

Structured Light Field Design for Correspondence Free Rotation Estimation

Ian Schillebeeckx

Washington University in St. Louis
1 Brookings Dr, St. Louis, MO 63130

ischillebeeckx@wustl.edu

Robert Pless

Washington University in St. Louis
1 Brookings Dr, St. Louis, MO 63130

pless@wustl.edu

Abstract

Many vision and augmented reality applications require knowing the rotation of the camera relative to an object or scene. In this paper we propose to create a structured light field designed explicitly to simplify the estimation of camera rotation. The light field is created using a lenticular sheet with a color coded backplane pattern, creating a lightfield where the observed color depends on the direction of the light. We show that a picture taken within such a light field gives linear constraints on the $K^{-1}R$ matrix that defines the camera calibration and rotation. In this work we derive an optimization that uses these constraints to rapidly estimate rotation, demonstrate a physical prototype and characterize its sensitivity to errors in the camera focal length and camera color sensitivity.

1. Introduction

Camera calibration is the first step in many Computer Vision and Photogrammetry applications because it defines the relationship between an image and the scene that the image depicts. Calibration characterizes the photometric and geometric properties of the camera, that define, respectively how pixels of the camera report color and intensity of the scene, and where scene elements appear on the image. This paper offers a new approach to partial geometric calibration of a camera. We design a passive calibration object that creates a specially structured light field.

The calibration object is based on a lenticular array. These are sheets of plastic which are comprised of a large set of parallel cylindrical lenses. These lenses focus light onto the back of the plastic sheet, and common children's toys exploit this by interleaving multiple pictures behind the lenticular array so that the apparent pattern changes as the array is rotated [18]. Some modern TVs use a lenticular layer in order to show different pictures in different directions in order to project 3D TV "without any glasses" [14, 10].

Here we explore a version of this idea, where the inter-




Figure 1. (Top Left) Viewed from the side, a lenticular array consists of a set of cylindrical lenses designed to focus parallel light rays onto a back-plane, here shown at the top of the figure. If this back plane has different colors arranged behind each lens, then the apparent color of the lenticular array changes based on viewing angle. (Top Right) This pattern is repeated at every location of the lenticular array, so we can think of each lenticule as a light source that has a different color projecting in each direction. (Bottom Left) This lenticular array creates a light field structured so that the color of each ray depends on its direction. (Bottom Right) The image from a pinhole camera in this light field will have colors that change across its field of view, and this paper derives simple constraints that allow this image to be used for camera calibration without solve for image correspondences.

leaved pictures behind each lenticular lens create a color spectrum, as shown in Figure 1. This creates a light field, and when a camera captures an image, the color of a pixel depends fundamentally on the direction of the ray viewed by that pixel. This is quite different than most imaging situations; for example, in Lambertian scenes, it is the set of rays that come from a single world point that have the same appearance. In our case, a camera that captures an image in this light field will see a color pattern that varies across the image.

These images have a clear relationship to the parameters of the camera's rotation and intrinsic calibration. Specifically, for a pinhole camera capturing an image as shown in Figure 1 (bottom right), the pixel that views a green color must be viewing the world in a vertical direction. The constraints are only slightly more complicated in a full three dimensional setting. The contributions of this paper are to formally derive those constraints, to give an initial exploration of the sensitivity to noise of these constraints, and to experimentally demonstrate their feasibility with a real prototype lenticular calibration object. We find several points particularly compelling:

- Lenticular sheets are cheap commodity items and the appropriate back-plane textures can be printed on commodity printers, so this approach is broadly feasible even in application domains where cost is a factor.
- There are linear constraints on a combination of intrinsic calibration and rotation based on the pixel location and observed color. They do not require correspondence between a pixel and a location on the calibration grid, so it is easy to use all observed pixels as measurements.

One explanation of why correspondence is not necessary for this process is that the light field created by the lenticular pattern is structurally different than light fields created by normal (Lambertian) objects. In particular, it is apparent from Figure 1 that when a camera is translated within this light field, the images it captures remain exactly the same.

