Catoptric Surface Characteristics and Visual Feedback Control

Samatha Kodali
Jeremy Manin
Lissette Torres-Escobedo
Run Zhang
Chandler Ahrens
Chris Gill
Roger D. Chamberlain

Samatha Kodali, Jeremy Manin, Lissette Torres-Escobedo, Run Zhang, Chandler Ahrens, Chris Gill, Roger Chamberlain, "Catoptric Surface Characteristics and Visual Feedback Control," in *Proc. of 19th IEEE Conference on Industrial Electronics and Applications* (ICIEA), August 2024. DOI: 10.1109/ICIEA61579.2024.10665117

Washington University in St. Louis

Catoptric Surface Characteristics and Visual Feedback Control

Samatha Kodali*, Jeremy Manin*, Lissette Torres-Escobedo*, Run Zhang*,
Chandler Ahrens, Chris Gill, Roger D. Chamberlain
Washington University in St. Louis
St. Louis, MO, USA
{kodalis,jmmanin,l.torres,run,cahrens,cdgill,roger}@wustl.edu

***Abstract*—The distribution of natural light in an interior space can be controlled via an array of mirrors (a catoptric surface), e.g., mounted on the interior of a large, southern-facing window. We describe the operational characteristics of a catoptric surface constructed on campus, including the use of computational vision techniques for active control of individual mirrors and for ensuring safe operation. We show that: (1) the approach we use can be effective, and (2) there is significant sensitivity to external factors requiring additional development for it to be robust.**

I. Introduction

The illumination of an environment with natural light provides tangible health benefits, economic benefits, and energy benefits [1], [2], [3], [4], [5], [6], [7], [8], [9]. By exposing individuals to more natural light, they have lower levels of cortisol as well as higher levels of melatonin at 10pm, both of which result in lower levels of depressive symptoms and an increased quality of sleep [10]. Furthermore, the use of natural light can alter one's spatial perception in the environment [11]. The beneficial effects of natural light supplies designers with another mechanism through which to alter the experience of occupants in an interior environment.

Figure 1 is an image of a catoptric surface installed in the atrium of Steinberg Hall on the campus of Washington Univ. in St. Louis. In the image, the camera is facing east, the southern exposure is to the right, and the building atrium is to the left. The surface (almost all of which is visible in the image) is comprised of just over 600 mirrors that are each individually under pan/tilt control. By judicious orientation of the mirrors, one can send natural light into the building's atrium, positioning it where desired.

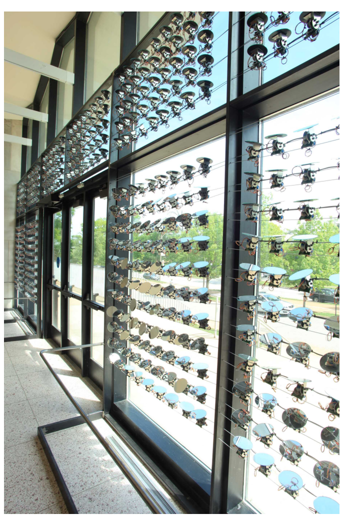

Fig. 1. Catoptric surface.

*Equal contribution; ordered alphabetically by last name.

The goals to be achieved with the control of natural light can be multi-fold [12], a number of which are quite relevant in an industrial environment. First, we can alter the illumination levels in the interior space so that individuals that desire more natural light (e.g., reading a paper report) can have a higher light level while individuals that desire less natural light (e.g., reading a laptop screen) can have a lower light level. This goal is relevant to an office workspace. Second, we can provide natural light to plants that are located in the space to facilitate their health. Third, we can position spots of light in specific locations on the ceiling or wall to create an artistic image in light. This goal is appropriate wherever we wish to improve the environment for the people present. Fourth, we can target excess light at components of the HVAC system (e.g., a heat exchanger) to harvest energy and thereby decrease heating costs. Finally, on a factory floor, it can be quite advantageous to fine-tune the ambient lighting, particularly with natural light, as has been documented for 100 years [13].

