Computational Imaging for Long-Term Prediction of Solar Irradiance

Leron K. Julian, Haejoon Lee, Soummya Kar, and Aswin C. Sankaranarayanan

Abstract—The occlusion of the sun by clouds is one of the primary sources of uncertainties in solar power generation, and is a factor that affects the wide-spread use of solar power as a primary energy source. Real-time forecasting of cloud movement and, as a result, solar irradiance is necessary to schedule and allocate energy across grid-connected photovoltaic systems. Previous works monitored cloud movement using wide-angle field of view imagery of the sky. However, such images have poor resolution for clouds that appear near the horizon, which reduces their effectiveness for long term prediction of solar occlusion. Specifically, to be able to predict occlusion of the sun over long time periods, clouds that are near the horizon need to be detected, and their velocities estimated precisely. To enable such a system, we design and deploy a catadioptric system that delivers wide-angle imagery with uniform spatial resolution of the sky over its field of view. To enable prediction over a longer time horizon, we design an algorithm that uses carefully selected spatio-temporal slices of the imagery using estimated wind direction and velocity as inputs. Using ray-tracing simulations as well as a real testbed deployed outdoors, we show that the system is capable of predicting solar occlusion as well as irradiance for tens of minutes in the future, which is an order of magnitude improvement over prior work.

Index Terms-Sky imaging, Catadioptric Systems, Solar Energy

1 INTRODUCTION

Solver the output of light energy from the sun measured at a location on Earth, is converted to usable energy through the use of photovoltaic devices. The amount of energy received is usually affected by cloud or aerosol occlusion which scatters, reflects, or absorbs solar irradiance [1], [2], [3], [4].

Predicting the amount of energy to be received is difficult due to the seemingly random shape and trajectory of clouds which are inherently influenced by a variety of complex factors. This, in-turn, makes it difficult to predict precisely when there will be a loss of power due to a reduction of solar irradiance being received [5], [6], [7]. The occurrence of these rapid fluctuations poses significant challenges for power grid operators. It leads to voltage and frequency fluctuations, limited time to adjust between energy sources, and ultimately energy disruptions [8], [9], [10], [11]. As a result, full integration of solar into the electricity grid poses challenges, in part, due to the difficulty in forecasting this intermittent natural phenomena. By imaging clouds and estimating their motion, we can formulate this into an imaging problem that can precisely forecast when a cloud will occlude the sun. By translating the overall problem into one of a trajectory forecasting, we can gain foresight into when there will be a cloud occlusion of the sun and predict the overall dip in the solar energy received at the ground.

Cloud motion around a particular localized region can be monitored using sky images captured periodically by a wide angle field of view (FoV) imager, usually a hemispherical mirror or a fisheye lens. By using various prediction methods, cloud motion for future time instances can be predicted using these sky images. However, wide FoV



Fig. 1. Our work presents a computational imaging system based on a catadioptric combination of mirrors and cameras. Seen above are images of the sky when using (left) a traditional hemispherical mirror and (right) the proposed hyperboloidal mirror. Overlaid on top are circles corresponding to different angular extents of the sky. The use of a hyperboloidal mirror enables increased. angular resolution near the horizon. We also propose a learning-based framework for predicting future solar irradiance. Together, these contributions enable a more accurate prediction over a longer time horizon than the hemispherical imagers.

imaging systems provide non-uniform spatial resolution of the sky with a higher detail and resolution at the zenith (or directly overhead) and significantly lower-resolution near the horizon. As a consequence, our ability to detect a cloud at the horizon is negatively influenced due to the lowered resolution; further, the lower resolution also implies that the estimates of cloud velocity is also poor for those at the horizon. This inability to detect clouds at the horizon and estimate their motion eventually limits the time horizon over which we can make precise predictions of solar occlusion by a cloud and the irradiance measured on the ground.

As a result of these limitations, traditional works have attempted to combat this issue by digitally undistorting the sky image to limit the non-uniform spatial resolution [12],

All authors are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213.
 E-mail: {ljulian,haejoonl,soummyak,saswin}@andrew.cmu.edu.

[13], [14] or by employing a multiple camera setup [15], [16], [17]. Digital warping does help alleviate some of the challenges underlying non-uniform flow estimation, but it is fundamentally limited by the loss of resolution at image formation. Adding more pixels by using more cameras or even a higher resolution sensor can be an effective approach, but comes with increased costs. Further, direct imaging of the sky needs to be done with some care, given that the potential damage to the sensor caused by a focused image of the sun. We instead pose a different question: is it possible to *optically redistribute* the pixels in a wide FoV camera so that resolution is uniform for a cloud as it traverses the field of a sensor?

The centerpiece of this work is a catadioptric system that *optically* warps the sky to provide a uniform spatial resolution of the sky (for each height), over the entire field of view of the camera. We achieve this by imaging the sky through a mirror whose shape is designed to provide the aforementioned property. This design also has the added benefit of making motion of the clouds equally perceptible, be it at the zenith or the horizon. As a result of using this mirror shape in a catadioptric setup, our ability to estimate cloud trajectory is improved over traditional methods even when a cloud is farther away. This improves long-term cloud evolution prediction and as a result, prediction of when a cloud will occlude the sun.

Contributions. Our method advances long-term forecasting of cloud evolution, enabling predictions that extend far beyond previous works, reaching into the tens of minutes. Our contributions are as follows:

- *Imaging system for whole sky imager*. We have designed and deployed a novel sky imager comprising a catadioptric system with an adapted hyperboloidal mirror [18] to capture and analyze sky images with the eventual goal of improved solar irradiance forecasting.
- Dataset of sky images. Using our imaging system, we have captured high dynamic range sky images across a period spanning many months. The dataset also includes time-synchronized ground solar irradiance values captured via a pyranometer.
- *Predicting from spatial-temporal slices.* We have a novel prediction algorithm that uses estimated wind velocity to identify an informative 2D space-time slice of the imagery; this allows us to ignore the clouds that are unlikely to occlude the sun at our vantage point. More importantly, it significantly simplifies the resulting prediction problem, which we address using a lightweight learned-network.

Impact. The contributions above have resulted in a system that provides precise prediction of sun occlusion and solar irradiance over a time horizon of tens of minutes (~30 minutes for the simulated data, and 10-20 mins for the real system). This *long-term prediction* is an order of magnitude improvement over previous works; for example, Julian and Sankaranarayanan [12] who rely on a similar premise, but with digital warping, show prediction results that only span 2-3 minutes. Finally, we also hope this work will continue the renewed interest in this problem that is at the intersection of imaging and solar prediction. To this effect, we have released the dataset and its associated code base at [19]. The

Limitations. Our current prediction system is based on an assumption that clouds move linearly, driven by wind. While this is true for the most part, at least over the halfhour or so of our prediction horizon, the model does not account for cloud genesis and extinction events within the field of view. It is very likely that such evaporation and condensation provides a fundamental limit on achievable prediction accuracy. Our real dataset includes such phenomena; this suggests the need for more robust statistical models along with additional data such as humidity and other atmospheric conditions to account for such events.

