

A Pot of Gold: Rainbows as a Calibration Cue

Scott Workman, Radu Paul Mihail, and Nathan Jacobs

Department of Computer Science
University of Kentucky
{scott, mihail, jacobs}@cs.uky.edu

Abstract. Rainbows are a natural cue for calibrating outdoor imagery. While ephemeral, they provide unique calibration cues because they are centered exactly opposite the sun and have an outer radius of 42 degrees. In this work, we define the geometry of a rainbow and describe minimal sets of constraints that are sufficient for estimating camera calibration. We present both semi-automatic and fully automatic methods to calibrate a camera using an image of a rainbow. To demonstrate our methods, we have collected a large database of rainbow images and use these to evaluate calibration accuracy and to create an empirical model of rainbow appearance. We show how this model can be used to edit rainbow appearance in natural images and how rainbow geometry, in conjunction with a horizon line and capture time, provides an estimate of camera location. While we focus on rainbows, many of the geometric properties and algorithms we present also apply to other solar-refractive phenomena, such as parhelia, often called sun dogs, and the 22 degree solar halo.

1 Introduction

Understanding natural outdoor scenes is challenging because a large number of physical factors affect the imaging process. Natural phenomena provide a variety of cues for estimating camera calibration and understanding scene structure. For example, image haze is a strong cue for inferring scene models [6], as are cloud shadows from a partly cloudy day [7], and the motion of shadows is a cue for camera calibration [33]. We explore another natural cue, the rainbow. Rainbows are a fascinating atmospheric effect—in addition to having strong symbolic meaning, they also have interesting geometric properties.

In particular, the location of a rainbow is exactly constrained by the relative geometry of the sun and the viewer. A rainbow is always centered around the antisolar point (the point exactly opposite the sun), and the outer radius of the rainbow is about 42° from the line connecting the viewer’s eye to the antisolar point. In Figure 1, the projection of the antisolar point is visible as the shadow of the photographer’s head. These strong geometric constraints make the rainbow a powerful calibration object.

In this paper, we explore these constraints both theoretically and practically. First, we characterize the minimal set of constraints necessary to capture the




Fig. 1: The rainbow is an ephemeral, but well-defined, geometric object that can be used to perform camera calibration and provides constraints on camera location. The figure above shows the final result of our calibration method, with the image of the antisolar point (yellow circle) and two color bands (red and blue) on the primary rainbow. (“Double Alaskan rainbow” by Eric Rolph - CC BY-SA 2.5)

relative viewing geometry of the camera and the sun in both calibrated and uncalibrated cases. Second, we introduce methods for estimating this geometry from an image of rainbow, including an image-based refinement technique. Finally, we evaluate the ability of these methods to calibrate a large dataset of real world images and present several use-cases: a data-driven approach to rainbow appearance modeling, rainbow editing and geolocation estimation.

From a practical standpoint, single-image calibration “in the wild” has become an important vision problem. Many cues have been proposed because not every scene has, for example, orthogonal vanishing points or coplanar circles. Rainbows and similar solar refractive phenomena are an important new cue for this problem. While rainbows are rare, there are many rainbow pictures and many webcams will eventually view a rainbow. Rainbows have advantages for calibration: are one of the few calibration cues suitable for “mostly sky” webcams, are easier to localize than the sun (which results in a large oversaturated image region), give strong constraints on the focal length and sun position from a single image, and have more distinctive appearance than sky color gradients.

We focus on rainbows but our geometric framework and analysis applies to other solar-refractive phenomena, e.g., sun dogs and halos, and gives a foundation for future work in using webcams to estimate atmospheric conditions using such phenomena.

1.1 Related Work

Our work introduces a new cue that provides constraints on intrinsic and extrinsic camera calibration. Typical approaches to intrinsic camera calibration rely on either reference objects, such as coplanar circles [3], camera motions, such as camera rotation [32], or both [34]. Extrinsic calibration approaches rely on matching to known, static scene elements [18]. However, when such objects do not exist, and the camera is in an unpopulated area without reference imagery, the problem is more challenging.

