

SparkleGeometry: Glitter Imaging for 3D Point Tracking

Abby Stylianou, Robert Pless
Washington University in St. Louis
{abby | pless}@wustl.edu

Abstract

We consider a geometric inference problem for an imaging system consisting of a camera that views the world through a planar, rectangular sheet of glitter. We describe a procedure to calibrate this imaging geometry as a generalized camera which characterizes the subset of the light field viewed through each piece of glitter. We propose an easy to construct physical prototype and characterize its performance for estimating the 3D position of a moving point light source just by viewing the changing sparkle patterns visible on the glitter sheet.

1. Introduction

We consider an extreme form of imaging geometry: a camera system that views the world through a surface covered with glitter. Glitter is highly reflective so this creates an imaging system that views the scene from a wide variety of narrow viewpoints. Once calibrated, this wide variety of different imaging viewpoints can be thought of as an extreme form of a generalized camera [6] or as a sample of the light field. In either case, this allows 3D information to be extracted from a single image of the glitter.

In 2006, Fergus et. al. proposed a random lens imaging system that included random reflective and refractive elements in lieu of the lens within the optical path of a camera [5]. They demonstrated the ability to reconstruct low-resolution images of the environment and also showed that point light sources at different distances from the camera result in different intensity patterns on the imaging plane. In 2015, Zhang et. al. proposed SparkleVision [13], which, like us, considers the problem of using a standard camera to view the world through a sheet of glitter. They state: “For a surface covered in glitter, it is difficult to build a proper physical model,” and they sidestep this issue by learning an elegant linear approach to calibrating the camera system in order to reconstruct planar patterns (such as images on a TV screen) that are scrambled when viewed through glitter.

This paper offers an approach to creating the physical model of the flat specular parts of a glitter plane, and there-

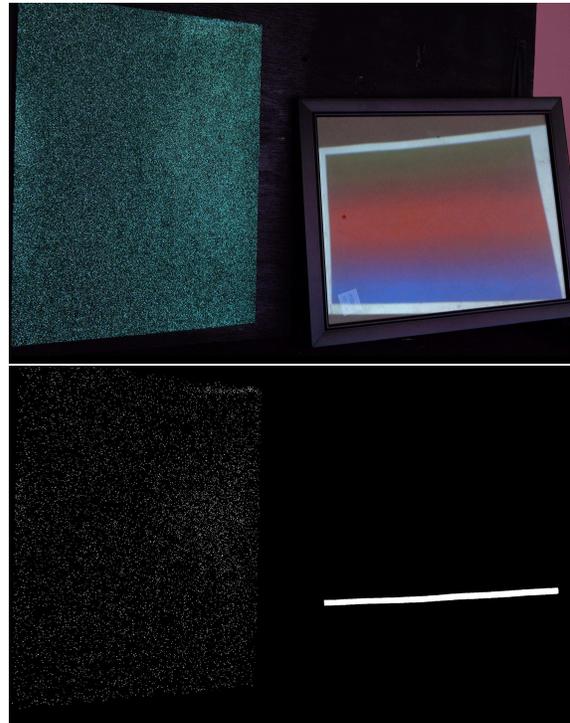


Figure 1: A sheet of glitter (top left) samples the world (shown, for convenience, reflected through glass in the top right) in a way that scrambles the appearance but samples a widely distributed part of the light field. In this paper we explore the geometric inference potential of an imaging system comprised of a camera and glitter plane. Showing controlled patterns (bottom right) and measuring when glitter pieces sparkle (bottom left, size and brightness enhanced for visibility) allows us to calibrate this system. We can then estimate the 3D position of an unknown point light source using only the pattern of visible sparkles.

fore create a calibrated imaging system that can solve for the 3D position of an unknown light source from only the pattern of visible sparkles. We explore the geometry of this system by considering a pinhole camera model where each