2 PRIOR WORK

Large Field-of-View Sky Imagers. Devices used to acquire images of the sky—often dubbed *sky imagers*—encompass a fisheye lens or a catadioptric combination of lenses and hemispherical mirrors [20], [21], [22]. These images are wide FoV RGB images captured in regular intervals.

A limitation of these imagers is the non-linear fisheye distortion introduced which affects clouds optical flow estimates for trajectory predictions. The further-out clouds attenuate evolution over time and their prior motion estimates attribute to longer forecasting horizons. To combat this issue many works have attempted to spatially warp these images to achieve uniform apparent motion and limit the fisheye distortion [13], [14], [23]. Julian and Sankaranarayanan [12] have even shown that spatially warping these images achieves longer forecasting horizons. Limiting the long-term prediction accuracy when using digital warping is the loss of resolution of pixels at the periphery [24], [25]. Think of it as digital zoom versus optical zoom. Periphery pixels are stretched out and then interpolated. Therefore, true pixel values at these locations are absent. For learning and optical flow based methods, these pixel values are essential for accurate long-term prediction.

Predicting Cloud Movement. Modeling cloud evolution through tracking and forecasting clouds solely using sky images are achieved using mainly two methods. Initial works utilized optical flow based methods [22], [26], [27], [28] using pairs of subsequent sky images to forecast cloud trajectory. Overall, this method is inadequate for long-term prediction due to the variability of cloud shapes and trajectory between image captures making it difficult to forecast.

More recent works have seen greater success using deep learning to predict a subsequent sky image for a future time instance [12], [23], [29], [30], [31]. Learning-based methods can be further improved for more accurate long-term forecasting by addressing the fisheye distortion of the scene introduced by traditional large-FoV imagers. In this work, we show that optical warping leads to even greater success when coupled with learning-based predictions.

Photovoltaic Power Output Prediction. Directly predicting photovoltaic power output has been a direction taken by previous works. These studies either take a statistical approach to predicting future irradiance values from past values [32], [33] or finding the relationship between an associated sky image and its irradiance value; also known as nowcasting [34], [35]. However, these works do not take the future state of cloud patterns into consideration which directly influences the amount of irradiance received at the ground. Therefore, by having an accurate method of predicting the future distribution of clouds, a better estimate of future irradiance can be obtained.

Computational Imaging for Atmospheric Tomography. The ideas in this paper are closely related to a rich body of work related to computational imaging of clouds, especially in the context of tomography reconstruction. This includes results that model the heterogeneous multi-scattering media and reconstruct the full volumetric field using distributed ground-based camera systems [36], [37], [38], [39], or airborne imagery [40], [41], [42], [43]. Such three-dimensional analysis of clouds also provides a pathway for prediction of solar irradiance over a large geographic region, provided we are able to recover the cloud distribution over the region; in contrast, we rely on simpler two-dimensional reasoning in our work given that we are making predictions at a single location that is collocated with our imaging instruments.

Catadioptric Imaging Systems. Catadioptric imaging which utilizes the reflective nature of mirrors during the acquisition process are designed such that their shape can achieve various tasks. Baker and Nayar [18] have extensively studied the family of these shapes. In particular, the hyperboloidal shape provides practical wide-angle imaging with minimal distortion and solves to the aforementioned challenges. This selected mirror shape will be further discussed in the subsequent sections.

3 PROBLEM DEFINITION

In this section, we introduce the problem of sky imaging by stating the desired specifications and discussing the gaps between these requirements and current wide FoV imagers.

3.1 Design Specifications

Prediction of cloud movement requires precise estimates of their velocities when they are at the periphery of the FoV. For example, a cloud with a typical height of 1 km and a velocity of 50 km/hr, will cover an angle of $\tan^{-1}(25/1) =$ 87° , in a camera's view, over half an hour. If the sun is at the zenith, to provide a reliable prediction half-an-hour in the future, we need a 174° field of view, as well as the ability to precisely sense motion when the cloud is near the horizon. This provides us with our design specifications: an imager with extremely large field of view approaching 180°, while providing the ability to detect and estimate cloud motion over the entire field. We interpret the second part of the specification as one of providing uniform spatial resolution for the cloud as it appears and traverses the field of view of the imager. While uniform resolution by itself is not a necessary condition (for example, we could ask for higher resolution at the horizon over the zenith), it allows for a robust solution that can also accommodate cloud creation events within the field of view.



Fig. 2. A rendering obtained using Blender to visualize the non-uniform resolution of a hemispherical mirror. (Top) We create a scene consisting of a checkerboard with a length and width of 50 km, placed 2 km high above the ground. Each square on the checkerboard has physical extent of 1 km. (Bottom-left) Our imaging system consists of a pinhole camera observing the sky or the checkerboard indirectly through a hemispherical mirror. (Bottom-right) The image observed on the camera has high resolution at the zenith of the image and significantly lower resolution at the periphery.

3.2 Gaps in Current Sky Imagers

Current wide FoV imagers can be built with a fisheye lens or more commonly with a catadioptric system where the sky is imaged through a hemispherical mirror. Such traditional sky imagers are not conducive for long-term prediction due to their lack of resolution at the periphery of the imager.

Specifically, in such systems, an object placed at the zenith of the sky will appear to have a larger total spatial extent as opposed to the same object at the horizon. Figure 2 visualizes this circumstance via a large checkerboard placed above a simulated hemispherical mirror. The checkerboard, which has a length and width of 50 km, is placed 2 km high above the mirror where each square is uniformly spaced at 1 km per square space. This hemispherical setup shows that squares at the zenith of the imager appear larger compared to squares at the periphery, despite their physical dimensions being the same. This compression of the squares at the horizon translates to poor localization of clouds that are located at the horizon in the world. Another related factor is our ability to estimate motion. In current sky imagers, motion of clouds appear to be non-uniform despite their physical speed being largely the same (since clouds are driven by wind), with large apparent motion at the zenith and significantly smaller ones at the horizon. In more practical terms, the imagery only allow for precise estimates of cloud velocity only after it is significantly away from the horizon; in turn, this limits the time horizon over which sun occlusions can be predicted.



Fig. 3. (Left) We design the mirror for a setup where a camera is placed 1 meter above the mirror. The shape is optimized so that the overall system scales the field of view of the pinhole camera from 3.58° to 170° uniformly. (Right) The resulting hyperboloidal mirror shape that we use in our setup.

3.3 Solution Outline

Our goal is to address the limitations of current sky imagers which achieve a large viewing angle at the cost of two issues that limit long-term cloud motion prediction: lack of resolution at the horizon, and non-uniformity of motion. How can the problem of non-linear motion and lack of pixel resolution be circumvented?

