

Recording and controlling the 4D light field in a microscope using microlens arrays

M. LEVOY*, Z. ZHANG* & I. MCDOWALL†

*Computer Science Department, Stanford University, Stanford, California, U.S.A.

†Fakespace Labs, 241 Polaris Ave., Mountain View, California, U.S.A.

Key words. Illumination, light field, light field microscope, microlens array, plenoptic function, synthetic aperture imaging, spatial light modulator, spherical aberration, Shack–Hartmann sensor.

Summary

By inserting a microlens array at the intermediate image plane of an optical microscope, one can record four-dimensional light fields of biological specimens in a single snapshot. Unlike a conventional photograph, light fields permit manipulation of viewpoint and focus after the snapshot has been taken, subject to the resolution of the camera and the diffraction limit of the optical system. By inserting a second microlens array and video projector into the microscope's illumination path, one can control the incident light field falling on the specimen in a similar way. In this paper, we describe a prototype system we have built that implements these ideas, and we demonstrate two applications for it: simulating exotic microscope illumination modalities and correcting for optical aberrations digitally.

Introduction

The light field is a four dimensional (4D) function representing radiance along rays as a function of position and direction in space. Over the past 10 years our group has built several devices for capturing light fields (Levoy, 2000; Levoy, 2005; Ng, 2005b; Wilburn *et al.*, 2002). In particular, Ng *et al.* is a handheld camera in which a microlens array has been inserted between the sensor and main lens. A photograph taken by this camera contains a grid of circular subimages, one per microlens. Each subimage records one point in the scene, and within a subimage each pixel records one direction of view of that point. Thus, each pixel in the photograph records the radiance along one ray in the light field.

Recently, we have shown that by inserting a microlens array at the intermediate image plane of an optical microscope, we can capture light fields of biological specimens in the

same way. We call this a light field microscope (LFM) (Levoy, 2006). From the image captured by this device, one can employ light field rendering (Levoy & Hanrahan, 1996) to generate oblique orthographic views or perspective views, at least up to the angular limit of rays that have been captured. Since microscopes normally record orthographic imagery, perspective views represent a new way to look at specimens. Figure 1 shows three such views computed from a light field of mouse intestine villi.

Starting from a captured light field, one can alternatively use synthetic aperture photography (Isaksen *et al.*, 2000; Levoy *et al.*, 2004) to produce views focused at different depths. Two such views, computed from a light field of Golgi-stained rat brain, are shown in Fig. 10. The ability to create focal stacks from a single input image allows moving or light-sensitive specimens to be recorded. Finally, by applying 3D deconvolution to these focal stacks (Agard, 1984), one can reconstruct a stack of cross-sections, which can be visualized using volume rendering (Levoy, 1988).

Summarizing, the LFM allows us to capture the 3D structure of microscopic objects in a single snapshot (and therefore at a single instant in time). The sacrifice we make to obtain this capability is a reduction in image size. Specifically, if each microlens subimage contains $N \times N$ pixels, then our computed images will contain N^2 fewer pixels than if the microlenses were not present. In return, we can compute N^2 unique oblique views of the specimen, and we can generate a focal stack containing N slices with non-overlapping depths of field (Levoy, 2006). Note that this trade-off cannot be avoided merely by employing a sensor with more pixels, because diffraction places an upper limit on the product of spatial and angular bandwidth for a given aperture size and wavelength, regardless of sensor resolution. Despite this limit, light fields contain much useful information that is lost when an object is photographed with an ordinary microscope.

While technologies for recording light fields have existed for more than a century, technologies for generating light




Fig. 1. Three oblique orthographic views from the indicated directions, computed from a light field captured by our LFM. The specimen is a fixed whole mount showing *L. monocytogenes* bacteria in a mouse intestine villus 7 h post infection. The bacteria are expressing green fluorescent protein, and the actin filaments in the brush border of the cells covering the surface of the villus are labelled with rhodamine-phalloidin. Scale bar is 10 μm . Imaging employed a 60 \times / 1.4NA oil objective and an $f/30$ microlens array (see section 'Prototype light field illuminator and microscope'). Contrast of the bacteria was enhanced by spatially masking the illumination as described in Fig. 7. The inset at left shows a crop from the light field, corresponding to the square in the first oblique view. In this inset, we see the circular subimages formed by each microlens. The oblique views were computed by extracting one pixel from each subimage – a pixel near the bottom of each subimage to produce the leftmost view and a pixel near the top to produce the rightmost view.

fields have been limited until recently by the low resolution and high cost of spatial light modulators (SLMs). With the advent of inexpensive, high resolution liquid crystal displays and digital micromirror devices, interest in this area has burgeoned. In microscopy, SLMs have been used to control the distribution of light in space (Hanley *et al.*, 1999; Smith *et al.*, 2000; Chamgoulov *et al.*, 2004) or in angle (Samson & Blanca, 2007). The former is implemented by making the SLM conjugate to the field of the microscope, and the latter by making it conjugate to the aperture. Systems have also been proposed for manipulating the wave front using an SLM placed in the aperture plane (Neil *et al.*, 2000). However, in these systems the illumination must be coherent, a significant limitation.

As an extension of our previous work, we show in this paper that by inserting a microlens array at the intermediate image plane of a microscope's illumination path, one can programmatically control the spatial and angular distribution of light (i.e. the 4D light field) arriving at the microscope's object plane. We call this a light field illuminator (LFI). Although diffraction again places a limit on the product of spatial and angular bandwidth in these light fields, we can, nevertheless, exercise substantial control over the quality of light incident on a specimen.

In particular, we can reproduce exotic lighting effects such as darkfield, oblique illumination, and the focusing of light at planes other than where the microscope is focused. In addition, by generating structured light patterns and recording the appearance of these patterns with our LFM after they have passed through a specimen, we can measure the specimen's index of refraction, even if this index changes across the field of view. In this application we are essentially using the LFI

as a 'guide star' and the LFM as a Shack–Hartmann sensor. Finally, we can use this information to correct digitally for the optical aberrations induced by these changes in index of refraction.

The four-dimensional light field

We begin by briefly reviewing the theory of light fields. In geometrical optics, the fundamental carrier of light is a ray (Fig. 2(a)). The radiance travelling along all such rays in a region of 3D space illuminated by an unchanging arrangement of lights has been dubbed the plenoptic function (Adelson & Wang, 1992). Since rays in space can be parametrized by three coordinates and two angles as shown in Fig. 2(b), the plenoptic function is 5D. However, if we restrict ourselves to regions of space that are free of occluders, then this function contains redundant information, because the radiance of a ray remains constant from point to point along its length. In fact, the redundant information is exactly one dimension, leaving us with a 4D function historically called the light field (Gershun, 1936).

Although the 5D plenoptic function has an obvious parametrization, the 4D light field can be parametrized in a variety of ways (Levoy, 2006c). In this paper, we parametrize rays by their intersection with two planes (Levoy & Hanrahan, 1996) as shown in Fig. 2(c). Although this parametrization cannot represent all rays – for example rays parallel to the two planes if the planes are parallel to each other – it has the advantage of relating closely to the analytic geometry of perspective imaging. Indeed, a simple way to think about a two-plane light field is as a collection of perspective views of




Fig. 2. Plenoptic functions and light fields. (a) The radiance L of a ray is the amount of light travelling along all possible straight lines through a tube of a given solid angle and cross-sectional area. The units of L are watts (W) per steradian (sr) per meter squared (m^2). (b) In 3D space, this function is 5D. Its rays can be parametrized by three spatial coordinates x , y , and z and two angles θ and ϕ . (c) In the absence of occluders, the function becomes 4D. Shown here is one possible parametrization, by pairs of intersections (u , v) and (s , t) between rays and two planes in general position.

the st plane (and any objects that lie beyond it), each taken from an observer position on the uv plane.

One important extension of the two-plane representation is that one of the planes may be placed at infinity. If for example we place the uv plane at infinity, then the light field becomes a collection of orthographic views of the st plane, where each view is composed of parallel rays having fixed direction uv . This parametrization is similar to the ray matrices used in geometrical optics (Halbach, 1964), where s and t represent position on some imaging surface and u and v represent ray direction, which is expressed as a plane intercept in the paraxial approximation. We will need this extension when we consider microscopes, which record orthographic rather than perspective imagery due to the presence of a telecentric stop (Kingslake, 1983) one focal distance behind the objective's second principal point.

Light field rendering and focusing

Whether one of the two planes is placed at infinity or not, we can treat the samples in a two-plane light field as an abstract 4D ray space indexed by u , v , s and t . If we extract from this space a 2D slice parallel to the st axes, we obtain a pinhole view of the scene taken from one of the original observer positions. For the easier-to-draw case of a 2D ('flatland') light field, this procedure is diagrammed in Fig. 3(a). More interestingly, if we extract a 2D slice that is not axis-aligned, we obtain a novel view, i.e. from an observer position not contained in the original data set (Fig. 3(b)). We call this technique light field rendering (Levoy & Hanrahan, 1996), and it has found numerous applications in computer graphics and computer vision.

Alternatively, if we compute a projection of the 4D space perpendicular to its st axes, i.e. by adding together samples of constant u and v , then we obtain a non-pinhole view. Such a view is focused on the st plane in 3D and has an aperture equal in extent to the uv plane (Fig. 3(c)). We call this *synthetic*

aperture photography (Levoy *et al.*, 2004). Finally, if we shear the 4D space parallel to its st axes before summing it (Fig. 3(d)), then we move the plane of focus perpendicular to the st plane (Isaksen *et al.*, 2000; Ng, 2005a). If we shear the array in other directions in 4D, we obtain a tilted focal plane (Vaish *et al.*, 2005). Curved focal surfaces are also possible, if we warp the space non-linearly.

Light field illumination

The algorithms described in the last two paragraphs for generating views from light fields can also be applied to shaping illumination. Consider the 4D light field falling on a wall after passing through an aperture. Let the wall be the st plane and the aperture be the uv plane. As before, indices u , v , s and t form an abstract 4D ray space. If we illuminate only those rays lying on a 2D slice parallel to the st axes, we produce the appearance of a point light source located somewhere inside the boundaries of the aperture. If the uv plane (hence the aperture) is placed at infinity, the light source becomes collimated, similar to Fig. 3(a). If the slice is not axis-aligned in 4D ray space, the light source becomes either converging or diverging (as in Fig. 3(b)).

Alternatively, if we assign a pattern to a slice that is parallel to the st axes, and we extrude the pattern in 4D perpendicularly to the st plane, then we produce a finite-aperture image of the pattern focused on the wall (as in Fig. 3(c)). Finally, if we shear or warp the 4D ray space after performing the extrusion, we can focus the pattern at a different depth (like 3(d)), or onto a tilted plane or curved surface – without moving any optics. We call this *synthetic aperture illumination* (Levoy *et al.*, 2004).