Recent interest in calibrating Internet imagery has led to the need for new techniques for intrinsic and extrinsic camera calibration. Much of this work has focused on problems associated with calibrating widely distributed cameras, such as webcams. In this domain, clear-sky appearance has been used to estimate orientation, focal length, and camera location [10,15,13,16]. Other work has explored the use of video from cloudy days for estimating focal length and absolute orientation [7,8]. In addition, photometric and shadowing cues have been used for geolocalization and calibration [19,11,12,9,27]. Our work is most closely related to work on calibration and localization from sky appearance, with the important differences being that rainbows provide much stronger, single-frame constraints on the focal length and have very consistent color properties.

Methods for outdoor appearance modeling are used for applications ranging from compression to scene understanding. The main focus is on modeling the effect of sun motion and weather conditions on scene appearance. Sunkavalli et al. [30] build a factored representation that directly models sun motion and shadows. Subsequent work in this area has sought to extract deeper scene information, such as surface material properties [29]. More recently, the focus has shifted to estimating 3D scene models using photometric cues [1,2]. While most work has focused on static scene elements, there has also been significant research in building models of sky appearance. Lalonde et al. propose to use webcam image sequences to estimate global lighting models [14] for object insertion. Shen et al. [28] estimate local weather conditions. Peng and Chen [23] propose a random field model to estimate per-pixel sky cloudiness. We extend this line of research by including the geometry and appearance of rainbows.

2 Rainbow Image Formation

We briefly describe the physical process which leads to rainbows and the geometric relationship between the camera calibration, sun position and the image location of the rainbow.

2.1 Physical Rainbow Formation

We present the basic aspects of the physical process that creates rainbows, see [22] for additional details. Rainbows are an atmospheric phenomenon prompted by the interplay of light and water droplets. Typically, rainbows begin with a

passing rain shower leaving behind water droplets suspended in the air. In some light paths through a droplet, the ray refracts upon entering, undergoes an internal reflection, and refracts again upon exiting. These light paths generate a rainbow. The amount that a ray of light bends is a function of its wavelength. For example, red light (longer wavelength) bends slightly less than blue and violet. The dispersion of light inside water droplets separates light into its component colors, resulting in a spectrum of light appearing in the sky.

The location of a rainbow depends entirely on the sun position. From the point of view of the observer, the outside of the rainbow (red) is at roughly a 42° angle relative to the antisolar point, the point directly opposite the sun. As we decrease the angular distance from 42° , the colors gradually change from red to violet. The rays from the rainbow to the observer form a conical surface (see Figure 2). While the precise relationship between angle and color depends somewhat on atmospheric scattering, particularly the size and shape of the water droplets [17], we assume spherical drops and use reference angles as computed in [26] (e.g. red = 42.3° and violet = 40.4°).

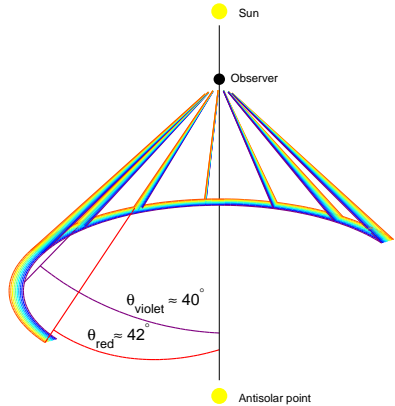


Fig. 2: Rainbows are visual phenomena with relatively simple geometric properties that result from the refraction of sunlight inside of millions of water droplets.

2.2 Geometry of Rainbow Images

We now define the relationship between the image of the antisolar point (IAP) and

a rainbow point with known color. Consider a world point, $P = [X, Y, Z]^T$, that projects to an image location, $\lambda p = [\lambda u, \lambda v, \lambda]^T = \mathbf{K}[\mathbf{R} \mid \mathbf{t}]P$, with extrinsic rotation, \mathbf{R} , and translation, \mathbf{t} . We assume the camera has zero skew, known principal point, square pixels, and is aligned with the world frame. This results in a simple pinhole camera model, $\lambda p = \mathbf{K}P = \text{diag}([f, f, 1])P$, with the focal length, f , as the only unknown.

We define the *absolute angle constraint* which relates the image point on the rainbow, p , and the IAP, s , as follows:

$$p^T \mathbf{K}^{-T} \mathbf{K}^{-1} s = \|\mathbf{K}^{-1} p\| \|\mathbf{K}^{-1} s\| \cos(\theta_p) \quad (1)$$

where, θ_p , is the angle between the rainbow ray, $\mathbf{K}^{-1} p$, and the antisolar point, $\mathbf{K}^{-1} s$. In practice, the camera calibration, \mathbf{K} , and the projection of the antisolar point, s , are unknown and we estimate θ_p from image data.

