

Shadow cameras: Reciprocal views from illumination masks

Sanjeev J. Koppal and Srinivasa G. Narasimhan
Robotics Institute, Carnegie Mellon University, Pittsburgh, USA
Email: (koppal,srinivas)@ri.cmu.edu

Abstract

Scene appearance from the point of view of a light source is called a reciprocal or dual view. Since there exists a large diversity in illumination, these virtual views may be non-perspective and multi-viewpoint in nature. In this paper, we demonstrate the use of occluding masks to recover these dual views, which we term shadow cameras. We first show how to render a single reciprocal scene view by swapping the camera and light source positions. We extend this technique for multiple views by building a virtual shadow camera array with static masks and a moving source. We also capture non-perspective views such as orthographic, cross-slit and a pushbroom variant, while introducing novel applications such as converting between camera projections and removing refractive and catadioptric distortions. Finally, since a shadow camera is artificial, we can manipulate any of its intrinsic parameters, such as camera skew, to create perspective distortions.

1. What are shadow cameras?

Of all the light rays emitted, scattered or reflected by a scene, only the light-field measured by a camera can be directly accessed. In addition, when a programmable source, such as a projector, illuminates the scene, the incident light-field can also be known and controlled. In this paper, we introduce a third set of computable rays: the shadow-field of a non-programmable light source. The shadow-field is captured by using an occluding mask to block rays from the light source. Therefore, our representation is purely geometric and not photometric, since incident light rays are detected by *not* measuring them. Ray reciprocity allows us to describe these rays as if they were viewed by a virtual camera, which we term a **shadow camera**.

Although specific instances of shadow cameras have been used for scene reconstruction (as in [2]), we develop a view synthesis framework for image-based rendering. Figure 1 lists traditional approaches to the IBR problem. Historically, view synthesis involves either camera motion or multiple cameras, while scene relighting requires illumination control (as in a light stage [21]). Dual methods, in contrast, exploit Helmholtz reciprocity and treat light sources as cameras and vice-versa. Shadow cameras utilize reciprocity to allow flexible control of camera pixels, rearranging them in geometries determined by the relative source-mask motion. In this sense, we extend the dual photography technique ([18]) beyond programmable light sources.

Research area	Traditional approaches	Dual approaches
View Synthesis	Move camera (or use multiple cameras)	Vary illumination
Scene Relighting	Vary illumination	Move camera

Figure 1. **Dual approaches for view synthesis and relighting:** Most view synthesis approaches require camera motion or multiple cameras. Similarly, traditional scene relighting techniques involve varying illumination. However, there also exist dual approaches that exploit the Helmholtz reciprocity between viewing and illumination rays. This paper describes a dual technique termed **shadow cameras**, which allows virtual scene views to be obtained by varying illumination with an occluding mask. In Section 6 we briefly introduce a second dual method called **relighting cameras**, to relight a scene using camera motion.

The recovered light ray geometry depends on the mask shape, the light source type and the relative motion between the two. The shadow camera is located at either the mask or the source, depending on which is stationary. In this paper, we focus on the use of linear masks (approximated in practice by thin, rigid wires). Two such masks, perpendicular to each other and moving with uniform motion, produce two distinct intensity minima at each scene point. The times at which these minima occur represent the horizontal and vertical image coordinates of a shadow camera centered at the source.

If a distant light source is used, the shadow camera is **orthographic** in nature and if a near light source is used the shadow camera is **perspective**. Multiple pairs of such perpendicular masks allow the creation of a virtual shadow camera array. Varying the mask speed/angle controls the shadow camera’s intrinsic parameters (skew and scaling). Depending on the type of viewing camera, it becomes possible to switch between orthographic and perspective views.

We can also create shadow cameras with multiple viewpoints. Traditional **cross-slit** cameras are created by imaging through two perpendicular thin slits separated by some distance. Similarly, a cross-slit shadow camera is created when the motion paths of the two linear masks do not intersect and are separated. When compared to mosaicing cross-slits from many perspective views (as in [29]), our technique requires only a single reciprocal pair.