Our approach relies on the insight that we can redesign the mirror used in a sky imager to spatially redistribute the pixels with the eventual goal of having the same spatial resolution on a cloud over the field of view of the imaging system—immaterial of whether the cloud is at the horizon or at the zenith. This allows for early detection of clouds, as well as simplifies the motion estimation problem since the clouds largely translate over the field of view. We discuss the mirror design problem in Section 4 as the associated testbed and dataset in Section 5.

The second part of our solution is an algorithmic technique for prediction over time-horizons of tens of minutes. Part of the challenge here is the high-dimensionality of the input image which makes any learning-based solution hard to implement due to compute and memory requirements, as well as the need for a large amount of input data. To simplify this problem, we argue that cloud motion due to wind is largely translational; hence, to predict the occlusion of the sun as well as solar irradiance at future time instants, it is sufficient if we look at a spatial slice through the sun that is parallel to the wind direction. While this likely misses out on predicting irradiance due to indirect skylight, it has all the relevant information for predicting direct sunlight which is the dominant term in the overall irradiance. We describe this algorithm in Section 6.

Finally, we evaluate our algorithms over this dataset, as well as a synthetic counterpart, in Section 7.

4 MIRROR DESIGN

We frame the problem of mirror design as one that "flattens" the sky image formed on the sensor. Figure 3 illustrates the relevant variables.

We delve into the derivation of the mirror shape profile, which share the same goals as [18], in the supplemental material. However, we concisely present the setup here. Our basic setup is that of a pinhole camera with a sensor size of



Fig. 4. Simulated case visualizing the uniform resolution of our hyperboloidal mirror. (Left) The same parameters as Figure 2 with the proposed mirror replacing the hemispherical mirror. (Right) Observe how the checkerboard resolution is spatially uniform—a consequence of the system acting as overall scaling operation.

w = 12.5mm, placed at a distance c = 1m from the mirror, with a field of view of $f_c = 200$ mm. These choices are based on design considerations for the final implementation where we need the camera to be sufficiently far away to avoid blocking a significant portion of the field of view. The long focal length also allows us to effectively mimic the pinhole camera with a lens-based counterpart.

The mirror has a shape $z = f(\rho)$, where ρ is the radial distance over the ground plane. We make an additional assumption that the cloud is at some height h; the exact height of the clouds do not play an actual role as we will assume that $h \gg c$ and so only the tangent of the angle subtended by the cloud at the mirror matters. With this, we formulate the mirror design as one of designing the profile $f(\cdot)$ such that the effective sky to sensor mapping is a scaling operation over the desired field of view. Effectively, we are scaling the FoV of the camera—which is $\theta_{cam} = 3.58^{\circ}$ —using the mirror by a constant spatial factor to achieve a target FoV $\theta_{target} = 170^{\circ}$.

To determine the mirror shape, we use a numerical procedure where we solve for the axial profile $f(\cdot)$ by densely ray tracing over the image plane. With this, we also have the constraint that the ray — after mirror reflection — behaves like a pinhole camera with the target field of view. This provides a constraint on the derivative of the f (since the surface normal is determined by the normal). Integrating this derivative provides us with the desired shape. A visualization of the resulting mirror shape is shown in Figure 3 (right). This solution falls under the family of shapes described by Baker and Nayar [18], and in particular it is a hyperboloidal shape.

In Figure 4, we use Blender to render an identical setup as Figure 2, but with the hyperboloidal shaped mirror. Our mirror achieves a uniform image of the checkerboard while maintaining a large FoV, showing that we are able to image the sky with uniform resolution. As is to be expected, this design also enables uniform motion estimates throughout the whole FoV of the imager.

5 TESTBED

We now describe our imaging setup in the context of our hyperboloidal-based mirror.



Fig. 5. Various images captures from the synthetic dataset. (Top) Captures from hemispherical setup. (Bottom) Captures from the hyperboloidal setup. Each column is captured at the same time instant

5.1 Simulation Setup

We initially evaluate and report results of our setup and methods on simulated data which achieves the idealized scenario of a real-world setup with known parameters. The platform used to develop our simulated data is Blender which uses the same setup as in our real-world data. In Blender, the Pure-Sky Pro package which simulates an array of cloud formations inspired by [44] is used. Although not modeled as mathematically in-depth as a large-eddy simulation [45], Pure-Sky Pro is accurate to the scale of this simplified simulated application. The package does allow for the modification of cloud dynamics such as how warm/cold air affects cloud evolution.

Using the computer generated hyperboloidal-mirror shape, as shown in Figures 2 and 4 placed with a reflective mirror material property, we capture simulated data on various cloud scenes with a sampling period of $T_0 = 30$ seconds. We also capture the same data using a hemispherical mirror with the same parameters. These cloud scenes include randomized cloud parameters across a 28-day period from 8AM to 5PM based on real-world factors such as wind, hot/cold air patterns, and sunlight. Images of simulated data for both mirrors are shown in Figure 5.

5.2 Hardware Prototype

To develop a physical prototype for our mirror, we fabricated the desired 3D surface of the hyperboloidal mirror using a Computer Numerical Control (CNC) machine. We used aluminum for the material due to the ease with which it could be polished; for the final mirror surface we used a chemical deposition process that is commonly used to produce highly reflective surfaces [46].

For image acquisition in our system, we utilize an RGB camera mounted on a cuboidal frame above the mirror. The mirror itself lies on the horizontal axis of the frame and coupled with a mini PC, captures sky images with a sampling period $T_0 = 30$ seconds. To minimize any nearby building occlusion, our imaging device is placed on a building roof and captures data continuously during daylight with a frequency of T_0 . We also included a second system with a hemispherical mirror for evaluating improvements of



Fig. 6. Our deployed testbed with a detailed visualization of mirror placement of the hyperboloidal mirror and hemispherical mirror.



Fig. 7. Field of view of the sky (in tangent of angle) as a function of radial distance from center for the hyperboloidal and hemispherical mirrors in our deployed system. The numbers to the right of the figure provide the corresponding half-FoV in degrees. Note how the hyperboloidal mirror provides a linear relationship to the radial pixel displacement, thereby leading to flattening of the sky.

our proposed work. To handle the large dynamic range of the sky, due to the sun, we capture images using exposure bracketing and fuse them to get a single HDR image. The top of our system also includes a pyranometer that measures solar irradiance in the form of global horizontal irradiance (GHI), defined as the total solar irradiance received at a location horizontal to the Earth's surface measured in the units of watts per meters-squared (W/m^2) . Our setup as described is shown in Figure 6.

Next, we characterize the angular mapping provided by both mirrors in our system in Figure 7. We placed a light source at various heights along the frame of our real setup, and used it to calculate the FoV observed at each mirror as a function of radial distance from the center of the mirrors, as observed in the acquired images. Figure 7 plots the half FoV, in its tangent, observed at each mirror across the image; for consistency, we normalize the radial distance of each mirror by the radius of the mirror for a fair comparison. Two key observations emerge. First, the hyperboloidal mirror pro-



Fig. 8. Real images captured from various dates and weather conditions. (Top) Images from the hemispherical setup. (Bottom) Images from the proposed hyperboloidal setup. Each column is captured at the same time instant. The cropped cloud in the red box in both images show the benefit of our mirror being able to image a cloud much further out.



