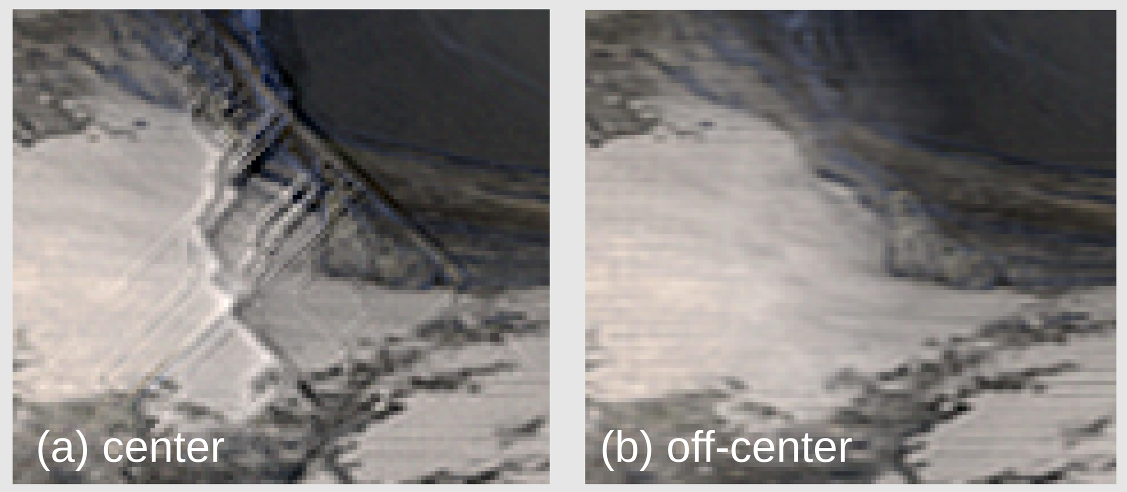

Figure 7: Seams are visible at the central perspective (a) and are distributed over a larger region in off-center perspectives (b).

In future, we want to investigate more advanced techniques for filtering critical seams and for reducing deconvolution errors.

[1] {firstname.lastname}@jku.at

More information at: www.jku.at/cg