2. Related Work

Geometric camera calibration is a very well studied problem, with widely shared practical toolboxes [21, 6], and well understood geometric constraints (e.g. [9, 19]). Most common calibration approaches start with an image of an object with known 3D geometry, or several images of an object with known 2D geometry, find correspondences between points on the object and points in the image, and use these to solve for the camera geometry.

There are some approaches that are not based on identifying exactly corresponding points in a scene. Calibration patterns that consist of patches or parallel lines can be used for intrinsic camera calibration and extrinsic calibration (including rotation and translation relative to other cameras). Approaches that do not require correspondence between specific lines are based on seeing the orientation and spacing of those parallel lines on the image including those based on the prismatic line constraint [4, 3], and an approach using the Radon transform as a filtering operator [13, 12]. In some circumstances there are scenes that have large numbers of parallel lines with a known orientation, such as vertical edges of buildings or plumb-lines;

the orientation and position of those lines in an image provide constraints for both intrinsic and extrinsic calibration even without matching the pixel to a specific line in the world [15, 5].

Other approaches use textures for calibration. The authors of [23] assume that natural urban textures have low rank (for example, grids of windows on the exterior of a building). Using this assumption, the authors solve for the calibration and lens distortion that minimizes the rank of the textures in the image, using all pixels, not just point locations of geometric points.

In our work we explicitly create a light field whose color depends on the direction of the light rays. Fascinating research has gone into formalizing representations of the light field [1, 8], and creating acquisition systems that can capture them [20, 17].

A few prior works have explicitly created fiducial markers whose relative appearance depends on the direction from which they are viewed. Agam fiducials [7], take advantage of properties of 3D stairstep like structures where the vertical part of all steps are painted a different color than the horizontal parts of each step. When viewed from afar, the darkness of this pattern relates to the angle at which the stairstep is viewed. A recent patent uses small lenticular markers that vary in color based on the orientation from which they are viewed for the purpose of pose-estimation [2]. More recently BoKodes [16] created a highly structured pattern of light projected away from one point in a scene. This structured pattern is based on thousands of small QR-codes. When a defocused camera takes a picture of a scene, this pattern is visible, and the identity of QR-codes in view indicates the relative direction of the camera from the camera to the bokode marker. Both the BoKode markers and the AGAM fiducials can be thought of as creating a structured lightfield, but in both cases that structure is limited to a pencil of rays that intersect the bokode marker or AGAM fiducial. We do not know of any previous work that explicitly designs a large scale light field for more than a pencil of rays for the purpose of rotation estimation or augmented reality.

3. Structured Light Field Patterns Using Lenticular Arrays

We define a structured light field based on a lenticular array with a back-plane pattern that varies in hue. This leads to an object whose apparent color smoothly varies as it is rotated. The observed hue depends on how far the viewing angle is rotated around the major axis of the lenticular array. Figure 2 describes notation that we will use; in the coordinate system of the lenticular patch, x is the direction of major axis, v_{hue} is an example ray observed at a given color and $n_{hue} = V_{hue} \times x$ is the surface normal to the




Figure 2. We define the major orientation to be the direction along the lenticular lenses. Depending on the angle at which the pattern is viewed, it appears to have a different hue. For notation, we define the major axis of the lenticular pattern to be x , and an example ray that would be seen for each hue as V_{hue} . Any ray that views the lenticular sheet in the plane spanned by x and V_{hue} has the same color, our derivations are written most simply in terms of n_{hue} , surface normal to that plane.




Figure 3. (left) The top surface of a lenticular array is comprised of parallel cylindrical parts that focus parallel rays onto particular rows of a backplane. (right) We use the hue spectrum, interleaved under each lenticule, as our back focal plane texture. This is an example of a hue encoder for an array of 5 lenticules. The picture is dramatically enlarged; when placed under the lenticular array each of the five repetitions of this color spectrum is scaled to fit under one lenticule, which in our experiments was 0.34mm wide.

plane containing all rays of that color.