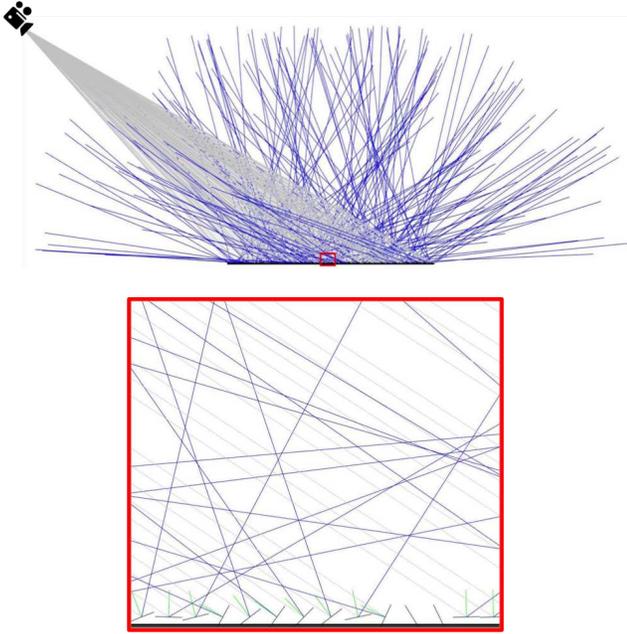


Figure 2: A pinhole camera that views a plane with randomly oriented glitter is transformed into a generalized camera that samples a wide variety of scene rays. The bottom shows an expanded version of the red rectangle near the center of the glitter plane, which emphasizes that some glitter may be oriented so that it is the glitter array itself which is reflected into the camera. Otherwise, the scene rays that are captured by the camera (in blue) provide a diverse sampling of the light-field that is incident on the glitter plane.

pixel views a ray in space. When those rays bounce off of glitter pieces, the pinhole camera model is transformed into a generalized imaging model [6], where neighboring pixels may observe the world along dramatically different rays. Figure 2 gives a visualization of this geometry.

As a light moves through the scene, different parts of the glitter plane sparkle. Those sparkles provide constraints on the light location. The sparkle pattern changes quickly as the light moves (those fast changes are the definition of a sparkly surface), so slightly different light positions may create substantially different observed images, as shown in Figure 3. Inferring the 3D position of the light is the geometric complement to previous work that reconstructs planar images, scrambled through random lens imaging.

Our contributions are:

- The geometric calibration of a glitter plane as a generalized camera that samples the scene along rays in space,
- a characterization of that sampling for a physical prototype, and

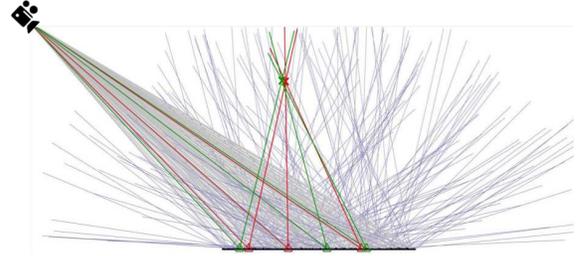


Figure 3: Each possible lighting position creates a set of sparkles that are visible by the camera. Because similar lighting positions (the red and green x's) cause very different pieces of glitter to sparkle in the camera view (the red and green triangles), there is the potential for accurate point location from glitter.

- an error analysis of an end-to-end system that tracks a point light source in 3D space.

2. Background

This work is directly inspired by two previous works that consider random lens imaging. Fergus et. al. proposed a camera system that includes random reflective and refractive elements within the optical path of a camera [5]. In this case each pixel may be affected by light from many parts of the scene. This paper proposes to model the scene as an image and searches for the stochastic transform relating known input images to the set of measurements taken by the camera. A corresponding probabilistic MAP algorithm is created to estimate an unknown image given just the scrambled picture captured by the camera. Results show recognizable reconstructions of 32 x 32 images. Final results show that a point light source at different distances from the random lens camera creates a different pattern of responses on the sensor, but no formal algorithm was proposed to estimate 3D position.

More recently, Zhang et. al. proposed SparkleVision [13]. This paper defined a spectrum of reflective imaging geometries based on the surface shape, spanning from imaging the world through a mirror (which maintains a pinhole imaging geometry), through a curved mirror (leading to various panoramic or catadioptric sensors) and through irregular mirrors. They focus on imaging with an array of irregular small mirrors (glitter) using a pinhole camera, and also explore the problem of reconstructing an image that is viewed through the glitter. Unlike Fergus, they use a pinhole camera; so the response of a particular pixel in this imaging system depends on the orientation, specularity and albedo of the glitter that pixel observes. They calibrate the SparkleVision sensor to reconstruct an image by solving for a linear transformation between the image pre-

