Title
A virtual sky imager testbed for solar energy forecasting

Authors
Kurtz, Benjamin
Mejia, Felipe
Kleissl, Jan

Publication Date
2017-12-01

DOI
10.1016/j.solener.2017.10.036

A Virtual Sky Imager Testbed for Solar Energy Forecasting

Benjamin Kurtz[a,*], Felipe Mejia[a], Jan Kleissl[a]

[a]*Center for Renewable Resources and Integration, Department of Mechanical and Aerospace Engineering, University of California, San Diego, United States*

Abstract

Whole sky imagers are commonly used for forecasting irradiance available for solar energy production, but validation of the forecast models used is difficult due to sparse reference data. We document the use of Large Eddy Simulations (LES) and a 3D Radiative Transfer Model to produce virtual clouds, sky images, and radiation measurements, which permit comprehensive validation of the sky imager forecast. We then use this virtual testbed to investigate the primary sources of sky imager forecast error on a cumulus cloud scene. The largest source of nowcast (0-minute-ahead forecast) errors is the converging-ray geometry implied by use of a camera, while longer-term forecasts suffer from overly-simplistic assumptions about cloud evolution. We expect to use these findings to focus future algorithm development, and the virtual testbed to evaluate our progress.

Keywords: whole sky imager, forecast, Large Eddy Simulations

1. Introduction

In recent years, whole-sky imagers have become popular for forecasting solar energy availability on short time horizons [1, 2, 3, 4, 5]. However, validation of these forecasts can be tricky; reference data is often limited to at most a few irradiance sensors, and even in the case where many sensors are present

*Corresponding author: B. Kurtz
Email address: `bkurtz@ucsd.edu` (Benjamin Kurtz)

over a large area, detailed validation data on the cloud field itself is uniformly unavailable. Under these circumstances, validation can determine the forecast accuracy, but apportionment of the forecast error to different components of the algorithm is difficult due to the lack of data about the actual state of the atmosphere and the resulting radiation field. Therefore prioritization of forecast development work is usually not well-informed and is unable to follow cost-benefit principles.

We propose to address some of these limitations by producing a virtual sky imager testbed, in which the configuration of the clouds and resulting irradiance is known. The purpose of this paper is to describe the setup of the virtual testbed and briefly illustrate its potential through a case study. The virtual testbed is used to design and test improvements to whole-sky imager forecast methodology developed at UC San Diego, but it is straightforward to adapt it to any other algorithm.

Simulating clouds is one of the grand challenges of atmospheric physics as it includes scales from micrometers (cloud condensation nuclei) to kilometers (cloud size), multiple phases (vapor, liquid, ice), and even chemistry (hydrophobicity of aerosol species). In terms of short-term (order of 10 minutes) cloud dynamics that are most relevant to sky imager solar forecasting, the multi-scale and multi-phase fluid dynamics need to be represented. In particular atmospheric turbulence plays a critical role in cloud formation (e.g. thermals) and cloud dynamics. Not only do clouds "live" in the turbulent atmospheric boundary layer flow field, but they also generate their own turbulence due to longwave radiative cooling at the cloud top and latent heat release. Large Eddy Simulation (LES) is a uniquely suited tool to simulate these boundary layer and cloud dynamics. In LES the large turbulent eddies that are responsible for most of the momentum, heat, and moisture transport are explicitly resolved and simulated faithfully based on the Navier Stokes equations. The small scales (less than about 10 meters) cannot be resolved due to computational cost and are parameterized through subfilter scale models [6]. LES also simulates all modes of heat transfer, water vapor transport and phase change, as well as cloud microphysics.

LES is a mature field in engineering and atmospheric science and the resolution, subfilter scale models, and microphysics models have been continually improved over the past decades [7, 8].

Virtual cloud fields will be produced using LES. Surface-level irradiance fields and simulated whole-sky images will be derived from a 3-dimensional radiative transfer model (3D RTM). These tools (LES and 3D RTM) are significantly more physically grounded and accurate than current sky imager forecast algorithms, so there is considerable scope for improving sky imager forecasts based on the virtual testbed. It is worth noting that the virtual testbed need not reproduce a given observed cloud field for this to be useful, so long as the virtual clouds behave similarly to real clouds. Why not just use the LES and 3D RTM for forecasting in the first place? First, while recent GPU-accelerated LES codes [9] approach the speeds necessary to produce operational forecasts, the computational requirements for LES and 3D RTM tools are currently too large to be feasible for short-time-horizon forecasting. Furthermore, even in those cases where LES has been run operationally on a wide variety of measured data [10, 11], the cloud fields are statistically accurate on timescales from tens of minutes to hours. To produce meaningful forecasts of individual clouds, LES would require input of a detailed state of the atmosphere including detailed humidity and velocity fields which, as noted, are generally unavailable. Even here, the virtual testbed is useful, as it allows improved testing of 3D cloud detection algorithms for whole-sky imagers, which could eventually be used as input to an LES-based forecast.