Prototype light field illuminator and microscope

To explore these ideas, we have built a system that combines a LFI and LFM. Our design is a standard optical microscope in

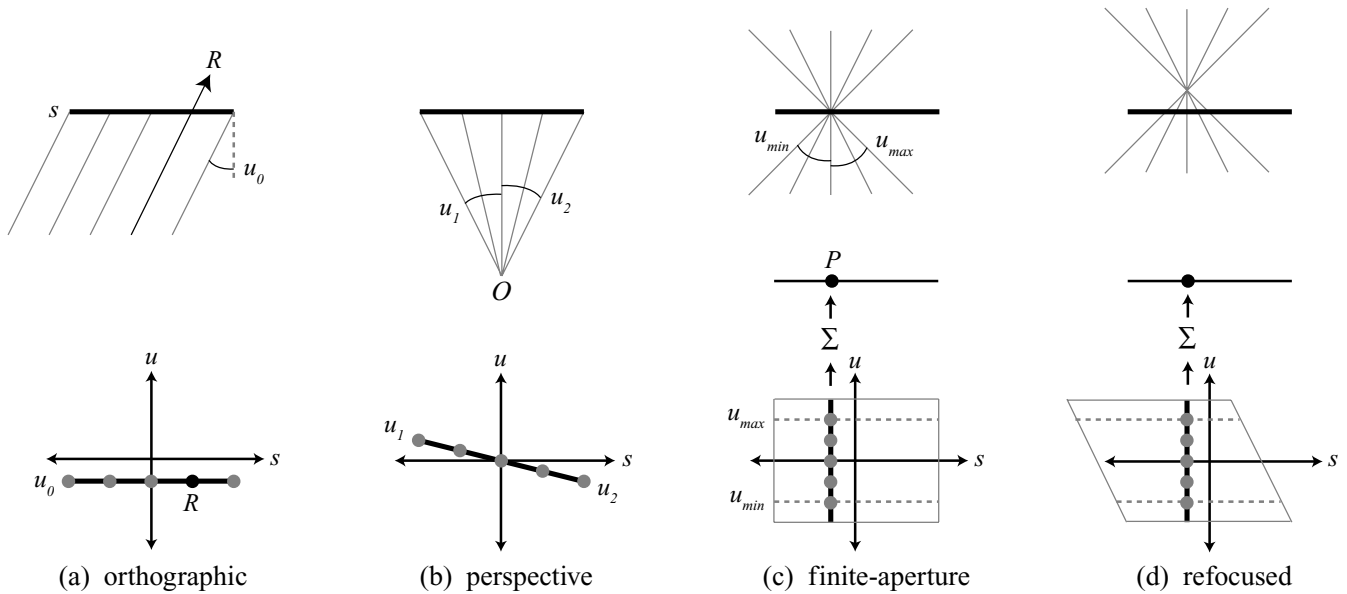


Fig. 3. Light field rendering and focusing in flatland. In a 2D world, rays (like r) are parametrized by their intersection with two lines u and s . To represent the situation in a microscope, we place the u line at infinity as described in the text. (a) A bundle of parallel rays with direction u_0 (upper diagram) constitutes a 1D pinhole orthographic view of the s line. In the abstract ray space defined by u and s (lower diagram), this view corresponds to a horizontal line at u_0 . Each point on this line denotes one ray. (b) Rays diverging from a point (like O) constitute a perspective view. These correspond to a tilted line in ray space, comprising rays having directions from u_1 to u_2 . (c) Summing part or all of a column in ray space produces one pixel (like P) in a non-pinhole view focused on the s line. Its aperture spans ray directions from u_{\min} to u_{\max} . (d) Shearing the ray space leftward in s with respect to u , then summing as before, we produce a view focused above the s line, or above the object plane if this were a microscope. A rightward shear would produce a view focused below the s line. In (c)–(d), the bundle of rays constitutes the PSF of a synthetic aperture. In a microscope this PSF is spatially invariant, i.e. the vertical integral for every s spans the same range of directions in u .

which microlens arrays have been inserted at the intermediate image planes of the illumination and imaging paths. In the illumination path this plane is imaged by a video projector and in the imaging path by a digital camera. Since we need optical quality in both paths, we chose epi-illumination for our prototype rather than transmitted light, since the latter would require building a custom microscope having objectives above and below the stage. Epi-illumination also facilitates fluorescence studies, an important application as we shall see. The microlens arrays in the illumination and imaging paths are similar but distinct. In fact, their optical recipes may differ, permitting different spatio-angular trade-offs in lighting and viewing.

Optical layout

Fig. 4 shows our optical layout. Our SLM is a Texas Instruments Digital Light Processor (DLP) having 1024×768 pixels, of which our prototype uses the innermost 800×600 . When imaged onto light field plane Lf_{LFI} in Fig. 4(b), this produces a $12 \text{ mm} \times 9 \text{ mm}$ field. The illumination microlens array at Π'_{LFI} is a standard part made by Adaptive Optics. Each microlens is a square-sided, planoconvex lenslet with a pitch of $300 \mu\text{m}$ and a focal length of 3.0 mm or 7.6 mm . This

gives us an addressable spatial resolution of $12 \text{ mm}/300 \mu\text{m} = 40$ tiles horizontally and $9 \text{ mm}/300 \mu\text{m} = 30$ tiles vertically. The angular resolution within each tile is $800 \text{ pixels}/40 = 20$ addressable directions horizontally and vertically.

The DLP is driven by a Discovery 1100 Controller (Digital Light Innovations), which is fed by a PC's DVI output using a custom interface (Jones *et al.*, 2007). In this interface, a 24-bit image drawn on the PC results in an endless loop of 24 binary images being displayed in rapid succession on the DLP. To produce grey, a fraction of the bits in each pixel are turned on. The resulting sequence of binary images integrates over time to approximate the desired intensity. Although this interface supports only 24 grey levels, it has proven sufficient for our applications.

Imaging is performed using a Retiga 4000R monochrome CCD camera (2048×2048 pixels, each $7.4 \mu\text{m} \times 7.4 \mu\text{m}$). When imaged onto light field plane Lf_{LFM} in Fig. 4(c), this corresponds to a field $15 \text{ mm} \times 15 \text{ mm}$. The imaging microlens array at Π'_{LFM} is a custom design manufactured to our specifications by Adaptive Optics and consists again of square-sided, planoconvex lenslets, this time with a pitch of $125 \mu\text{m}$. This gives us an addressable spatial resolution during imaging of $15 \text{ mm}/125 \mu\text{m} = 120$ tiles horizontally and 120 tiles vertically. The angular resolution within each tile




Fig. 4. Prototype LFI and microscope. (a) The base microscope is a Nikon 80i. The LFI (yellow outline) consists of a video projector *A* and microlens array *B*. The LFM (red outline) is a camera *C* and second microlens array *D*. (b) Light from a 120 W metal halide source (X-Cite 120, EXFO) is conducted through a 1.5 m liquid light guide (3 mm diameter, 0.59NA, EXFO) to a rectangular hollow reflective light pipe (Discovery Visible Optics Kit, Brilliant Technologies). From there it is relayed to a TIR prism and focused onto a DLP-based SLM. The DLP's image is focused by a 1:1 telecentric relay lens (Brilliant Technologies) onto light field plane Lf_{LFI} (also shown in (c)). Solid lines are the chief rays for the DLP, and dashed lines are the marginal rays for one micromirror. (c) Light field plane Lf_{LFI} is the front focal plane of the illumination microlens array, which lies at a conjugate Π'_{LFI} of the microscope's intermediate image plane. Plane Π'_{LFI} is imaged by a Nikon imaging tube lens TL_{LFI} , beam splitter *BS*, and standard microscope objective *Ob*, which focuses it onto object plane Π . Returning light passes through the objective and the microscope's native tube lens TL_{LFM} , which focuses it onto the imaging microlens array at the native intermediate image plane Π'_{LFM} . The back focal plane Lf_{LFM} of this array is imaged by a 1:1 relay lens system (not shown), which consists of two Nikon 50 mm *f*/1.4 AIS manual focus lenses mounted front-to-front, and is recorded by a digital camera. Solid lines are the chief rays for a single microlens, and dotted lines are the marginal rays for that microlens. To simplify the drawings, only three microlenses are shown in each array. Further details on the optical recipes of LFMs can be found in Levoy (2006b).

is $2048 \text{ pixels}/120 = 17$ addressable directions horizontally and vertically.

The optics connecting the illumination and imaging paths consists of microscope objective *Ob* in Fig. 4(c) and two tube lenses TL_{LFI} and TL_{LFM} . We used a Nikon imaging tube lens at TL_{LFI} rather than a Nikon illuminator tube lens because we believe the former to have higher optical quality, which we need when making focused images. The effective focal length of this tube lens matches that of the microscope's native tube lens at TL_{LFM} , thereby making the two intermediate image planes Π'_{LFI} and Π'_{LFM} conjugates. As a result, we can treat the two microlens arrays as if they were superimposed. Comparing their fields, we see that the projector somewhat underfills our camera's field of view. More importantly, the camera has a higher pixel count than the projector, reflecting the current state of the art of these two technologies at roughly comparable component costs. By our selection of microlens pitches (125 μm versus 300 μm), we have chosen to devote those extra camera pixels to spatial rather than angular resolution.

To complete the optical layout, we must specify the focal lengths of the microlenses. To optimize use of our projector and camera, we want the telecentric stop D_{Ob} in Fig. 4(c) to map through the microlens arrays to arrays of imaging circles that exactly fill light field planes Lf_{LFI} and Lf_{LFM} . Ideally these circles should abut, neither overlapping nor leaving a gap between them. (Our square-sided microlenses produce a rectilinear grid of circles, which wastes additional space; a hexagonal grid would be more efficient.) This mapping is achieved if the NA of the microlenses matches the image-side NA of the objective. Expressing this constraint as an *F*-number *N* gives $N \approx M/(2NA)$. For example, for a $40 \times / 0.95$ NA objective, we should employ *f*/20 microlenses for both illumination and imaging. The LFI microlenses in our prototype have a pitch of 300 μm , so their focal length should be $300 \mu\text{m} \times f/20 = 6.0 \text{ mm}$ for this objective. Similarly, the LFM microlenses with a pitch of 125 μm should have a focal length of $125 \mu\text{m} \times f/20 = 2.5 \text{ mm}$.

At this point it is instructive to look back at our discussion of light field parametrizations. In Fig. 3, we represented

orthographic light fields abstractly by placing the u line at infinity. In Fig. 4 the st plane can be treated as lying at the microscope's object plane Π (or any conjugate, such as image planes Π'_{LFI} or Π'_{LFM}), and the uv plane can be treated as lying at the objective's back focal plane, i.e. at telecentric stop D_{ob} . Since this stop is the microscope's aperture, we say formally that $-1 \leq (s, t) \leq +1$ gives a position in the microscope's field of view, and $-1 \leq (u, v) \leq +1$ gives a position in its aperture or exit pupil.

From the foregoing it should now be clear that each circular subimage on illumination light field plane Lf_{LFI} controls the light incident at one position in the microscope's field of view, and each pixel within a subimage controls the light passing through one part of microscope's aperture, hence one direction of light for that position. Similarly, each circular subimage on imaging light field plane Lf_{LFM} records the light reflected from one position in the field of view, and each pixel within a subimage records one part of the aperture, hence one direction of view for that position.