Traditional **push-broom** cameras are created by imaging through a translating slit. Similarly, we concatenate scene points shadowed by a single, translating linear mask to create a ‘shadow-pushbroom’ camera. We also consider scene views either reflected by an unknown mirror surface or refracted through transparent solids. Given a point source illuminant, the computed dual view removes any distortions in the non-single viewpoint dioptric/catadioptric image.

The shadow-camera technique is simple, easy to implement and requires only a background calibration plane to estimate the relative source-mask motion. It is widely applicable since reliable shadow detection is possible for scenes with complex BRDFs. In addition, it is not necessary to store the complete video and instead, we can efficiently detect the moving shadow edge using only a window of a few frames. Finally, the shadow camera resolution depends on the edge detector quality and not on the shadow width.

1.1. Related work

Illumination masking follows a ‘less is more’ trend in vision, where coded camera apertures have been used to capture additional light-field information ([14]). Shadows can also be used for spatio-temporal correspondences across cameras ([5]). Other related work includes:

Light-field rendering and Mosaicing: Light-field rendering ([11],[7]) obtains new views by interpolating between densely sampled images, with applications such as view synthesis ([26]), all-focus images ([13]), seeing past obstacles ([22]) and mosaicing both multi-viewpoint images ([17]) and new cameras (cross-slit ([29]), pushbroom ([25])). Although any shadow camera could be rendered by capturing the entire light field, our approach is more efficient since we only use a single reciprocal pair.

Shadow-based approaches: These techniques have been popular since they are invariant to material properties ([15], [16], [9]). Linear masks have been previously used for scene point triangulation ([2], [3],[10]), while shadow-gram methods reconstruct intricate objects ([24]) and multi-flash approaches obtain occluding edges ([23]). Although we enjoy the advantages of shadows, we differ from these methods since our goal is not scene reconstruction.

Dual views from programmable light sources: Dual photography ([18], [6]) exploits Helmholtz reciprocity ([8], [27]) between every projector-camera pixel pair to render the scene from the projector’s point-of-view. Since the full light transport is measured, any scene relighting technique can be applied. In contrast, our method obtains only a dual scene view. However, it has the advantage of allowing non-programmable illumination and can create multi-viewpoint cameras from a *single* reciprocal pair. To create similar images using [18] we would need a ‘dual video’ created while the projector moves smoothly in a line or plane, which could be prohibitive in terms of time.





Figure 2. **Dual views for non-programmable illumination:** Traditional image-based rendering techniques (IBR) seek to capture the entire light field so that any desired virtual view can be rendered, as in (I). Dual photography (II) reduces the number of images required by placing a projector at the desired viewing location. However, for a single projector position, dual photography is restricted to a perspective view. Our method (III) recovers the dual view by using an illumination mask. Depending on the relative source-mask motion we create shadow cameras that are non-perspective and multi-viewpoint in nature.

2. Light-source centric image-based rendering

Let a scene be illuminated by a source located at $L = (L_x, L_y, L_z)$, and imaged by a pin-hole camera C moving along on a plane Π_1 whose location is given by (u, v) (Figure 2(I)). The moving camera samples the 4D light field at rays specified by their points of intersection, (u, v, s, t) , on Π_1 and a parallel plane Π_2 . To render any new scene view we select light rays that describe the virtual camera’s *caustic*, which is a curve in space that all the light rays must be tangent to ([20]).

For the virtual pin-hole camera C' in the figure, the caustic degenerates to a point in space (the camera center) which lies on a third plane Π_3 parallel to Π_1 , at position (u', v') . The key advantage of traditional camera-centric IBR is that a desired caustic can be created without knowing scene shape. The trade-off is that the entire light field has to be captured, requiring many images. This can be reduced by making scene BRDF assumptions ([19]) or by doing more work to find correspondences across fewer samples by statistical modeling ([28]).