To help reduce the cost associated with so many mirrors, the pan/tilt control is implemented using inexpensive stepper motors that do not incorporate shaft encoders for positive feedback of their position. As a result, there is a need for an alternative mechanism to provide positional feedback to the control system that is managing the mirrors' orientation. This paper investigates the application of computer vision techniques to provide this feedback. Given positional feedback, the control system can then ensure that there is never too much light at any individual position. Our goal is not to develop new computer vision algorithms, but rather to assess the applicability of existing algorithms to the application of providing positional feedback to the control system of the catoptric surface.

The outline of the paper is as follows. Section II provides the operational characteristics of the catoptric surface and its components. This motivates the use of computer vision in the system. Section III articulates the algorithms used and effectiveness of using images of the mirrors to assess their current orientation. Section IV follows with a discussion of using images of the reflected light for fine-tuned pointing control. Section V concludes and provides direction for the further work needed to enable the system to operate safely and effectively.

II. Operational Characteristics

The initial operational question for the catoptric surface is, does it work? I.e., does it provide enough difference in light levels to make a difference. Figures 2 and 3 show images taken between 11am and 1pm on a sunny day in the building atrium, with the camera facing north and positioned in the center of the south wall (in the doorway visible in Figure 1, pointing to the left). In Figure 2, the mirrors have been commanded to direct light away from the ceiling (to lower the illumination level in the room), and in Figure 3, the mirrors have been commanded to illuminate the ceiling (to raise the illumination level).

Fig. 2. Non-illuminated ceiling.

Fig. 3. Illuminated ceiling.

A light meter was positioned on a table between the support columns, and Figure 4 shows the results of those measurements. Clearly, the room was significantly brighter at the position of the light meter when the mirrors were illuminating the ceiling right above it. This effectiveness was corroborated by the informal perceptions of the individuals in the room. The space was perceived to be noticeably brighter when the ceiling was being illuminated using the catoptric surface.

The next set of experiments were aimed at quantifying the ability to perform fine-tuned positioning of individual mirrors. Here, the test setup was a 2×2 prototype system in which the 4 mirrors were mounted using the same cabling system as used in the full installation. The experiment only used one of the four mirrors. An LED pointer was aimed at the center of the mirror and the mirror oriented so that (initially) the LED light was normally incident to the mirror surface. The mirror was then panned fully to the left, reaching the hard stop provided by the mount.

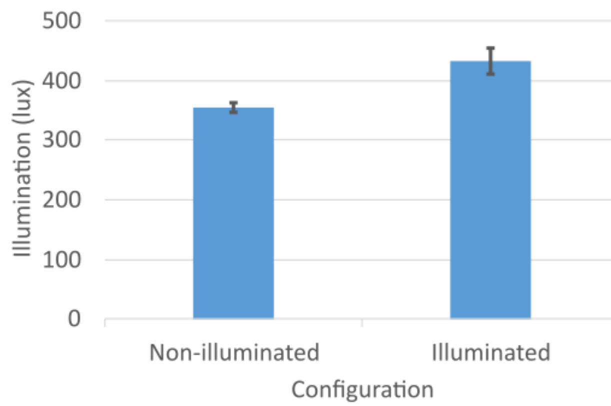

Fig. 4. Illumination data.

Figure 5 shows the raw data representing reflected LED light positions on a wall 2 m in front of the mirror as the mirror was panned first to the right (10 steps, pause, record data) and then back to the left (10 steps, pause, record data). The grid on the wall is marked with ⅛-in. spacing. In the figure, the blue dots represent measurements moving left-to-right and the orange dots represent measurements moving right-to-left. If the mechanical subsystem were free of hysteresis, the blue and orange dots would be at the same positions. Clearly this is not the case, there is a noticeable amount of hysteresis in the mirror movement mechanics.