It costs a few thousand dollars to convert a microscope with a digital camera into a LFM, and about the same to convert a Nikon epi-fluorescence module (Nikon Instruments, Japan) and Digital Light Innovations projector into a LFI. Although we replaced the illumination optics in the Nikon module to ensure optical quality as noted earlier, we continue to use the module's rotating turret of beam splitters and fluorescence filters, as well as the Nikon trinocular for switchable observation or photography. In the future, we envision a system having ports for the camera and projector and two turrets of microlenses. These turrets would permit

the user to optimize the microlens arrays' pitches and focal lengths for each microscope objective. This would permit the user to trade-off spatial and angular resolutions, and to do so separately for illumination and imaging, as needed for their particular application.

Alignment and calibration

No attempt is made in our prototype to physically align the microlens arrays with the camera or projector chips. Instead, we calibrate our system in software. The coordinate systems we use are defined in Fig. 5. We first find the mapping from camera pixels to camera microlenses, and from there to positions and directions of imaging rays. Then we find the mapping from projector pixels to camera pixels. Combining this mapping with the first one gives us the mapping from projector pixels to illumination rays, since they pass through the same microscope optics. The entire calibration procedure, including image capture, takes about 1 min on a personal computer.

To calibrate the imaging path, we place a blank slide on the microscope stage, illuminate it with transmitted light, focus the microscope for Koehler illumination, reduce the condenser aperture to a pinhole, and capture a light field. This light field will contain a grid of circular spots each a few pixels in diameter as shown in Fig. 5(a). We find the centroids (x_c, y_c) of these spots in camera coordinates (with subpixel accuracy), form a correspondence between the grid of centroids and points (s_c, t_c) on a 2D integer lattice (grey grid in Fig. 5(b)) representing the camera microlens array, and solve for a cubic warp mapping lattice points to camera coordinates. In addition to translation,




Fig. 5. Definition of coordinate systems. (a) The image returned by our camera or displayed on our projector. The pixel grid (black lines) is denoted using coordinates (x, y) . With the microlens arrays in place, these images can be treated as an array of circular subimages (grey circles), one per microlens. This array is generally not aligned to the pixel grid, as shown. Given a pixel in one of these images (red dot in (a)), the cubic warps described in the text tell us the integer coordinates (s, t) of the microlens it lies in (red circle in (b)), as well as the position in rectangular coordinates (u, v) or polar coordinates (r, φ) of the pixel (red square in (c)) relative to the centroid of its microlens. From r and subimage radius r_{NA} , we can apply Eq. (1) to obtain declination angle θ of the pixel's ray bundle (pink tube in (d)) with respect to the microscope's optical axis. Polar coordinate φ in (c) directly gives the ray's azimuthal angle in (d). Finally, microlens coordinates s and t (from (b)) tell us the spot in the field of view (red ellipse in (d)) where the ray bundle intersects object plane Π . The optics reflect θ and φ around the origin, as shown in the diagrams, but not s and t . In the text, camera coordinates bear the subscript c , and projector coordinates bear the subscript p .

scaling, rotation, and shear, this warp corrects for the radial distortion in our relay lenses. We assume that the microlens array, although not aligned, is not itself distorted.

Assuming the microlenses are themselves distortion-free¹, a pixel at distance r from the centroid of a microlens subimage (Fig. 5(c)) can be mapped to ray declination angle θ relative to the microscope's optical axis (Fig. 5(d)) using the Abbe sine condition (Oldenbourg, 2008)

$$\sin \theta = \frac{r}{r_{\text{NA}}} \cdot \frac{\text{NA}}{n}, \quad (1)$$

where r_{NA} is the radius of the microlens subimage produced by an objective having a numerical aperture NA, and N is the refractive index of the medium at the object plane.

To calibrate the illumination path, we place a mirror on the stage, we display on the projector a sequence of patterns consisting of black-and-white stripes, and for each pattern we capture a light field. First, we divide the light field by a previously captured image with every projector pixel on. This performs radiometric correction, to account for vignetting across the field of view and aperture. (The former affects the entire image; the latter affects each microlens.) We then subtract from each pixel a fraction (40%) of the average intensity in its subimage, clamping at zero. This partly corrects for light scattered from the seams between microlenses. This scattering is not completely understood and may include diffraction from the edges of the microlenses; characterizing it is a topic of ongoing work. Finally, thresholding the intensities seen in each camera pixel on each frame of the sequence, we produce a binary code per pixel. We perform this procedure first using horizontal stripes, then vertical stripes. These codes tell us (to the nearest pixel) that projector pixel (x_p, y_p) was seen by each camera pixel.

The stripe patterns we employ are Gray codes (Bitner *et al.*, 1976). These codes have the property that if a stripe boundary falls inside a pixel, making it hard to determine if the pixel should be zero or one, it will do so only in once in the sequence of patterns. Thus, the error we suffer as a result of making the wrong determination will be no more than ± 1 projector pixel. The number of patterns required (horizontal and vertical) is $2 \log_2 W$, where W is the width of the projector in pixels. For example, for $W = 1024$ pixels we must display 20 patterns, producing a 20-bit code for each pixel. Such a code is shown in Fig. 6(a).

Although Gray codes are commonly used to calibrate projector-camera pairs (Rusinkiewicz *et al.* 2002), the microlens arrays in our system make the task harder. In particular, no attempt is made to align our arrays to one another, nor to match the pitches of their microlenses. This

¹Although our microlenses are uncorrected, we are using them at long conjugates (between $f/10$ and $f/30$) to produce small images (≤ 20 pixels on a side). We believe that aberrations or distortions are smaller than one pixel and can be ignored in a first-order approximation.

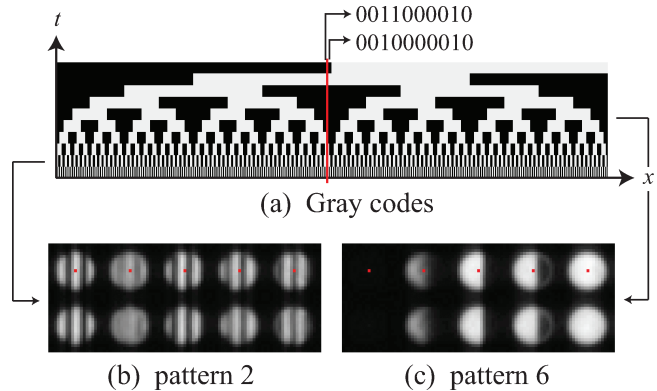


Fig. 6. Calibrating our projector using Gray code sequences. (a) This plot contains 10 rows, one for each frame in the sequence. Each row is 1024 pixels long, the width of our video projector. A row is extruded vertically to make the pattern displayed on the projector for that frame. Within row t , each pixel x is either black or white. Taking a vertical slice through the plot gives the 10-bit Gray code for that pixel. Two adjacent codes, on either side of the red line, are shown. Note that they differ in only one digit, in keeping with the Gray code's unique property. At bottom are the images captured by our camera for the second (b) and sixth (c) patterns. The specimen is a first-surface mirror with cover slip. Imaging employed a Nikon $60\times/1.2$ NA water objective, an $f/25$ illumination array, and an $f/20$ imaging array. See text for explanation of the red dots.

produces a mapping from projector pixels to camera pixels that folds back on itself repeatedly. This in turn permutes the sequence of Gray codes observed across along a row (or column) of camera pixels, destroying its error-tolerant property. To avoid this problem, we first discard any camera microlenses whose subimage is blurry, since this represents a fold in the mapping. We evaluate blurriness by thresholding the sum of the absolute values of the pixel gradients in X and Y , integrated across a microlens subimage. For example, we would discard the second column of microlenses in Fig. 6(b). Fortunately, the projector microlenses in our prototype are more than twice as wide as the camera microlenses, so we can afford to be conservative. We then transpose the light field, producing a 2D array of 'imagelets' each of which represents a single position (u_c, v_c) in the microscope's exit pupil, hence a constant ray direction. Two imagelets can be seen in Fig. 17(h)–(i). Adjacent pixels in an imagelet come from corresponding positions in adjacent microlenses, for example the red dots in Fig. 6(b)–(c). In these pixels, the Gray codes are guaranteed not to be permuted with respect to camera pixel position, preserving their error-tolerant property.

Once we have assembled an imagelet, which may be sparse because we have discarded some of its pixels, we form a correspondence between the remaining pixels, whose Gray codes give the coordinates (x_p, y_p) of the projector pixel they saw, and points (s_c, t_c) on a 2D integer lattice representing the camera microlens array. We then solve for a quadratic warp mapping projector coordinates to lattice points. Concatenating

this warp to the mapping we already know from the camera microlens lattice to ray positions and directions tells us the position and direction of ray produced by turning on any projector pixel. To illuminate a ray of fractional position and direction, we turn on a weighted set of pixels. If we assume linear interpolation, then to turn on a single ray we partially activate 16 projector pixels (the two neighbours in each dimension of the 4D light field). As a result of this interpolation, the resulting ray bundle will be broader in space and angle than we would like, but this is the best we can do given our discrete sampling of the light field.

Shaping microscope illumination

The control provided by our LFI over the lighting in a microscope can be described formally as

$$L'(u, v, s, t) = L(u, v, s, t) \times M_{LFI}(u, v, s, t), \quad (2)$$

where L is the light field that would arrive at the object plane if we performed no modulation, and M_{LFI} is the attenuation provided by our SLM. Although the SLM is 2D, the presence of a microlens array allows us to define M_{LFI} as a 4D function.

Unfortunately, Eq. (2) is rather general, making it difficult to appreciate which attenuation functions might be useful. In this paper, we consider two specializations of this equation: separable spatio-angular control and synthetically focused illumination.

Separable spatio-angular control

In a microscope whose illumination path provides physical access to conjugates of the field and aperture planes, substantial control can be exercised by inserting masks at one or both planes. This control can be described as

$$M_{LFI}(u, v, s, t) = M_a(u, v) \times M_f(s, t), \quad (3)$$




Fig. 7. Spatial control over illumination. The specimen and imaging arrangement are described in Fig. 1. Illumination employed an $f/25$ microlens array. Scale bar is 10 μm . (a) The GFP channel under brightfield illumination. The bacteria are visible, but contrast is low due to autofluorescence from surrounding structures. (b) A video projector pattern designed to restrict the illuminated field to a circular mask 18 μm in diameter. The circle spans 29 microlenses. Each microlens is represented in the mask by a solid disk, indicating that within the 18 μm circle the microscope aperture is filled with light. (c) The GFP channel under the masked illumination. Contrast between the bacteria and its background, computed using $(I_1 + I_2) / (I_1 - I_2)$, is improved 6 \times over (a). Figure 1 shows three views of (c) (visualized in green) together with the surrounding structures (visualized in red).

where M_a and M_f are attenuation functions of 2D, and (u, v) and (s, t) denote lateral positions in the field and exit pupil, respectively. In this formulation M_f controls the spatial character of the illumination, and M_a controls its angular character. To control spatial or angular character alone, one merely sets the other function to unity.