In section 2, we present the virtual testbed and whole-sky imager forecast. Section 3 compares the results of the sky imager forecast to those of the virtual testbed, paying special attention to the newfound ability to determine errors of difficult-to-measure quantities such as wind speed aloft and 3D cloud structure. Differing geometrical perspectives and cloud field dynamics constitute the largest sources of error in the current forecast, with geometry playing a larger role at short forecast horizons, and cloud evolution dominating the error for further-ahead forecasts. Discussion and conclusions are provided in Section 4.

2. Virtual Testbed Components

2.1. Large Eddy Simulation

LES are carried out using the UCLA LES [12, 13, 14], which has been thoroughly validated and tested for a number of cases including continental cumulus [15], raining cumulus [8], and stratocumulus clouds [13]. The UCLA LES uses the Smagorinsky sub-gridscale model, and parameterizes cloud microphysics following Stevens and Seifert [8]. Interactive radiation is implemented via a Monte Carlo version [16] of the delta-four-stream model [17]. Cloud droplet radius for both radiation and microphysics is modeled by assuming a fixed cloud droplet mixing ratio.

A single 14.5 hour simulation was carried out using example input data modeled for continental cumulus clouds, following the base case in [18], which is itself based on a detailed LES study of measurements taken at the Southern Great Plains (SGP) site of the Atmospheric Radiation Measurement (ARM) program [15]. Following prior simulations [18], precipitation was disabled in the microphysics model, leaving cloud liquid water diagnosed as the total water mixing ratio in excess of the saturation mixing ratio, and with the fixed cloud droplet mixing ratio of 70×10^6/kg. Initial profiles of atmospheric temperature and humidity, as well as input surface fluxes are shown in Figure 1. Small volumetric forcings are applied as in [15] in order to represent observed large-scale advection in the periodic simulation domain. This day represents typical formation of a convective boundary layer due to surface heating, with cumulus clouds forming at the top of the (initially clear) boundary layer. As the day progresses, the cloud base rises from 1000 m to around 1500 m, with maximum cloud thickness of around 1250 m. Both the boundary layer and the clouds continue to deepen until late afternoon when solar radiation has decreased significantly. Typical horizontal cloud size is 400 m. Hemispherical cloud cover peaks just above 65% around solar noon; Figure 6 later shows hemispherical cloud cover over the course of the day.

LES grid cells are 50 meters across in both horizontal dimensions and 40

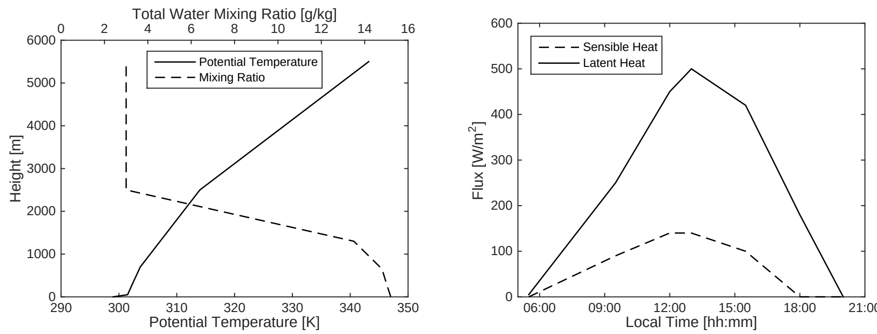

Figure 1: Profiles of temperature and humidity at simulation start, along with surface convective heat fluxes during the simulation.

meters high, spanning a 6.4 km domain that is 5.1 km deep. Periodic boundary conditions are used in the horizontal dimensions. A 10-cell thick sponge layer is used at the top of the domain to prevent wave reflection, while the lower surface uses a no-slip boundary with roughness length of 0.035 m, representative of long grass.