In the remainder of this section we provide details on constructing a physical prototype. In the next section we derive constraints that relate an image of a lenticular pattern to its rotation and the camera calibration.

3.1. Prototype Creation

Our experiments are based on lenticular fiducial markers we created using the EcoLens Lenticular Array, purchased from Pacur, Inc., with a design that includes lenticules that are elliptical instead of cylindrical because that gives several optic advantages [11]. We use arrays where each lenticule was 0.34 mm wide, and the thickness of the array was 0.4 mm. Interleaved behind the back focal-plane

we use a texture that includes 12 discrete samples of a hue spectrum, sampled every 30 degrees around a color wheel. Our pattern is shown in Figure 3. A pattern with 12 colors aligned along the major axis every 0.34 mm requires printing a texture at ≈ 900 dpi to get a stripe of each color 1 dot wide. In reality, the printing process mixes adjacent colors to a certain degree, and the lenses do not provide a perfect focus, leading to an observed hue that changes continuously as the lenticular pattern is rotated.

3.2. Experimental Measure of the Hue-Angle Relationship

We create a large lenticular array with our back-focal plane texture. To measure the apparent color we use two motorized rotation stages from ThorLabs to move the lenticular array in a controlled way around its major axis and, orthogonal but coplanar to the major axis, the minor axis. Images were captured with a Nikon D60 and a 300mm lens.

In the first experiment, we rotate the mounted lenticular array around its major axis at 1 degree increments ($\theta \in \{-40^\circ, \dots, 40^\circ\}$) and imaged the pattern at each angle. This set of angles is repeated with the lenticular array tilted at 5 rotations around its minor axis ($\phi \in \{0^\circ, 10^\circ, 20^\circ, 30^\circ, 40^\circ\}$). Figure 4(a) shows the setup for the experiment: a large lenticular array was mounted square on a right angled bracket attached to a motorized stage to vary θ , while another motorized stage could be tilted to vary ϕ .

For each of the images, we average a small region (5 lenticules wide) at the center of the stage to capture a representative hue value for each viewpoint. Figure 4(b) shows that for θ in the range of $[-35, 35]$, all five sets of images show an approximately linear HRF giving a one to one relationship between the hue and the angle. Furthermore, the HRF does not vary greatly between different rotations around the minor axis (ϕ).

For values of θ outside the range of $[-35, 35]$, the optics of the lenticular array focus the rays hitting one lenticule on the back-focal plane texture of the adjacent lenticule. Since the hue spectrum is repeated under each lenticule on the BFP texture, the hue appearance of the lenticular array “rolls over” on the hue color wheel (and jumps from 0 to 1) on the standard mapping of hues to numbers from 0 to 1).

The next experiment explores the question: how small of angular changes can the lenticular patterns detect? Here we set $\phi = 0$ and change θ in increments of $\frac{1}{10}^\circ$. Figure 4(c) shows the observed hues for θ values between 10 and 14. The blue line interpolates measurements at each integer degree, and the more jagged red line interpolates measurements between every $\frac{1}{10}^\circ$. These were done in two different trials, so the difference between the red and blue lines gives an indication of the variance.

The experiments in the remainder of the paper are de-




Figure 4. (a) We experimentally measure the HRF in the laboratory using a precision rotation controller. The observed hue (b) has a nearly linear response as a function of rotation around the major axis θ , even for different rotations around the minor axis (shown for 5 different angles ϕ). The drastic jumps in hue at wide angles of θ are a consequence of the design of our BFP texture and is explained in Section 3.2. (c) The HRF resolution with our current prototype suggests that accuracy is possible to about half a degree. At smaller scales, noise in the measurement process makes the hue/angle function non-monotonic.

signed so that the lenticular fiducial marker is always viewed from an angle within 35° of straight on so there is a one-to-one mapping. The estimate of n_{hue} is computed by converting Figure 4(b) into a lookup table that maps observed hues into values for θ .