We first consider spatial control. Microscopists know that by restricting illumination to a small region of the field, in-scattered light from objects outside this region is reduced, improving contrast. Indeed, this is the basis for confocal microscopy. Since we can programmatically mask illumination and imaging, we can implement confocal imaging. However, since we are partitioning our projector's pixels between spatial and angular resolution, and we have chosen to restrict spatial resolution to 40×40 tiles, our confocal performance would be poor. In particular, it would be worse than using the projector to perform only spatial modulation (Hanley *et al.*, 1999; Smith *et al.*, 2000; Chamgoulov *et al.*, 2004).

Nevertheless, our spatial resolution is high enough to improve the contrast of ordinary fluorescence imaging. Fig. 7 demonstrates this idea for the difficult case of imaging *Listeria monocytogenes* bacteria in the intestinal epithelium. By restricting the illumination to a small circle interactively positioned to enclose only the bacteria, contrast was improved 6 \times , allowing us to see them clearly even though they are embedded 40 μm into the tissue sample.

We now consider angular control. Textbooks on optical microscopy typically devote an entire chapter to exotic ways of illuminating the condenser aperture, including brightfield, darkfield, oblique illumination, optical staining and so on (Pluta, 1988, vol. 2). As currently practiced, each of these techniques requires a special apparatus. It has been previously shown that by placing a SLM conjugate to this aperture, many of these effects can be simulated programmatically (Samson & Blanca, 2007). Since we are partitioning our projector's pixels among spatial and angular resolution as previously




Fig. 8. Angular control over illumination using the LFI alone. The specimen is a blond human hair. Scale bar is 100 μm . At top are the patterns displayed on our projector. As they pass through objective *Ob*, these patterns produce the illumination shown in the middle row. At bottom are photographs captured by a colour camera (Canon 5D) using a $10\times/0.45\text{NA}$ objective and no imaging microlens array. The illumination microlens array was $f/10$. (a) Brightfield. M_a (see Eq. (3)) is a full-aperture disk convolved with a Gaussian ($\sigma = \text{aperture diameter} \times 0.1$), which apodizes M_a to reduce diffraction. (b) Quasi-darkfield. M_a is a full-aperture annulus apodized as before. Note the enhanced visibility of scattering inside the hair fibre. (c) Headlamp, i.e. brightfield with a reduced aperture. M_a is a Gaussian ($\sigma = \text{aperture diameter} \times 0.4$). This produces a specular highlight, whose undulations arise from scales on the fibre surface. (d) Oblique, produced by shifting (c) to the edge of the aperture. See text for explanation of the arrow.

mentioned, our angular resolution is relatively low (20×20 directions).

Nevertheless, we too can simulate these effects, as demonstrated in Fig. 8. The patterns at top differ only within the circular subimage destined for each illumination microlens. These differences represent different aperture attenuation functions (M_a in Eq. (3)). The oblique illumination in 8(d) is particularly interesting. Some of the light reflects from the top surface of the hair fibre, producing a specular highlight, but some of it continues through the fibre, taking on a yellow cast due to selective absorption by the hair pigment. Eventually this light reflects from the back inside surface of the fibre, producing a second highlight (white arrow), which is coloured and at a different angle than the first. This accounts for the double highlight characteristic of blond-haired people (Marschner *et al.*, 2003).

Strictly speaking, darkfield imaging requires the condenser to overfill the objective back aperture. This allows light to be cast on the specimen from directions not seen by the objective. Since we are illuminating and imaging through the same objective, Fig. 8(b) does not satisfy this condition. However, if we capture a light field instead of an ordinary photograph, and we compute a synthetic aperture photograph using only the centre of each microlens subimage, then we can correctly simulate darkfield imaging, albeit with a reduced aperture. This idea is shown in Fig. 9. Under brightfield (a) the pixels in the yellow rectangle at bottom include a specular reflection from the marble slab. This reflection makes the background appear bright in the view at top. Under quasi-darkfield illumination (b) light arrives obliquely, reflects specularly from the marble, and leaves obliquely. Thus, it fills only the periphery of each microlens; the centre remains dark.




Fig. 9. Using the LFI and LFM together to obtain true darkfield imaging. At top are views computed from light fields captured by our LFM. The specimen is L-Glutamic acid crystals on a slab of polished black marble, imaged with a $20\times/0.75\text{NA}$ objective. The illumination microlens array was $f/10$, and the imaging array was $f/13$. Scale bar is $100\ \mu\text{m}$. The middle row are excerpts from the light fields, showing the microlenses used to compute the pixels inside the red rectangles at top. Visualized in yellow are the rays summed to compute one pixel. The bottom row contains blowups of one microlens. The yellow rectangle shows the rays used.

Extracting the centre then produces a true darkfield image, shown at top.

Synthetically focused illumination

If we relax the requirement in Eq. (3) that M_a and M_f be functions of (u, v) and (s, t) alone, we obtain a formulation with more generality,

$$M_{LFI}(u, v, s, t) = M_a(T_a(u, v, s, t)) \times M_f(T_f(u, v, s, t)), \quad (4)$$

where T_a and T_f are arbitrary mappings from 4D to 2D. For example, if $T_a = (u, v)$ and $T_f = (s + au, t + av)$, then we obtain a shear of st with respect to u and v , where a controls the magnitude and sign of the shear. As shown in Fig. 3, shearing a light field is equivalent to refocusing it (Ng, 2005). Thus, using T_f to shear M_f corresponds to creating a spatial illumination pattern focused on a plane parallel to the microscope's natural object plane but displaced from it. As another example, if $T_a = (u + bs, v + bt)$ and $T_f = (s, t)$, then the bundle of directions from which light arrives at the object plane, i.e. the illumination

PSF, will vary as a function of field position. This allows us to simulate point sources at arbitrary 3D positions relative to the specimen. These capabilities differ fundamentally from existing programmable-array illumination systems.

What does it look like to shear a light field? Fig. 10 shows two views from a light field of Golgi-stained neurons. To produce a view focused on the microscope's natural object plane, it suffices to sum the rays in each microlens subimage. (The pixel at the crosshair lies between microlenses, so rays were drawn from four microlenses.) To focus higher in the specimen, we shear the light field. This is equivalent to summing the central rays from the central microlenses and the peripheral rays from nearby microlenses, with the latter taken from the periphery furthest from the central microlenses. If we wished to focus lower, we shear the light field the other way. This causes us to draw the latter rays from the periphery closest to the central microlenses.

Shearing an illumination light field works the same way. If we replace the camera in Fig. 10 with a projector, and we turn on pixels corresponding to the yellow areas in (a) or (b), we would produce a spot of light at the lateral position indicated by the crosshair but focused at different depths. The minimum lateral extent of this spot would equal the spacing between microlenses in object space ($3\ \mu\text{m}$ for this example), and its minimum axial extent would equal the depth of field given




Fig. 10. Synthetic refocusing by shearing a light field. The specimen is a $100\ \mu\text{m}$ slice of Golgi-stained rat brain, imaged with transmitted light and a $40\times/1.3\text{NA}$ oil objective. At top are views computed from a single light field captured by a variant of the optical layout in Fig. 4. Scale bar is $100\ \mu\text{m}$. At bottom are excerpts from the light field. Visualized in yellow are the rays summed to compute the pixel under the crosshair at top. (a) Finite-aperture view focused on the microscope's natural object plane. (b) Focused $10\ \mu\text{m}$ higher. Note that rays are drawn from more microlenses than in (a).

by Eq. (5.2) in Levoy (2006) ($4\ \mu\text{m}$ for this example). These shears can be performed in real time using a modern graphics card.

Experimental demonstration

To demonstrate our ability to refocus light programmatically, we use our LFI to throw a light pattern onto a partially reflective slide (see Fig. 11). We then move the microscope stage axially, throwing the slide out of focus and the illumination doubly out of focus, since its light passes twice through this increased distance. Shearing the incident light field in 4D, we can refocus it back to the plane of the slide. Shearing the captured light field instead, we can refocus our view back to the same plane. Doing both at once brings the slide and illumination back into focus – without moving any optics.

Let us analyze these results. Applying Sparrow's criteria the minimum resolvable spot on the intermediate image plane (Levoy, 2006, Eq(2)) is $6.9\ \mu\text{m}$ for our $20\times/0.75\text{NA}$ objective, assuming $\lambda = 550\ \text{nm}$. On the imaging side, this gives us 18 spots within each $125\ \mu\text{m}$ -wide microlens. Our camera resolution is slightly worse than this spot size (pixel size = $7.4\ \mu\text{m}$), so in practice we resolve only 16.8 spots. Taking this as N_{ii} , the depth of field of a synthetically focused view (Levoy,

2006, Eq. 5.2) is $18.4\ \mu\text{m}$, and the axial range over which these views can be focused before their circle of confusion exceeds one microlens (Levoy, 2006, Eq. 5.3) is $139\ \mu\text{m}$.

On the illumination side, we have 43 resolvable spots below each $300\ \mu\text{m}$ -wide microlens by Sparrow's criteria. Our projector pixel size ($13.7\ \mu\text{m}$) is much larger, so in practice we can address only 21.8 spots. Taking this as N_{ii} , the depth of field of a synthetically focused pattern is $11.6\ \mu\text{m}$, and the axial range over which these patterns can be focused before their circle of confusion exceeds one illumination microlens is $233\ \mu\text{m}$. However, if we wish the pattern to appear sharp in the captured light field as well as the incident light field, then the axial range must be reduced by the ratio of the two microlens pitches, yielding a maximum range of only $97\ \mu\text{m}$. Thus, we can expect to refocus the illumination by only $40\ \mu\text{m}$ before it becomes noticeably blurry in views of the light field. Indeed, the words in Fig. 11(f) appear slightly less sharp than in (b).

What is synthetically focused illumination good for? First, it allows us to move a focused spot of light axially without moving the microscope optics. Combined with our synthetically refocusable imaging, this facilitates the illumination and viewing of rapidly moving organisms. Second, it allows us to create masked illumination at multiple depths simultaneously, by summing several patterns each sheared by a different amount. In ongoing work, we are using this technique to create steerable 'follow spots' (by analogy to theatrical lighting) to stimulate fluorescence in multiple *L. monocytogenes* as they move through live tissue, and minimize autofluorescence from surrounding structures. Finally, it allows us to implement a low-resolution scanning confocal microscope capable of acquiring a 3D volume without moving the microscope optics.