Fig. 9. We show examples of adverse weather conditions captured from our real setup of rain (top) and snow (bottom).

vides a linear mapping between the radial distance observed in the image plane and the tangent of the observed half FoV; this is a key design property of the mirror that allows a fixed plane above the ground to be mapped to the image plane of the camera. Second, notice how the hemispherical mirror devotes most of the image plane to a small central cone at the zenith of the sky. Specifically, the central cone of 45° around the zenith of the sky occupies more than 50% of the radial distance with the hemisphere as opposed to less than 10% with the hyperboloid.

5.3 Dataset Collection

Figure 8 shows a gallery of real images captured from our setup. These images are captured from our hyperboloidal mirror and a hemispherical mirror at the same time instant. Similar to the simulated data, the real images achieve the benefits of using the hyperboloidal shaped mirror. We are able to see more clouds within a single capture and the motion is more translational through time. Our dataset for this work consists of imagery collected from October 20th, 2023 to March 5th, 2024. We excluded days that were entirely cloudy or completely clear skied, so as to remove scenarios where GHI is nearly constant over the entire day. This left us with 76 days worth of data with most days having partly cloudy conditions. In Figure 9, we show some adverse conditions from the captured datasets.

5.4 Pre-processing

Before we can apply learning-based techniques on this dataset, we need to perform certain operations on it. In particular, knowledge of the sun as well as the wind velocity at each frame is helpful for the algorithms we describe next.

Sun localization. As a crude pre-processing step, we use the shortest exposure in our HDR stack to estimate the location of the sun. However, this technique fails when the sun is occluded by clouds. To get a robust estimate, we pool the data across multiple days of maximal saturation and reject outliers using RANSAC. This provides a sun estimate as a function of daytime where occluded sun estimates are filled by fitting a polynomial function over sun locations identified by maximal saturation. Of course this will fail for incorrect predictions; therefore manual identification is required for some cases.

On an aside, the location of the sun in absolute angular coordinates with respect to the zenith of the sky can be analytically computed given the latitude and longitude of the testbed. We can in principle map such elevation and azimuthal position of the sun to the image plane coordinates using a calibration procedure; details of such a procedure can be found in prior work on cloud imaging [39].

Wind velocity estimation. Another useful information for the learning-based formulation that we will present next is the direction of wind velocity. A challenge here is that clouds are largely featureless, which makes traditional optical flow techniques fragile. Further, there are features in our field of view which are constant, for example buildings at the periphery and the frame used to hold the cameras. These static features bias the optical flow estimates especially since they are also high contrast ones.

To overcome these effects we use the mask to suppress the static regions and run the optical flow technique proposed by Liu [47] with very a strong weight associated with the spatial regularization term. Finally, we use an aggressive temporal median filter on the estimated optical flow across frames to ensure a smooth flow field.





Fig. 10. We show how space-time-slices are extracted (top) Images of the sky from the hyperboloidal mirror at the start and end times of a sequence. We track the sun location and slice the image along the direction of wind motion, as marked in each image. (bottom) The spacetime image produced by the hyperboloidal mirror on the top and the hemispherical mirror on the bottom.

6 ALGORITHMS FOR LONG TIME-HORIZON PRE-DICTION

The sudden rise and fall of received solar irradiance at the ground within a short period of time is crucial information for electricity grid operators to mitigate disruptions in power output. This event, also called a ramp event (RE) [48], [49], is influenced by the occlusion state of the sun by a cloud. This is necessary to forecast and can be achieved through prediction of cloud trajectory. Thus, we show the benefits of utilizing our hyperboloidal imaging system and focus on a suite of algorithms that can exploit these benefits to better predict RE's.

6.1 Space-Time-Slice Image

The key benefit of having uniform apparent motion of clouds with the hyperboloidal mirror is that we can linearly back-trace the trajectory of clouds through time to the cloud's projected path toward the sun. We can interestingly use this fact and state that the only part of the image that is important is the sun and clouds that are moving towards it. This corresponds to a linear slice of the image along the wind velocity as shown in Figure 10. By collecting such slices over time, centered around the sun as it moves, we can build a *space-time-slice image* that effectively summarizes the relevant cloud movement.

We briefly describe how we create the space-time image and utilize it for inference. For each time instant for a single day, the x and y coordinate of the sun is initially identified. Next, the general direction of cloud motion $(\hat{\theta})$ through time is obtained. We take a sun-centered slice of the



Fig. 11. We show resulting space-time-slice images from captured from our real setup. (Top) a result of selecting $\theta = 65^{\circ}$ while the bottom is $\theta = 85^{\circ}$.

image in the direction of cloud motion for each time instant and horizontally concatenate these images through time to formulate the space-time image.

For the case of the simulated sky images, we are benefited by having the ground truth sun location, direction of cloud motion, along with the binary sun occlusion state. Therefore, we have the necessary parameters for an ideal space-time image. Figure 10 shows this space-time image and compares the image obtained from our system to the hemispherical-based system. Our system clearly achieves the desired linear apparent motion and is the ideal case for predicting sun occlusion discussed in Section 6.2.

However, for the real images, ground-truth parameters are not given and are estimated. Sun location and general cloud direction is identified using the methods described in Section 5.4. In the real-world environment, cloud motion direction is not static and changes through time. Therefore, we estimate cloud motion over a whole day using estimated optical flow and apply an aggressive temporal median filter to ensure a smooth flow field. We present sample space-time slice images for the real case in Figure 11.

6.2 Non-Learning Occlusion Prediction

Given that the ground-truth parameters for the simulated images are available, we are able to perform non-learning based sun occlusion and show the benefits of our setup.

6.2.1 Back Projected Sun Occlusion Prediction

Looking at the space-time image produced by our hyperboloidal shaped mirror, the linear streaks of clouds whose trajectory through time occludes the sun at the center of the space-time image, can easily be seen. Due to the nonlinearity of the apparent motion within the hemispherical images, these space-time images do not have the same effect. We can use this and make the assumption that a cloud that occluded the sun at a time instant T is the same cloud that is at a location v on the image such that:

$$v = \tau \tan \theta, \tag{1}$$

where $\tau = (T - t)$ is the time displacement from *T* along the horizontal x-axis. If we sweep over a range of τ and obtain the slice warped to the current location of the sun at *T* based on Eq. (1), this results in a new image which, although seemingly meaningless, actually obtains information about the future sun occlusion states.



Fig. 12. (Top) Space-time image. (Bottom: left-to-right) We visualize the non-learning steps for occlusion prediction. We take some window out of the original space-time image, warp it based on the optimal $\hat{\theta}$, convert it to a new color space based on [50] and plot the mean through time. Red points show ground truth occlusion states.