4. Derivation of Rotation Constraints from Lenticular Patterns

Consider a situation where a pinhole camera views a lenticular pattern. Let the rotation of that lenticular pattern relative to the camera be R , and the calibration matrix of the camera be K . The basic constraint is that a ray observing a hue must be perpendicular to n_{hue} . In the coordinate system of the camera, a pixel \vec{p} captures light traveling along a ray \vec{r} that depends on the calibration matrix K as:

$$\vec{r} = K^{-1}p$$

In the coordinate system of the lenticular pattern, this ray is $RK^{-1}p$, so when observing a particular hue, we must satisfy the constraint:

$$(RK^{-1}p) \cdot \vec{n}_{hue} = 0. \quad (1)$$

We write this dot-product in matrix notation,

$$(RK^{-1}p)^\top \vec{n}_{hue} = 0,$$

rewrite the transpose to get

$$p^\top K^{-1} R^\top \vec{n}_{hue} = 0,$$

collect terms to get:

$$(p^\top K^{-1})^\top R^\top \vec{n}_{hue} = 0.$$

In this equation p is a pixel location where a color is measured, K is the calibration matrix that we assume is known and \vec{n}_{hue} is a vector defined as normal to the plane that contains all the rays of a particular hue only when viewed from some directions. Thus, each pixel measurement gives a constraint on the unknown rotation R .

Furthermore, the constraint is based only on the hue of the pixel and the pixel location; this constraint does not require finding a correspondence between the image and a particular point on the pattern. However, if all observed points are from a single lenticular sheet with the same major orientation, then the n_{hue} vectors will all lie along a pencil of planes (all planes that include the major axis direction x). In Sections 5.1.3 and 5.2 we explore prototypes such as one lenticular pattern surrounded by a lenticular pattern with a perpendicular orientation. In these cases we need to know “which part” of the lenticular pattern a pixel observes, but not the correspondence between a pixel and a particular location on the lenticular pattern.

To optimize over the rotation matrix R , we parameterize it with a Rodrigues vector ρ and define $R(\rho)$ as the rotation matrix for a particular Rodrigues vector. Then, we can solve for these variables minimizing the squared error function, summed over all pixels i :

$$g(\rho) = \sum_i \|p_i^\top K R(\rho) \vec{n}_{hue_i}\|_2^2.$$

5. Experimental Results

The following section shows simulated and empirical results of using lenticular arrays to estimate the rotation of a camera. First, we explore the sensitivity of our geometric

constraints to noise, number of sample points, and calibration object design. Second, we show rotation estimate error rates for a prototype on captured video. Then, we introduce using rotation estimation in an augmented reality application. After, we show empirical results of this prototype to estimate camera rotation across different cameras and focal lengths. Lastly, we show that our method works for even out of focus images.

5.1. Simulated Camera Calibration

We create a simulator to explore the geometric constraints, allowing us to change parameters that control:

- the amount of noise in the measured hue of a simulated lenticular array
- the number of measurements taken from a simulated calibration object
- the orientation and relative positioning of various lenticular arrays that make up a simulated calibration object

This simulator randomly generates a known position and orientation for the virtual calibration object that is modeled to be 110 mm tall and 70 mm wide. This object can appear anywhere in the field of view of a virtual camera from 150 to 850 mm away. In addition, the virtual calibration object cannot be rotated more than 30° around any of its axes. With a randomly generated position and orientation, the simulator projects the object from a camera’s 3d space onto an image plane. We model the pinhole geometry of an iPhone sensor (a 1/3.2” format image sensor at a default 4.1 mm focal length). This image is used to optimize the measured geometric constraints to get a rotation estimation. These estimations are compared against the true simulated conditions to gauge the performance of the derived geometric constraints.