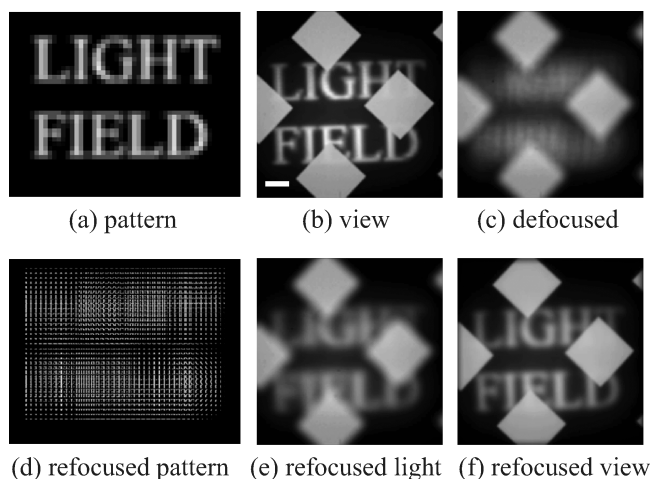


Fig. 11. Refocusing the illumination in a microscope. The specimen is a micrometer slide (Edmunds M53–713) consisting of diffuse diamonds on a mirror background, imaged using a $20\times/0.75\text{NA}$ objective. The illumination microlens array was $f/10$, and the imaging microlens array was $f/13$. Scale bar is $100\ \mu\text{m}$ (a) The word 'LIGHT FIELD' was rasterized with antialiasing and displayed on our video projector. (b) A view computed from a light field, focused on the microscope's object plane as in Fig. 10(a). The diamonds and words are sharp. (c) Lowering the stage by $50\ \mu\text{m}$ defocuses our view of the slide and doubly defocuses our view of the words. (d) Modifying the projected image as in Fig. 10(b) refocuses the illumination downwards by $50\ \mu\text{m}$. (e) The words are now focused at the plane of the slide, but our view of them is still blurry. (f) Synthetically refocusing our view downwards by $50\ \mu\text{m}$ brings the diamonds and words back into focus.

Digital correction of optical aberrations

Aberrations arise in an optical system when light leaving a point on the object do not reconverge to a point on the image plane. Since 4D light fields separately record each ray passing through the system, it is possible in principle to compute focused views free of aberrations – without altering the optics – by capturing and resampling a light field. This idea is shown geometrically in Fig. 12. Whether one can accomplish these corrections in practice depends on capturing rays at fine enough resolution in space and angle, and accurately characterizing the rays actually captured. In previous work in our laboratory, Ng (2006) showed that non-chromatic aberrations in photographic lenses whose optical recipe he knew could be reduced in this way.

Like any imaging system, microscopes are sensitive to optical aberrations. Minimizing these aberrations requires controlling the index of refraction of the immersion medium, and for a non-dipping objective, the thickness of the cover slip. If the specimen has a different index of refraction than the immersion medium, then the objective can only be well




Fig. 12. Aberrations in a generalized plenoptic imaging system, having an object plane Π , objective lens Ob , image plane Π' with microlenses, and light field plane Lf . (a) The subimage formed on Lf by one microlens records the different directions of light (grey paths) leaving a point (B) on the object plane. In the presence of aberrations, light field sample A on the periphery of that subimage maps (solid black path) to a laterally displaced point C . Adding this sample to the others in its subimage during synthetic focusing would produce a hazy image. What we actually wanted was the dashed path, which was recorded (at D) in the subimage formed by another microlens. (b) By selecting and summing light field samples from several microlenses, we can synthesize a full-aperture, aberration-free image. Drawing is not to scale.

corrected for a single depth. Researchers have proposed numerous ways to reduce aberrations, including introducing additional optics (akin to the cover slip corrector found on some objectives), changing the index of refraction of the immersion medium, or employing computer-controlled deformable mirrors (Albert *et al.*, 2000; Booth *et al.*, 2002; Sherman *et al.*, 2002; Kam *et al.*, 2007; Wright, 2007). However, these techniques offer only a limited ability to correct for aberrations that vary spatially. When imaging thick biological specimens the index of refraction may vary both laterally and axially, making full correction difficult.

In a LFM, aberrations affect us two ways. First, they cause light field samples to deviate from the nominal position and direction we expect for them, as shown in Fig. 12(a). We can correct for these deviations by resampling the light field as shown in 12(b). The paths in 12(a) arise if the medium between Π and Ob has a lower refractive index than expected, causing rays to bend excessively as they pass through the objective. In this case, D is to the left of A , and D is closer to the centre of its microlens subimage than A is to the centre of its subimage. If the medium had a higher refractive index than expected, D would be to the right of A and further from the centre of its subimage than is A . In fact, it may fall outside the objective's

exit pupil and be irretrievable. This unavoidably reduces the angular spread of rays captured.

The second effect of aberrations is that the light we record for each sample may itself be aberrant. Since a sensor pixel records an integral over space and angle, in the presence of aberrations our samples may represent integrals that are asymmetric with respect to their nominal position and direction or that overlap those of adjacent samples.

In this paper, we measure and correct for the first effect, but we ignore the second. Since aberrations are proportional to aperture size, and the aperture of a single light field sample is small, this simplification is justifiable. However, for severely aberrant imaging the intrasample aberrations become significant. We will return to this issue later.

There are several ways we could measure the deviations of light field samples from their expected position and direction. Here are two methods we have tried:

Method 1: Tracking guide stars

In terrestrial telescopes, aberrations due to refraction in the atmosphere can be reduced by projecting a laser guide star into the sky, observing it with a Shack–Hartmann wave front sensor (1971), and physically modulating the shape of the telescope's main reflector to improve convergence of the rays – a technique called adaptive optics. A Shack–Hartmann sensor is simply a microlens array focused on a sensor. In microscopy Beverage *et al.* (2002) have used a Shack–Hartmann sensor the same way, employing as a guide star a single fluorescent microsphere placed at the microscope's object plane. In a related approach, Pavani *et al.* (2008) measures the deformation of a structured illumination pattern as it passes through a phase object.

In our method, we create a grid of widely spaced 'guide stars' by reflecting a sparse grid of spots produced by our LFI off a mirror or glass slide. After two passages through the specimen, we record images of these spots using our LFM. Fig. 13 shows the light field captured for a single moving guide star under normal and aberrant conditions. By measuring the shift in rows 3–4 of 13(b) relative to 13(a), we can correct digitally for the aberrations using the ray selection strategy shown diagrammatically in Fig. 12.

Although similar in some respects to Beverage *et al.*, our method differs in several ways. First, we do not characterize the complete PSF; we measure only lateral shifts. However, this simplification is justified by the relatively low resolution of images we can produce from our light fields. Second, our method permits aberrations to be measured simultaneously in each part of the field of view, although we have not yet implemented this. Third, rays deviate twice in our approach due to the mirror. This makes it twice as sensitive as observing a fluorescent microsphere.

Although the analogy with guide stars is alluring, this method of measuring aberrations has several problems. First,

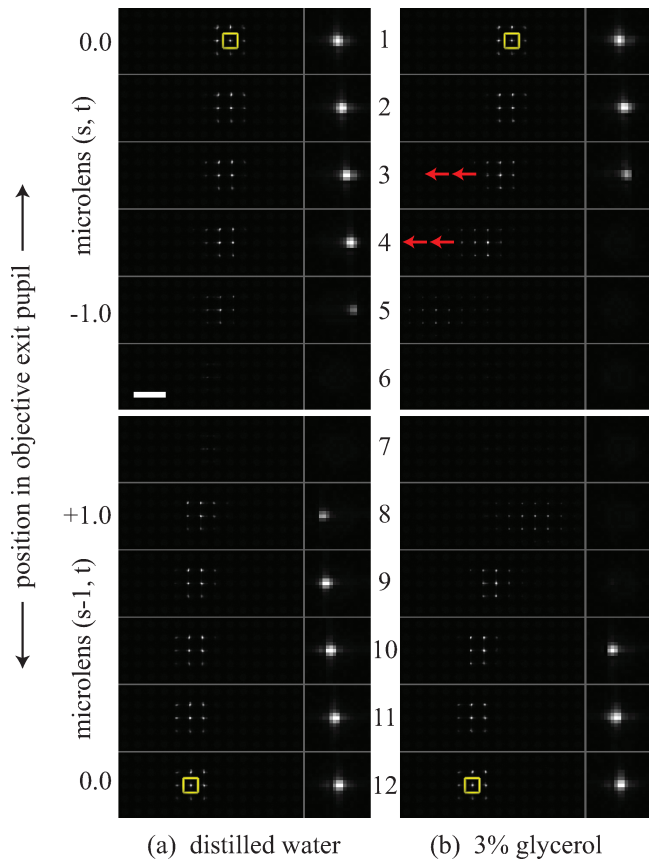


Fig. 13. Lateral shifts due to aberrations. The specimen is a mirror, imaged using a Nikon $60\times/1.0\text{NA}$ water dipping objective. The immersion medium is distilled water (a) or 3% glycerol (b). The illumination microlens array is $f/25$, and the imaging array is $f/30$. Scale bar is $5\ \mu\text{m}$. In row 1 we illuminate the centre pixel (pupil position 0.0) in one projector microlens. This produces a ray parallel to the optical axis, which reflects from the mirror and is recorded by several imaging microlenses (due to their smaller pitch). A blowup of one microlens (yellow rectangle) is shown at right. As the illuminated pixel moves leftward, the illumination angle increases. This induces an opposite tilt in the reflection, causing it to move rightward in the imaging microlenses (rows 2–5). Eventually the angle exceeds the pupil diameter (row 6), and the light is blocked. Continuing leftward, we pass into an adjacent projector microlens (rows 7–12), producing a similar sequence in reverse. In glycerol the paraxial rays (rows 1–2) behave as in (a), but peripheral rays (rows 3–4) start shifting left (red arrows).

it requires tracking the lateral motion of a reflected spot over a potentially long sequence of frames. In the presence of severe aberrations, the reflection can shift substantially, causing it to be confused with the reflections from adjacent spots. More importantly, the reflection from a single projector pixel is weak, so its SNR is poor, especially if the medium exhibits scattering or attenuation.

Method 2: Analyzing Gray code sequences

The calibration procedure in section ‘Alignment and calibration’ depends on finding correspondences between

projector pixels and camera pixels after reflection from a mirror. In the presence of aberrations, axial rays will be unaffected, but marginal rays will be shifted laterally, leading them to strike the ‘wrong’ camera pixel. Since we can predict the ‘right’ pixel from the axial rays, we can measure these shifts, allowing us to solve for the refractive index of the medium as before.

This method has the important advantage over the guide star method that it can be performed simultaneously at every camera pixel with no danger of confusion. This provides robustness and good SNR. It also allows us to solve for the refractive index no matter how fast it changes across the field, limited only by the resolution of our camera. Of course, if we perform this measurement independently in every pixel, we trade-off spatial resolution against SNR. In this paper, we demonstrate only the ability to measure the refractive index of a homogeneous medium; mapping a spatially varying index of refraction is left as future work.