3 Rainbow to Calibration and Sun Position

The image of a rainbow provides strong constraints on the calibration of the camera and the position of the sun. We begin with an analysis of the constraints in different settings and then describe several alternative calibration methods that build upon them.

3.1 Constraint Analysis

We describe inherent ambiguities and minimally sufficient sets of constraints for two scenarios: one with a fully calibrated camera and one with an uncalibrated camera.

Calibrated Camera When the camera calibration matrix, \mathbf{K} , is known, the location of the rainbow in image space is entirely dependent on the IAP, s . We show that, in the calibrated case, three points at known angles are necessary and sufficient to uniquely identify s .

Consider a set of image points with known angles relative to the antisolar point. With a single point, p , there is a circle of possible solutions, on the view sphere, for the antisolar point, $\mathbf{K}^{-1}s$, which make the required angle, θ_p , with the pixel ray, $\mathbf{K}^{-1}p$. With two distinct points at known angles, there are at most two solutions where the respective circles intersect. Intuitively, this is because the image of the rainbow could be “bent” in two different directions. Therefore, three distinct points at known angles are necessary to uniquely estimate the antisolar point. This minimal set of constraints is visualized in Figure 3. It shows that, in the ideal case, the circle of possible solutions for each point all intersect at a single location, the antisolar point.




Fig. 3: Absolute Angle Constraint: (left) A synthetic rainbow image (FOV = 60°) with image points annotated on different color bands. (right) For each annotated point there is a circle on the view sphere where the antisolar point could be. In the calibrated case, three points at known angles are necessary to solve for an unambiguous antisolar point.

Given these constraints, numerous algorithms could be used to estimate the antisolar point, and hence the sun position, using three, or more image constraints. We have developed analytic methods and nonlinear optimization methods, but we omit them here and instead focus on the uncalibrated case.

Uncalibrated Camera We now consider the case of an uncalibrated camera in which both the IAP, s , and the focal length, f are unknown. We initially focus on defining angular constraints and show that at least three points, at two distinct angles relative to the antisolar point, are required to fully constrain a solution for s and f . In the calibrated case, three points at known angles, which are not necessarily distinct, are needed to guarantee a unique solution for the antisolar point. When the focal length is unknown, we find that to eliminate an ambiguity caused by infinite focal lengths, points must be at two *distinct* angles. This is demonstrated visually in Figure 4.

Given a single point at a known angle, there is a circle of possible solutions on the view sphere, at every possible focal length, that make the required angle with the corresponding pixel ray. With two points at the same angle, there is an interval $[x, \infty)$ of focal lengths where the circles will intersect. At infinite focal length, the pixel rays lie along the optical axis and the circles converge. This shows we need at least two distinct angles to remove infinity as a solution. Therefore, with unknown focal length, the minimal configuration necessary to get an unambiguous solution for f and s is three points and two distinct angles.

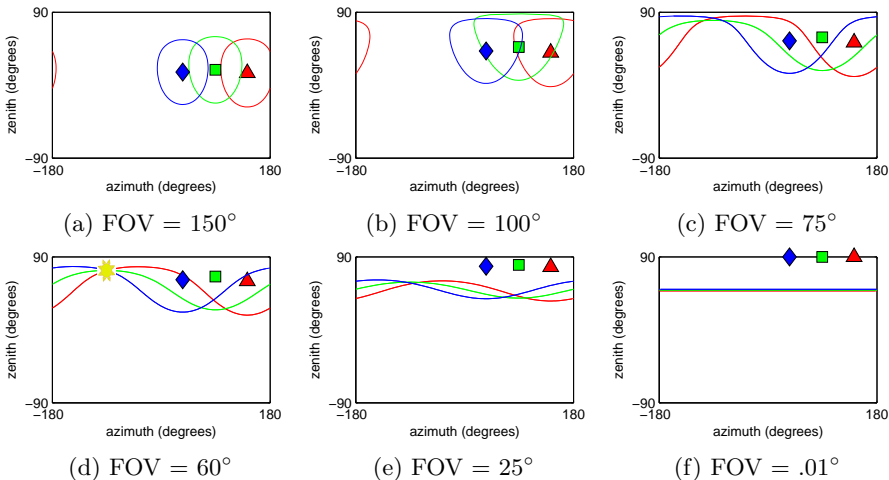


Fig. 4: Absolute Angle Constraint: In the uncalibrated case, the constraints for each annotated point (Figure 3, left) change as the focal length changes. At the correct focal length, all constraints are satisfied at the true antisolar point.

