

An experimental method to merge far-field images from multiple longwave infrared sensors for short-term solar forecasting

Andrea Mammoli^{a,*}, Guillermo Terren-Serrano^b, Anthony Menicucci^a, Thomas P. Caudell^b, Manel Martínez-Ramón^b

^a Department of Mechanical Engineering, The University of New Mexico, Albuquerque, NM 87131, United States

^b Department of Electrical and Computer Engineering, The University of New Mexico, Albuquerque, NM 87131, United States



ARTICLE INFO

Keywords:

Solar forecasting
Sky imaging
Longwave infrared
Image stitching

ABSTRACT

In a system used for short-term forecasting of solar irradiance, multiple longwave infrared sensors are used to acquire an image of a large, continuous section of the sky dome. The field of view of each sensor is directed at a particular, fixed portion of the sky dome, with some overlap between adjacent edges of each field of view to ensure that a continuous image can be acquired. Because of unavoidable imperfections in the optical components and in the alignment of the sensors, and because of the complex optics, it is difficult to pre-determine adjustable parameters in the geometric transformations required to merge the multiple images into a single one. Instead, it is possible to do this experimentally, by rotating the imaging sensor array via a two-axis rotating platform, using the sun itself as a convenient far-field reference object, and using the images collected to generate relations that map each sensor pixel into an altitude-azimuth direction, α and ϕ respectively. The experimental proof of concept of this method is described here.

1. Introduction

Merging the output of multiple imaging sensors is common practice in many applications, including collection and archiving of panoramic ground-based imagery (e.g. Google Streetview), digital maps from satellite photos (e.g. Google Earth), automatic stitching of multiple images to form panoramas on digital cameras or smartphones, and medical imaging. As a consequence of the wide array of applications, image stitching techniques have received much attention. In general, image stitching is used to construct an image with a larger field of view than that could be obtained with a single photograph. Common problems that are addressed by various algorithms are preservation of the dynamic range of each of the images, automatic alignment, finding and merging overlapping regions, and distorting subsections of individual images to remove visible seams and other artifacts.

These problems have been addressed by several studies. Brown and Lowe (2007) propose an algorithm for fully automated multi-row stitching, using invariant local features present in each image for ordering and alignment. Automatic gain compensation and straightening are also part of the algorithm. Eden et al. (2006) present a method that allows for large exposure differences, and is thus suitable for high-dynamic-range (HDR) applications, and that is capable of dealing with

rapidly moving objects. Kekre and Thepade (2009) concentrate on the problem of identifying the extent of overlap regions, and propose methods to blend images so that artificial edges do not appear in the final stitched image, while also minimizing brightness differences. Juan and Oubong (2010) propose algorithmic improvements using the SURF (Speeded Up Robust Features) in combination with an image blending algorithm based on RANSAC (Random Sample Consensus), resulting in rapidly-generated large panoramas with invisible seams. Processing speed is addressed by Xiong and Pulli (2010), who propose a fast algorithm to create high-resolution images from long sequences of source images with large color and brightness differences. The algorithm uses little memory and is suitable for use in mobile phones. The problem of stitching images that do not satisfy assumptions of planarity or simple rotation, which are often violated in practice, is addressed by Zaragoza et al. (2013). To correct for this, as-projective-as-possible warps are proposed, which approximate projectiveness globally but allow for local non-projective warps and result in better alignment in overlap regions. Solar irradiance forecasting is becoming important because the penetration of photovoltaic energy generation, at scales ranging from a few kW to hundreds of MW, is increasing fast and can lead to reduced power quality on distribution networks and even instability on the transmission grid. While methods for intra-hour forecasts, based on

* Corresponding author.

E-mail address: mammoli@unm.edu (A. Mammoli).

satellite imagery, are well-established, short-term forecasting is still relatively undeveloped. One of the main reasons for this is that it is difficult to obtain ground-based images of the cloud field with a wide field of view, high resolution, and with a clear distinction between the solar disk, the clear sky and clouds. In the present application, we seek to acquire images of the sky dome to use as input for a learning algorithm that can forecast solar irradiance in the short-term, with time horizons on the order of minutes, as previously described by Mammoli et al. (2013). Successive images for this algorithm could be required at a rate on the order of 1 Hz. Many approaches to solar forecasting using sky imaging use optics (e.g. a fisheye lens) to capture an image of the sky dome with sufficiently large solid angle. This is the case, for example, for the work by Sayeef and West (2014), who use relatively inexpensive fisheye security cameras that have been repurposed to acquire sequences of images for use as input to forecasting algorithms. An alternative approach to a fisheye lens is taken by Chow et al. (2011), whose ‘Total Sky Imager’ uses a camera mounted vertically above a curved mirror surface to obtain an image of the entire sky dome, with a shadowband blocking the sun, to avoid overwhelming the dynamic range of the camera sensor.