Experimental demonstration

To apply this method, we need only run our calibration procedure, then compare the results in axial rays and marginal rays. For a specimen consisting of air bubbles caught between a mirror and cover slip, Fig. 14 illustrates the procedure. When aimed at the bubbles, marginal rays undergo a strong lateral shift. Squinting at one of the microlens subimages in (b), one

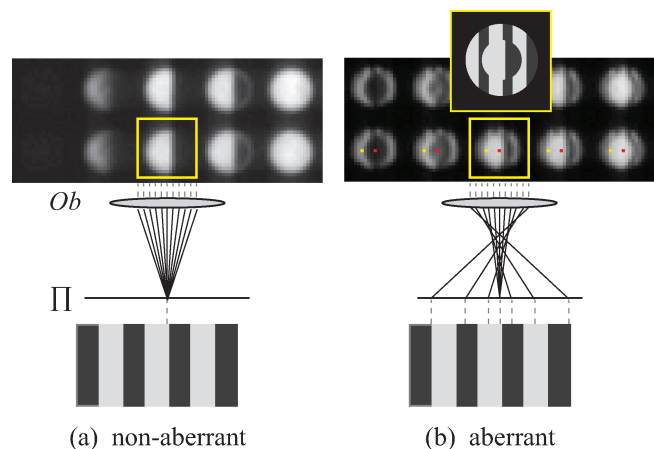


Fig. 14. The effect of aberrations on calibration. The specimen is a mirror to which a cover slip has been adhered using water, but leaving air bubbles in parts of the field. Illumination consists of the same Gray code as Fig. 6(c). Imaging is performed with a Nikon $60\times/1.2\text{NA}$ water objective. The illumination microlens array is $f/25$, and the imaging array is $f/20$. (a) Non-aberrant imaging. Each camera microlens focuses at a point on the object plane Π . In particular, the microlens outlined in yellow sees the boundary between a black and white stripe, as diagrammed below (not to scale). (b) Aimed at an air bubble. The core of paraxial rays still focuses on the stripe boundary, but marginal rays strike Π at shifted positions, causing them to see other stripes, as shown in the diagram. The dots and inset image is explained in the text.

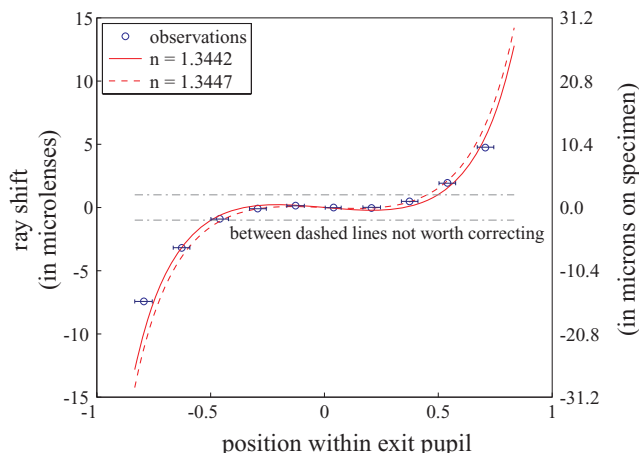


Fig. 15. Measuring aberrations using Gray codes. The specimen and imaging is described in Fig. 13, but with 10% glycerol instead of 3%. The blue dots give the lateral ray shift observed in pixels taken from a transect through the microlens subimages, averaged across the field of view and plotted as a function of pixel position in the pupil. The dots are collectively positioned vertically so that the centre dot (pupil position 0.0) is assigned zero shift on the assumption that axial rays are aberration-free. Moving away from this pixel, ray shifts increase, then decrease to zero again at pupil positions -0.3 and $+0.3$. These secondary zeros arise from our attempt to maximize sharpness by refocusing the microscope. Beyond these positions, ray shifts accelerate rapidly. The red curves are fits to these observations, as explained in the text.

can imagine its annulus as a separate image, focused higher in z than the object plane Π . A cartoon depicting this effect is inset at top. In reality no sharp boundary exists between the core and annulus.

Applying this procedure to the optical arrangement in Fig. 13 allows us to measure the ray shift seen by each camera pixel. These shifts are plotted as blue dots in Fig. 15. In this experiment, we employed only vertical stripes, and observations are made using pixels from a horizontal transect through the microlenses; similar results are obtained for horizontal stripes and a vertical transect. The geometry underlying these ray shifts is diagrammed in Fig. 16. Summarizing the formulas in that figure's caption, the predicted shift x'_a for rays of a given ideal angle θ is

$$x'_a = z_a \tan \theta - (z_a - \Delta z) \tan [\sin^{-1} (n/n_a \sin \theta)]. \quad (5)$$

Applying this equation to the present imaging situation, we obtain the dashed red curve in Fig. 15.

Alternatively, assuming we do not know the refractive index n_a of the medium, we can solve for it using unconstrained non-linear optimization. The function we minimized is the weighted sum of squared differences between the shifts predicted by Eq. (5) and the measured shifts (blue dots). The weighting we employed was a Gaussian tuned to match the observed falloff in light intensity towards the pupil boundary for this optical arrangement. Optimization was performed using Matlab's




Fig. 16. A dipping objective used with the wrong immersion medium. M is the objective's first glass surface and Π is the object plane. (a) Normally, rays leaving the objective converge at point a on Π . If the immersion medium has index of refraction n_a instead of n , then a marginal ray, which normally makes an angle θ with the optical axis, will instead make an angle $\theta_a = \sin^{-1} (n/n_a \sin \theta)$. This causes it to strike Π at B instead of A . In this example, $n_a > n$, so $\theta_a < \theta$. If z_a is the height of the immersion column, then the shift (from A to B) is $x_a = z_a \tan \theta - z_a \tan \theta_a$. (b) To measure these shifts empirically, we must focus the microscope as best we can despite the aberrations. If for example we move the stage down by Δz , so that paraxial rays (dotted lines) come to an approximate focus at C , then our marginal ray strikes Π at D instead of B , and its shift is reduced by $\Delta x = \Delta z \tan \theta_a$. If n_a were less than n , the signs of x_a , Δz , and Δx would be reversed.

fminsearch function. For this arrangement the minimum occurred at $n_a = 1.3442$, plotted as the solid red curve in Fig. 15. This agrees closely (within 3%) with the known value of 1.3447. The entire procedure, including image capture, takes about 1 min on a personal computer.

Despite this agreement, the observed data (blue dots) do not agree well with the model curves in the extreme periphery of the exit pupil. We ascribe these discrepancies to incomplete subtraction of scattered light. In addition, we believe there is slight miscalibration of projector pixels relative to camera pixel for axial (non-aberrant) rays. We represent this error on the figure as horizontal error bars of length ± 0.5 pixels, or equivalently ± 0.05 of pupil diameter.

Whether we measure the refractive index of a specimen experimentally or we know it *a priori*, we can employ the ray selection strategy diagrammed in Fig. 12 to correct digitally for aberrations. As an example, we replaced the mirror used in Figs. 13 and 15 with a micrometer slide, focused the microscope on the scale markings to maximize image sharpness, captured a light field using our camera, and used the ray shifts given by the dashed curve in Fig. 15 to improve the quality of focus synthetically.

The results of this experiment are shown in Fig. 17. Note the microlens subimages in the visualization for (b); they are

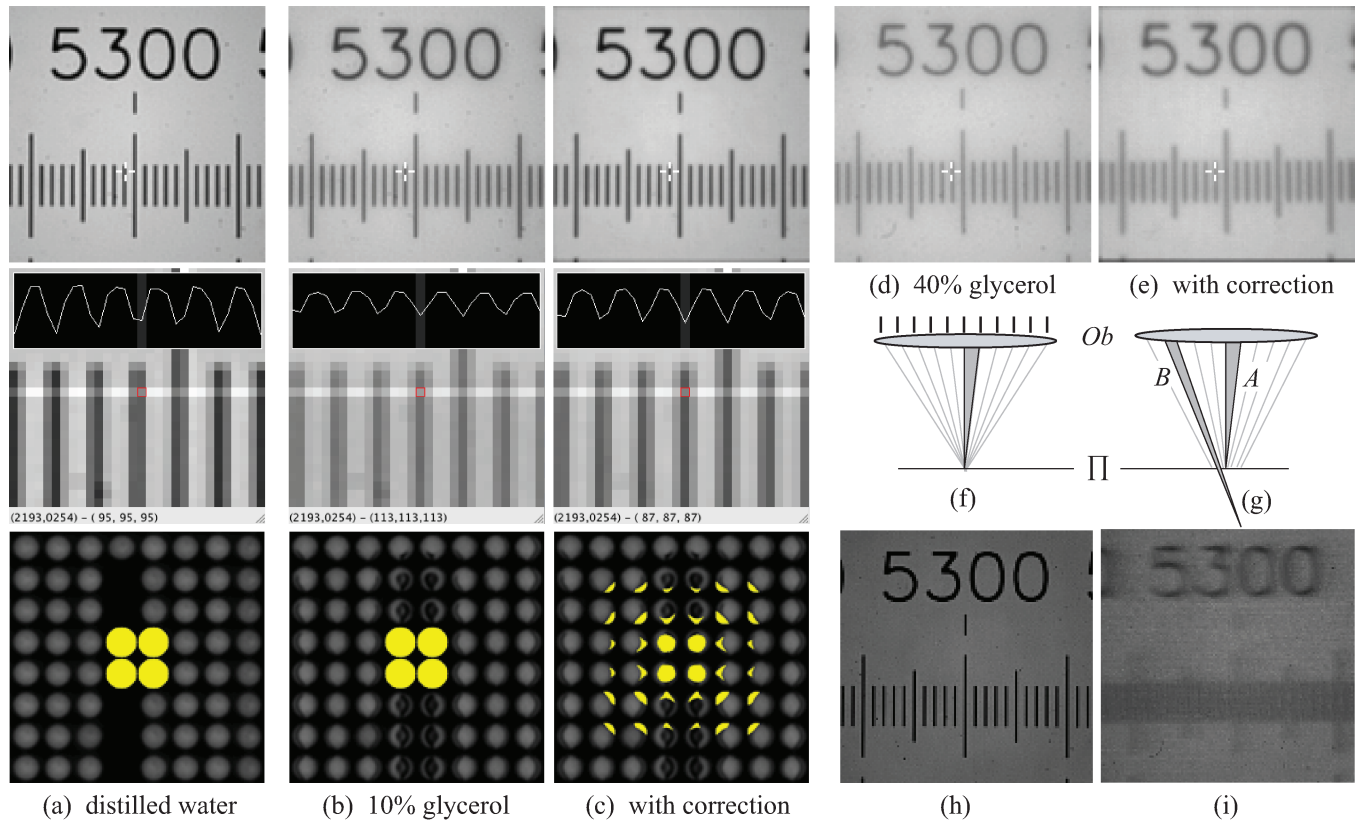


Fig. 17. Digital correction of aberrations. The specimen is a micrometer slide (Edmunds M53–713), imaged as in Fig. 13. The minor scale ticks are 10 μm . The synthetic views at top are focused on the ticks. At middle are blowups and line plots. At bottom are portions of the light field, showing which rays (yellow) were summed to compute the pixel under the crosshair at top. (a) Distilled water. The ticks exhibit a contrast (as defined in Fig. 7) of 54%. (b and c) 10% glycerol ($n_a = 1.3447$), without and with correction for aberrations. Contrast is 25% and 39%, respectively (an improvement of 56%). The visualization at bottom shows that rays have been drawn from several microlenses to form the corrected focus in (c). (d and e) A failure case: 40% glycerol ($n_a = 1.3895$), without and with correction for aberrations. Contrast is 12% and 10%, respectively. See text for explanation of (f–i).

distorted as in Fig. 14(b). Comparing the yellow highlighting in (c) to the diagram in Fig. 12(b), one can see how rays are being borrowed from nearby microlenses to improve the focus. In this case, rays were borrowed from the inner edges of microlenses. If the index of refraction had been lower rather than higher than expected, these rays would have been borrowed from the outer edges instead. Note also that the ray borrowing is half as extensive laterally as the ray shifts in Fig. 15, because those shifts reflect two passages through the aberrant medium rather than one. Finally, note that the sampling required to correct aberrations is quite different than the sampling required to perform focusing in the absence of aberrations (Fig. 10(b)). In focusing, relatively few samples are drawn from any one microlens, whereas in aberration correction most samples are drawn from the central microlenses. These samples are the paraxial rays, which suffer little aberration.