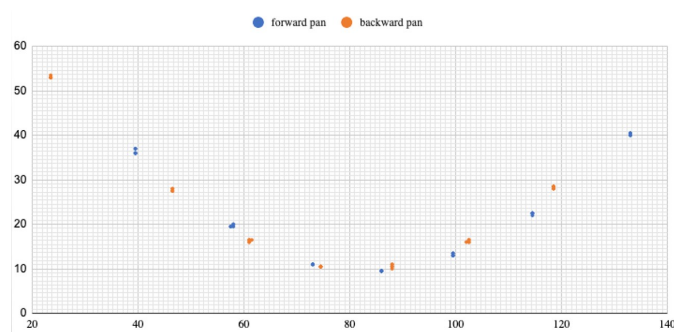

Fig. 5. Raw data illustrating hysteresis.

The implications of this are significant. To accurately control the mirrors' orientation, a feedback mechanism is needed to indicate the actual orientation of each mirror. Even more importantly, the lack of knowledge (on the part of the control software) of the precise orientation of each mirror implies that the system is inherently unsafe. If too many mirrors point sunlight at a specific position, it is quite possible to generate a heating hazard that could be damaging to whatever is at that position.

We address these issues with a pair of applications of computer vision techniques. First, we image the mirrors in the catoptric surface so as to ascertain their orientation via direct observation. We use this information for course-grained control of each mirror's orientation. Second, we image the reflected surface to enable a fine-tuning of the orientation. Our approaches to computer vision for each of these applications are described in the following two sections.

III. Imaging the Mirrors

We first consider the application in which the mirrors themselves are in the image, and we are trying to discern the orientation of each mirror. Here, we only consider the problem of orientation discernment, leaving the problem of image segmentation for later work (which should be straightforward given that the geometry of the surface is fixed and well known).

As has already been reported, computer vision is made more difficult with mirrors in the field of view [14], [15]. We illustrate this in Figure 6, which shows three images of the same mirror in three distinct orientations. What happens, as one would expect, is that the pixels within the mirror's surface are dramatically impacted by the mirror's orientation.

Fig. 6. Example mirror orientations.

While the fact that we are imaging mirrors is a challenge, we are helped, however, by the fact that the objects we are seeking to identify and characterize are constrained in size and shape. We know the distance to the mirrors and also know that the actual geometry of the mirror itself does not change, just its orientation. We will exploit these advantages in the image analysis described next.

A. Image Analysis

Given the complications that can ensue when attempting to understand reflected imagery in a mirror, our approach is to avoid analysis of the interior of the mirror as much as is reasonable. To do this, we search for ovals within the image, since independent of orientation the mirror's boundary will present as an oval, and then limit our consideration to the boundaries of the oval.

The basic steps in the image analysis pipeline we attempt when imaging the mirrors are edge detection, connected components, and random sample consensus. As stated above, we are not seeking to develop new computer vision algorithms, but rather to investigate the use of existing algorithms for our application. The steps in the pipeline are illustrated in Figure 7, which has the detected oval boundary superimposed on the original image on the far right and the output of the various steps moving left to right. This image is of a prototype pan/tilt unit mounted on a wooden base.

Edge detection is a fundamental step in many computer vision tasks, including object detection, image segmentation, and feature extraction. It involves identifying points in an image where there is a significant change in intensity or color, which often correspond to object boundaries or other important image features. Popular edge detection methods include the Canny edge detector [16], Sobel operator [17], and the Roberts operator [18], which have been widely used in computer vision research and applications.

For our application, the Canny edge detector was quite effective. The Canny edge detector is a multi-stage algorithm that involves smoothing the image with a Gaussian filter, calculating the gradient magnitude and direction, suppressing non-maximum edges, and finally, applying hysteresis thresholding. The result of the Canny edge detector is a binary image where the edges are white and the non-edges are black. In Figure 7, the leftmost image is the output of edge detection.