Dual photography solves these issues by replacing the source at (L_x, L_y, L_z) with a camera and placing a projector at the desired virtual camera location (u', v') , as shown in Figure 2(II). If the camera image is given by $I(x, y)$ and a virtual image at the projector is denoted by $I'(x', y')$, then Helmholtz reciprocity relates these as $I'(x', y') = I(x, y)$ where $(x, y) \leftrightarrow (x', y')$ are corresponding projector-camera pixels. Surface normals, required in Helmholtz stereopsis, are not computed since the expression for pixel irradiance $I(x, y)$ contains both illumination and viewing foreshortening (see Appendix at [18]).

Dual photography's limitation is that, for a given projector position, it recovers a single *perspective* view of the scene. To create virtual views for other camera caustics, we have to imitate the traditional IBR setup in Figure 2(I), with a translating projector instead of a moving camera, capturing the virtual image at each projector location. Instead, we wish to more fully and efficiently exploit Helmholtz reciprocity by creating dual views directly for non-perspective, non-programmable light-sources. The major challenge is finding correspondences between the real view and the dual view without the control a projector affords.

2.1. Ray correspondence through shadows

Consider now a non-programmable point light source placed at (u', v') , as in Figure 2(III). We move an opaque mask in front of the light source, creating a shadow that falls on the scene. A scene point P , located in the camera image I at pixel (x, y) , displays a minima in its measured intensity at time t if it is occluded by the mask. Let R_P be the set of all time instances when such intensity minima occur at P . If, for every pair of scene points P and Q , $R_P \cap R_Q = \emptyset$, then R_P uniquely identifies a ray from the light source to P . Therefore each incident light ray at the scene is assigned a unique ID that corresponds to the locations in time when it was occluded.

If we know the caustic of the illuminant (defined similarly to a camera caustic) then we can map the identifier R_P to a pixel on the virtual image I' , $R_P \rightarrow (x', y')$. The intensity at this virtual pixel is found by exploiting Helmholtz reciprocity such that $I'(x', y') = I(x, y)$ for the correspondence $(x, y) \leftrightarrow (x', y')$ between real camera pixels and the virtual camera pixels, *exactly* as in dual photography.

Finding the mapping R_P can be done in two ways. The first is to design the mask motion such that the minima location gives the virtual image pixel directly. Since the pixel location has two degrees of freedom, we need at least two minima locations in time to do this. In Figure 3(I) we show a ray diagram of the shadow cast by a linear mask translating in front of a static point light source. Since the shadow hull of a line mask is a plane, the locations in time of two intersecting perpendicular plane shadow hulls specifies the horizontal and vertical coordinates of a virtual image pixel.






Figure 3. **Virtual perspective view using linear masks:** In (I) we show a ray diagram for a perspective shadow camera. Using linear masks we obtain the input images, (III), which create a dual view in (IV). Note the foreshortening effects of looking down at the object (shortened legs, extended scales), that specularly locations do not change and that shadows in (II) are occluded in (IV).

The second method to find R_P generalizes beyond linear masks but requires a simple calibration step. The experiment must be repeated twice, first with the actual scene and the second time with a plane placed at the location where we wish the virtual image plane to be. If the plane is visible to the camera and the motion of the mask is identical in each experiment, then we can map each scene point to a coordinate on calibration plane.

The minima detection is independent of both BRDF variation and intensity-fall off and allows cast/attached shadow disambiguation. The virtual camera resolution depends on the shadow edge detector and not the shadow width and is theoretically only limited by the camera resolution. Our method also handles multiple sources since these would create many intensity minima scene points, each of which gives a different view point. Similarly, we can detect a ‘minima interval’ for area sources corresponding to a continuous set of viewpoints.

3. Perspective shadow cameras

Our experiments were performed using a 12 bit Canon XL2 video camera running at 30fps. We selected a 1mm diameter piano wire for the mask, held by a Manfrotto tripod that allows controlled height adjustment and our scenes were illuminated by a Lamina ceramics DK4 LED. We placed a plane behind the scene, to estimate the shadow locations and a uniform interpolation of these obtains proper sampling for the virtual image.