In our approach, we capture longwave infrared (IR) images rather than visible light images. There are several advantages associated with IR imaging, including reduced forward Rayleigh scattering of incoming direct normal radiation (see e.g. Howell et al. (2010)), a much smaller dynamic range, and the ability to clearly distinguish clouds from clear sky, as noted by Liu et al. (2011), even in regions of the image very close to the solar disk. The reduced scattering produces a better discrimination between the solar disc and the circumsolar regions. This is an advantage over visible radiation cameras in very short term irradiance forecast. Indeed, the present direct radiation of the sun does not provide useful information for the prediction of the future received irradiance from the sky, and then a common practice is simply to apply a mask to exclude the sun from the prediction data. Conversely, because of the strong forward-scattering, it is very difficult to resolve cloud edges when using visible radiation methods.

While information in the circumsolar region is essential, a wide field of view is also necessary, especially in the case of fast-moving clouds. In the visible range, high-resolution sensors combined with wide-angle optics satisfy both requirements. Unfortunately, conventional glass optics do not transmit IR, and materials such as Germanium must be used instead. Custom Germanium optics are much more expensive than the sensor itself, and not widely available. While the ‘Total Sky Imager’ approach would potentially work well (metal surfaces reflect IR well), the resolution of our imaging sensor is too low to capture the entire sky dome while retaining essential image features. The ‘Total Sky Imager’ approach was indeed used by Liandrat et al. (2017) and by Redman et al. (2018) in the IR range, using higher resolution imaging equipment. We propose an alternative solution which is more compact in volume and does not involve any moving parts, by generating of a composite image from multiple IR fixed sensors at different orientations. This solution is not unlike the biologically-inspired approach taken by Dunkel et al. (2015) and by Brückner et al. (2010), who fabricate small wafer-level camera lenses inspired by insect compound eyes.

Our system consists of a number of imaging sensors, connected to a single-board computer that is responsible for grabbing images from each sensor, and for subsequent image pre-processing. Because of the relatively large throughput of information, and the relatively small computing capacity of a low-cost single-board processor, the algorithm for image stitching should not be computationally burdensome. As a result, conventional image-stitching techniques that are based on recognizing common features in overlap regions may not be suitable. Moreover, the presence of distinctive features in a sky-imaging application is not a given - consider, for example, the case of a clear or uniformly overcast sky. Fortunately, the far-field nature of the images of interest in our system (clouds and sun) make it possible to implement

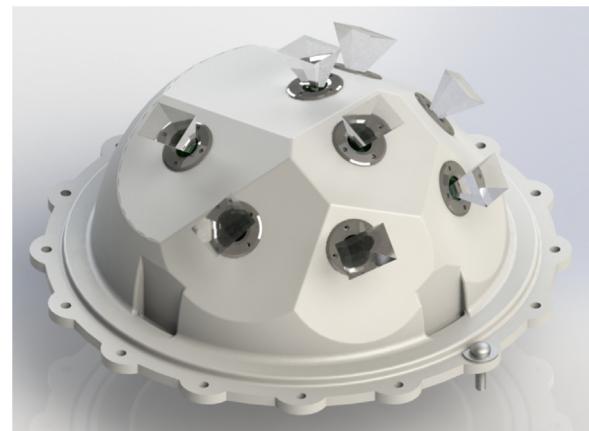


Fig. 1. Rendering of the eight-sensor imaging array. Each sensor images one of eight overlapping sectors of the sky dome, as shown by the field-of-view projections emanating from the sensors.

static mappings that account for the fixed (although not explicitly known) optical features of the system. Thus, the main contribution of the present work is the development of an experimental procedure to map the location of a pixel on a specific sensor to the corresponding location on a master image.

In what follows, first we present the experimental setup and data collection procedures; we then provide the core of the paper, namely the algorithms used to map local coordinates on individual sensors to global altitude-azimuth coordinates; we then present results consisting of the stitching of three adjacent images with overlap; we finally provide concluding remarks.

2. Experimental hardware setup and data collection

The imaging array consists of eight IR sensors, mounted on a support designed so that the collective image spans the region of the sky dome around the path of the sun, when the imager is mounted in a fixed horizontal position. This is illustrated in Fig. 1.

A typical image collected by the system is shown in Fig. 2. Clear regions of overlap are evident. The purpose of the present work is to develop an experimental technique that results in a single, seamless image with uniform gain. For simplicity, only the top three sensors are considered henceforth, since the procedure to stitch all eight sensors is similar.