LES requires on the order of an hour of simulation time to properly "spin-up" the turbulent flow and cloud field. After spin-up, the 3D state of the atmosphere (velocity, temperature, pressure, humidity, and liquid water content) is saved every 60 seconds of simulation time for input into the 3D RTM and reference against the sky imager forecast results.

2.2. 3D Radiative Transfer Model

The Spherical Harmonic Discrete Ordinate Method (SHDOM) [19] is used to solve the 3D Radiative Transfer Equation. SHDOM is the most computationally intensive portion of the virtual testbed, requiring over half of the approximately 5000 CPU-core-hours used for the run presented here. SHDOM inputs are derived from the liquid water content output by UCLA LES, combined with the aerosol loading shown in Figure 2, which is based on the `nauru19990707` data file included with SHDOM adjusted to match the observed annual-average aerosol concentration, and effective radius at the ARM SGP AERONET site in 2013. This rapid decrease in aerosol concentration with height matches the

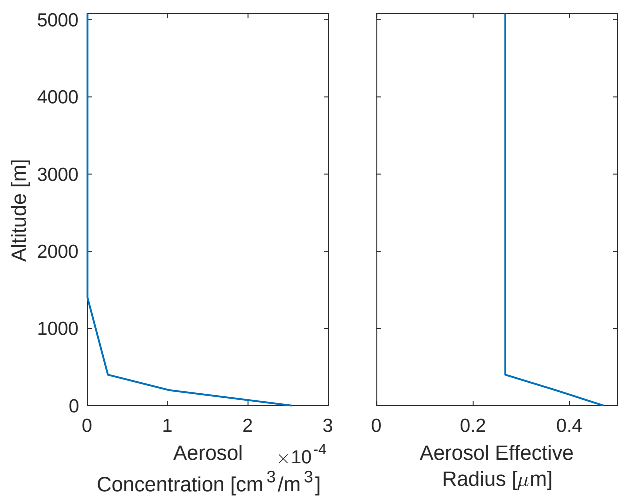

Figure 2: Aerosol loading and effective radius used to produce blue sky in SHDOM.

exponential decay proposed in [20]. SHDOM also uses atmospheric temperature when computing scattering properties; input vertical temperature profiles were derived from LES outputs. In order to simplify interpretation of the results, SHDOM is run with a constant sun position (solar zenith angle of 45°) for the entire simulation time period; this avoids changing clear sky irradiance and geometric perspectives.

At each time step, SHDOM produces a map of surface global horizontal irradiance (GHI) across the simulation domain. In addition, it produces one or more simulated sky images (essentially a map of radiance versus direction at a single location) that can be fed into the sky imager forecast routines. SHDOM results at three different wavelengths (450 nm, 550 nm, and 670 nm) are combined to produce full-color images, and are averaged to approximate broadband GHI. As in the LES, periodic boundary conditions are used.

Figure 3 shows an example of clouds from the LES and the corresponding virtual sky image from SHDOM.

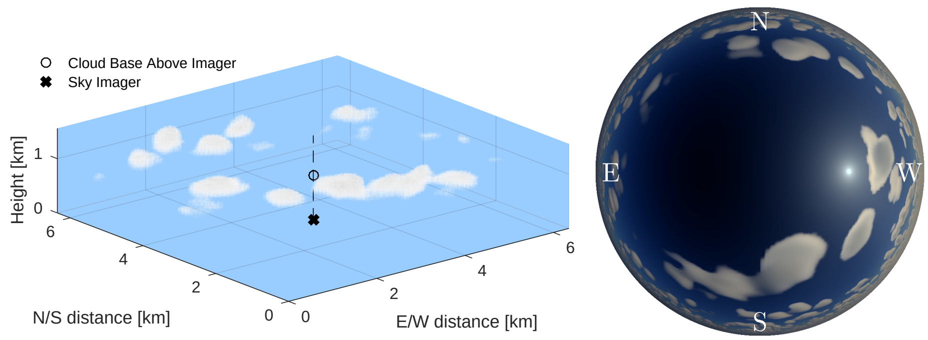

Figure 3: Example LES clouds and virtual sky image at 10:43 local time.

Cloud Map		Cloud Motion		Radiation	
Sky Image	(si)	Pixel Motion	(pix)	k_t Histogram	(kthist)
LES Converging Ray	(conv)	LES Layer Mean	(llm)	Per-class Mean	(ktmean)
LES Zenith Parallel Ray	(zen)			No Quantization	(noquant)
LES Sun Parallel Ray	(sun)			k_t Advection	(ktadv)

Table 1: Naming shorthands for modified versions of the forecast algorithm. The standard forecast is si-pix-kthist.