In Figure 3 (I) we show a ray diagram for a perspective shadow camera. I(a) and I(b) show two translating linear masks occluding P at times i and j respectively. The intersection of the shadows associated with the two masks is a virtual ray, shown at I(c). This shadow ray passes through the light source, and is associated with a pixel (i,j) in a virtual image at I(d). In (II) we show a non-convex plastic object. Using linear masks we obtain the input images, shown at (III). In Figure 3 (IV) we show a dual view for a non-convex plastic toy obtained using a set of two linear masks. Note the viewing foreshortening effects since the light source is higher than the camera position, such as shortening of the legs and the extension of the scales. In addition, the shadows of the dinosaur’s tail are occluded in the dual image and vice-versa.

Multiple shadow cameras are possible by moving the source and performing experiments at the new light locations. We also demonstrate a different setup, where the mask is static and the light source moves. For example, in Figure 4 (I) we show a set of horizontal and vertical linear masks that are illuminated by a light source moving in a line. The shadows of the mask can be seen in Figure 4 (II) and horizontal and vertical experiments are performed separately to enable easy minima detection. By recording the times at which the horizontal and vertical masks occlude every scene point, we can recover virtual shadow cameras at the intersection of the horizontal and vertical masks.

Extra steps are required when moving the source instead of the mask. Since the scene intensities vary, shadow detection becomes harder. We also need additional images taken with the light source placed at each of the mask’s grid intersections. However, the advantage is that a single experiment produces multiple viewpoints: more precisely, if m and n

are the number of vertical and horizontal masks, we obtain an mn virtual array of shadow cameras. Six such images are shown in Figure 4 (III) and (IV), each of which are rectified horizontally and vertically since the wires in the mask array are perpendicular to each other.

4. Multi-view and non-perspective shadow cameras

An indication of the potential of shadow cameras is that even for the limited case of linear masks and point light source discussed in this paper, we demonstrate many multi-view and non-perspective cameras. Figure 5 (I) shows that a ray diagram describing orthographic virtual view can be obtained from a perspective viewing camera by using a distant light source. In Figure 5 (II) we show a real view of plastic toy shark under distant (orthographic) illumination. Note the specularities on the shark and the shadows are correctly demonstrated in the dual view in Figure 5 (III). To validate this view, we co-located the light-source and a camera using a half-mirror. Note the image in Figure 5 (IV) has no shadows. The shape of the shark is qualitatively similar in Figure 5 (III) and (IV). However, the appearance of the object is different since co-location is not the same as switching the source and camera positions. Despite this, the comparison demonstrates the correctness of the dual view.

Shadow cameras also allow us to demonstrate the opposite effect by switching from an orthographic view of the scene to a perspective view as shown in Figure 6. The perspective distortions are clear in the dual views, such as the change in angles of colored squares away from 90 degrees and the foreshortening when looking down at the octopus. The other aspect of the switch is that the illumination also changes: for example, in the dual view of the toy octopus the shadows are smaller, due to the orthographic source.

A multi-perspective cross-slit image of the scene can be created if the two linear masks do not intersect at a point, and are instead shifted by some amount, as in Figure 7 (II). Mathematically, this is identical to having an image plane intersected by rays passing through two slits if the light source is moved between experiments. In Figure 7 (III) we show such an image: note that the head of the octopus appears ‘unwrapped’ as if we are viewing all around it while the colored squares are curved.

We also introduce a new camera view inspired by the real pushbroom camera, which is created by imaging a translating slit. In our case we take the scene points corresponding to the shadow of a single translating linear mask and concatenate them together, as illustrated in Figure 7 (IV). The image, shown in Figure 7 (V), appears to be lit with sources at both the camera and the light source positions. Note the double shadows of each of the octopus tentacles. We call this camera the ‘shadow pushbroom’ camera.