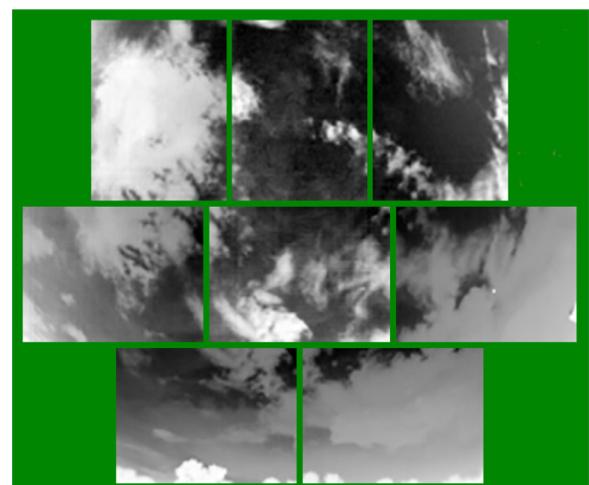


Fig. 2. Set of eight IR images collected by the imaging system. Note the position of the sun, the small bright circle in the rightmost image in the middle row.

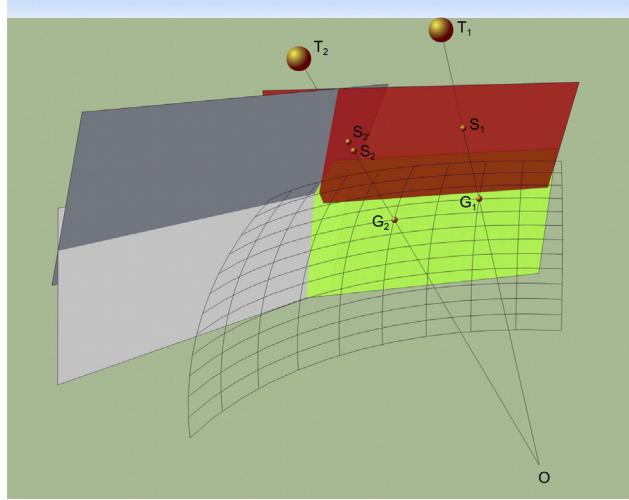


Fig. 3. Schematic of experimental setup to match sensor pixel coordinate location with location on a master grid.

A far-field master grid of altitude-azimuth coordinate pairs was used to calibrate the system for the conditions of interest, namely far-field images. This is illustrated in Fig. 3. The sensors, all located, for practical purposes, at the origin 'O', each have a field of view (red, green, dark grey and light grey) that spans a portion of the master grid. The line from the sensor location at the origin to a far-field target 'T' intersects the master grid at 'G' and the sensor field of view at 'S'. The location 'S' corresponds to pixel coordinates on the sensor itself. A specific far-field target, corresponding to a given location on the master grid, could be seen by one or more sensors. For example, target 'T₁' is seen by the 'red' sensor only, at pixel location S₁, and intersects the master grid at G₁. In contrast, target 'T₂' is seen by both the 'red' sensor at pixel location S₂, and by the 'dark grey' sensor at pixel location S₂, and intersects the master grid at a single location G₂. In the areas of field-of-view overlap, pixel locations on different sensors corresponding to the same target naturally map to a single point on the master grid. Once the mappings from each sensor pixel to the master grid are known, image stitching is a straight-forward operation. To obtain such mappings, images of the target should be collected for target locations coinciding with points on a known altitude-azimuth grid.

To accomplish this, first a suitable far-field target must be selected. There are two requirements: (1) the target must be large enough so that it spans several pixels, in a way that an accurate centroid may be calculated; (2) the target must be far enough away from the sensor array. Thus, the position of the sensors can be considered as a single point. The

sun meets both criteria, with only one difficulty, namely that it cannot be placed in an arbitrary position. This difficulty can be overcome by mounting the sensor array assembly on a pan-tilt tracker (FLIR PTU-E46-17 pan-tilt, 0.013° location accuracy). Rather than moving the target with respect to a fixed sensor array and altitude-azimuth grid, the sensor array and the attached altitude-azimuth grid are moved with respect to the target (the sun, whose position can be calculated based on well-known equations given a geographic location, date and time, see e.g. [Duffie and Beckman \(2013\)](#)), so that a line between the target and the sensors corresponds to points on the altitude-azimuth grid.

The pan-tilt unit is controlled in such a way that the center of the grid (the direction of the central imaging sensor) is at 0° azimuth and 101.25° altitude (measured from the 0° azimuthal direction, meaning that the direction of the sensor is to the North), when in the home position. The tracker is then rotated between the positions of -74° to 74° pan and -22° to 22° tilt, in increments of 4° in each direction, with respect to the home position. The resulting master grid then has dimensions of 38 × 11 in the pan and tilt directions respectively. To enable this, a controller calculates the position of the sun at all times during the course of the experiment, and uses this to set the angular direction for the home position in real-time.

The calibration experiment that produced the results described here was performed on October 25, 2017 (a clear day), beginning at 19:35 Universal Time and ending at 20:54 Universal Time. At each grid position, an image was collected from each sensor. The images were transferred to a laptop computer and stored for post-processing.