2.3. Sky Imager Forecast

The sky imager forecast [1] investigated here models clouds as occurring in a single plane at the height of the cloud base. Current cloud positions are detected based on the color of the input image, and future positions are forecast using the "frozen cloud advection" assumption, which assumes that the entire cloud field moves in a uniform direction without changing shape. Inputs to the sky imager forecast are a sky image, cloud base height usually derived from lidar (Light Detection and Ranging) data, and recent measured GHI—used to estimate average cloud optical thickness, which is difficult to determine from the image. Figure 4 illustrates data flow through the sky imager forecast algorithm, along with inputs from the virtual testbed. In addition, several variations of the algorithm are discussed as part of the virtual testbed; naming conventions for these variations are given in Table 1.

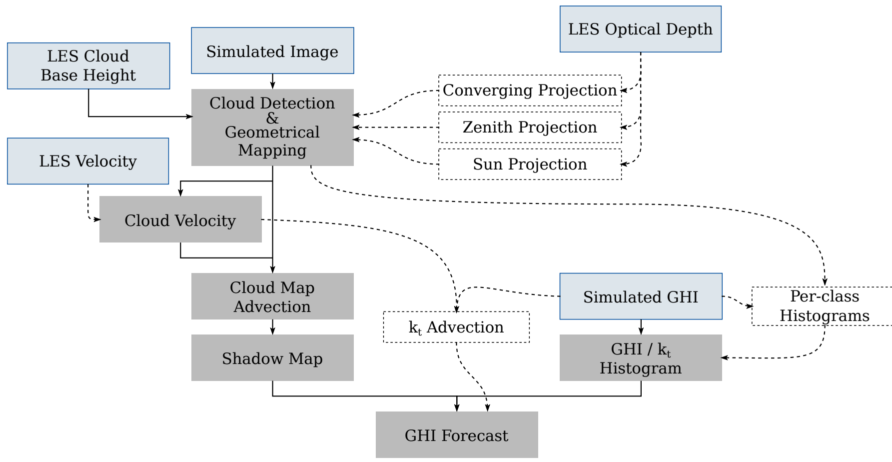

Figure 4: Data flow through sky imager forecast algorithms with inputs from virtual testbed. Solid arrows indicate the standard flow of data through the algorithm, while dashed lines show where "correct" data from the virtual testbed can be used in place of a step in the forecast algorithm. Outputs of LES or SHDOM are shown with a thin solid outline, while derived results have a dashed outline; steps in the basic sky imager forecast have no outline.

2.3.1. Cloud Detection and Geometrical Mapping

In the virtual sky imager testbed, cloud base height is determined based on the first grid cell to have significant liquid water content. As lidar point measurements of cloud base height are generally accurate, the "correct" LES-derived cloud height is used directly for forecasting. In practice, errors would be introduced in the process of interpolating point measurements of cloud height into an accurate height for an entire layer, particularly in the presence of topography or heterogeneous land surface and over larger areas. In the interest of brevity, we do not address these errors here.

Cloud detection operates on the virtual sky images in the same manner as real sky images, and classifies each pixel of the input image as clear sky, thin cloud, or thick cloud, by applying thresholds to the difference between the red-blue ratio (RBR) of the image being analyzed and RBR of a clear sky. Pixels with $\mathrm{RBR} - \mathrm{RBR}_{\mathrm{clear}} \geq 0.4591$ are considered thick cloud, while those with $0.4591 > \mathrm{RBR} - \mathrm{RBR}_{\mathrm{clear}}$ and $\mathrm{RBR} - \mathrm{RBR}_{\mathrm{clear}} \cdot \mathrm{HCF} \geq 0.3044$ are considered

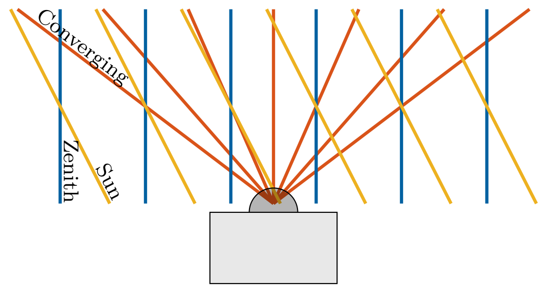

Figure 5: Ray geometry for cloud projections. All three methods are used for computing reference cloud maps from LES. Converging-ray projection is also implicitly used when producing cloud maps from sky images, and “sun projection” is used to map from the cloud plane to shadows on the ground.

thin cloud. These thresholds generally vary with camera and location, and these values were manually selected specifically for use with the virtual testbed based on five images. HCF is the haze correction factor, and helps distinguish thin clouds from background haze. It is iteratively determined for each frame so that the average RBR of portions of the image detected as clear matches the RBR of the haze-corrected clear sky.