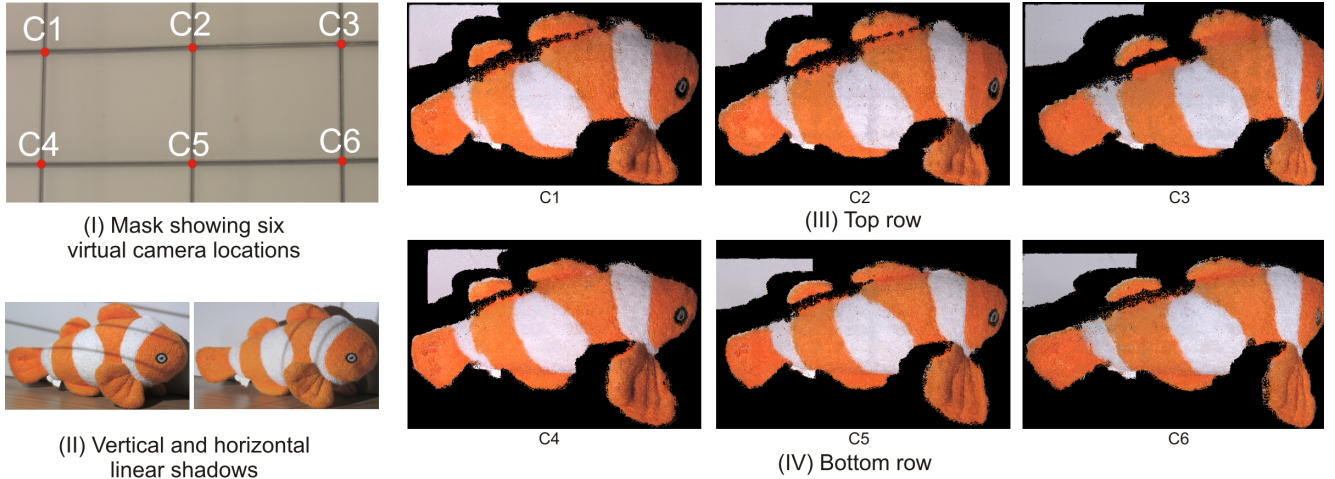


Figure 4. **Virtual camera array from multiple linear masks:** In (I) we show a grid mask with six intersections. We collect data using the horizontal and vertical masks separately, as shown in (II). We then place light sources at the six intersections of the mask, capturing six images. We swap these images using Helmholtz reciprocity and the minima locations of the mask shadows. The result is a virtual camera array, and in (III) and (IV) we show the two rows of the array. Note the images are stereo rectified both horizontally and vertically.

Since shadow cameras are completely virtual, we can change their intrinsic parameters such as image skew or scaling. In Figure 8 we describe some of these perspective distortions applied to a scene. In Figure 8 (I) we show the dual view when the linear masks are perpendicular and move with uniform velocity, showing no distortions. We performed a third shadowing experiment with a slanted mask, which created virtual images with a non-zero pixel skew as shown at the left of Figure 8 (II). We also repeated the experiment with a horizontal mask moving slower than the vertical mask, resulting in higher horizontal sampling and a stretched image at the right of Figure 8 (II), which can be seen in the left dinosaur.

5. An imaging application: Seeing through refraction and reflection

Optical elements, such as reflective surfaces and refractive solids, bend light and are designed into many imaging systems to create multiple new viewpoints (as in catadioptric cameras ([20]) and fish-eye lenses). While these cameras are useful tools for many applications, the images produced are considered distorted when viewed directly. Similar distortions are also present in images of everyday scenes with objects like water drops and glossy surfaces inadvertently behaving as lenses and mirrors. Unwarping these effects in images without any prior knowledge of the geometry of the optical elements is almost impossible.


However, consider a source that illuminates the scene such that the light only travels through free space, without passing through any transparent occluder or reflector. Applying our shadow camera method removes all distortions rendering a warp free image of the scene since the shadow

correspondences map the viewing rays bent by the optical elements to the light rays incident on to the scene. This mapping is very similar to recovering the light-transport as an environment matte ([4]).