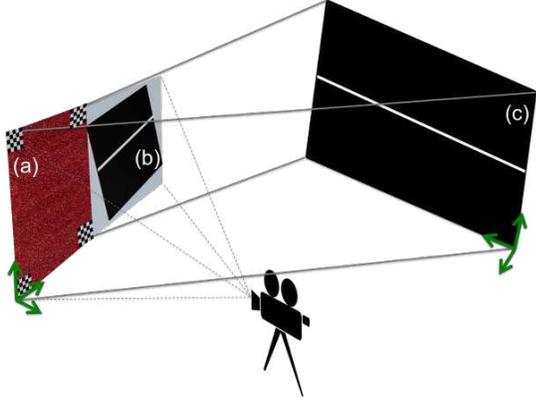


Figure 4: A sequence of lines are projected onto the calibration plane (c). They smoothly move across the calibration plane in several directions. This calibration sequence is reflected to the observing camera both by a sheet of glitter (a) and a small piece of glass (b). “Lit” glitter locations are correlated with the calibration sequence to determine their orientation, as shown in Figure 5.

sented to the sensor and the image observed by the camera. They also show recognizable reconstructions of low resolution images, and demonstrate that for images consisting of a point light source, small motions of that light lead to very different output images. Other related work includes systems that use micro-mirror arrays that can be programmed on the fly to modify the optical system [8] and work that tries to infer the shape of specular surfaces from the motion of specular reflections [11, 2].

This paper is inspired by both previous works that notice large changes from small motions in a point light source. We seek to formalize and quantify the estimation of 3D light positions from sparkle appearance in images. Under more standard imaging geometries, this Geometric Point Light Source Calibration problem has been extensively studied. Early work estimates isotropic lighting position from the shading of a known Lambertian 3D object [12]. Other approaches use fish-eye lens cameras [4], specular reflections off shiny spheres in the image [1], the reflections of the front and back of a clear hollow sphere [3], or from specular reflections seen in multiple viewpoints or reflecting off multiple planar, reflective surfaces [7]. Recent work has considered the problem of solving for the position and orientation of an anisotropic light source (like a flashlight) based on the pattern of intensity variations observed on nearby planar Lambertian surfaces [9].

3. Geometric Representation and Calibration

We create a calibration and testing configuration that is similar to previous work [13], but geared towards geomet-

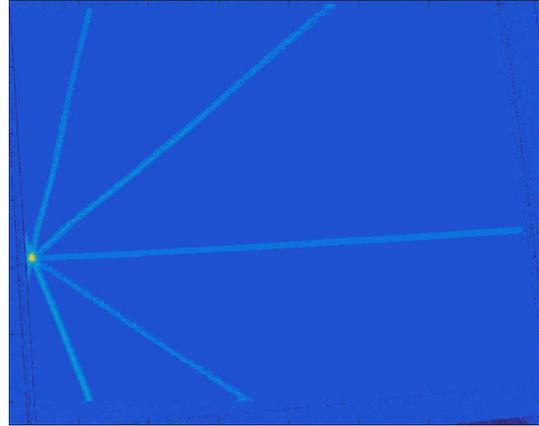


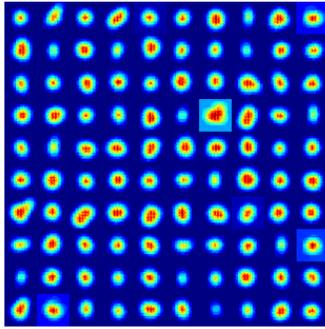
Figure 5: The calibration process illustrated in Figure 4 sweeps lines across the calibration plane in 5 directions, leading a piece of glitter to sparkle 5 times. The orientation of that piece of glitter is modeled as the ray from the sparkle location to the location on the calibration plane where the most correlated lines from the five different calibration directions intersect (the center of the brightest region in the image above).

ric calibration and inference. The glitter covered planar surfaces used in our experiments are vinyl sheets densely covered in small hexagonal glitter (less than one millimeter per side). These sheets are available from many craft supply stores. We affix these sheets to a stable vertical surface and position this surface in front of a projector screen. The configuration of the glitter, projector screen and the observing camera are approximately shown in Figure 4. A small piece of glass with an opaque backing is placed by the glitter. This glass reflects to the camera the patterns that are displayed on the projector screen during the calibration process.

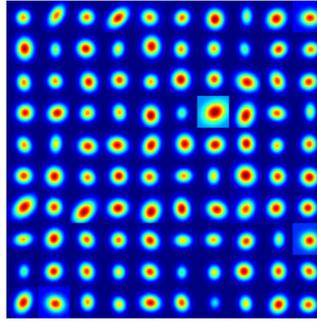
The calibration protocol displays on the projector screen a sequence of smoothly moving lines in multiple directions across the screen. We record the reflections off of both the glitter and the small piece of glass during this sequence using a Sony RX10-II capturing video at 30 frames per second.