We take this warped image, and then convert it to a new color space defined as the ratio (B - R)/(B + R) where B and R are the intensities in the blue and red channels of the image, respectively. This space allows for a simple contrast-based delineation of cloud versus non-cloud pixels [50]. The mean of this image along the y-axis is computed and used to identify the binary sun occlusion state. As shown in Figure 12, dips in the plotted mean relate to sun occlusion states.

It should be noted that to find the optimal $\hat{\theta}$, we sweep over a range of empirically chosen $\hat{\theta} = [60^\circ, 85^\circ]$. Instead of computing the mean, we use the standard deviation and attribute the lowest value of the standard deviation to the optimal $\hat{\theta}$ and warp the image by that value using Eq. (1).

To obtain metrics in terms of accuracy of the sun occlusion state, we employ the receiver operating characteristic (ROC) curve to obtain the optimal threshold for deciding the occlusion state. We then look at accuracy through time for future time steps using the area under the curve (AUC) which provides an accumulated measurement of performance across all classification thresholds.

6.3 Learning-Based Occlusion Prediction

We believe that tasking a learning-based system to understand the dependencies of the space-time image to predict a future sun occlusion state will yield better results than a non-learning approach.

6.3.1 Neural Occlusion Prediction For Simulated Images

Non-learning based prediction of sun occlusion is limited by the dimensionality of the space-time image. As a result, the time prediction to T+N of a sun occlusion is capped at a certain value of N which is based on $\hat{\theta}$. Using a simple CNN-MLP, we are able to show that a learning-based method can learn the dependency between spatial cloud locations at τ along with $\hat{\theta}$ to predict the sun occlusions state for a future time instant. Our model is fed the warped image consisting of τ space-time-slices and produces a $(1 \times T + N)$ vector which are the binary occlusion state predictions from T + 1 to T + N. Our model is trained end-to-end using binary cross entropy as the loss function.



Fig. 13. We show sample GHI values based on cloud conditions.

6.3.2 Neural Occlusion Prediction For Real Images

For the real images, we utilize a different approach. Rather than directly predicting the binary occlusion state of the sun, we predict the solar irradiance, in the form of GHI, at a future time instant which directly correlates to sun occlusion. As shown in Figure 13, a decrease in GHI directly correlates to the occlusion of the sun by a cloud and therefore can be used as a prediction method for real images. Overall, GHI is predicted as opposed to the direct sun occlusion state due to the fact that we do not have the ground truth occlusion state for the real images. However, GHI can be deemed an improved occlusion metric because in a binary state, there is no indication of *how* occluded the sun really is. GHI better captures this information and is therefore a more informed prediction value for sun occlusion.

Our learning pipeline for forecasting GHI is 2-fold and very similar to [51] being that we first pre-train our model on masked-input reconstruction followed by fine-tuning for prediction. Our model architecture employs a transformer encoder [52], and instead of a transformer-based decoder, we use a simple, lightweight reconstruction head for pretraining and a forecasting head for fine-tuning. Both of these heads are small multi-layer perceptrons (MLPs) consisting of linear and dropout layers.

Pre-training. During pre-training, we use information from our space-time slices along with their associated GHI values as input $\{\tau; G_{\tau} \dots G_T\}$. The input space-time slice image is sent to a small image encoder consisting of 5 convolution layers where each layer is followed by a ReLU and 2D Maxpool; except for the last layer. This results in a latent embedding (K) consisting of 2-channels that encapsulates the information from the space-time slice image. Concurrently, for the associated GHI values, we employ masked-input reconstruction where, during training, 25% of the input is masked-out and replaced with a learnable mask embedding. Theoretically, we treat this input GHI as a time series data where information about the current cloud conditions are added via the space-time slice images. The masked GHI is fed into a patch embedding layer similar to [53] and the resulting latent embedding (I) is concatenated with K and passed into the transformer encoder. The encoder passes its learned output to the reconstruction head which reconstructs the original masked input.

Overall, the goal of this pre-training is to allow the model to learn a representation of the original GHI with cloud



Fig. 14. We forecast future irradiance values via a 2-step end-to-end training pipeline similar to [51]. We first pre-train on reconstructing the input irradiance followed by fine-tuning for irradiance forecasting.

information present in the space-time slice image. Figure 14 presents a visual of the described steps.

Fine-tuning for forecasting. The goal of fine-tuning the model to the task of forecasting is to use the pre-trained learned representation of a space-time slice image and its associated GHI value. During the fine-tuning for forecasting, we replace the reconstruction head with a forecasting head. The forecasting head is again a lightweight MLP consisting of a dropout and linear layer. Every other weight parameter of the model is frozen during fine-tuning except for the forecasting head. The model takes the same input of the space-time slice image and its associated GHI values $\{\tau; G_{\tau} \dots G_T\}$. However, instead of reconstructing the original input, the model predicts the GHI at future time instances: $[G_{T+1} \dots G_{T+N}]$. The goal of this is to limit the amount of trainable parameters for the task of forecasting all while using the high-level features and learned weights from the encoder. Both pre-trained and fine-tuned models utilize mean-squared error (MSE) as loss functions.

7 EVALUATION

In this section, we present results from the above algorithms for both the simulated and real data.

7.1 Simulations

For experiments on the non-learning based sun occlusion prediction for the simulated data, we use 100 minutes ($\tau = 200$) of past data to predict 30 minutes (N = 60) of sun occlusion state values, in 30 second intervals. Using the algorithms expressed in section 6, we achieve promising results.



Fig. 15. We compare learning and non-learning based approaches for binary sun occlusion states on the simulated data. Using 90% as an AUC value for confident predictions, this plot clearly visualizes the stark difference in prediction that we are able to obtain. Looking at the learning-based approaches, for the hemispherical-based mirror, we are limited to a confident forecasting window of around 3 minutes. Compared to hyperboloidal mirror which is able to achieve confident forecasts around 18 minutes into the future.

As seen in Figure 15 we are able to obtain reliable occlusion predictions up to ≈ 18 minutes into the future compared to the hemispherical mirror that is only able to maintain solidified predictions to \approx 3 minutes. Learningbased methods for occlusion prediction provide the best results, providing greater improvement over the back projected method. For the hyperboloidal mirror we benefit with even greater accuracy, longer through time, all while still substantially outperforming predictions obtained using the hemispherical mirror. For comparison, we experimented using a Transformer-based architecture on the simulated data with the same inputs. Simulated results clearly show the benefit of using a hyperboloidal mirror setup coupled with a learning-based system for long-term prediction of sun occlusion by a cloud. We now present real-world results and show the benefits of using our system.