3.2 Estimating Sun Position and Calibration

Three rainbows points with at least two distinct angles are sufficient to guarantee a unique camera calibration solution. We now define several alternative objective functions for solving the calibration problem given the image of a rainbow.

Absolute Angular Error Given a set of image points, $\{p_i\}$, at known angles, $\{\theta_i\}$, relative to the IAP, s , we formulate the following objective function:

$$\operatorname{argmin}_{f,s} \sum_{i=1}^M \left| \cos^{-1} \left(\frac{p_i^\top \mathbf{K}^{-\top} \mathbf{K}^{-1} s}{\|\mathbf{K}^{-1} p_i\| \|\mathbf{K}^{-1} s\|} \right) - \theta_i \right| \quad (2)$$

derived from the absolute angle constraint (1). We first grid sample 20 focal lengths, f , between 1 and 10 image widths, and for each, optimize over the IAP, s . The minimum error configuration is used to initialize a Nelder-Mead simplex search [21] to estimate the focal length and sun position.

Figure 5 shows the shape of this objective function with different numbers of known points and distinct angles. As described in Section 3.1, three points and two distinct angles is the minimal configuration necessary to ensure a unique global minima.




Fig. 5: The solution space of the absolute angular error objective function (2) as the number of points and distinct angles is varied.

Absolute Pixel Error In addition to directly minimizing error based on angular constraints, we can also optimize over image-space distances. Consider once again a set of image points, $\{p_i\}$, at known angles, $\{\theta_i\}$, relative to the IAP, s . We define the following optimization problem:

$$\operatorname{argmin}_{f,s,\{\tilde{p}_i\}} \sum_i \|p_i - \tilde{p}_i\|^2 \quad (3)$$

where \tilde{p}_i is the projection of the image point, p_i , onto the rainbow and where the IAP, s , and the projection of p_i onto the estimated rainbow, \tilde{p}_i , are constrained by (1). We use Nelder-Mead simplex search [21] to estimate the focal length and sun position. We initialize the optimization by fitting a circle, with radius, r , and center, c , to the clicked points: $f_0 = r / \tan(41^\circ)$, $s_0 = c$, and $\{\tilde{p}_i\}_0 = \{p_i\}$.

As compared to the angular error method, this approach has the advantage of functioning correctly with image points at a single fixed angle, it does not have a trivial solution at $f = \infty$ and the objective function more accurately models typical user errors in clicking points.

3.3 Automatic Calibration Refinement

Since it can sometimes be difficult to identify the precise angle for a particular point on a rainbow, we propose a method to automatically refine our manual calibration estimates by maximizing the correlation between the observed rainbow and the expected appearance of a rainbow.

We first estimate the expected appearance of a rainbow from a set of rainbow images with known focal length, f , and IAP, s . For each image, we compute the angle, θ_p , relative to s for each pixel p using (1):

$$\theta_p = \cos^{-1} \left(\frac{p^T \mathbf{K}^{-T} \mathbf{K}^{-1} s}{\|\mathbf{K}^{-1} p\| \|\mathbf{K}^{-1} s\|} \right). \quad (4)$$

We define a rainbow signature as the average color change in the L*a*b* color space, for a given antisolar angle, θ , as we move radially away from the IAP. We use the radial derivative because it is much less dependent on the scene behind the rainbow than the raw intensity. We construct a rainbow signature by quantizing θ (we use 200 bins between $[38^\circ, 44^\circ]$) and averaging the radial image gradients across the image regions that contain a rainbow. To model the expected appearance of a rainbow, we average the signatures for each image in our dataset and obtain an average rainbow signature, $E[\frac{\partial L}{\partial \theta}]$. To reduce blurring in the rainbow signature due to imperfect manual calibration, we sequentially align individual signatures [4] to the average signature until convergence.