3. Surface fitting algorithm

At a given grid point in the experiment, an image of the sun (the target) may be present in one or more of the sensor images. The dynamic range of the imaging sensors is 0 to 45,000. The image of the sun saturates the corresponding pixels at 45,000, while the background is close to 25,000. To determine the position of the sun, each sensor image was thresholded, with pixel values lower than 35,000 set to zero, and left to their original value otherwise. When the sun is in the field of view of the sensor, its location is clearly visible, as shown in Fig. 4. Thresholding allows the simple calculation of the sun object position. The centroid of the sun object (s_i, s_j) is calculated by weighting based on the pixel values:

$$s_i = \frac{\sum_{k=1}^I \sum_{l=1}^J k p_{(k,l)}}{\sum_{k=1}^I \sum_{l=1}^J p_{(k,l)}}, \quad (1)$$

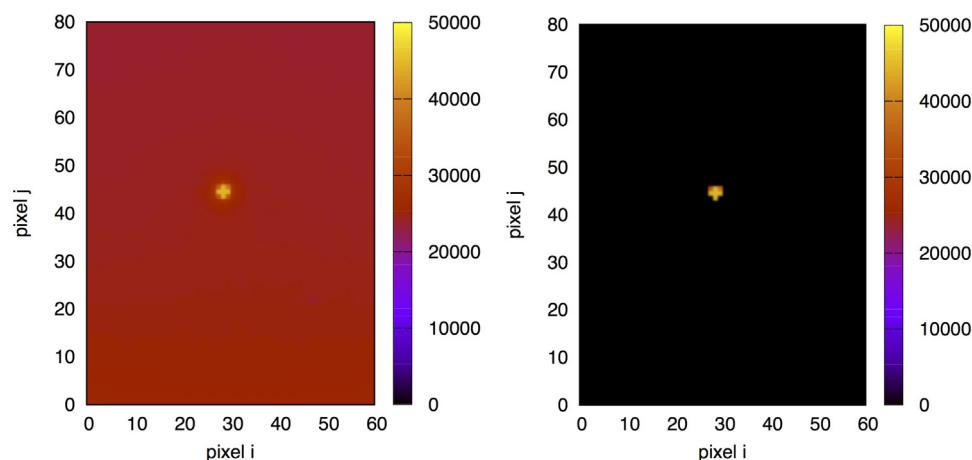


Fig. 4. Image collected by a sensor with the sun in its field of view, before and after thresholding.

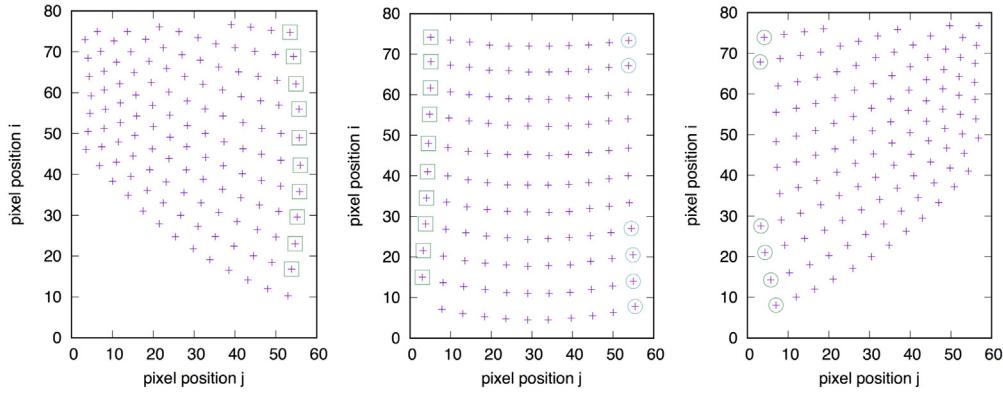


Fig. 5. Array of solar object centroid positions for three adjacent sensors. Overlapping solar positions are highlighted by the square and the circle symbols for the left and right overlaps respectively. Note the absence of solar object centroids in the immediate vicinity of the sensor edges.

$$s_j = \frac{\sum_{k=1}^I \sum_{l=1}^J l p_{(k,l)}}{\sum_{k=1}^I \sum_{l=1}^J p_{(k,l)}}, \quad (2)$$

where I and J are the pixel dimensions of the sensor array, in this case 60 and 80 pixels respectively, and $p_{(k,l)}$ is the intensity present at pixel (k, l) of the sensor. The position of the target is well-approximated by the position of the centroid of the sun object. This procedure allows the position of the target to take non-integer values, achieving sub-pixel accuracy. The positions of solar objects that were too close to the edge of the sensor were discarded, because of the possibility of an incomplete pixel set.

The sun positions for the entire experiment are shown in Fig. 5. Visual inspection indicates that the solar locations are determined with high accuracy. Rows of overlapping solar positions are shown highlighted.