We demonstrate unwarping of both reflective and refractive distortions with no prior scene knowledge. In Figure 9 (II), objects are reflected off a spherical mirror and the straight lines of the colored squares appear curved. As the light source is perspective, these distortions are removed in the dual views in Figure 9 (III). In Figure 10 we show a planar poster viewed through thick glass objects. Note the heavy distortion, especially in the glass on the right. Figure 10(II) shows the undistorted view, with readable text.

6. Discussion and Future work

We have demonstrated a reciprocal approach by creating novel views from the relative motion between non-programmable sources and occluding masks. When compared to light-transport methods like [18], we only recover the matrix’s diagonal as the dual view. For future work, we wish to look at the intensity minima at a scene point created by occluding other parts of the scene. We believe this would allow the recovery of more than just the light-transport diagonal and is related to fast separation of global illumination ([12]), where shadows are used to extract off-diagonal information in the light-transport matrix. Finally, the inverse of our technique could relight scenes with static light sources and moving cameras. For example, a translating camera can generate epipolar plane images ([1]) with dense correspondences. We can relight any image in the sequence by replacing its pixels with corresponding pixels from other frames. From Helmholtz reciprocity, such images are identical to those created by moving the light source. We term



(II) Original perspective view with orthographic light-source




(III) Dual orthographic view illuminated by perspective source



(IV) Real orthographic view by co-locating source and camera

Figure 5. Validating the dual view by perspective to orthographic conversion: In (I) we show how the dual view due to a distant light source is orthographic. In (II) we show the perspective view of a plastic shark illuminated by a distant light source. Note the perspective effects since the shark's snout is close to the camera. In (III) we show the dual orthographic view. This compares reasonably with (IV) where we co-located the light source and the camera using a half-mirror (note the lack of shadows). The co-location does not move the light-source, explaining the difference in specularities between (III) and (IV).



(I) Orthographic view/near light to perspective view/distant light

Figure 6. Orthographic to Perspective conversion: In left column we show the original images viewed by an orthographic camera. In the right column, the dual views have foreshortening effects associated with perspective cameras.

this dual method as **relighting cameras** and show a planar example in Figure 11, where we change a specularity location with a translating camera. The challenge here will be finding correspondences for non-linear camera motions.

7. Acknowledgements

This work was supported in parts by ONR grants N00014-08-1-0330 and DURIP N00014-06-1-0762 and an NSF CAREER Award IIS-0643628.

References

- [1] R. Bolles, H. Baker, and D. Marimont. Epipolar-plane image analysis: An approach to determining structure from motion. *IJCV*, 1987. 5
- [2] J. Bouguet and P. Perona. 3d photography on your desk. *ICCV*, 1998. 1, 2
- [3] Y. Caspi and M. Werman. Vertical parallax from moving shadows. *CVPR*, 2006. 2
- [4] Y. Chuang, D. Zongker, J. Hindorff, B. Curless, D. Salesin, and R. Szeliski. Environment matting extensions: Towards higher accuracy and real-time capture. *SIGGRAPH*, 2000. 5
- [5] J. Davis, D. Nehab, R. Ramamoorthi, and S. Rusinkiewicz. Spacetime stereo : A unifying framework for depth from triangulation. *CVPR*, 2003. 2
- [6] G. Garg, E.-V. Talvala, M. Levoy, and H. P. A. Lensch. Symmetric photography: Exploiting data-sparseness in reflectance fields. *EGSR*, 2006. 2
- [7] S. Gortler, R. Grzeszczuk, R. Szeliski, and M. Cohen. The lumigraph. *SIGGRAPH*, 1996. 2
- [8] H. V. Helmholtz. Treatise on physiological optics. *Dover*, 1925. 2