Given this average rainbow signature, and a new unseen rainbow, we refine our estimate of the focal length and the IAP by maximizing the average correlation between the signature of each radial line and the average signature using Nelder-Mead simplex search [21]. In practice, directly optimizing over focal length and the IAP failed to converge to a globally optimal solution due to coupling between the parameters. To reduce this coupling, we reparameterize the problem by replacing the IAP, s , by a point on the rainbow at 41 degrees, on the line from the sun position to the principal point. The practical result is that focal length can change without requiring a change in the sun position to keep the rainbow nearby in roughly the same image location. This small change significantly improved our automatic refinement results.




Fig. 6: Rainbow Localization: (left) input rainbow image, (middle) probability of rainbow pixel, and (right) predicted antisolar angle.

3.4 Fully Automatic Calibration

We describe a discriminative approach to rainbow localization and camera calibration that eliminates the need for manually clicking rainbow points. This is challenging because rainbows are transparent, often highly transparent, and their appearance varies due to atmospheric conditions [20], camera optics, CCD sensor properties and software post-processing.

For an uncalibrated rainbow image, we first use random forest regression to estimate two per-pixel labels: the likelihood that the pixel contains a rainbow and the most likely angle the backprojected pixel ray makes with the antisolar point. We use a 7×7 HSV patch and the eigenvalues and first eigenvector of the 2D structure tensor of the hue channel for each pixel as low-level features. We build a training set by manually calibrating a set of images, assigning a label of rainbow or not (manually filtering non-rainbow pixels between 40 and 42 degrees), and computing the antisolar angle for each pixel using (4). An example image and corresponding per-pixel labels can be seen in Figure 6.

Using the per-pixel predictions, our approach is as follows: 1) select the most probable rainbow pixels (top 5 %, selected empirically to filter out false positives), 2) randomly sample three points and use the optimization from Section 3.2, assuming the estimated antisolar angle is correct, to get an estimate of sun direction and focal length, 3) use our image-based refinement technique from Section 3.3. We repeat this process multiple times and use the configuration with lowest error as the final calibration.

4 Applications

We use the geometric properties and algorithms we derived for several applications.

4.1 Calibration of Internet Imagery

We demonstrate the effectiveness of our calibration approaches on a dataset of 90 images we collected from a popular photo sharing site (<http://flickr.com>). We only include images from the iPhone 4, a popular camera phone, because it has a

fixed focal length and has been used to capture and share many rainbow images. This is a realistic and challenging dataset containing many small, often faint rainbows, some barely visible (see Figure 7 for sample images). The dataset, including the results of our methods on all images, is freely available online (<http://cs.uky.edu/~scott/projects/rainbows>).

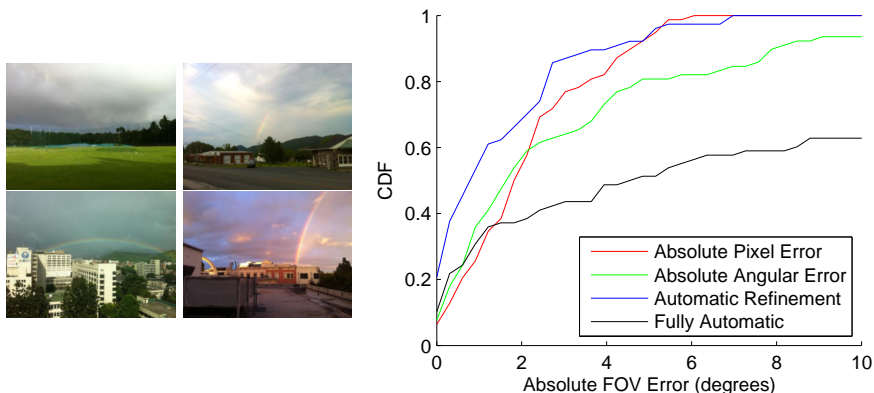


Fig. 7: (left) Example images from our rainbow image dataset. (right) The distribution of errors for various calibration approaches. The method based on pixel error produces fewer outliers. The blue curve shows the result of applying our refinement technique to the output of the pixel error method.