Each of the centroids shown in Fig. 5 is associated with a unique global coordinate, determined using the pan and tilt angles of the tracker with respect to current the home position, which is time-varying, due to the earth's rotation during the sampling time. Consider a point on sensor i that is defined by sensor coordinates r_i and s_i (units of pixels). The corresponding position in the global altitude-azimuth coordinate system is found using:

$$\varphi = f_i(r_i, s_i), \quad (3)$$

$$\alpha = g_i(r_i, s_i). \quad (4)$$

The purpose of the present experiment is thus to find suitable approximations to the functions f_i and g_i that map a coordinate pair on sensor i to the global altitude-azimuth coordinates. Each of the sensors is then associated with a mapping. In turn, the mapping inherently merges the images, by placing pixel coordinates in overlapping regions of the sensors on the same global coordinate system. The entire process is experimental, with no need to identify common features in post-processing, as is done in the majority of image stitching algorithms, by taking advantage of the fact that the relative angular position of the sensor arrays with respect to far-field objects is fixed.

The mapping can be done, for example, by fitting models of the functions f_i and g_i to the dataset consisting of the (r_i, s_i) pairs on sensor i and the corresponding altitude azimuth global coordinates. The transformations required to find the global altitude-azimuth coordinates that correspond to the known tracker pan-tilt position can be found by referring to the diagram shown in Fig. 6.

The problem is to find the position of the solar object in altitude-azimuth coordinates (α', φ') in a local coordinate system attached to the sensor array, that is rotated by some known amount. The angular position of the sun for a particular date and time in a coordinate system

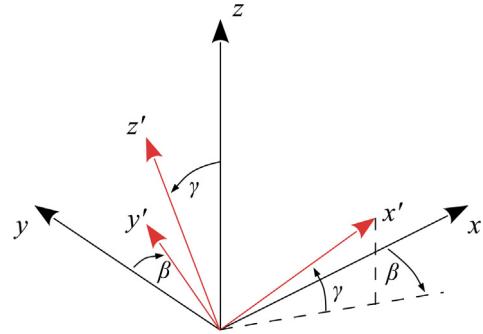


Fig. 6. Angles used for the transformation of the sun object direction from global coordinates (the black axes (x, y, z)) and the coordinate system attached to the sensor array (the red axes (x', y', z')), which in turn is attached to a pan-tilt tracker, whose position is defined by the pan coordinate β and the tilt coordinate γ .

with the z axis normal to the ground (global coordinates) can be found using well-known equations (e.g. [Duffie and Beckman \(2013\)](#)). Let the coordinate system (x, y, z) represent global rectangular coordinates, with x due south, y due east and z vertical with respect to the horizontal direction for the particular geographic location. Also let the coordinate system (x', y', z') represent local rectangular coordinates, with x' along the major axis of the sensor array, y' along the transverse axis of the sensor array, and z' vertical with respect to the base of the sensor array. The angular position of the coordinate system (x', y', z') mounted on the sensor array is defined by the angular position of the tracker. Specifically, the y' axis is rotated by angle β (the pan angle) with respect to the y axis, the z' axis is rotated by angle γ (the tilt angle) with respect to the z axis. Also, the angle between the projection of the x' axis on the global horizontal plane (the xy plane) and the x axis is β , and the angle between the x' axis and the global horizontal plane is γ . Thus, the components of the solar object direction in the global rectangular coordinate system are:

$$s_x = \cos(\alpha)\cos(\varphi) \quad (5)$$

$$s_y = -\cos(\alpha)\sin(\varphi) \quad (6)$$

$$s_z = \sin(\alpha). \quad (7)$$

The components of the sensor array axes x' , y' and z' with respect to the global rectangular coordinate system x , y , z are:

$$(x'_x, x'_y, x'_z) = (\cos(\beta)\cos(\gamma), -\sin(\beta)\cos(\gamma), \sin(\gamma)) \quad (8)$$

$$(y'_x, y'_y, y'_z) = (\sin(\beta), \cos(\beta), 0) \quad (9)$$

$$(z'_x, z'_y, z'_z) = (-\sin(\beta)\cos(\gamma), \sin(\gamma)\sin(\beta), \cos(\gamma)) \quad (10)$$

The projection of the solar direction on the rotated coordinate system is then the inner product of the sensor array axes and the direction of the sun object, which can be written in the global rectangular coordinates as:

$$s'_x = s_x x'_x + s_y x'_y + s_z x'_z \quad (11)$$

$$s'_y = s_x y'_x + s_y y'_y + s_z y'_z \quad (12)$$

$$s'_z = s_x z'_x + s_y z'_y + s_z z'_z \quad (13)$$

Finally, the altitude and azimuth of the solar object in the rotated coordinate system can be retrieved using a four-quadrant inverse tangent:

$$\alpha' = \arctan(-s'_y, s'_x) \quad (14)$$

$$\varphi' = \arctan(s'_z, \sqrt{(s'_x)^2 + (s'_y)^2}). \quad (15)$$

The local altitude and azimuth, in polar coordinates, of all the solar object centroids imaged in the experiment are shown in Fig. 7. The mapping process is then to find the relationship between a pixel coordinate on a given sensor, to its associated global coordinates. This can be achieved simply by fitting a surface to the dataset of sun object positions, as shown in Fig. 5, to the corresponding altitude and azimuth shown in Fig. 7.