Reference cloud maps are derived from LES optical depth, with optical depths greater than 1.5 considered thick clouds and any smaller non-zero optical depth considered thin cloud. Optical depth is the integral of extinction coefficient μ along the rays of the projection, normalized by ray orientation.

$$\text{Optical Depth} = \int \mu \frac{dz}{ds}\, ds \tag{1}$$

$$\mu = \frac{3}{2} \frac{\text{LWC}}{\rho_l r_e}, \tag{2}$$

where LWC is the liquid water concentration in kg/m^3, ρ_l is the density of water, and r_e is the effective droplet radius, here fixed at 8 μm.

As there is no obviously “correct” way to compress a 3D cloud into a plane, reference optical depth maps are computed using three different geometries (illustrated in Figure 5): zenith projection, sun projection, and converging-rays projection. Zenith and sun projected cloud maps compute the cloud optical depth along parallel rays, while the converging-ray projection computes cloud optical depth along rays emanating from the location of the camera. Because it uses the same projection function as the camera, the converging-ray projection is representative of the best results we can expect to achieve with a pixel-by-pixel cloud detection on a sky image, while the sun projection is most relevant to the actual irradiance received at ground level. The zenith projection is similar to the view from a satellite positioned directly overhead.

2.3.2. Cloud Velocity and Cloud Map Advection

The sky imager forecast computes cloud speeds based on pixel motion between adjacent frames. Motion vectors are determined for small regions of the image, and then clustered and averaged to produce a single wind vector that will be used to advect the entire cloud field. Assuming that clouds travel on the background flow, reference wind vectors can be obtained directly from the LES as the vector average wind at the cloud base height.

2.3.3. Shadow Mapping and GHI Forecast

The final step of the forecast is to place cloud shadows and estimate $\mathrm{GHI}(x, y, t)$. The correct way to estimate surface GHI is to run a 3D RTM on a 3D field of extinction coefficients, which accounts for attenuation of the direct beam and 3D photon transport for diffuse radiation. Sky imager forecasts require simplifications both because 3D fields are not available and due to the computational complexity of 3DRTM. At present (kthist in Table 1), effects on direct and diffuse irradiance are lumped by assigning a clear-sky index k_t (fraction of clear-sky GHI that will be received) to each cloud class:

$$\mathrm{GHI}(x, y, t) = \mathrm{GHI}_{\mathrm{csk}}(t) \times k_t(\text{cloud class}(x, y, t)) \tag{3}$$

with cloud classes projected from the cloud plane to the ground using "sun projection" geometry from Figure 5. The k_t for each cloud class is selected by finding three peaks (modes) in the histogram of measured GHI data from the past 2 hours. If fewer than three peaks are found, defaults of 0.42, 0.70, or 1.06

(for thick, thin, and clear respectively) are used. "Correct" k_t for each class is determined by averaging the SHDOM GHI of pixels located in the shadows of each class.

In addition to reference GHI computed in SHDOM, we also compare several other radiation schemes, designed to illuminate the errors that arise in the existing forecast model. 1. Following the current sky imager forecast method, but using the "correct" k_t for each class (ktmean). 2. Converting directly from optical depth (Eq. 1, any projection) to k_t at each point via an exponential model fit at each time step, without quantizing into cloud classes (noquant). 3. k_t advection, i.e. $k_t(x, y, t) = k_t(x - ut, y - vt, 0)$, for clouds moving with velocity (u, v), without reference to detected clouds (ktadv). 4. Persistence, i.e. $k_t(x, y, t) = k_t(x, y, 0)$. Method (1) removes errors in the k_t assignment, while (2) removes errors due to quantization. Methods (3) and (4) are initially perfect, and are included primarily to illustrate model performance as the cloud field changes. We note that methods (2) and (3) require more detailed information about the cloud field than is generally available outside the virtual testbed.