For each image, we manually click points along different color bands and estimate the calibration parameters using the optimization methods from Section 3.2. On average, we annotate 20 points per image (max 39, min 8, $\sigma = 8.4$) on two color bands. See Figure 7 for the results of this experiment, shown in terms of field of view for easier interpretation. We find that the absolute angular error approach gives more accurate results on some images, but the absolute pixel error is more robust. The figure also shows that our image-based refinement technique (Section 3.3), when applied on the output of the absolute pixel error method, improves the calibration estimates relative to calibration based on manual clicks alone. Our fully automatic method (Section 3.4), trained using approximately 300 other rainbow images from Flickr (20 trees, 32 iterations), performs better than the fully manual approaches on some images, but fails dramatically on others. This highlights the difficulty of automatically labeling rainbow pixels in our challenging dataset.

Our implementation, running on a standard desktop PC, takes on average less than a second to perform manual calibration, and 30 seconds for refinement, for a single image.

4.2 Data-Driven Rainbow Appearance Modeling

We use the images in our dataset to construct a rainbow appearance model for the primary rainbow. To our knowledge this is the first attempt to build such a model in a data-driven manner. Previous work has focused on physics-based models [26]. These approaches, while very successful at rendering extreme rainbows, fail to capture the relative distribution of typical rainbow appearance. We show that a data-driven approach can capture this typical appearance, and, as we show in the following section, can be used to exaggerate or diminish the appearance of real rainbows.

We build upon our approach, described in Section 3.3, for estimating the expected rainbow appearance from color changes relative to the antisolar point. Instead of an average image, we estimate a linear basis which captures typical rainbow signature variation. To build this model, we collect rainbow signatures, $\frac{\partial L}{\partial \theta}$, from all images in our dataset. We use the result of our automatic refinement method to estimate the antisolar point and focal length. In Figure 8, we show the marginal distribution of radial gradients, $P(\frac{\partial L}{\partial \theta}|\theta)$, by aggregating radial color derivatives for all radial lines that contain rainbows (manually filtered). This shows the characteristic color changes of a rainbow overlaid with rainbow signatures from two different images. From these we can see that individual rainbows vary in saturation and intensity and that these changes covary from angle to angle. This motivates the use of a Probabilistic Principal Component Analysis (PPCA) model [31] to describe rainbow signatures.

For all images in our dataset, we compute the rainbow signature from our refined calibration estimates, vectorize these signatures and aggregate them into a matrix. From this, we estimate the PPCA decomposition of rainbow appearance. In Figure 8, we show ten rainbow images randomly sampled from our PPCA model. In the following section we show one possible use of this rainbow appearance model.

4.3 Rainbow Editing

We use a gradient-domain editing approach [24], coupled with our PPCA-based rainbow color model, to exaggerate or diminish rainbows in images. We first calibrate the camera then compute the radial and tangential image gradients. For each radial line, we estimate the parameters of the PPCA model that best describe the color derivatives and then subtract these changes from the radial derivatives. We then solve the Poisson equations to find the image that best fits the updated derivatives. Figure 9 shows several examples of rainbows we attempted to edit, including one failure case due to poor initial calibration.

4.4 Video Geolocation

We show how to estimate the geolocation of a static camera from a video containing a rainbow. We build on an existing algorithm [25] that computes sun position (zenith/azimuth angle) for a given time and location. Our approach is

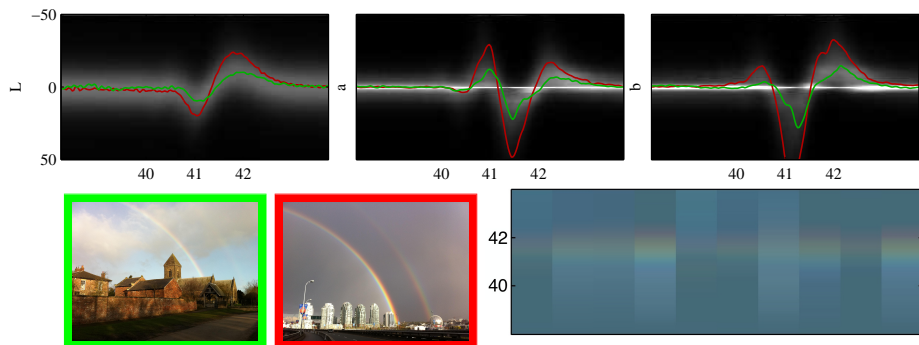


Fig. 8: (top) Conditional color derivative distributions for the $L^*a^*b^*$ color space. (bottom, left) Two rainbow images with rainbow signatures that correspond to two lines in the distribution plots. The green (red) line is the rainbow signature of the left (right) image. (bottom, right) Ten images randomly sampled from our empirical rainbow appearance model overlaid on a blue background image.