The process is illustrated using the sensor corresponding to the left panel in Fig. 5. Quadratic surfaces are fitted to the data triplets (r_i, s_i, α) and (r_i, s_i, φ) using the Gnuplot ‘fit’ function, that in turn uses an implementation of the nonlinear least-squares (NLLS) Marquardt–Levenberg algorithm (Marquardt, 1963). Visually the data fit appears excellent, as apparent from Fig. 8.

4. Results

To demonstrate the effectiveness of the procedure, the mapping is applied to two sets of images, both of which were obtained on a different day from the calibration experiment, with the sensor array at its regular fixed position. One image set is of a clear sky with the sun located in overlapping regions of two sensors, while the other set is of a cloudy sky. The unstitched sensor images and the stitched images are shown alongside each other for both cases in Fig. 9. We also note that while small gain adjustments were made manually to match the pixel magnitudes in the overlap regions here, such adjustments could be easily be made automatically in a production setting.

For the clear sky case, the sun is exiting the field of view of one

sensor and entering the field of view of the other. A slight misalignment in the unstitched images is noticeable, along the long axis of the sensors. In the stitched image, the misalignment is completely removed. The two sun images are merged seamlessly into a single object. The region of overlap, where higher pixel density exists, is also clearly visible. The nonlinear mapping of pixel coordinates into altitude-azimuth coordinates is also clearly visible. Moreover, the expectation of imperfect alignments is confirmed by the fact that the overlap of the two sensor combinations is not exactly symmetric. This is precisely the strong point of the algorithm proposed here, that is completely agnostic to such small misalignments.

For the case of the cloudy sky, overlap regions are clearly visible in the unstitched images. The stitched images overlap seamlessly, in location and in magnitude, confirming the ability of the algorithm to fully account for the optics and the geometry of the multi-sensor image capture. This is further confirmed by visual inspection of the 2-D projection of the stitched images after interpolation of the unstructured pixel array onto a structured grid, as shown in Fig. 10. In both the clear sky and the cloudy sky cases, the images appear as if taken by a single imaging sensor, rather than three.

5. Conclusions

We have demonstrated a simple procedure for calibrating the optics and geometry of a multi-sensor imaging array using longwave infrared images. Although the procedure was demonstrated for three out of the eight sensors in our imaging apparatus, the results unequivocally show that the procedure provides very accurate overlaps. Furthermore, the mappings are produced only by optics and geometry, with no artifacts from commonly used image-stitching algorithms that use only features in adjacent images to produce the stitching.

For a full set of experiments with solar objects located in all of the eight sensors of the array, a tracker with more degrees of freedom than the pan-tilt tracker available for our experiment is desirable. This would allow more freedom in locating the solar object uniformly over all sensor locations. The limitations of the two-axis tracker system are evident even in the present experiment, in which sensor locations in the bottom left and bottom right of the left and right sensors respectively were not populated by sun objects. Despite this, the algorithm is robust enough that the pixel coordinate to altitude-azimuth mappings were still able to produce accurate stitching.

Finally, the mappings produced by rotating the tracker with respect to the sun position could also be produced by simply leaving the tracker in a fixed position and waiting for the sun to cover all possible