2.3.4. Error Calculations

Comparison of each of these intermediate forecast quantities to the reference values can obviously be done directly, but it is also beneficial to compare the relative effects of errors at each step. For example, it is not clear how a cloud-speed error of 1 m/s relates to an error in cloud-cover of 10%. For this purpose, we compare the final forecast errors that result from substituting various reference values into subsequent forecast steps. For example, we might calculate forecast cloud positions and shadows using the "correct" sun projection reference cloud map rather than the cloud map derived from the sky image (corresponding to sun-pix-kthist in Table 1). This and other varying paths through the forecast algorithm are drawn in Figure 4. Naming conventions for variations are summarized in Table 1.

Note that domain-average GHI is nearly constant over short periods of time, so errors are computed for all points, rather than for the domain average. Errors thus obtained are representative of validating sky imager forecasts against point measurements at weather stations. Forecasts for power plants exhibit reduced random error magnitudes due to spatial averaging. Forecasting and error reporting commence 15 minutes before the formation of the first clouds and extend through the end of the simulation.

When comparing the error E of different methods to a baseline case, it is also useful to define forecast skill,

$$\text{forecast skill} = 1 - \frac{E}{E_{\text{ref}}}, \tag{4}$$

which is small positive number (up to 1 for a perfect forecast) if a method performs better than the baseline, and a negative number if the method under consideration is worse.

3. Results and Discussion

3.1. Errors in Intermediate Quantities

Time series of cloud cover, cloud velocity, and k_t results are illustrated in Figure 6 and demonstrate the forecast's ability to match overall atmospheric conditions. During the simulation run, the sky imager forecast had errors (RMS) of 2.0 m/s and 1.7 degrees for the detected cloud velocities compared to LES wind at the cloud base height. Considering multiple cloud classes, 83% of pixels were correctly classified, with 7% that were classified as a cloud of the wrong class, and the remaining 10% classified as clear when they should have been cloud or vice versa. Detected k_t values from the existing histogram-based method were also relatively reliable, with errors (RMS) of 0.033, 0.078, and 0.079 for clear, thin, and thick categories.

Based purely on these error numbers, only the cloud speed error appears large enough to be of concern; the following sections consider the relative importance of these different errors to the GHI forecasts. Errors at short time horizons will mainly be influenced by the cloud mapping and radiation models, while longer forecasts rely significantly on the ability to predict the evolution of the cloud field.

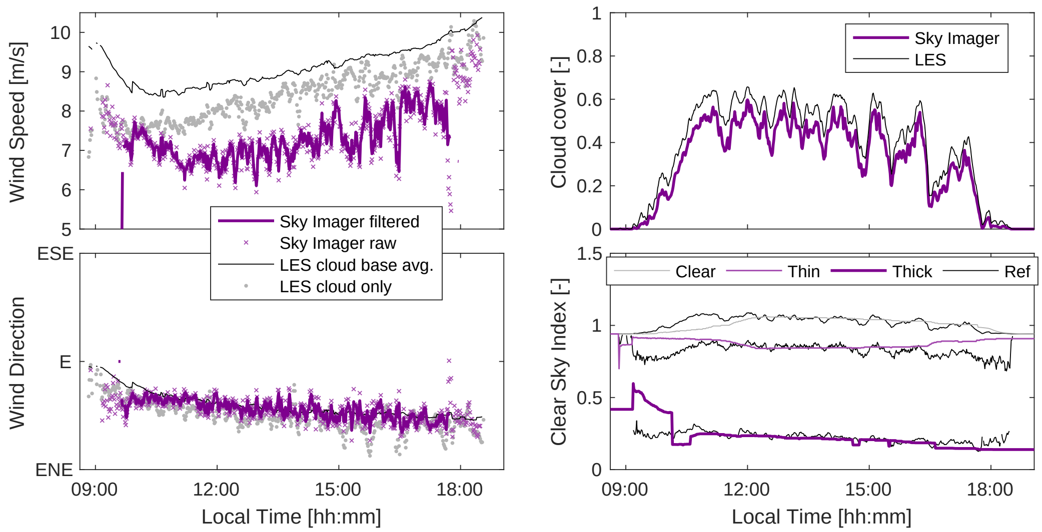

Figure 6: Time series comparison of cloud motion, hemispherical cloud cover, and per-class k_t values against the references derived from LES. Cloud motion filtering smooths data and removes points where cloud cover < 0.05. Clear-sky k_t exceeds 1 during much of the simulation due to cloud enhancement.