Our formal calibration protocol starts by measuring the 3D coordinate system of the glitter plane and the projection screen that are shown in Figure 4. Based on fiducial markers affixed to the corner of the glitter plane, and a set of four markers projected onto the screen, we solve for two homographies. The first maps pixel locations that view the glitter onto coordinates on the glitter plane. The second maps pixel locations that view the reflective glass onto coordinates on the projector screen, which is parallel to the glitter.

We then analyze the video that shows lines sweeping in five directions across the projector screen. Pieces of glitter



(a)



(b)

	Eccentricity	Size
Mean	0.007	7.024
St.D.	0.198	0.7736

(c)

Figure 6: Each image in (a) shows the measured receptive field for 100 pieces of glitter. Each image in (b) shows the corresponding Gaussian fit for each piece of glitter. Each block in (a) and (b) is 29mm x 29mm. Warmer regions correspond to higher intensity responses for the piece of glitter – that is to say, the piece of glitter lights of brightest when a light is at the center of the modeled Gaussian, and that intensity fades as the light moves away from that peak. The table in (c) shows statistics on the average Gaussian used to model the receptive field each glitter piece. Eccentricities closer to 0 are closer to being round, while more elongated Gaussians have eccentricities closer to 1. Size is measured as the ratio of the major axis to the minor axis (in mm).

that reflect a region of the screen back to the camera have an intensity time-series that is mostly near zero except for the 5 times that a line is sweeping past them. For pixels that view glitter, we solve for the correlation between their time-series and the time-series of every pixel of the reflected calibration sequence. Figure 5 shows an example correlation map for a glitter piece captured in our real prototype shown later; the correlation is highest at the screen location imaged by that piece of glitter. All pixels on the glitter field whose correlation is less than 0.8 are thrown out. We also throw out pixels whose correlation is not larger than the correlation of neighboring pixels. In our prototype, this results in 27,616 calibrated sparkle locations. For each location, we use our homographies to calculate the 3D line that goes from the location on the glitter plane to the location on the projection plane. These lines are used to represent the orientation of each piece of glitter. The validity of using these lines, and not a richer model of the glitter orientation (such as a cone), is discussed in Section 4.

4. Characterizing Glitter Imaging Geometry

In order to understand the accuracy and error modes of the glitter imaging geometry, we solve for several parameters of how it samples the world. Each piece of glitter is viewed by a pixel from the camera. The pixel intensity depends on what part of the world a piece of glitter reflects towards that pixel (which we call the receptive field of the piece of glitter), and the brightness of the light and reflectance of the glitter. In this section we characterize the

receptive field of the glitter, motivate a simple thresholding approach to make pixel measurements binary and ignore the question of the reflectance and explore the density with which our glitter samples the world.

4.1. Glitter Receptive Field

In order to probe the receptive field of a particular piece of glitter, we project a 5mm square white light onto a screen 1000mm in front of the glitter. We move this light in an overlapping raster pattern over a 100mm square region of the screen and record video of the glitter as the light moves. We use the same calibration approach as described in Section 3 to determine glitter pixel locations.

We then model the receptive field for each piece of glitter by finding the frames in the calibration sequence where intensity of the glitter location surpassed a threshold (for this experiment we use a threshold of 30, but the results are robust to a wide variety of choices). We sum all frames where this threshold was surpassed in order to get a map of which screen locations affect that piece of glitter.

We computed this model of the receptive field for 100 pieces of glitter in order to understand the average area to which each of our glitter pieces responds. We found that the receptive field of each piece of glitter can be approximately modeled with a Gaussian of varying location, eccentricity and size. Examples of these receptive fields and statistics of the Gaussian models are shown in Figure 6.

This experiment demonstrates that an individual piece of glitter is maximally responsive (or brightest) when a light is

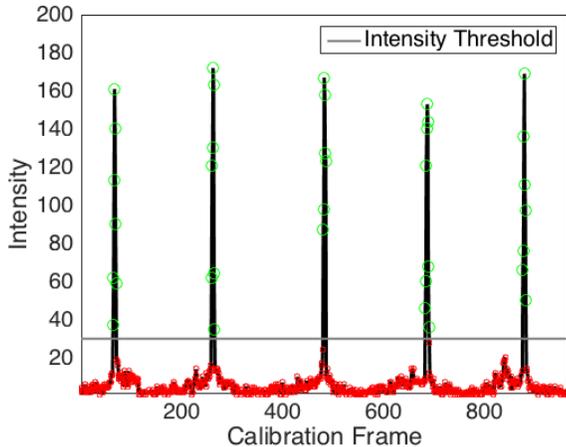


Figure 7: During the calibration procedure described in Section 3, lines are scanned smoothly across the calibration plane in a variety of directions. Shown here is how one piece of glitter responds in image intensity as the scanning line moves into and then out of the area that that piece of glitter reflects. The reflection from the glitter in most intense when the light is directly at the center of the glitter piece’s receptive field. We model the orientation of each glitter piece as the direction from the glitter to the location on the calibration plane. This model yields more accurate position estimation if we threshold our experiments to only include “lit” glitter pieces that are at least as bright as an intensity threshold that excludes pieces of glitter when they’re being lit from too distant of locations.