5.1.1 Sensitivity to Noise

The only measurement of our calibration system is the hue observed from the lenticular array at a given pixel location. Therefore, a source of error could be in measuring an incorrect hue. In terms of our geometric constraints in Section 4, this error manifests as an improper direction of \vec{v}_{hue} , and thus \vec{n}_{hue} . Therefore, to simulate the geometric effects of measurement noise, we introduce normally distributed aberrations to the direction of \vec{v}_{hue} . These aberrations are created by randomly choosing a 3D vector from a Gaussian distribution with a given standard deviation, adding that vector to \vec{v}_{hue} , and re-normalizing to again get a unit vector.

We start with zero noise as a baseline and add a maximum of 0.2 standard deviation of noise to the unit vector. Figure 5(b) gives insight into the practical effects of adding




Figure 5. As more noise is introduced to \vec{n}_{hue} , the a) rotation error for all axes increases and is less consistent. b) shows the quartiles of angular error introduced to \vec{n}_{hue} for a given standard deviation of noise.

noise to \vec{v}_{hue} and then computing \vec{n}_{hue} , by showing the angular error in the geometric constraint (the computed direction of \vec{n}_{hue}) as a function of the standard deviation of the added noise. The left of Figure 5 shows the sensitivity to noise in estimation rotation as a function of the amount of noise.

In Figure 5(a), we show the 1st, 2nd (median), 3rd quartiles of the errors in rotation estimation. We display the angular error for each axis of a rotated local reference frame of the calibration object. The angular error for each axis is measured as the difference (in degrees) of a coplanar ground truth projected axis and the estimated projected axis. For all three axes, the trend has higher median amounts of rotational error with wider distributions for increasing amount of noise. The x and y axes have a small amount of error more than the z axis. This is due to the fact that our lenticular arrays are directly measuring rotation *around* the x and y axes and not around the z axis. Thus, error in rotation around the z axis manifest as error in the angular error of the x and y axes. Even at very high noise levels, the median estimated rotation has less than 2 degrees of error.




Figure 6. Using more measurements, and thus geometric constraints, in the optimization results in a lower and more consistent rotation error.

5.1.2 Sensitivity to Number of Measurements

Because our calibration approach does not need point correspondences, it is easy to use a large number of measurements to provide redundancy. Thus, we analyze how the number of measurements of the calibration object increases calibration performance.

We ran 300 trials of randomly generated calibration object poses with 0.08 standard deviation in noise and used an increasing amount of measurements sampled evenly across the calibration object for optimization. Results are shown in Figure 6.

We show the 1st, 2nd (median), and 3rd quartiles of the errors in rotation estimation. Rotation error reduces and the estimations become more consistent as more measurements are used in the optimization.

Practically, a measurement represents the hue at one pixel of the image. One can get more measurements of the calibration object by having the calibration object fill more of the image — bringing it closer or by using a higher resolution camera. For subsequent simulation experiments we use about 30,000 measurements (200 x 150 pixels), which are feasible to capture with commodity cameras at reasonable distances.

5.1.3 Sensitivity to Orientation and Relative Position of Lenticular Arrays

The constraints created by observing a single lenticular array are not sufficient to solve for the camera rotation. To get a system of equations for 4 that is not rank deficient, we need to include observations of a lenticular sheet of a different orientation. Thus, our structured light field object must have 2 lenticular arrays, which have major axis in different directions. Beyond this, there is also the design consideration of relative positioning of the differently oriented lenticular arrays. We explore the different decisions by sim-

ulating various designs with the same simulation system as the previous section.

We assess how the relative orientation and placement of lenticular sheets affect the estimation accuracy by creating a large set of designs. Each design is depicted in Figure 7(a-g). These show the layout and orientation of the lenticular arrays. For each design we run 400 simulations with varied position and rotation, adding 0.2 standard deviation in appearance noise (the maximum tested in Section 5.1.1) to 30,000 measurements. We measure the rotation estimation error for each axis.