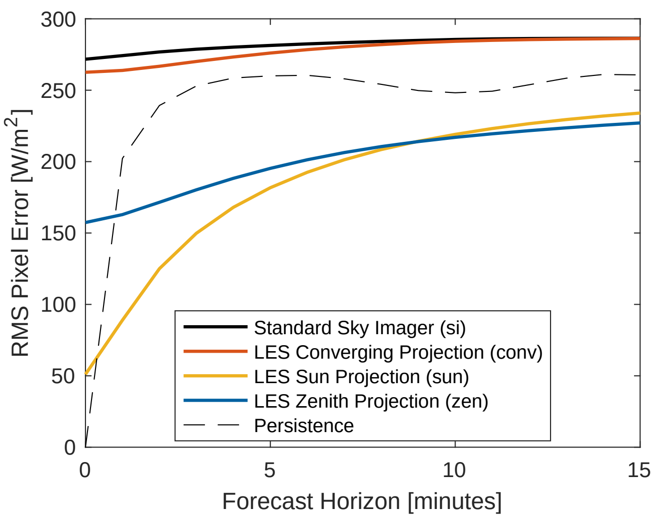

Figure 7: GHI forecast errors for several methods (X-pix-kthist) of mapping 3D clouds onto a horizontal plane at the cloud base height. All methods increase in error as the forecast horizon grows, however the methods with converging rays (Sky Imager and LES Converging) are unable to beat a persistence forecast. For reference, the typical range of GHI at any given time is around 670 W/m^2.

3.2. Projection

Figure 7 illustrates the difference between the different cloud projection schemes. The standard sky imager forecast errors and persistence forecast errors follow the trend observed in previous work involving real-world data [1]. The converging-ray reference cloud map produces slightly better short term forecasts, but does no better at longer time horizons. Most notable, however, is the significant improvement that comes from using one of the parallel-ray projections, particularly at short time horizons. The sun projection method works best for short forecasts because it best matches the actual path light takes through the atmosphere, while zenith projection seems to work better at longer time horizons. We suspect this is because cumulus clouds form convectively, and as a result are more dynamic in the vertical dimension (which is hidden in the zenith projection) than the horizontal dimensions. Converging-ray projection was generally known (e.g. [21]) to cause some degree of perspective error, but the authors had not previously realized just how much of the error (over 2/3 at the shortest time horizons) was a result of this. The remaining error at zero time horizon ("nowcast") is due to cloud detection (thresholding of optical depth) and the complex 3D diffuse irradiance field that is not captured by the k_t assignment; this error is further investigated in the following section. The inadequacy of the frozen cloud advection hypothesis and to a lesser extent, cloud speed errors (Figure 6), result in all the methods having larger errors at long time horizons.

3.3. Radiation

To investigate the remaining nowcast errors, we consider the radiation component of the forecast algorithm. The current algorithm makes two significant approximations. First, it treats GHI as depending only on the value in the 2D cloud map at a single point. This is accurate for the direct beam, but not at all representative of how diffuse irradiance propagates. Secondly, as a result of this single-point approximation and our quantized cloud map, the cloud shadows are also quantized. To assess the performance implications of these assumptions, the results of relaxing each of these assumptions are demonstrated in Figure 8. The sun projection is used for this comparison as it is most physically representative, and performs best (Figure 7) at short time horizons.

Nowcast errors are independent of cloud motion and therefore reveal the radiation model errors. Choosing the optimal (mean observed at zero horizon) k_t for each class (red line) results in modest (around 12%) improvements in the radiation model. However, even eliminating the quantization (blue line) leaves over 40% of the nowcast error. The remainder requires properly dealing with diffuse irradiance and 3D cloud structure.

At longer time horizons, the difference between the various methods decays as the advected static cloud map becomes less representative of the real cloud field. The k_t advection scheme uses the initial measured $k_t(x, y)$, and is thus perfect initially, but by 5 minutes is hardly any better than the standard algorithm. Interestingly, the mean k_t method actually performs better at long time