placed at the center of the Gaussian model of its receptive field. In our experiments, we imagine there is a ray from the glitter location to the center of this receptive field, as discussed in more detail in Section 3. The glitter, however, also responds less intensely to light approximately 6mm away from this center point at a distance of 1000mm between the screen and the glitter, or approximately 0.34 degrees around that center line.

4.2. Sparkle Intensity Thresholding

Because of the variable responsiveness of the glitter pieces as a light moves over their receptive field, we use only the brightest possible sparkles when estimating the position of the light (see Section 5 for more details on position estimation). In order to appropriately set that threshold, we evaluate the intensity of each glitter piece during the calibration sequence.

In Figure 7, we show a standard profile for a piece of glitter as the lines from our calibration sequence scan over the center of its receptive field. As a line approaches that location, the glitter becomes slightly brighter. When the

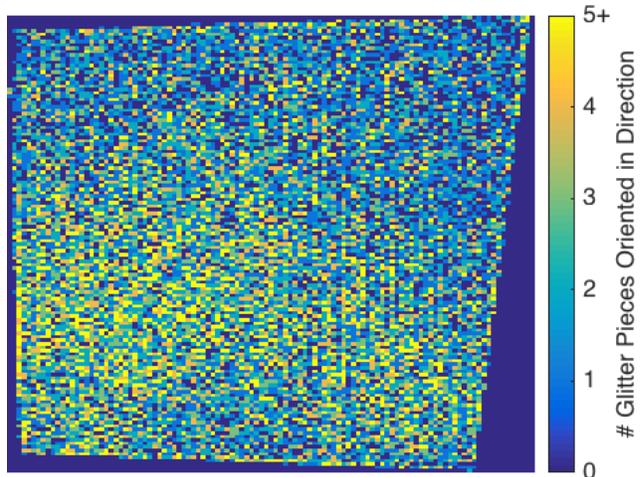


Figure 8: Shown here are the number of sparkles oriented towards different areas of the calibration plane. The calibration plane is displayed at 5mm x 5mm resolution. This resolution is slightly smaller than the receptive field of a single piece of glitter, as described in Section 4, and about the size of the dot that results from shining a laser on the calibration plane, as in the experiments described in Section 5.

line overlaps the glitter’s receptive field the intensity spikes then dims again as the line passes.

As discussed in Section 3, we choose to represent the surface orientation of each piece of glitter as a ray from the glitter center to approximately the center of its receptive field. The maximally correlated location from the calibration procedure may vary slightly from the precise center of the receptive field, as the scan lines used in the calibration process are wider than the points used to determine the receptive field. In order to achieve the best accuracy when determining the position of a light from an image of the glitter sheet, we want to choose only those pieces of glitter which are at their brightest – that is to say, we want to select those pieces of glitter that are being lit by a light closest to the center of their receptive field. We experimentally determine a threshold value of 30, which is shown in Figure 7.

While it would be possible to compute an ideal threshold for each piece of glitter separately, our evaluations demonstrated that a single threshold for every glitter piece is appropriate to retain a sufficient number of sparkles in every frame of position estimation while achieving the improvement in position accuracy that comes from only taking the most responsive sparkles (those which are being lit by a light closest to the center of their receptive field).

4.3. Sparkle Density

Figure 8 shows the number of sparkles that point to different parts of the calibration plane at a 5mm x 5mm resolu-