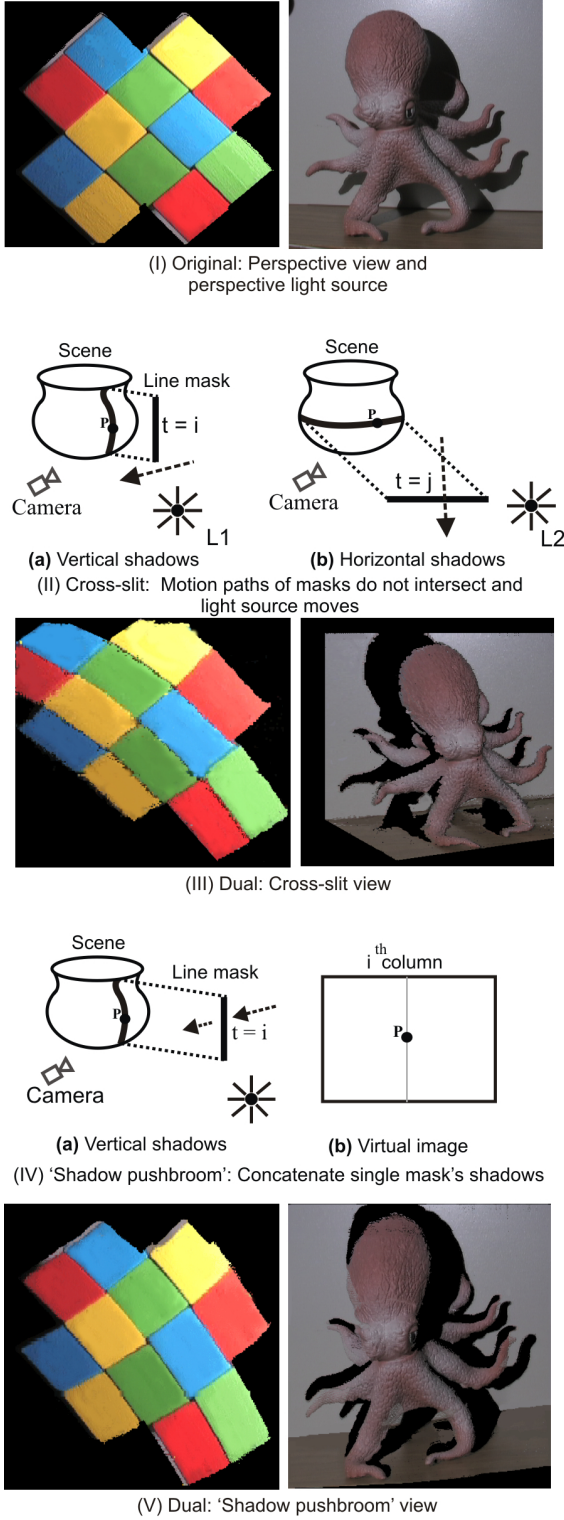


Figure 7. **Non-perspective and multi-view cameras:** In (II) we show a cross-slit view created when the motion vectors of the linear masks do not intersect and the light source is moved between experiments. We demonstrate the distortion effects in (III). In (IV) we create a pushbroom-like view using the shadow from a single moving mask. Note the octopus in (V) appears illuminated by two light sources and each tentacle has two shadows.






Figure 8. **Controlling intrinsic parameters of shadow cameras:** Since shadow cameras are virtual, we have control over the locations of the pixels in space. In (I) we show a dual view created by two linear masks moving with uniform motion at 90 degrees to each other. Varying the speed of these two perpendicular masks or adding a third linear mask that is not perpendicular to either of the first two results in skewed and stretched images as seen in (II).