7.2 Real Data

For experiments using real data, we pass 30 minutes ($\tau = 60$) of data to predict 30 minutes (N = 60) of GHI, in 30 second intervals. Although as not ideal as in the simulated case, results from the real case still achieve GHI prediction performance over the hemispherical mirror. For our accuracy metric, we use the normalized root mean-squared error (nRMSE):

$$RMSE = \sqrt{\frac{\sum_{n=1}^{N} (\hat{y}_n - y_n)^2}{N}}.$$
 (2)

$$nRMSE = RMSE / \sqrt{\frac{\sum_{n=1}^{N} y_n^2}{N}}.$$
 (3)

where \hat{y}_n , y_n is the predicted GHI and true GHI, respectfully. A lower value equates to a more accurate prediction.

We present results of prediction values in Figure 16 which shows that we are able to predict GHI longer into

We present comparison results on various models trained and tested on the same datasets. The Transformer model is the proposed architecture in Section 6.3.2. "Combined" concatenates the space-time-slice image from the hyperboloidal and the hemispherical mirror in the channel dimension. The Seq2Seq model exploits recurrence in the form of gated recurrent units (GRUs) to model the sequential aspect of the space-time-slice images along with the associated GHI. The values are nRMSE between ground truth and predicted GHI values at future time instants.

Model	Data	1 min	5 min	10 min	15 min	20 min	30 min
Persistence	GHI alone	0.191	0.267	0.300	0.323	0.340	0.336
Seq2Seq	GHI alone	0.169	0.233	0.265	0.287	0.304	0.307
Seq2Seq	Hemispherical + GHI	0.185	0.227	0.239	0.255	0.271	0.270
Transformer	Hemispherical + GHI	0.192	0.236	0.250	0.258	0.273	0.284
Transformer	Hemispherical Warped + GHI	0.190	0.236	0.253	0.294	0.310	0.341
Seq2Seq	Hyperboloidal + GHI	0.179	0.227	0.248	0.260	0.277	0.277
Transformer	Hyperboloidal + GHI	0.191	0.227	0.230	0.246	0.266	0.289
Transformer	Hyperboloidal Warped + GHI	0.189	0.233	0.266	0.294	0.303	0.319
Transformer	Combined + GHI	0.183	0.223	0.239	0.252	0.259	0.268

the future with lower error compared to the hemispherical mirror. As a baseline comparison, we also compare both imaging setups to the persistence model which states that the irradiance value will remain unchanged over the forecasting horizon $(\widehat{GHI}_{t+\Delta t} = GHI_t)$.

From Figure 16, we observe that for very short time horizons of 2 minutes (or less), persistence is competitive; this can be attributed to occlusions and disocclusions of the sun, where the GHI changes abruptly, are relatively sparse events. Looking further, we see a significant gap between methods that use visual imagery and persistence; for example, at 14 minutes, the error in the predictions from hyperboloidal imagery is nearly half of that of persistence. We also observe that hyperboloidal predictions are better than hemispherical ones for the first 27 minutes or so. The flip in performance after 27 minutes can likely be attributed to the relatively small size of the dataset. For example, we see dips in the performance of persistence (which is a rather simple non-learning predictor) around the 15 and 28 minutes, hinting at some bias in the underlying data. Further, for the learning-based methods, predictions at longer time horizons also have higher variance, making the small dataset issues worse.

7.3 Ablation

In Table 1, we present an ablation study that shows how different model architectures and different data sources affect the prediction results. For the models, we used two models: Transformer: the proposed architecture in Section 6.3.2, and Seq2Seq: a RNN architecture comprised of gated recurrent units (GRUs) that model the sequential aspect of the spacetime-slice images along with the associated GHI. For the data sources, we have access to the GHI measurements from the pyranometer, and the imagery from the hemispherical and hyperboloidal cameras. We trained models on various combinations as shown in the Table. Finally, as a reference, we report results on a commonly-used baseline called 'persistence' where the latest GHI value is used as the prediction for all future time instants. The results in Table 1 suggest that, except for the shortest time horizon, models trained with hyperboloidal data outperforms the rest. For the shortest time horizon of 1 minute, GHI alone model



Fig. 16. This figure shows the outperformance of the hyperboloidal mirror. Our mirror is able to achieve lower forecasting error longer into the future compared to the traditional hemispherical mirror. Both models are trained using the model described in section 6.3.2 which is fed 30 minutes of past data to predict 30 minutes of GHI. As a baseline comparison, we compare both imaging setups to the persistence model which states that the irradiance value will remain unchanged over the forecasting horizon ($\widehat{GHI}_{t+\Delta t} = GHI_t$). Our model still outperforms the persistence model longer into the future.

works the best; given that all models have access to the GHI data, we believe this is likely a manifestation of the small amount of data compared to the variability of the sky imagery.

8 DISCUSSIONS

This paper argues for a novel system that brings core computational imaging techniques to a compelling problem in renewable energy. Our work provides a pathway to improve the time horizon over which we can reliably forecast solar irradiance; specifically, over conventional wide FoV systems, we can improve predictions from minutes to tens of minutes. We expect such a prediction framework to be of wide interest in the solar photovoltaics community, where resource allocation and energy dispatch is often done in the absence of such predictive analytics. Finally, on a broader scale, we hope the techniques suggested in this paper continue to incite the interest in applications that lie at the intersection of imaging and climate change.

Beyond hyperboloidal mirrors. Our choice of a hyperboloidal shaped mirror provides a single viewpoint system while preserving uniform resolution over the sky plane. A natural question to consider is whether either of these assumptions are need for solar forecasting. For example, we could consider alternate designs that place a larger emphasis on resolution at the horizon at the cost of even reduced resolution at the zenith. Similarly, there is no critical reason why designs like conic mirrors that do not offer a single viewpoint [18] are not applicable here. Effectively, we could enable a framework where the solar forecasting guides the design or optimization of the mirror profile in an end-to-end manner.

Cloud formation and disappearance. One of the factors that we fail to consider in this paper is that clouds form and disappear based on changes in humidity, temperature, and pressure. This violates the slicing model used in this work, in part because clouds can appear in the middle of the field of view, or disappear as it traverses the field. In our testbed, this happens frequently at a particular spot that is over a water body, a few miles from the deployed system. The spatial consistency of the cloud formation suggests that statistical models that have a better understanding of how clouds form and terminate, augmented with other sources of data such as weather, humidity and the geographic layout of the surrounding regions, might have a better chance in handling the effects.

Self-occlusion by clouds. Another factor that we fail to consider is that clouds have vertical extent, and hence a cloud closer to the camera may block one that is further away. This shows itself in the form of radial streaks in Figures 5 and 8. This is a hard problem to resolve in the absence of additional view points. It is likely that a multicamera version of our system with a baseline in kilometers will be able to reason such occlusions and handle them effectively.

Incorporating other data sources. The techniques proposed in this paper will also benefit from other richer sources of data such as satellite imagery, weather prediction, and humidity measurements. Such sources of data are often publicly available; however, each of them have unique features that need to be accounted for. For example, satellite data that provides very large spatial extent, has very poor temporal resolution, often in minutes, if not hours. Further, the ground spatial resolution of such data is in meters, which may not be sufficient for the kind of prediction we envision. Wind velocity, which is often available from weather data, is something we can benefit from. However, such measurements are often made at ground level, and at very sparse locations, which limits their utility, since atmospheric wind velocities differ from ground measurements. Yet, the role that temperature, humidity, and, more broadly, the weather play in forecasting cannot be denied. Translating such models to near-future time horizons and higher precision that is demanded for solar forecasting is an interesting approach for subsequent research.