Discussion

Although this experiment proves that we can measure and correct for aberrations, the improvement was modest. Why

is this? When a user manually refocuses a microscope in an aberrant situation to maximize image sharpness, this reduces the magnitude of ray shifts. This reduction appears in Fig. 15 as a downward shear of the observed data for positive pupil positions and upward for negative positions. In mildly aberrant situations, this shear causes rays in the inner part of the pupil to shift by less than one microlens, indicated by the zone between the horizontal dashed lines. Since each microlens corresponds to one pixel in the final view, shifts smaller than a microlens are barely worth correcting. The outer part of the pupil is still worth correcting, and we have shown the utility of doing so. Discarding this part of the pupil after recording the light field is not a good alternative, since it reduces SNR. However, most of the energy passing through a microscope objective is paraxial, so the improvement we obtain is modest.

In severely aberrant situations, we must recall that light field samples are not abstract rays; they are integrals over space and direction. The PSF of a synthetically focused view, like that of the underlying microscope, is hourglass-shaped with its waist at the plane of focus. If this plane coincides with the microscope's natural object plane, we have the situation

in Fig. 17(f). The upper half of this PSF is depicted by a conical bundle of grey lines. The half-PSF of one light field sample (shaded area) is the same shape, but slenderer by a factor equal to the number of pixels behind each microlens (tick marks). In an aberrant situation (Fig. 17(g)), paraxial sample *A* is unaffected, but the focus of peripheral sample *B* is shifted both laterally and axially. The procedure described in this paper can correct only for the lateral shift; the axial shift causes light field samples to see a blurry view of the object plane Π .

To illustrate this secondary blur, Fig. 17(h) is a view of the same light field as in (d), but constructed using only the middle pixel in each microlens, i.e. the red dots in Figs. 14(b), and (i) is constructed using only pixels from pupil position $(-0.44, 0)$, the yellow dots in 14(b). These views were called 'imagelets' in section 'Alignment and calibration'. Note the horizontal blur in (i); these samples will be of little use in correcting aberrations. To reduce this blur, one must increase the angular resolution of the light field. However, the space-angle bandwidth product is limited by diffraction, so this requires a sacrifice in spatial resolution, which in turn widens the zone (dashed lines in Fig. 15) inside that it is not worth performing correction. It may be fruitful to investigate aspherical microlenses, which could provide finer angular discrimination in the periphery of the aperture without sacrificing spatial resolution.

In addition to improving image quality, correcting aberrations should increase the accuracy of volume data reconstructed from light fields using synthetic focusing followed by 3D deconvolution (Levoy, 2006). Discarding aberrant marginal rays is likely to be a markedly inferior solution in this application, because 3D deconvolution quality is known to depend critically on high numerical aperture. Indeed, our proof in (Levoy, 2006) that synthetic focusing followed by 3D deconvolution is equivalent to limited-angle tomography underscores the importance of maximizing the angular range of rays.

Finally, our correction methodology could lead to new designs for microscope objectives, in which optical performance is relaxed with respect to aberrations that can be corrected using microlenses, thereby allowing performance to be optimized in other respects, as suggested by Ng (2006).

Conclusions and future work

In this paper, we have described a system for controlling microscope illumination in space and angle. Coupled with our LFM, we have demonstrated two applications for this system: programmable illumination and correction of aberrations. The major limitation of our system is that spatial resolution is sacrificed to obtain angular resolution. This limits the precision with which we can direct and focus illumination, and it limits the degree to which we can correct for aberrations. A secondary limitation is the difficulty of obtaining a light source that is uniform in both space and angle. In our prototype,

we used a reflective light pipe, but this was not entirely successful.

Regarding future work, we have measured the improvement in contrast obtained by restricting illumination to a spot at the microscope's natural object plane; we now need to measure the improvement obtainable for other planes. Similarly, we have corrected for aberrations arising from a mismatch in the index of refraction that is homogeneous across the field, and we have argued that our technique also works if the index varies locally in *X* and *Y*; we now need to verify this argument. (It is less clear whether our technique can be used to measure 3D changes in refractive index.) Finally, we need to build a model predicting for which objectives, microlenses and specimens one can obtain worthwhile correction.

Since inserting microlenses into the illumination path of a microscope has not previously been done, it is not surprising that many potential applications remain untapped. Here are three ideas:

Scattering light field microscope

The bidirectional reflectance distribution function gives reflectance as a function of incident and reflected light direction (Nicodemus *et al.*, 1977). Kaplan *et al.* (1999) has recorded 2D slices of this 4D function by imaging an objective's back focal plane. Our system can generate 4D incident light fields and record 4D reflected light fields. Thus, we can measure reflectance as a function of incident and reflected position and direction. This 8D function is called the bidirectional surface scattering distribution function (BSSRDF). Although our spherical gantry (Levoy, 2002) can measure this function for macroscopic objects, ours is the first system that can measure it microscopically. We have not yet tried this, but Fig. 18 shows how it might work – using an LFM and a mockup of the angular control provided by an LFI to examine the spatio-angular dependence of scattering from iridescent squid skin.

Reconstructing the 3D shape of opaque objects

Although most objects are partially transparent at the microscopic scale, thick specimens viewed at low magnification are largely opaque. At these scales shape cannot be reconstructed using 3D deconvolution; instead we must employ methods from computer vision. One example is shape-from-focus (Nayar & Nakagawa, 1990; Noguchi & Nayar, 1994). Since our LFI can illuminate specimens with sheets of light from any angle, and our LFM can compute views from any direction, we should be able to digitize 3D shape using structured-light triangulation range finding (Levoy, 2000). Combining range data with BSSRDFs could lead to new appearance models for complex materials such as bird feathers.

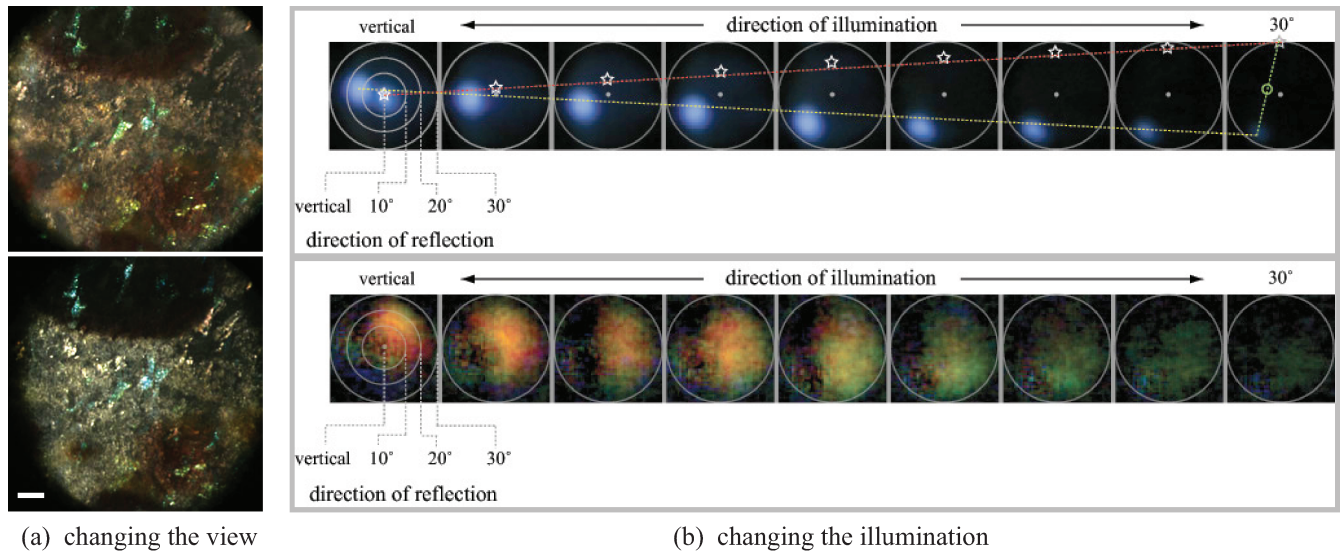


Fig. 18. Scatterometry using a LFM. The specimen is a live skin patch from longfin inshore squid *Loligo pealeii*. Imaging was performed using a $20\times/0.5\text{NA}$ objective and an $f/20$ microlens array. Scale bar is $100\ \mu\text{m}$. (a) Oblique views 15° (top) and 30° (bottom) from the optical axis, computed from a single light field captured under reflected light. Note that the centremost iridescent spot (iridiphore) turns from green to blue. (b) Moving a physical pinhole across the condenser aperture, one can change the direction of incident illumination. The polar plots at top show the angular distribution of reflected light for nine incident light directions (white stars and red line) ranging from axial (at left) to 30° (at right). Each plot was computed by summing several microlens subimages in a non-iridescent part of the skin. The yellow line is a linear fit to the path taken by the specular highlight. The green line and circle in the rightmost image shows the estimated surface normal, halfway between the light and highlight. The surface thus points slightly to the left. The plots at bottom show the light reflected from the centremost iridiphore in (a) after highlight subtraction and contrast expansion. At normal incidence, the iridescence is red-orange and preferentially to the right. As the illumination declines to the north, the iridescence becomes, then blue, and rotates towards the highlight direction.

General spatio-angular control over illumination

The ability to independently specify field and aperture masks, and to refocus illumination, gives us great flexibility when designing incident light fields. Fig. 19 provides an intuition for what this flexibility looks like – using the LFI to create a light pattern and the LFM to visualize its passage through a fluorescent liquid. What might this flexibility be good for? Using an array of video projectors, we previously demonstrated the ability to illuminate an object hiding behind foliage whereas not illuminating the foliage itself (Levoy *et al.*, 2004). An analogous application in microscopy is photostimulation of a single neuron whereas avoiding stimulating nearby neurons. Alternatively, one can imagine an interactive 4D paint program, in which the user outlines regions of interest in the field of view, then paints for each region a mask specifying from which directions light should arrive at that region. In both applications, one should be able to correct for aberrations whereas planning the illumination, using the methods described in this paper.

In closing, we note that there is a large body of literature on the use of coherent or partially coherent illumination to measure 3D structure (Poon *et al.*, 1995), phase (Barty *et al.*, 1998), or refractive index (Charrière *et al.*, 2006; Choi *et al.*, 2007), to correct for aberrations (Neil *et al.*, 2000), or to create illumination patterns in 3D (Piestun & Shamir,

2002; Nikolenko *et al.*, 2008). The techniques described in (Levoy, 2006) and this paper employ microlens arrays and incoherent light for similar aims. A comparison needs to be undertaken between these approaches. A useful starting point for this comparison might be to study the relationship between the Wigner distribution, a wave optics representation, and the light field (Zhang & Levoy, 2009). Our expectation is that each approach will prove to have advantages and disadvantages.