Connected components [19] are groups of pixels or image regions that are connected to each other based on certain criteria, such as having the same intensity value or being adjacent to each other. Since the output of our edge detection algorithm is a binary image, the criteria we use is adjacency. In the figure, the center image shows the output of the connected components step.

The Random Sample Consensus (RANSAC) algorithm, proposed by Fischler and Bolles in 1981 [20], is a widely used robust estimation technique for fitting models to noisy data. RANSAC is particularly suitable for applications where the data may contain outliers or noise, such as in our context. RANSAC iteratively selects random subsets of data points, fits a model to each subset, and then evaluates the goodness-of-fit of the model. The process is repeated multiple times, and the best-fitting model is selected as the final result. In the example of Figure 7, the rightmost image shows the output of RANSAC in red superimposed on the original image.

B. Results and Discussion

Figure 8 shows a pair of images (left and center) in which the algorithm does a very good job of identifying the oval that bounds the mirror. The detected oval is shown in blue, superimposed on the original image. This is exactly what we desire. Even though there are significant artifacts in the mirror reflection, the algorithm is robust enough to not be misled by them.

While this result happens for a good fraction of the images tested, that is not the complete story. The right image in the figure illustrates what can happen in extreme circumstances. Here, the mirror is tilted sufficiently high that the face of the mirror is not visible in the image, and rather than identifying the oval that represents the back of the mirror, another arbitrary oval is detected instead.

Upon further investigation, it becomes clear that the effectiveness of the mechanism depends significantly on the parameterization of the constituent algorithms, especially in

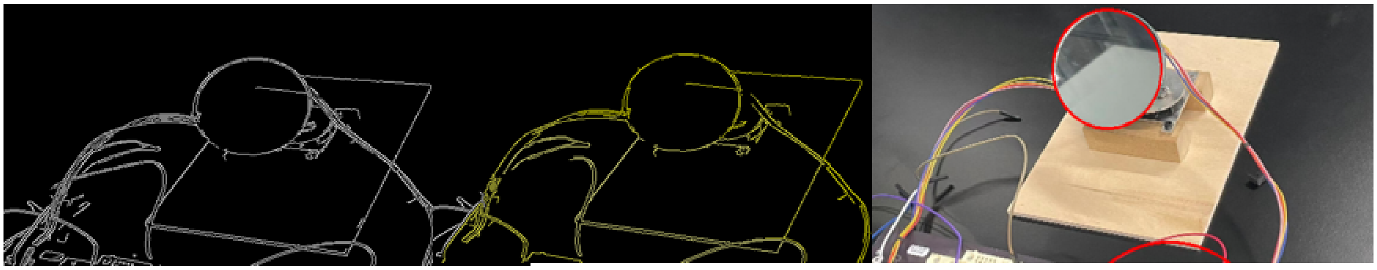

Fig. 7. Set of images illustrating the output of edge detection, connected components, and RANSAC.

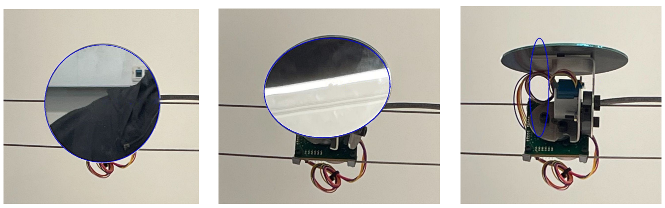

Fig. 8. Good (left, center) and poor (right) identifications.