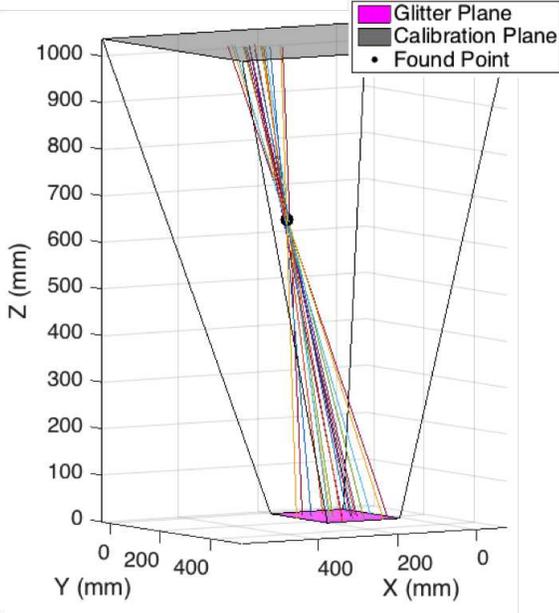


Figure 9: We model the orientation of pieces of glitter as a ray from the glitter piece to the part of the calibration screen that lit that piece of glitter, as described in Section 3. To determine the 3D position of an unknown light source given the set of rays from the pieces of glitter lit by that light, we optimize to find the location closest to all of those rays.

tion. While this figure shows that our calibration procedure yields sparkles pointing in most directions around the calibration plane, especially at the scale of a small laser dot, like those used in our experimental evaluation and shown in Figure 10, the density of the represented sparkle directions can be low. This is a potentially limiting factor in the accuracy of our position estimation.

5. Experiments

In order to determine the position of an unknown light source, we observe which pixel locations on the glitter plane are “lit” above the intensity threshold determined in Section 4. We select from those possible glitter locations the locations that were indexed during the calibration procedure detailed in Section 3. Each of those locations represents a ray from the real world location of the glitter to the location on the calibration plane that lit that piece of glitter during the calibration procedure. We then solve for the 3D position closest to all of these rays, solving a non-linear optimization using Cauchy’s function [10] as a robust variant of least-squares error that is tolerant to outliers. This position is our estimate of the 3D location of the point light source.

In order to evaluate this approach, we shine laser in a known pattern on both the calibration plane and other arbitrary

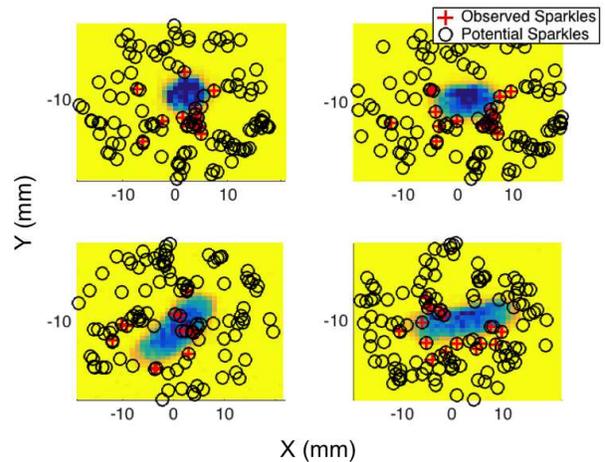


Figure 10: During one of our experiments, we shine a laser on the calibration plane. This laser dot is reflected in the piece of glass used during our calibration procedure to the camera. The background of each image in this figure shows that laser dot projected into real world units on the calibration plane, and the end points of lines from pieces of glitter to nearby locations on the calibration plane. We show both those sparkles which were observed during the frame where the laser was reflected, and those sparkles discovered during the calibration procedure which could have reflected light to the camera but did not.

planes in the region that has been calibrated (the calibrated region is the area between the corners of the glitter and the corners of the part of the calibration plane reflected to the camera by the piece of glass shown in Figure 4).

Figure 10 shows the end points of the rays from the observed glitter, as well as the nearby end points of the rays from every potential calibrated glitter location. These are shown on top of the reflected laser point for four frames where we shined a laser on the calibration plane. The observed points are clustered closely around the laser location. There are some sparkles predicted by our calibration that we do not observe, likely due to intensity thresholding. It is also possible, though less common, for there to be glitter locations that are lit that were not discovered during our calibration process. This can result from pieces of glitter that were too dim during the calibration process, or from locations with multiple pieces of glitter stacked on top of each other, reflecting light in different directions for the same pixel location, which were purposefully eliminated during our calibration procedure.