- [9] M. S. Langer, G. Dudek, and S. W. Zucker. Space occupancy using multiple shadowimages. *IROS*, 2005. 2
- [10] D. Lanman, R. Raskar, A. Agrawal, and G. Taubin. Shield fields: modeling and capturing 3d occluders. *Siggraph Asia*, 2008. 2
- [11] M. Levoy and P. Hanrahan. Light field rendering. *SIGGRAPH*, 1996. 2
- [12] S. Nayar, G. Krishnan, M. D. Grossberg, and R. Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *Siggraph*, 2006. 5
- [13] R. Ng. Fourier slice photography. *Siggraph*, 2005. 2
- [14] R. Raskar, A. Agrawal, and J. Tumblin. Coded exposure photography: Motion deblurring using fluttered shutter. *TOG*, 2006. 2
- [15] D. Raviv, Y. Pao, and K. Loparo. Reconstruction of three-dimensional surfaces from two-dimensional binary images. *Transactions on Robotics and Automation*, 1989. 2
- [16] S. Savarese, M. Andreetto, H. Rushmeier, and F. Bernardini. 3D reconstruction by shadow carving: Theory and practical evaluation. *IJCV*, 2005. 2
- [17] S. Seitz. The space of all stereo images. *ICCV*, 2001. 2
- [18] P. Sen, B. Chen, G. Garg, S. R. Marschner, M. Horowitz, M. Levoy, and H. P. A. Lensch. Dual photography. *SIGGRAPH*, 2005. 1, 2, 3, 5
- [19] H. Shum, S. Chan, and S. B. Kang. Image-based rendering. *Springer*, 2007. 2
- [20] R. Swaminathan, M. Grossberg, and S. K. Nayar. Caustics of catadioptric cameras. *ICCV*, 2001. 2, 5
- [21] J. Unger, A. Wenger, T. Hawkins, A. Gardner, and P. Debevec. Capturing and rendering with incident light fields. *EGRS*, 2003. 1



(I) Undistorting catadioptric camera reflections:
Dual view is perspective and distortion-free




(II) Original images: Distorted view
reflected on curved mirror

(III) Duals: Perspective view
with no distortions

Figure 9. **Unwarping catadioptric distortions:** (I) shows the ray diagram for unwarping mirror distortions. In (III) we show unwarped results: note the colored squares' lines are straight. The non-shadow holes are due to stereo occlusion.


- [22] V. Vaish, R. Szeliski, C. L. Zitnick, S. B. Kang, and M. Levoy. Reconstructing occluded surfaces using synthetic apertures: Stereo, focus and robust measures. *CVPR*, 2006. 2
- [23] D. Vaquero, R. Feris, M. Turk, and R. Raskar. Characterizing the shadow space of camera-light pairs. *CVPR*, 2008. 2
- [24] S. Yamazaki, S. G. Narasimhan, S. Baker, and T. Kanade. Coplanar shadowgrams for acquiring visual hulls of intricate objects. *ICCV*, 2007. 2
- [25] J. Yu and L. McMillan. General linear cameras. *ECCV*, 2004. 2
- [26] C. Zhang and T. Chen. A self-reconfigurable camera array. *EGSR*, 2004. 2
- [27] T. Zickler, P. Belhumeur, and D. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *ECCV*, 2002. 2
- [28] C. Zitnick, S. B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski. High-quality video view interpolation using a layered representation. *Siggraph*, 2004. 2
- [29] A. Zomet, D. Feldman, and S. Peleg. Mosaicing new views: The crossed-slits projection. *PAMI*, 2003. 1, 2



(I) Undistorting transparent object refractions:
Dual view is perspective and distortion-free



(II) Original image showing distortions



(III) Dual view with readable text

Figure 10. **Unwarping refractive distortions:** In (I) we show how the dual view of an object looking through refractive elements can be warp-free. This occurs when the light-source illuminates the object without being occluded by the refractive elements. In (II) we show an image of planar poster through two thick glass objects. (III) shows the dual view, which has undistorted, readable text.



(I) Uniform vertical
camera motion

(II) Original view
from I(b)

(III) Relit with
illumination from I(a)

Figure 11. **Relighting cameras:** In (I) we show the first and last images of a camera moving vertically. In (III), using the EPI images of this object, we relight the image at I(b) with pixels from the image taken at I(a). From Helmholtz reciprocity, this is equivalent to moving the light source. Note the position of the specularly in (III) has changed from its original location in (II).