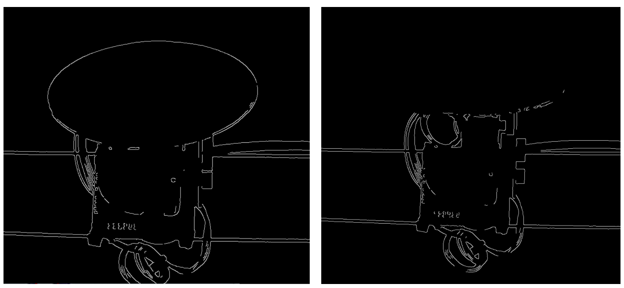

Fig. 9. Two examples of output from edge detection.

edge detection and RANSAC. This is illustrated nicely in Figure 9, which shows two results of edge detection with different tilt angles for the mirror. In the image on the left, the oval that represents the boundary of the mirror is readily evident, while in the image on the right, it is completely lost. In this example, the tilt angle on the right resulted in a lighter reflection in the mirror, and the threshold parameters in the edge detection did not support the identification of the mirror's boundary.

One optimization we did include is to incorporate some geometric constraints on the detected ovals. We restrict the eccentricity of the ellipse that defines the oval, which does reject a number of negative results. Given knowledge of the positioning of the camera relative to the mirrors, it should be possible to add additional geometric constraints to the pipeline. These might include absolute constraints on the dimensions of the ellipse (e.g., the major axis length) in addition to the eccentricity of the ellipse.

In summary, the computer vision techniques we have investigated do a reasonably good job identifying the orientation of mirrors the majority of the time. Interestingly, this is arguably sufficient for the purpose of ensuring safe operation, as it takes the majority of mirrors pointed at an individual point for unsafe operation, so as long as the vision algorithm works more often than not, safety can be ensured. However, there are examples in which the results are poor, necessitating additional investigation before we could consider the approach sufficiently robust for acceptable deployment for orientation feedback in our application.

IV. Imaging the Reflected Light

Concurrently with the investigation above, we also assessed the application of computer vision techniques when imaging the spots of reflected light. Observing Figure 3, it is apparent that the reflected natural light shows up on the ceiling primarily as circles of light with significant contrast relative to the portions of the ceiling that are not illuminated. Our goal here is to use these reflected natural light images to guide the fine-tuned control algorithm, enabling the positioning of light as precisely as the physical apparatus allows. The notion is that the imaging of the mirrors will enable the control system to orient the mirrors so that the reflected spot of light is close to the desired target, and the imaging of the reflected light will enable the final positioning of the spot of light.

The experimental setup in the laboratory consists of one of the mirrors in the 2×2 prototype system, a high-intensity artificial light source, and a black background that contains a target for the desired light position. An example image is shown in Figure 10, which has the reflected spot of light near the bottom of the image, the target just above the center of the image, and a red line superimposed showing the difference in position between the two.

A. Image Analysis

At a high level, the imaging approach is to mask the target and identify the center of the reflected light spot. The target is masked so that it doesn't interfere with identification of the reflected light spot, which is possible since the target's position is static and known beforehand. The target's location in the image is determined using scale-invariant template matching.

The initial imaging problem addressed here is the identification of the brightest spot in the image. For this, we use the technique described by Rosebrock [21], which includes a blurring stage as a pre-processing step. This need for blurring is illustrated in Figure 11, where the image on the left has no blurring and the image on the right has been blurred. In both images, the brightest point detected is marked with a blue

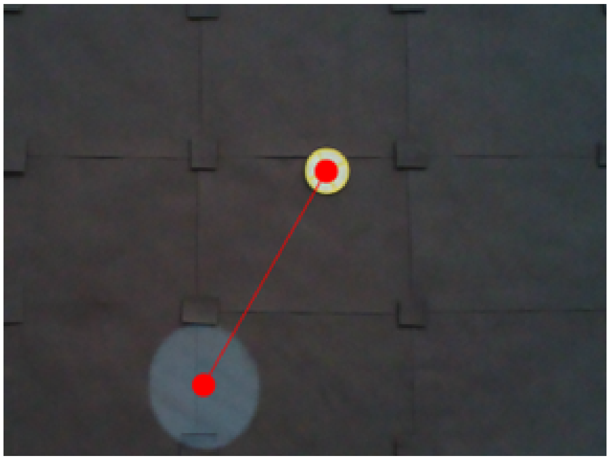

Fig. 10. Reflected-light control algorithm, initial condition.

circle. Note that in the zoomed region in the left image the algorithm focused on a single, abnormally bright point on the edge of the scene as opposed to the reflected light.