ACKNOWLEDGMENTS

This work was supported, in part, by the NSF EAGER award 2235063. Julian acknowledges support from a GEM Fellowship and the Fritsch Family Fellowship. Lee was supported in part by a TCS Presidential Fellowship.

REFERENCES

- L. Alados-Arboledas, J. Vida, and F. J. Olmo, "The estimation of thermal atmospheric radiation under cloudy conditions," *Intl. Journal of Climatology*, vol. 15, no. 1, pp. 107–116, 1995.
- [2] B. A. Wielicki, R. D. Cess, M. D. King, D. A. Randall, and E. F. Harrison, "Mission to planet earth: Role of clouds and radiation in climate," *Bulletin of the American Meteorological Society*, vol. 76, no. 11, pp. 2125 – 2154, 1995.
- [3] D. Serrano, M. Marín, M. Núñez, S. Gandía, M. Utrillas, and J. Martínez-Lozano, "Relationship between the effective cloud optical depth and different atmospheric transmission factors," *Atmospheric Research*, vol. 160, pp. 50–58, 2015.
- [4] "The effect of clouds on surface solar irradiance, based on data from an all-sky imaging system," *Renewable Energy*, vol. 95, pp. 314–322, 2016.
- [5] I. P. on Climate Change, Climate Change 2013 The Physical Science Basis: Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, 2014.
- [6] B. K. Sovacool, "The intermittency of wind, solar, and renewable electricity generators: Technical barrier or rhetorical excuse?" *Utilities Policy*, vol. 17, no. 3, pp. 288–296, 2009.
- [7] S. Sayeef, S. Heslop, D. Cornforth, T. Moore, S. Percy, J. Ward, A. Berry, and D. Rowe, "Solar intermittency: Australia's clean energy challenge. characterizing the effect of high penetration solar intermittency on australian electricity networks," CSIRO, p. 195, 2012.
- [8] M. Diagne, M. David, P. Lauret, J. Boland, and N. Schmutz, "Review of solar irradiance forecasting methods and a proposition for small-scale insular grids," *Renewable and Sustainable Energy Reviews*, vol. 27, pp. 65–76, 2013.
- [9] V. Kumar, A. Pandey, and S. Sinha, "Grid integration and power quality issues of wind and solar energy system: A review," in *Intl. Conf. Emerging Trends in Electrical, Electronics & Sustainable Energy Systems*, 2016.
- [10] S. Impram, S. V. Nese, and B. Oral, "Challenges of renewable energy penetration on power system flexibility: A survey," *Energy Strategy Reviews*, vol. 31, p. 100539, 2020.
- [11] D. Infield and L. Freris, Renewable energy in power systems. John Wiley & Sons, 2020.
- [12] L. Julian and A. C. Sankaranarayanan, "Precise forecasting of sky images using spatial warping," in *ICCV Workshops*, 2021.
- [13] W. Richardson, H. Krishnaswami, R. Vega, and M. Cervantes, "A low cost, edge computing, all-sky imager for cloud tracking and intra-hour irradiance forecasting," *Sustainability*, vol. 9, no. 4, 2017.
- [14] R. A. Rajagukguk, R. Kamil, and H.-J. Lee, "A deep learning model to forecast solar irradiance using a sky camera," *Applied Sciences*, vol. 11, no. 11, 2021.
- [15] M. A. Brown and D. G. Lowe, "Automatic panoramic image stitching using invariant features," *Intl. Journal of Computer Vision*, vol. 74, pp. 59–73, 2007.
- [16] Y. Dong, M. Pei, Y. Wu, and Y. Jia, "Stitching images from a conventional camera and a fisheye camera based on nonrigid warping," *Multimedia Tools and Applications*, vol. 81, 05 2022.
- [17] E. Nghonda Tchinda, M. K. Panoff, D. Tchuinkou Kwadjo, and C. Bobda, "Semi-supervised image stitching from unstructured camera arrays," *Sensors*, vol. 23, no. 23, 2023.
- [18] S. Baker and S. K. Nayar, "A theory of single-viewpoint catadioptric image formation," *Intl. Journal of Computer Vision*, vol. 35, no. 2, pp. 175–196, 1999.
- [19] https://github.com/Image-Science-Lab-cmu/SkyCam].
- [20] D. F. Victor, Morris, "Total sky imager (tsimovie)," atmospheric Radiation Measurement (ARM) user facility.
- [21] T. Mathanlal and J. Martin-Torres, "Design and development of an ultraviolet all-sky imaging system," *Sensors*, vol. 23, no. 17, 2023.
 [22] Z. El Jaouhari, Y. Zaz, and L. Masmoudi, "Cloud tracking from
- [22] Z. El Jaouhari, Y. Zaz, and L. Masmoudi, "Cloud tracking from whole-sky ground-based images," in *Intl. Renewable and Sustainable Energy Conf. (IRSEC)*, 2015, pp. 1–5.