Looking at the results in Figure 7, we see that the design b is the worst in estimating rotation. This is because the non-orthogonal lenticular orientations give less complementary cues about rotation angle. Most of the patterns are similar in estimating the rotation angle. However, designs e, f, and g are much more difficult to build in practice.

5.2. Evaluation of Physical Prototype

To test our structured light field in practice, we created a prototype that was 110 mm tall and 70 mm wide. It is comprised of two lenticular arrays, one cut smaller and placed directly on top of the other at a perpendicular orientation. We choose to implement design c in Figure 7 because of its simplicity.

5.2.1 Camera Rotation Estimation Accuracy

We first assess the rotation estimation accuracy when using our prototype. We use a standard checkerboard grid placed coplanar with our prototype to provide ground truth. The ground truth rotation is estimated using Zhang’s method [22] implemented in the MATLAB 2014a Camera Calibration Toolbox. We estimate the rotations for 345 frames of a 13 second video captured with an iPhone 5 that shows the prototype at various orientations. We omit the frames that we were unable to ground truth.

In Figure 8(a), we show the rotation error in degrees for each local axis across every frame of the video. Often, there are frames where the error for the z axis is noticeably lower than that of the x and y axes. This is the case when our rotation estimate has nearly correctly estimated the surface normal of the prototype, but has rotated the x and y axis around the surface normal or the z-axis

In Figure 8(b), we show the summary statistics for the rotational error for the entire video. Each box indicates the 1st, 2nd (median), and 3rd quartiles, the most extreme values as the feet, and outliers as red crosses. Rotation estimations using our prototype have around 5 degrees of error for all axes. There is less error in estimating the z axis or the surface normal of the prototype because the two lenticular arrays that comprise the prototype directly measure the rotation around the x and y axis, but not the z axis. This result




Figure 7. Various designs (a-g) change the orientation and relative positions of lenticular arrays to make a calibration object. Estimations achieved for design g) have the lowest h) rotation error, however it is an unrealistic design. c) is a more reasonable design.

corroborates the same phenomena seen in simulated results.

5.2.2 Augmented Reality Application

A possible application of our prototype is in Augmented Reality. With a rotation estimation, one gains knowledge of the real world orientation. With a planar object, then one could overlay any graphics or digital object in the image to give the appearance of the digital object being in the real world on top of the object. We demonstrate this application in the same video data used in the above experiment. In supplementary material, we show the video with overlaid 3d arrows to denote the local axes of the prototype structured light field and a ground truthing checkerboard. 6 random frames of this video can be seen in Figure 9. This augmented reality application is a good way to visualize the rotation estimation of the prototype, so we use it in future figures to make it easy to quickly assess the accuracy of the rotation estimation.

5.2.3 Camera Rotation Estimation at Different Focal Lengths

We test our prototype with two different cameras at various focal lengths. First, we determine rotation with an iPhone (4.1 mm focal length) at various orientations. Second, we determine rotation with a Nikon D90 SLR camera at various orientations and focal lengths. Results are shown in Figure 10. For all the images, we show the image used to calibrate the camera with the estimated rotation visualized as the local axis of the calibration object in addition to the




Figure 8. For a test video with varying camera orientations, we test the rotation accuracy of our prototype. We show the rotation error for each frame of the video in a) and summary statistics in b).




Figure 9. We show random frames from a video used to assess the rotation accuracy of our prototype structured light field. As an Augmented Reality Application, we overlay representations of the local axes. We do this for both our prototype and the ground truthing checkerboard to visually assess rotation estimation accuracy.

focal length of the image in the title. In some images, like Figure 10(a) (left), a direct reflection of overhead fluorescent lights is visible. Systematic errors in the measured hue, caused by reflections or white balance corrections from the cameras are not modeled in the current implementation of our constraints and may affect rotation estimation results.

5.2.4 Rotation Estimation of out of Focus Camera

Using calibration grids requires that all pictures be in focus to minimize error in locating point correspondences. For close focal lengths, it becomes challenging to keep the entire grid in the depth of field, as a variety of orientations and positions of the grid pattern is needed to achieve strong calibration results. This is not a concern for our system, as we can calibrate using out of focus images.