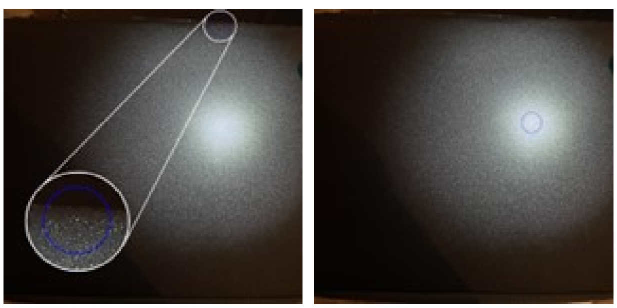

Fig. 11. Brightspot detection.

This method is useful in finding the single brightest point on an image. However, when considering a scene with a reflected light on it, the single brightest point is not necessarily the center point of the reflected light. This is because the brightness of the reflected light is not uniform across the entire region of reflected light. The remainder of the development on bright region detection is focused on finding the center point of the reflection using the brightest spot as a basis.

Region detection proceeds by thresholding the image with the threshold set to a fixed fraction of the brightest point. This is then morphologically closed [22] to address gaps in the region of interest. The center point of the largest resulting contour is then identified as the center of the region of interest.

A set of images showing how the bright region detection algorithm works is provided in Figure 12. Shown from left to right are the original scene; the blurred, greyscale scene; the masked scene; the masked scene after morphological closing has been applied; and the result showing the centers of the reflected light and goal marker.

The above technique works quite reliably when the ambient light in the room is at a fixed level, which is easy to maintain in the laboratory. It is less robust, however, as the ambient light

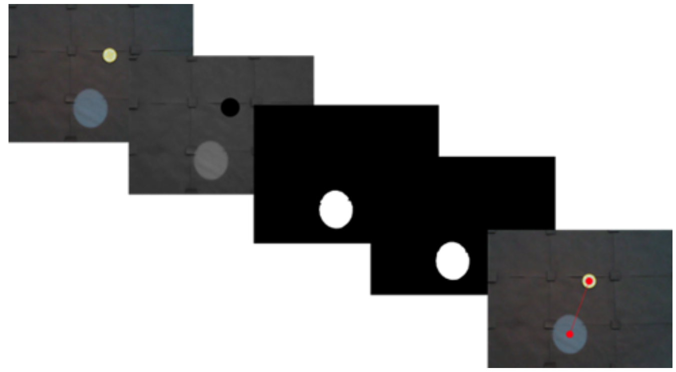

Fig. 12. Bright region detection algorithm steps.

varies, which will be the case in the actual installation. There must be sufficient contrast between the reflected light spot and the target surface for the spot to be identified. Continuing in the laboratory setting, we next describe the control of the mirror itself.

B. Mirror Control

For initial testing, a simple control algorithm is adopted. Once the location of the reflected light is known, the distance between the target and the center of the reflected spot is computed. The axis with the greatest absolute distance (in pixels) is determined, and the mirror is commanded to move in that axis a fixed amount in the direction that diminishes the distance. While the axes that are present in the image do not correspond precisely to the control axes for the mirror, this simple approximation of correspondence works quite well in practice.

We next illustrate the control algorithm in action. Starting from the initial condition illustrated in Figure 10, the sequence of positions (one per control iteration) is shown in Figure 13 (with each bright region represented by a red dot and a red line indicating the separation from the target, in yellow). The initial position is closest to the bottom of the image. Subsequent positions move up and to the left; then up and to the right; followed by down and to the left.