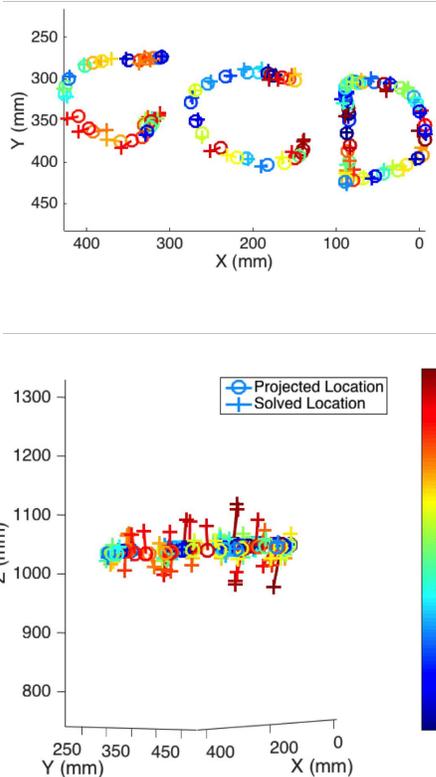


Figure 11: We measure error (in mm) as the distance between the projected location of the laser point (reflected in the calibration glass) and the location solved for from glitter using the approach detailed in Section 5. The average error is 19.6mm, with a median error of 15.2mm and a standard deviation of 15.7mm. Most of the error is along the Z-axis, with significantly smaller errors in the X- and Y-directions, which is consistent with attempting to reconstruct the depth of a point with measurements from a limited baseline.

6. Results

To evaluate our 3D position estimation, we use a laser to create a bright spot on two planes, drawing “CCD” on the plane that was used in the calibration process and “2016” on another plane. The 3D position of the laser was reconstructed in each frame just from the sparkles observed on the glitter sheet. Figure 12 shows two views of the reconstructed points. A fronto-parallel view shows that the text can easily be read, and we show a perpendicular plane designed to visualize the Z-axis error.

We assess the metric error of our technique by using the homography that warps image coordinates to the calibration plane to determine the real world, per-frame position of the laser point reflected from the calibration plane, off of the piece of glass shown in Figure 4, to the camera. Over all points in Figure 12a, the average error is 19.6mm, with a median error of 15.2mm and a standard deviation of

15.7mm. For reference, 15mm is approximately the width of a finger nail. Figure 11 visualizes the reconstruction error of all points, showing the position of the laser dot on the calibration plane connected to the point determined from the glitter reconstruction. We see that the dominant error is along the Z-axis, which is consistent with trying to reconstruct the depth of a point with measurements from a limited baseline.

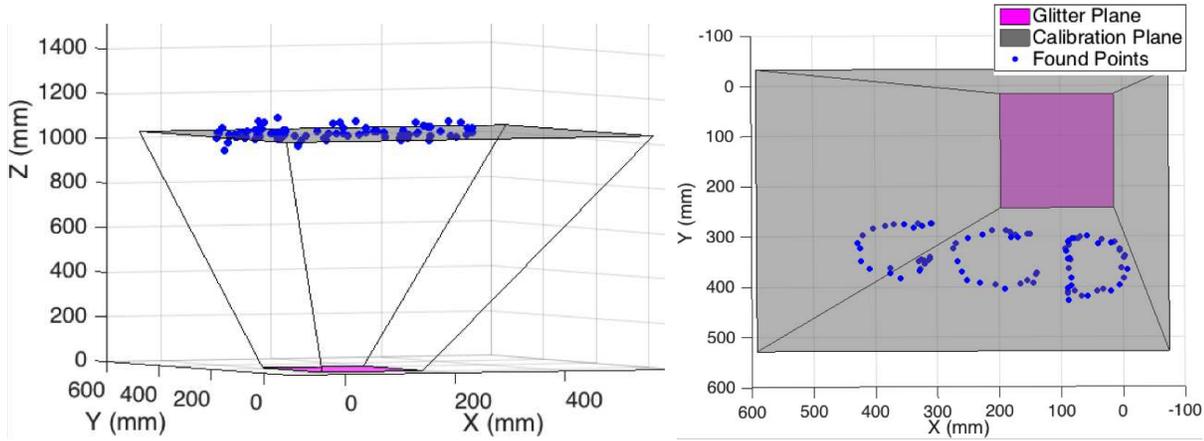
7. Discussion and Conclusions

This paper describes the design, implementation and testing of a geometric imaging system that views the world through glitter for the purpose of geometric inference. With a physical prototype system, we characterized the geometric properties of components of the system including the receptive field for each pixel viewing a piece of glitter and the density with which the glitter field samples a plane in space. For one possible end-to-end approach to estimating the 3D position of a point light source, we show reasonable results and characterize the position estimation error.