Using a Nikon D90, we focus 3 images in front of, at, and beyond the prototype. Results are shown in Figure 11. As before, we show each image with focal length and the rotation estimation. For all images, our method is able to estimate the orientation using our prototype. Out of focus images have blurring, which we believe gives more accurate estimates of the hue because the colors of neighboring pixels are averaged together.

5.2.5 Exploring Free Parameters

We explore the relationship between the focal length and the constraints on rotation. In this experiment, we optimize for the rotation parameters, but provide a focal length via the intrinsic parameter matrix K . We use the same images from




Figure 10. We estimate camera rotation from various poses and focal lengths. In the first row, we show calibrated images taken by an iPhone for different orientations of the calibration object. In the second row, we show calibrated images of images taken by a Nikon D90 DLSR with varying focal lengths. For each image, we show the focal length in the title and the rotation of the local axis of the calibration object superimposed on the image.

Figures 10(a) (the right) and 11(b), and show the results in Figure 12.

We introduce error in the focal length to understand the effects of bad camera calibration on rotation estimation of




Figure 11. Our rotation estimation object works even for out of focus images. For each image, we show the rotation estimate as a set of axis. We show the results for images a) focused in front, b) focused on, and c) focused behind the calibration object.

the structured light field object. We show the angular error of the rotation estimations as compared to the ground truth rotation. The ground truth is the rotation estimation for the correct focal length.

These error plots are sensitive to both the design and orientation of the calibration pattern, but we find it compelling that the estimate of the rotation (shown in 12(a) and 12(b)) is relatively insensitive to the focal length.

6. Discussion and Conclusion

In this paper we have introduced a new type of object that uses lenticular arrays to create structured light fields. We derived constraints on rotation for images captured of this object, and started an initial exploration of the sensitivity to noise of these constraints and some possible designs. We give initial experiments with a physical prototype that show that our approach works with real images, but also highlights a number of issues that should be explored further.

The first issue is optimizing the calibration pattern. This calibration pattern must be comprised of lenticular arrays of different orientations, but the best layout and configuration of those arrays is not yet clear. It would be interesting to explore if it suffices to have measurements along the border of a square instead of in the middle.

The second issue is understanding why the errors in our physical prototype are so much larger than our simulation. We believe the major factor is that the noise we see in our prototype is systematic, and not the random, independent




Figure 12. We explore the optimization of the rotation separately for two images (from Figures 10(a) (the right) and 11(b)). For both images, we vary the focal length of the camera calibration matrix used in our constraints and observe the effects on rotation estimation. Even for relatively large focal length errors, the rotation estimation has less than 5 degrees of error.

noise at every pixel that was in our model. One cause of this may be imprecise manufacturing of the lenticular array. A very small mismatch of the spacing of the interlaced color pattern under the lenticular lenses, or slight mis-alignment of the interlaced pattern would lead to a color-to-angle mapping that is not consistent across the array. Another source of systematic error arises from well-known challenges in interpreting color from commodity cameras; our algorithm does not model the effects of white-balance and color gamut corrections.

Despite these issues, we believe that the idea of creating a structured light field where the color of an observed ray relates to its angle is a very powerful idea. It leads to a set of constraints for camera calibration that are linear and do not require direct correspondence between pixels and specific locations on the calibration object, allowing calibration to be based on measurements at very large numbers of pixels. This may be particularly interesting for future work in calibrating omni-directional or catadioptric cameras that deform an image so that features are hard to match.

Acknowledgements

We greatly thank the reviewers and the area chairs for their thoughtful comments and suggestions which greatly improved this paper. In addition, we thank Dr. Gruev for use of his lab's motorized rotation stages. This research was supported in part through the following grants: NSF IIS-1111398, NSF EF-1065734, NSF IIA-1355466.