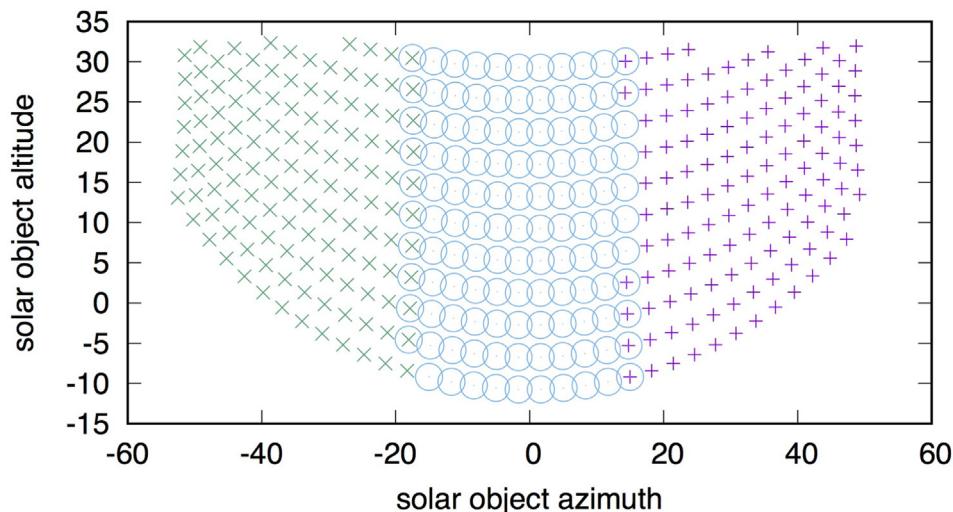


Fig. 7. Positions of all solar object image centroids, plotted in local altitude and azimuth coordinates. Note that the overlapping points identified in the individual sensor images shown in Fig. 5 now actually overlap, as expected. Centroids resulting from individual sensors are keyed differently.

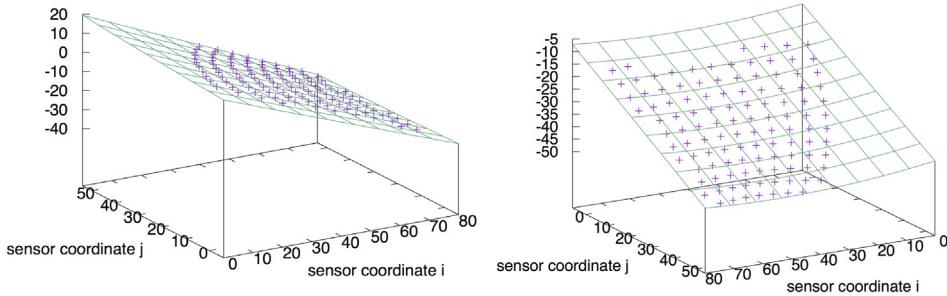


Fig. 8. Surface fits for the sensor pixel coordinate to altitude and sensor pixel coordinate to azimuth, left and right panels respectively. The variance of residuals (reduced chisquare) was 0.13173 and 0.0625206 for the altitude and azimuth fits respectively.

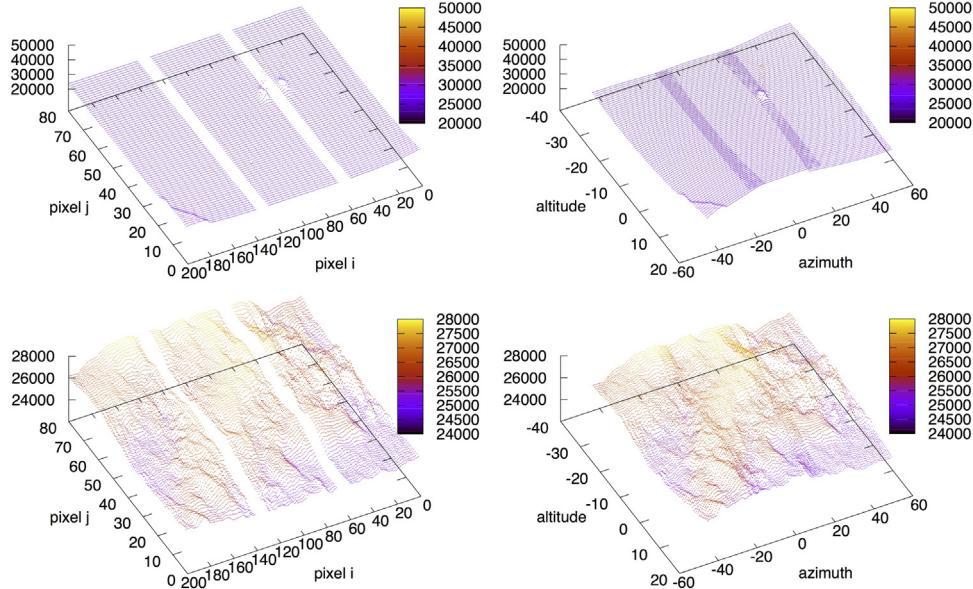


Fig. 9. Comparison of unstitched (left panels) and stitched (right panels) images for a clear sky (top panels) and for a cloudy sky (bottom panels).