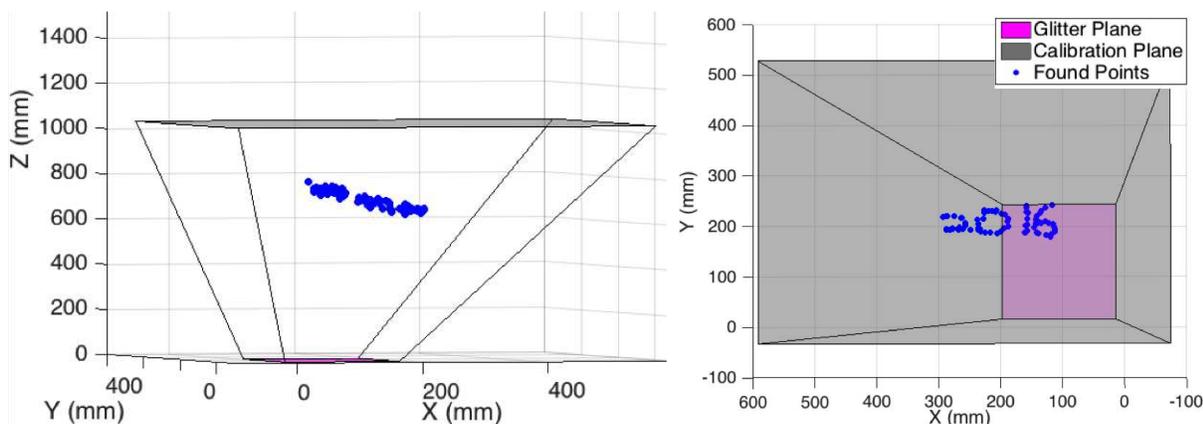
There are several directions of future work. First, the position estimation process now includes only information from lit sparkles, but sparkles that are dark also provide constraints about where a light is not. Second, while we treat sparkles as binary measurements in this work, there is a smooth (if rapid) transition of sparkle brightness when a light enters its field of view, and this could be exploited for additional accuracy. Third, it is possible to constrain the position of multiple lights simultaneously which offers a path towards pose estimation, or to determine the pose of non-discrete light sources such as lines of light.

References

- [1] J. Ackermann, S. Fuhrmann, and M. Goesele. Geometric point light source calibration. In *VMV*, pages 161–168, 2013.
- [2] Y. Adato, Y. Vasilyev, T. Zickler, and O. Ben-Shahar. Shape from specular flow. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(11):2054–2070, 2010.
- [3] T. Aoto, T. Taketomi, T. Sato, Y. Mukaigawa, and N. Yokoya. Position estimation of near point light sources using a clear hollow sphere. In *ICPR*, pages 3721–3724, 2012.
- [4] F. Einabadi and O. Grau. Discrete light source estimation from light probes for photorealistic rendering. In *Proceedings of the British Machine Vision Conference 2015*. BMVA Press, 2015.
- [5] R. Fergus, A. Torralba, and W. T. Freeman. Random lens imaging. Technical report, 2006.



(a) Results from shining a laser in a pattern writing "CCD" on the calibration plane.



(b) Results from shining a laser in a pattern writing "2016" on an angled plane closer to the glitter.

Figure 12: We recover the 3D positions of a laser shined on different surfaces. These positions are found by observing which pieces of glitter sparkle in a frame, and finding the point closest to all of the rays from those pieces of glitter to the locations on the calibration plane that lit those pieces during the calibration procedure.

- [6] M. D. Grossberg and S. K. Nayar. A general imaging model and a method for finding its parameters. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, volume 2, pages 108–115. IEEE, 2001.
- [7] A. Morgand, M. Tamaazousti, and A. Bartoli. Reconstruction d’une source lumineuse modélisée par une quadrique à partir d’images multiples. In *Journées francophones des jeunes chercheurs en vision par ordinateur*, 2015.
- [8] S. K. Nayar, V. Branzoi, and T. E. Boulton. Programmable imaging: Towards a flexible camera. *International Journal of Computer Vision*, 70(1):7–22, 2006.
- [9] J. Park, S. N. Sinha, Y. Matsushita, Y.-W. Tai, and I. S. Kweon. Calibrating a non-isotropic near point light source using a plane. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 2267–2274. IEEE, 2014.
- [10] W. Rey. *Introduction to robust and quasi-robust statistical methods*. Springer Science & Business Media, 2012.
- [11] S. Roth and M. J. Black. Specular flow and the recovery of surface structure. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 1869–1876. IEEE, 2006.
- [12] M. Weber and R. Cipolla. A practical method for estimation of point light-sources. In *BMVC*, pages 1–10, 2001.
- [13] Z. Zhang, P. Isola, and E. H. Adelson. Sparkle vision: Seeing the world through random specular microfacets. *arXiv preprint arXiv:1412.7884*, 2014.