References

- [1] E. H. Adelson and J. R. Bergen. The plenoptic function and the elements of early vision. In *Computational Models of Visual Processing*, pages 3–20. MIT Press, 1991. 2
- [2] anonymous. Tracking head position and orientation. US Patent 8922644, 2014. 2
- [3] P. Baker and Y. Aloimonos. Structure from motion of parallel lines. In *Computer Vision-ECCV 2004*, pages 229–240. Springer, 2004. 2
- [4] P. T. Baker and Y. Aloimonos. Calibration of a multicamera network. In *Computer Vision and Pattern Recognition Workshop, 2003. CVPRW'03. Conference on*, volume 7, pages 72–72. IEEE, 2003. 2
- [5] M. Bansal and K. Daniilidis. Geometric urban geolocalization. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 3978–3985. IEEE, 2014. 2
- [6] J.-Y. Bouget. Camera calibration toolbox for matlab. http://vision.caltech.edu/bougetj/calib_doc/index.html. Accessed: 2013-10-09. 2
- [7] A. M. Bruckstein, R. J. Holt, T. S. Huang, and A. N. Netravali. New devices for 3d pose estimation: Mantis eyes, agam paintings, sundials, and other space fiducials. *International Journal of Computer Vision*, 39(2):131–139, 2000. 2
- [8] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen. The lumigraph. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 43–54. ACM, 1996. 2
- [9] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2003. 2
- [10] M. Hirsch, G. Wetzstein, and R. Raskar. A compressive light field projection system. *ACM Transactions on Graphics (TOG)*, 33(4):58, 2014. 1
- [11] R. B. Johnson and G. A. Jacobsen. Advances in lenticular lens arrays for visual display. In *International Society for Optics and Photonics*, 2005. 3
- [12] A. Makadia, C. Geyer, and K. Daniilidis. Correspondence-free structure from motion. *International Journal of Computer Vision*, 75(3):311–327, 2007. 2
- [13] A. Makadia, C. Geyer, S. Sastry, and K. Daniilidis. Radon-based structure from motion without correspondences. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 796–803. IEEE, 2005. 2
- [14] W. Matusik and H. Pfister. 3d tv: a scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes. In *ACM Transactions on Graphics (TOG)*, volume 23 (3), pages 814–824. ACM, 2004. 1
- [15] R. Melo, M. Antunes, J. Barreto, G. Falcao, and N. Gonçalves. Unsupervised intrinsic calibration from a single frame using a” plumb-line” approach. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 537–544. IEEE, 2013. 2
- [16] A. Mohan, G. Woo, S. Hiura, Q. Smithwick, and R. Raskar. Bokode: imperceptible visual tags for camera based interaction from a distance. *ACM Transactions on Graphics (TOG)*, 28(3):98, 2009. 2
- [17] R. Ng, M. Levoy, M. Brdif, G. Duval, M. Horowitz, and P. Hanrahan. Light field photography with a handheld plenoptic camera. Technical report, Stanford University Computer Science, 2005. 2
- [18] T. Okoshi. *Three-dimensional imaging techniques*. Elsevier, 1976. 1
- [19] R. Szeliski. *Computer Vision: Algorithms and Applications*. Springer, 2010. 2
- [20] J. C. Yang, M. Everett, C. Buehler, and L. McMillan. A real-time distributed light field camera. In *Proceedings of the 13th Eurographics workshop on Rendering*, pages 77–86. Eurographics Association, 2002. 2
- [21] Z. Zhang. Flexible camera calibration by viewing a plane from unknown orientations. In *Proc. IEEE International Conference on Computer Vision*, pages 666–673, 1999. 2
- [22] Z. Zhang. A flexible new technique for camera calibration. *Pattern Analysis and Machine Intelligence, IEE Transactions on*, 22(11):1330–1334, 2000. 6
- [23] Z. Zhang, Y. Matsushita, and Y. Ma. Camera calibration with lens distortion from low-rank textures. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 2321–2328. IEEE, 2011. 2