Fig. 9: Rainbow image manipulation. Several examples of editing an original image (left), by exaggerating (middle) and diminishing (right) the rainbow using our empirical rainbow appearance model. Poor results occur when we have sub-optimal initial calibration (bottom).

similar to previous work on sun-based localization [5] which is founded upon the relationship between time, sun position and geolocation. Unlike this previous work, our method does not require extensive pre-calibration of the camera. All we need for our method is the image of a rainbow, the capture time of the video, and an estimate of the horizon line.

Given an image, we use our calibration methods to solve for the focal length, f , and IAP, s . We then compute the sun direction, $S = -\mathbf{R}_{\phi\psi}^{-1}\mathbf{K}^{-1}s$, where $\mathbf{K} = \text{diag}([f, f, 1])$ and $\mathbf{R}_{\phi\psi}$ is a rotation matrix that encapsulates camera roll and tilt and is computed directly from the horizon line. Solving for this rotation allows us to compute the solar zenith angle of S relative to the world coordinate system. To estimate the geolocation, we first compute the true sun position for the image capture time [25] on a dense grid of possible geolocations. For each location, we assign a score that reflects how close the image-based estimate of the sun zenith angle, z_{est} , is to the true zenith angle, z , for that location. We use the absolute difference between these values, $|z - z_{est}|$, as our score. We average this score across multiple images from the same video and choose the geolocation with the minimum value.

In Figure 10, we show several localization score maps generated from a time-lapse video, captured in the Czech Republic, that contains both a rainbow and the moon (which provides constraints on geolocation that are very similar to the sun). In the video, the rainbow is visible for only fourteen minutes and the moon for thirty. We sample four frames of each, using our automatic refinement technique to estimate the focal length and antisolar point, and hand labeling the centroid of the moon and the horizon line. Because the moon is in the frame for longer, it provides stronger constraints on location. But, it requires the focal length we estimated from the rainbow to interpret. By combining the localization scores from all frames for both cues, we get a final localization result that clearly highlights the true camera location. Geolocating outdoor cameras is challenging and often requires combining multiple cues. To our knowledge, this is the first work on using rainbows (and potentially other solar-refractive phenomena) for camera localization.

5 Conclusion


We derive constraints and demonstrate methods that allow rainbows to be used for camera geolocation, calibration, and rainbow-specific image editing. These methods exploit the strong geometric cues that rainbows offer through the physics of their formation. This adds to a growing body of work termed “lucky imaging” that exploits occasional and transient atmospheric effects to simplify various image analysis challenges.

Acknowledgements


We gratefully acknowledge the inspiration, support and encouragement of Robert Pless. This research was supported by DARPA CSSG D11AP00255.




(a) Rainbow Frame




(b) Moon Frame



(c) Sun Cue



(d) Moon Cue



(e) Sun and Moon

Fig.10: Localization result from a time-lapse video containing a rainbow and the moon. (a-d) Examples of single frame localization scores. (e) The final localization result obtained by combining scores from individual frames. (Video courtesy of Martin Setvak)