- [23] "Eclipse: Envisioning cloud induced perturbations in solar energy," *Applied Energy*, vol. 326, p. 119924, 2022.
- [24] C. Ishii, Y. Sudo, and H. Hashimoto, "An image conversion algorithm from fish eye image to perspective image for human eyes," in *IEEE/ASME Intl. Conf. on Advanced Intelligent Mechatronics (AIM* 2003), 2003.
- [25] C. Eising, M. Glavin, E. Jones, and P. Denny, "Review of geometric distortion compensation in fish-eye cameras," in *IET Conf. Publications*, 2008.
- [26] M.-C. Chang, Y. Yao, G. Li, Y. Tong, and P. Tu, "Cloud tracking for solar irradiance prediction," in *IEEE Intl. Conf. on Image Processing* (*ICIP*), 2017.
- [27] V. Jayadevan, J. J. Rodríguez, V. P. A. Lonij, and A. D. Cronin, "Forecasting solar power intermittency using ground-based cloud imaging," 2012.
- [28] Y. Ai, Y. Peng, and W. Wei, "A model of very short-term solar irradiance forecasting based on low-cost sky images," in AIP Conf. Proceedings, 2017.
- [29] S. Sun, J. Ernst, A. Sapkota, E. Ritzhaupt-Kleissl, J. Wiles, J. Bamberger, and T. Chen, "Short term cloud coverage prediction using ground based all sky imager," in *IEEE Intl. Conf. on Smart Grid Communications*, 2014.
- [30] L. Wei, T. Zhu, Y. Guo, C. Ni, and Q. Zheng, "Cloudprednet: An ultra-short-term movement prediction model for ground-based cloud image," *IEEE Access*, vol. PP, pp. 1–1, 01 2023.
- [31] Y. Nie, E. Zelikman, A. Scott, Q. Paletta, and A. Brandt, "Skygpt: Probabilistic short-term solar forecasting using synthetic sky videos from physics-constrained videogpt," 2023.
- [32] M. Sharifzadeh, A. Sikinioti-Lock, and N. Shah, "Machine-learning methods for integrated renewable power generation: A comparative study of artificial neural networks, support vector regression, and gaussian process regression," *Renewable and Sustainable Energy Reviews*, vol. 108, pp. 513–538, 2019.
- [33] A. Alzahrani, P. Shamsi, C. Dagli, and M. Ferdowsi, "Solar irradiance forecasting using deep neural networks," *Procedia Computer Science*, vol. 114, pp. 304–313, 2017.
- [34] Z. Zhen, J. Liu, Z. Zhang, F. Wang, H. Chai, Y. Yu, X. Lu, T. Wang, and Y. Lin, "Deep learning based surface irradiance mapping model for solar pv power forecasting using sky image," *IEEE Transactions on Industry Applications*, vol. 56, no. 4, pp. 3385–3396, 2020.
- [35] J. Zhang, R. Verschae, S. Nobuhara, and J.-F. Lalonde, "Deep photovoltaic nowcasting," *Solar Energy*, vol. 176, pp. 267–276, 2018.
- [36] D. Veikherman, A. Aides, Y. Schechner, and A. Levis, "Clouds in the cloud," in ACCV, 2014.
- [37] V. Holodovsky, Y. Y. Schechner, A. Levin, A. Levis, and A. Aides, "In-situ multi-view multi-scattering stochastic tomography," in IEEE Intl. Conf. on Computational Photography, 2016.
- [38] F. A. Mejia, B. Kurtz, A. Levis, Íñigo de la Parra, and J. Kleissl, "Cloud tomography applied to sky images: A virtual testbed," *Solar Energy*, vol. 176, pp. 287–300, 2018.
- [39] A. Aides, A. Levis, V. Holodovsky, Y. Y. Schechner, D. Althausen, and A. Vainiger, "Distributed sky imaging radiometry and tomography," in *IEEE Intl. Conf. on Computational Photography*, 2020.
- [40] D. J. Diner, S. W. Boland, M. Brauer, C. Bruegge, K. A. Burke, R. Chipman, L. D. Girolamo, M. J. Garay, S. Hasheminassab, E. Hyer, M. Jerrett, V. Jovanovic, O. V. Kalashnikova, Y. Liu, A. I. Lyapustin, R. V. Martin, A. Nastan, B. D. Ostro, B. Ritz, J. Schwartz, J. Wang, and F. Xu, "Advances in multiangle satellite remote sensing of speciated airborne particulate matter and association with adverse health effects: from MISR to MAIA," *Journal of Applied Remote Sensing*, vol. 12, no. 4, p. 042603, 2018.
- [41] J. Martonchik, D. Diner, R. Kahn, T. Ackerman, M. Verstraete, B. Pinty, and H. Gordon, "Techniques for the retrieval of aerosol properties over land and ocean using multiangle imaging," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 36, no. 4, pp. 1212–1227, 1998.
- [42] D. J. Diner, F. Xu, M. J. Garay, J. V. Martonchik, B. E. Rheingans, S. Geier, A. Davis, B. R. Hancock, V. M. Jovanovic, M. A. Bull, K. Capraro, R. A. Chipman, and S. C. McClain, "The airborne multiangle spectropolarimetric imager (airmspi): a new tool for aerosol and cloud remote sensing," *Atmospheric Measurement Techniques*, vol. 6, no. 8, pp. 2007–2025, 2013.
- [43] R. Ronen, Y. Y. Schechner, and E. Eytan, "4d cloud scattering tomography," in *Proceedings of the IEEE/CVF International Conference* on Computer Vision (ICCV), 2021.
- [44] https://cloudatlas.wmo.int/fr/descriptions-of-clouds.html.

- [45] P. J. Mason, "Large-eddy simulation: A critical review of the technique," *Quarterly Journal of the Royal Meteorological Society*, vol. 120, no. 515, pp. 1–26, 1994.
- [46] https://angelgilding.com/media/wysiwyg/PDF/How_To_ Silver_3-D_Glass_Objects.pdf.
- [47] C. Liu, "Beyond pixels: exploring new representations and applications for motion analysis," Ph.D. dissertation, Massachusetts Institute of Technology, 2009.
- [48] "Characterizing and analyzing ramping events in wind power, solar power, load, and netload," *Renewable Energy*, vol. 111, pp. 227–244, 2017.
- [49] T. Godfrey, S. Mullen, D. W. Griffith, N. Golmie, R. C. Dugan, and C. Rodine, "Modeling smart grid applications with cosimulation," in *IEEE Intl. Conf. Smart Grid Communications*, 2010.
- [50] Q. Li, W. Lyu, and J. Yang, "A hybrid thresholding algorithm for cloud detection on ground-based color images," *Journal of Atmospheric and Oceanic Technology*, vol. 28, pp. 1286–1296, 10 2011.
- [51] M. Goswami, K. Szafer, A. Choudhry, Y. Cai, S. Li, and A. Dubrawski, "Moment: A family of open time-series foundation models," in *Intl. Conf. on Machine Learning*, 2024.
- [52] C. Raffel, N. M. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *J. Mach. Learn. Res.*, vol. 21, pp. 140:1–67, 2019.
- [53] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," *ICLR*, 2021.



Leron Julian received his B.S. in Computer Science at Morehouse College in 2019 and is a Ph.D candidate at Carnegie Mellon University. His research aims to increase the efficiency of solar power by leveraging computer vision, computational photography, and learning to capture and forecast the relationship between atmospheric conditions and received solar irradiance at the ground. He is also the recipient of a GEM fellowship and the Fritsch Family fellowship.



Soummya Kar is the Buhl Professor of Electrical and Computer Engineering at Carnegie Mellon University. He received a B.Tech. in electronics and electrical communication engineering from the Indian Institute of Technology, Kharagpur, India, in May 2005 and a Ph.D. in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, in 2010. His research interests include decision-making in large-scale networked systems, stochastic systems, multiagent systems and data science, with applica-

tions in cyber-physical systems and smart energy systems. He is a Fellow of the IEEE.



Aswin C. Sankaranarayanan is a professor in the ECE department at CMU, where he is the PI of the Image Science Lab. His research interests are broadly in compressive sensing, computational photography, signal processing and machine vision. His doctoral research was in the University of Maryland where his dissertation won the distinguished dissertation award from the ECE department in 2009. Aswin is the recipient of best paper awards at SIGGRAPH 2023, CVPR 2019 and ICCP 2021 & 2022, the CIT

Dean's Early Career Fellowship, the Spira Teaching award, the NSF CA-REERaward, the Eta Kappa Nu (CMU Chapter) Excellence in Teaching award, and the Herschel Rich Invention award from Rice University.