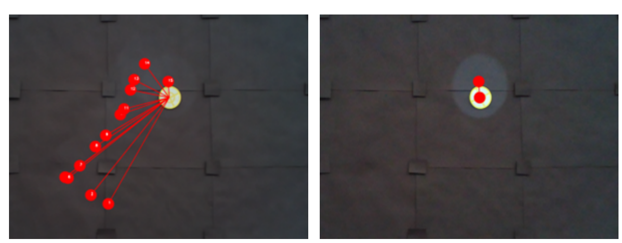

Fig. 13. Control algorithm, sequence of positions (left) and end condition (right).

As can be seen in the figure, commands to the mirror to move up (tilt up) also cause the reflected light to move horizontally as well. Similarly, commands to the mirror to move right (pan right) also cause the reflected light to move vertically as well. Nonetheless, the resulting orientation of the mirror has the reflected spot of light overlapping the target. The final image, after convergence of the control algorithm, is shown in on the right.

C. Results and Discussion

Under controlled lighting conditions, the techniques described here worked quite well at analyzing the reflected light imagery and controlling the mirror's orientation. The challenge going forward will be to make the procedure robust to more dynamic ambient lighting.

For bright region detection, an iterative approach to the current method could be explored to increase its robustness. In this iterative approach, a weighted score based on the area of the largest region and the distance from the known brightest point could be used as the threshold is gradually increased to find the optimal threshold value.

While the simple control algorithm employed was successful, one improvement worthy of investigation is to have a dynamic step size in the mirror's motion. This could substantially decrease the time required to orient the mirror.

V. Conclusions and Future Work

This paper investigates the use of existing computer vision techniques to the application of mirror orientation feedback for a catoptric surface. Imagery of the mirrors and imagery of the reflected light are both considered, and the conclusions are similar for both circumstances. In both cases, the techniques are effective, but not yet robust. When imaging the mirrors, there is a strong sensitivity to the parameter settings within the algorithms, especially in edge detection and in random sample consensus. When imaging the reflected light, there is strong sensitivity to the ambient light in the room. Further investigation is needed to address both of these issues.

Also, both imaging techniques are studied here for the case of a single mirror or a single reflected spot of light. They need to be generalized to the multiple mirror and multiple spot cases. Given the fixed position of the mirrors relative to the camera, segmenting an image into single-mirror subimages should be straightforward. Generalizing to multiple reflected spots will be more of a challenge.

When imaging the mirrors, we plan to investigate autotuning mechanisms for parameter setting, incorporating additional geometric constraints into the algorithms that are invariant for our problem context. For example, given the fixed size of the mirrors, there is a minimum dimension for the long axis of the edge ellipse that is independent of orientation. By refusing to accept detected ovals that don't meet this criterion, we can adjust algorithm parameters until the constraint is met.

When imaging the reflected light, we plan to investigate edge detection techniques for recognition of the bright region. In addition, more sophisticated control algorithms are needed.

A final concern that is relevant to both sets of imagery is that when using a camera in a public space, it is highly likely that individuals will at least occasionally end up in the field of view. This clearly can be a privacy concern, and we plan to explore the use of the Viola-Jones recognition approach [23] to identify and mask any faces detected in any of the imagery.

References

[1] I. Acosta, M. Á. Campano, R. Leslie, and L. Radetsky, "Daylighting design for healthy environments: Analysis of educational spaces for optimal circadian stimulus," *Solar Energy*, vol. 193, pp. 584–596, 2019.

[2] M. Bodart and A. De Herde, "Global energy savings in offices buildings by the use of daylighting," *Energy and Buildings*, vol. 34, no. 5, pp. 421–429, Jun. 2002.

[3] L. Edwards and P. Torcellini, "Literature review of the effects of natural light on building occupants," National Renewable Energy Lab., Golden, CO, USA, Tech. Rep. NREL/TP-550-30769, Jul. 2002.