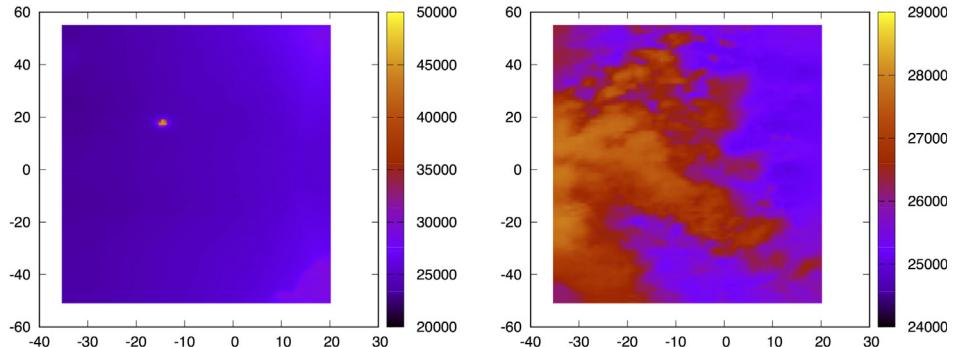


Fig. 10. Two-D stitched image of clear sky with sun object (left panel) and of cloudy sky (right panel). Note the absence of edge artifacts resulting from the stitching process. The unstructured stitched data in Fig. 9 were interpolated on a regular grid using the Gnuplot ‘dgrid3d’ function with parameters 100,100,6.

positions. However, the experiment described here can be performed in a matter of minutes, depending on the speed of the tracker, while the alternative method would require at least six months. However, the two methods are not mutually exclusive. In fact, because it is possible that the sensors could move slightly from their factory-aligned position when deployed in the environment over long periods of time, continuous re-mapping is likely desirable, and could be performed easily by processing of image data after the sensor array is deployed. In this case, the experimental calibration would simply serve to provide initial mappings, which would then be updated in an online fashion. Large deviations of the mapping coefficients could also serve as indicators of a

problem with the sensor array.

Acknowledgment

This material is based in part upon work supported by the National Science Foundation under Grant No. 1648751.

References

Brown, M., Lowe, D.G., 2007. Automatic panoramic image stitching using invariant features. *Int. J. Comput. Vision* 74 (1), 59–73.
 Brückner, A., Duparré, J., Leitel, R., Dannberg, P., Bräuer, A., Tünnermann, A., 2010.

Thin wafer-level camera lenses inspired by insect compound eyes. *Opt. Express* 18 (24), 24379–24394.

Chow, C.W., Urquhart, B., Lave, M., Dominguez, A., Kleissl, J., Shields, J., Washom, B., 2011. Intra-hour forecasting with a total sky imager at the UC San Diego solar energy testbed. *Sol. Energy* 85 (11), 2881–2893.

Duffie, J.A., Beckman, W.A., 2013. *Solar Engineering of Thermal Processes*. John Wiley & Sons.

Dunkel, J., Wippermann, F., Reimann, A., Brückner, A., Bräuer, A., 2015. Fabrication of microoptical freeform arrays on wafer level for imaging applications. *Opt. Express* 23 (25), 31915–31925.

Eden, A., Uyttendaele, M., Szeliski, R., 2006. Seamless image stitching of scenes with large motions and exposure differences. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 2. IEEE, pp. 2498–2505.

Howell, J.R., Menguc, M.P., Siegel, R., 2010. *Thermal Radiation Heat Transfer*. CRC Press.

Juan, L., Oubong, G., 2010. Surf applied in panorama image stitching. In: 2010 2nd International Conference on Image Processing Theory Tools and Applications (IPTA). IEEE, pp. 495–499.

Kekre, H., Thepade, S.D., 2009. Scaling invariant fusion of image pieces in panorama making and novel image blending technique. *Int. J. Imaging Robot.* 1 (A08), 31–46.

Liandrat, O., Cros, S., Braunv, A., Saint-Antonin, L., Decroix, J., Schmutz, N., 2017. Cloud cover forecast from a ground-based all sky infrared thermal camera.

Liu, L., Sun, X., Chen, F., Zhao, S., Gao, T., 2011. Cloud classification based on structure features of infrared images. *J. Atmosph. Oceanic Technol.* 28 (3), 410–417.

Mammoli, A., Ellis, A., Menicucci, A., Willard, S., Caudell, T., Simmins, J., 2013. Low-cost solar micro-forecasts for PV smoothing. In: 2013 1st IEEE Conference on Technologies for Sustainability (SusTech). IEEE, pp. 238–243.

Marquardt, D.W., 1963. An algorithm for least-squares estimation of nonlinear parameters. *J. Soc. Ind. Appl. Math.* 11 (2), 431–441.

Redman, B.J., Shaw, J.A., Nugent, P.W., Clark, R.T., Piazzolla, S., Apr 2018. Reflective all-sky thermal infrared cloud imager. *Opt. Express* 26 (9), 11276–11283.

Sayeef, S., West, S., 2014. Very short-term solar forecasting using inexpensive fisheye camera sky-imagery. In: Proceedings of the 52nd Annual Conference of the Australian Solar Council.

Xiong, Y., Pulli, K., 2010. Fast panorama stitching for high-quality panoramic images on mobile phones. *IEEE Trans. Consum. Electron.* 56 (2).

Zaragoza, J., Chin, T.-J., Brown, M.S., Suter, D., 2013. As-projective-as-possible image stitching with moving dlt. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2339–2346.