References

1. Abrams, A., Hawley, C., Pless, R.: Heliometric stereo: Shape from sun position. In: European Conference on Computer Vision (2012)
2. Ackermann, J., Langguth, F., Fuhrmann, S., Goesele, M.: Photometric stereo for outdoor webcams. In: IEEE Conference on Computer Vision and Pattern Recognition (2012)
3. Chen, Q., Wu, H., Wada, T.: Camera calibration with two arbitrary coplanar circles. In: European Conference on Computer Vision (2004)
4. Cox, M., Sridharan, S., Lucey, S., Cohn, J.: Least squares congealing for unsupervised alignment of images. In: IEEE Conference on Computer Vision and Pattern Recognition (2008)
5. Cozman, F., Krotkov, E.: Robot localization using a computer vision sextant. In: International Conference on Robotics and Automation (1995)
6. He, K., Sun, J., Tang, X.: Single image haze removal using dark channel prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2011)
7. Jacobs, N., Bies, B., Pless, R.: Using cloud shadows to infer scene structure and camera calibration. In: IEEE Conference on Computer Vision and Pattern Recognition (2010)
8. Jacobs, N., Islam, M., Workman, S.: Cloud motion as a calibration cue. In: IEEE Conference on Computer Vision and Pattern Recognition (2013)
9. Jacobs, N., Miskell, K., Pless, R.: Webcam geo-localization using aggregate light levels. In: IEEE Workshop on Applications of Computer Vision (2011)
10. Jacobs, N., Roman, N., Pless, R.: Toward fully automatic geo-location and geo-orientation of static outdoor cameras. In: IEEE Workshop on Applications of Computer Vision (2008)
11. Jacobs, N., Satkin, S., Roman, N., Speyer, R., Pless, R.: Geolocating static cameras. In: IEEE International Conference on Computer Vision (2007)
12. Junejo, I.N., Foroosh, H.: Estimating geo-temporal location of stationary cameras using shadow trajectories. In: European Conference on Computer Vision (2008)
13. Lalonde, J.F., Narasimhan, S.G., Efros, A.A.: What does the sky tell us about the camera? In: European Conference on Computer Vision (2008)
14. Lalonde, J.F., Efros, A.A., Narasimhan, S.G.: Webcam clip art: Appearance and illuminant transfer from time-lapse sequences. In: *ACM Transactions on Graphics* (2009)
15. Lalonde, J.F., Narasimhan, S.G., Efros, A.A.: Camera parameters estimation from hand-labelled sun positions in image sequences. Tech. rep., CMU Robotics Institute (2008)
16. Lalonde, J.F., Narasimhan, S.G., Efros, A.A.: What do the sun and the sky tell us about the camera? *International Journal of Computer Vision* (2010)
17. Lee, R.L.: Mie theory, airy theory, and the natural rainbow. *Applied Optics* (1998)
18. Li, Y., Snavely, N., Huttenlocher, D., Fua, P.: Worldwide pose estimation using 3d point clouds. In: European Conference on Computer Vision (2012)
19. Lu, F., Cao, X., Shen, Y., Foroosh, H.: Camera calibration from two shadow trajectories. In: International Conference on Pattern Recognition (2006)
20. McCartney, E.J., Hall, F.F.: Optics of the atmosphere: Scattering by molecules and particles. *Physics Today* (1977)
21. Nelder, J.A., Mead, R.: A simplex method for function minimization. *The Computer Journal* (1965)
22. Nussenzveig, H.M.: The theory of the rainbow. *Scientific American* (1977)

23. Peng, K.C., Chen, T.: Take your favorite cloud with you. In: IEEE International Conference on Computer Vision (2013)
24. Pérez, P., Gangnet, M., Blake, A.: Poisson image editing. *ACM Transactions on Graphics (TOG)* (2003)
25. Reda, I., Andreas, A.: Solar position algorithm for solar radiation applications. *Solar energy* (2004)
26. Sadeghi, I., Munoz, A., Laven, P., Jarosz, W., Seron, F., Gutierrez, D., Jensen, H.W.: Physically-based simulation of rainbows. *ACM Transactions on Graphics (TOG)* (2012)
27. Sandnes, F.E.: Determining the geographical location of image scenes based on object shadow lengths. *Journal of Signal Processing Systems* (2011)
28. Shen, L., Tan, P.: Photometric stereo and weather estimation using internet images. In: IEEE Conference on Computer Vision and Pattern Recognition (2009)
29. Sunkavalli, K., Romeiro, F., Matusik, W., Zickler, T., Pfister, H.: What do color changes reveal about an outdoor scene? In: IEEE Conference on Computer Vision and Pattern Recognition (2008)
30. Sunkavalli, K., Matusik, W., Pfister, H., Rusinkiewicz, S.: Factored time-lapse video. *ACM Transactions on Graphics (SIGGRAPH)* (2007)
31. Tipping, M.E., Bishop, C.M.: Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* (1999)
32. Wang, L., Kang, S.B., Shum, H.Y., Xu, G.: Error analysis of pure rotation-based self-calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2004)
33. Wu, L., Cao, X., Foroosh, H.: Camera calibration and geo-location estimation from two shadow trajectories. *Computer Vision and Image Understanding* (2010)
34. Zhang, Z.: A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2000)