[4] M. Figueiro, M. Kalsher, B. Steverson, J. Heerwagen, K. Kampschroer, and M. Rea, "Circadian-effective light and its impact on alertness in office workers," *Lighting Research & Technology*, vol. 51, no. 2, pp. 171–183, 2019.

[5] M. G. Figueiro and M. S. Rea, "Office lighting and personal light exposures in two seasons: Impact on sleep and mood," *Lighting Research & Technology*, vol. 48, no. 3, pp. 352–364, 2016.

[6] P. Ihm, A. Nemri, and M. Krarti, "Estimation of lighting energy savings from daylighting," *Building and Environment*, vol. 44, no. 3, pp. 509–514, Mar. 2009.

[7] R. P. Leslie, "Capturing the daylight dividend in buildings: why and how?" *Building and Environment*, vol. 38, no. 2, pp. 381–385, 2003.

[8] A. Mohammadi Tabar, M. Moradi, and R. Fayaz, "Effects of interior architecture for optimal use of natural light and electrical energy saving," *Energy Sources, Part A: Recovery, Utilization, and Environmental Effects*, vol. 44, no. 1, pp. 2415–2429, 2022.

[9] A. Pellegrino, S. Cammarano, V. R. L. Verso, and V. Corrado, "Impact of daylighting on total energy use in offices of varying architectural features in Italy: Results from a parametric study," *Building and Environment*, vol. 113, pp. 151–162, Feb. 2017.

[10] F. Harb, M. P. Hidalgo, and B. Martau, "Lack of exposure to natural light in the workspace is associated with physiological, sleep and depressive symptoms," *Chronobiology International*, vol. 32, no. 3, pp. 368–375, 2015.

[11] C. Ahrens, R. Chamberlain, S. Mitchell, A. Barnstorff, and J. Gelbard, "Controlling daylight reflectance with cyber-physical systems," in *Proc. of 24th Int'l Conf. on Computer-Aided Architectural Design Research in Asia*, vol. 1, Apr. 2019, pp. 433–442.

[12] R. D. Chamberlain, C. Ahrens, C. Gill, and S. A. Mitchell, "Work-in-progress: Hierarchical control of a catoptric surface," in *Proc. of International Conference on Embedded Software*, Oct. 2018.

[13] L. Gaster, "The illumination of factories," *The Journal of State Medicine (1912-1937)*, vol. 32, no. 12, pp. 574–581, 1924.

[14] J. Lin, G. Wang, and R. W. Lau, "Progressive mirror detection," in *Proc. of IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Jun. 2020, pp. 3694–3702.

[15] X. Yang, H. Mei, K. Xu, X. Wei, B. Yin, and R. Lau, "Where is my mirror?" in *Proc. of IEEE/CVF International Conference on Computer Vision*. IEEE, Oct. 2019, pp. 8808–8817.

[16] J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-8, no. 6, pp. 679–698, Nov. 1986.

[17] N. Kanopoulos, N. Vasanthavada, and R. L. Baker, "Design of an image edge detection filter using the Sobel operator," *IEEE Journal of Solid-state Circuits*, vol. 23, no. 2, pp. 358–367, Apr. 1988.

[18] L. G. Roberts, "Machine perception of three-dimensional solids," Ph.D. dissertation, Massachusetts Institute of Technology, 1963.

[19] C. Ronse and P. A. Devijver, *Connected Components in Binary Images: The Detection Problem*. John Wiley & Sons, Inc., 1984.

[20] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, Jun. 1981.

[21] A. Rosebrock, "Finding the brightest spot in an image using Python and OpenCV," https://pyimagesearch.com/2014/09/29/finding-brightest-spot-image-using-python-opencv/, Sep. 2014.

[22] R. M. Haralick, S. R. Sternberg, and X. Zhuang, "Image analysis using mathematical morphology," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, no. 4, pp. 532–550, Jul. 1987.

[23] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," in *Proc. of Conference on Computer Vision and Pattern Recognition*. IEEE, Dec. 2001.