High-resolution images, light fields, and videos related to this work can be found at <http://graphics.stanford.edu/papers/lfmicroscope/>.

Acknowledgements

This research was supported by the National Science Foundation under award 0540872, the Stanford University Office of Technology Licensing, and the Stanford University School of Engineering. We thank Matthew Footer and Mickey Pentecost for several of the specimens used in this paper, Shinya Inoué for the rat brain, Oliver Hamann for the L-Glutamic acid, and Lydia Mathger for the squid. Finally, we thank Mark Horowitz, Ren Ng, Julie Theriot, Matthew Footer, Stephen Smith, Mark Schnitzer, and Rudolf Oldenbourg for many useful discussions, and our anonymous reviewers for several useful suggestions.




Fig. 19. Visualizing incident light fields. The specimen is a chamber filled with 100 μm fluorescein isothiocyanate in 1 M tris, imaged using a $20\times/0.75\text{NA}$ objective. The illumination microlens array was $f/10$, and the imaging microlens array was $f/13$. Scale bar is 100 μm . (a) The field mask at bottom (M_b in Eq. (3)) is three vertical stripes, and the aperture mask at top (M_a) is two pinholes. (b) The resulting pattern displayed on our projector. (c) An oblique view from the captured light field. The three fluorescing shapes correspond to the three stripes, and the two sheets per stripe show that light is arriving from two directions, corresponding to the pinholes. (d) The pattern was altered to refocus the illumination 25 μm lower. (e) The sheets now intersect inside the chamber, as indicated by the superimposed drawings.

References

- Adelson, T. & Wang, J.Y.A. (1992) Single lens stereo with a plenoptic camera. *Trans. Pattern Analysis and Machine Intelligence* **14**(2), 99–106.
- Agard, D.A. (1984) Optical sectioning microscopy: cellular architecture in three dimensions. *Ann. Rev. Biophys. Bioeng.* **13**, 191–219.
- Albert, O., Sherman, L., Mourou, G. & Norris, T.B. (2000) Smart microscope: an adaptive optics learning system for aberration correction in multiphoton confocal microscopy. *Optics Letters* **25**(1), 52–54.
- Barty, A., Nugent, K.A., Paganin, D. & Roberts, A. (1998) Quantitative optical phase microscopy. *Optics Letters* **23**(11), 817–819.
- Beverage, J., Shack, R. & Descour, M. (2002) Measurement of the three-dimensional microscope point spread function using a Shack-Hartmann wavefront sensor. *J. Microsc.* **205**, 61–75.
- Bitner, J.R., Erlich, G. & Reingold, E.M. (1976) Efficient generation of the binary reflected Gray code and its applications. *CACM* **19**(9), 517–521.
- Booth, M.J., Neil, M.A.A., Juskaitis, R. & Wilson, T. (2002) Adaptive aberration correction in a confocal microscope. *Proc. Nat. Acad. Sci. (PNAS)* **99**(9), 5788–5792.
- Chamgoulov, R.O., Lane, P.M. & MacAulay, C.E. (2004) Optical computed-tomography microscope using digital spatial light modulation. *SPIE* **5324**, 182–190.
- Charrière, F., Marian, A., Montfort, F., Kuehn, J., Colomb, T., Cuche, E., Marquet, P. & Depeursing, C. (2006) Cell refractive index tomography by digital holographic microscopy. *Optics Letters* **31**(2), 178–180.
- Choi, W., Fang-Yen, C., Badizadegan, K., Oh, S., Lue, N., Dasari, R.R. & Feld, M.S. (2007) Tomographic phase microscopy. *Nature Methods* **4**, 717–719.
- Gershun, A. (1936) The light field. Translated by P. Moon and G. Timoshenko. *Mathematics and Physics* **18**, 51–151.
- Halbach, K. (1964) Matrix representation of Gaussian Optics. *Am. J. Phys.* **32**(2), 90–108.
- Hanley, Q.S., Verveer, P.J., Gemkow, M.J., Arndt-Jovin, D. & Jovin, T.M. (1999) An optical sectioning programmable array microscope implemented with a digital micromirror device. *J. Microsc.* **196**, 317–331.
- Isaksen, A., McMillan, L. & Gortler, S.J. (2000) Dynamically reparameterized light fields. *Proc. Siggraph* **2000**, 297–306.
- Jones, A., McDowall, I., Yamada, H., Bolas, M. & Debevec, P. (2007) Rendering for an interactive 360° light field display. *ACM Trans. on Graphics (Proc. SIGGRAPH)* **26**(3), 40.
- Kam, Z., Kner, P., Agard, D. & Sedat, J.W. (2007) Modelling the application of adaptive optics to wide-field microscope live imaging. *J. Microsc.* **226**, 33–42.
- Kaplan, P.D., Trappe, V. & Weitz, D.A. (1999) Light-scattering microscope. *Applied Optics* **38**(19), 4151–4157.
- Kingslake, R. (1983) *Optical System Design*. Academic Press, New York.
- Levoy, M. (1988) Display of surfaces from volume data. *Comp. Graphics App.* **8**(3), 29–37.
- Levoy, M. & Hanrahan, P. (1996) Light field rendering. *Proc. SIGGRAPH* **1996**, 31–42.
- Levoy, M. *et al.* (2000) The digital Michelangelo project: 3D scanning of large statues. *Proc. SIGGRAPH* **2000**, 131–144.
- Levoy, M. (2002) *The Stanford Spherical Gantry*. Available at: <http://graphics.stanford.edu/projects/gantry/>.
- Levoy, M., Chen, B., Vaish, V., Horowitz, M., McDowall, I. & Bolas, M. (2004) Synthetic aperture confocal imaging. *ACM Trans. on Graphics (Proc. SIGGRAPH)* **23**(3), 825–834.
- Levoy, M., Ng, R., Adams, A., Footer, M. & Horowitz, M. (2006) Light field microscopy. *ACM Trans. on Graphics (Proc. SIGGRAPH)* **25**(3), 924–934.
- Levoy, M., Optical recipes for light field microscopes. (2006) Stanford computer graphics laboratory technical memo 2006–001.
- Levoy, M. (2006) Light fields and computational imaging. *IEEE Computer* **39**(8), 46–55. August 2006.
- Marschner, S.R., Jensen, H.W., Cammarano, M., Worley, S. & Hanrahan, P. (2003) Light scattering from human hair fibers. *ACM Trans. on Graphics (Proc. SIGGRAPH)* **22**(3), 780–791.
- Nayar, S.K. & Nakagawa, Y. (1990) Shape from focus: an effective approach for rough surfaces. *Proc. International Conference on Robotics and Automation (ICRA)*, 218–225.
- Neil, M.A.A., Wilson, T. & Juskaitis, R. (2000) A wavefront generator for complex pupil function synthesis and point spread function engineering. *J. Microsc.* **197**, 219–223.
- Ng, R. (2005) Fourier slice photography. *ACM Trans. on Graphics (Proc. SIGGRAPH)* **24**, 735–744.
- Ng, R., Levoy, M., Brédif, M., Duval, G., Horowitz, M. & Hanrahan, P. (2005) Light field photography with a hand-held plenoptic camera. *Stanford Technical Report CTSR 2005–02*.
- Ng, R. & Hanrahan, P. (2006) Digital correction of lens aberrations in light field photography. *SPIE* **6342**.
- Nicodemus, F., Richmond, J., Hsia, J., Ginsberg, I. & Limperis, T. (1977) Geometrical considerations and nomenclature for reflectance. National Bureau of Standards.
- Nikolenko, V., Watson, B.O., Araya, R., Woodruff, A., Peterka, D.S. & Yuste, R. (2008) SLM microscopy: scanless two-photon imaging and

- photostimulation with spatial light modulators. *Front Neural Circuits* **2**(5). DOI: 10.3389/neuro.04.005.2008.
- Noguchi, M. & Nayar, S. (1994) Microscopic shape from focus using active illumination. *Proc. IAPR International Conference on Pattern Recognition (ICPR) A*, 147–152.
- Oldenbourg, R. (2008) Polarized light field microscopy: an analytical method using a microlens array to simultaneously capture both conoscopic and orthoscopic views of birefringent objects. *J. Microsc.* **231**, 419–432.
- Pavani, S.R.P., Libertun, A.R., King, S.V. & Cogswell, C.J. (2008) Quantitative structured-illumination phase microscopy. *Applied Optics* **47**(1), 15–24.
- Piestun, R. & Shamir, J. (2002) Synthesis of three-dimensional light fields and applications. *Proc. IEEE* **2**(90), 222–244.
- Pluta, M. (1988) *Advanced Light Microscopy (3 vols)*. Elsevier, Amsterdam.
- Poon, T.-C., Doh, K.B., Schilling, B.W., Wu, M.H., Shinoda, K. & Suzuki, Y. (1995). Three-dimensional microscopy by optical scanning holography. *Optical Engineering* **34**(5), 1338–1344.
- Rusinkiewicz, S., Hall-Holt, O. & Levoy, M. (2002) Real-time 3D model acquisition. *ACM Trans. on Graphics (Proc. SIGGRAPH)* **21**(3), 438–446.
- Samson, E.C. & Blanca, C.M. (2007) Dynamic contrast enhancement in widefield microscopy using projector-generated illumination patterns. *New Journal of Physics* **9**(363). DOI: 10.1088/1367-2630/9/10/363.
- Shack, R.V. & Platt, B.C. (1971) Production and use of a lenticular Hartmann screen. *J. Opt. Soc. Am.* **61**, 656.
- Sherman, L., Ye, J., Albert, O. & Norris, T. (2002) Adaptive correction of depth-induced aberrations in multiphoton scanning microscopy using a deformable mirror. *J. Microsc.* **206**, 65–71.
- Smith, P.J., Taylor, C.M., Shaw, A.J. & McCabe, E.M. (2000) Programmable array microscopy with a ferroelectric liquid-crystal spatial light modulator. *Applied Optics* **39**(16), 2664–2669. Vol. 39, No. 16, 1 June 2000.
- Vaish, V., Garg, G., Talvala, E., Antunez, E., Wilburn, B., Horowitz, M. & Levoy, M. (2005) Synthetic aperture focusing using a shear-warp factorization of the viewing transform. *Proc. CVPR* 2005.
- Wilburn, B., Joshi, N., Vaish, V., et al. (2005) High Performance Imaging Using Large Camera Arrays. *ACM Trans. on Graphics (Proc. SIGGRAPH)* **40**(3), 765–766.
- Wright, A.J., Poland, S.P., Girkin, J.M., Freudiger, C.W., Evans, C.L. & Xie, X.S. (2007) Adaptive optics for enhanced signal in CARS microscopy. *Optics Express* **15**(26), 222–228.
- Zhang, Z. & Levoy, M. (2009) Wigner distributions and how they relate to the light field. *IEEE International Conference on Computational Photography (ICCP)*.