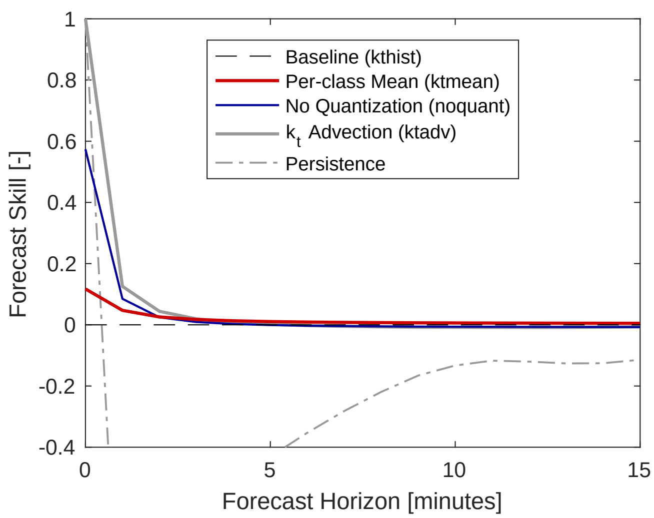

Figure 8: Forecast skill for various methods of modeling radiation. The baseline (reference) method is sun-pix-kthist, and other illustrated methods are sun-pix-X. The mean k_t observed for each class at the time of forecast gives best results for a quantized, single-point radiation model. Results are also shown for a single-point model without quantization and the full (3D with diffuse) radiation model run at the time of forecast. For comparison, a persistence forecast (constant k_t at each point) is also shown.

horizons, presumably because localized fluctuations about the mean values tend to change more quickly with time and smoothing forecast fields therefore tends to reduce errors.

It should also be noted that this cloud scene contains only medium-thickness fair-weather cumulus clouds which probably tends to improve the performance of the baseline radiation model compared to conditions with a mix of thin and thick clouds. In particular, the algorithm would likely have more difficulty selecting the correct peaks from a more complicated k_t histogram.

3.4. Cloud Evolution

To address errors at longer forecast horizons, additional comparisons were run using the nominal average wind vector from LES. As illustrated in Figure 9, using the nominal wind vector from LES results in less than 4% improvement in

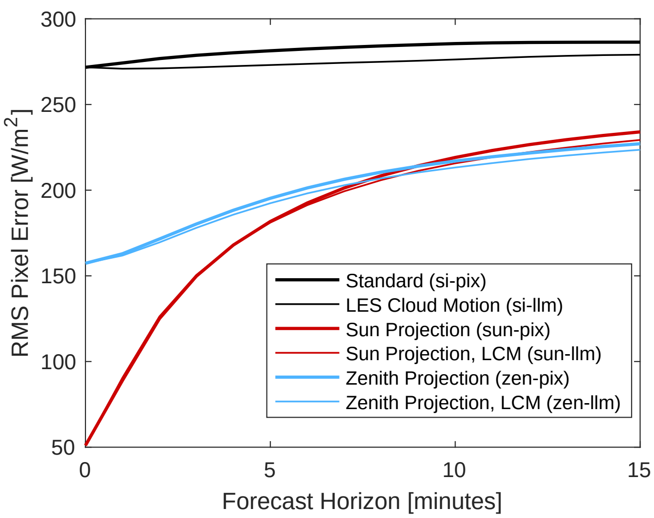

Figure 9: Forecast errors for frozen cloud advection compared with reference motion vectors from LES. Algorithm variations shown are X-Y-kthist. Similar behavior is observed for other forecast variants not shown.

forecast accuracy. For the sun and zenith projections, these improvements are relatively small (median 1.1% and 1.6% respectively across forecast horizons) in comparison to the overall increase in error with forecast horizon, suggesting that the current sky imager forecast's motion vector algorithm works well (at least, for this simple, one-layer cloud case), and that we have essentially saturated the capabilities of the frozen cloud advection model; further improvements would require a more dynamic model for cloud development. After a forecast horizon of 5 minutes, a forecast that assumes constant k_t thoughout the domain (not shown) outperforms all other forecast variants. Thus, 5 minutes can be considered to be the decorrelation time scale of this cloud field and an upper bound for the validity of the frozen cloud assumption; the decorrelation time scale is expected to vary with atmospheric conditions.

Some additional attention is required to the motion estimation algorithm as applied to the sky image or converging cloud map. In Figure 6 previously, a significant deviation was observed between the detected cloud speed and the LES reference speed—the pixel motion estimation consistently under-predicts speed. While the contribution to overall error is still always less than 4% (median 2.9%) in this case, approximately half of the forecast-horizon-dependent error is attributable to this velocity under-prediction. This under-prediction appears to be related to the vertical geometry of the cloud, as Figure 10 shows that the

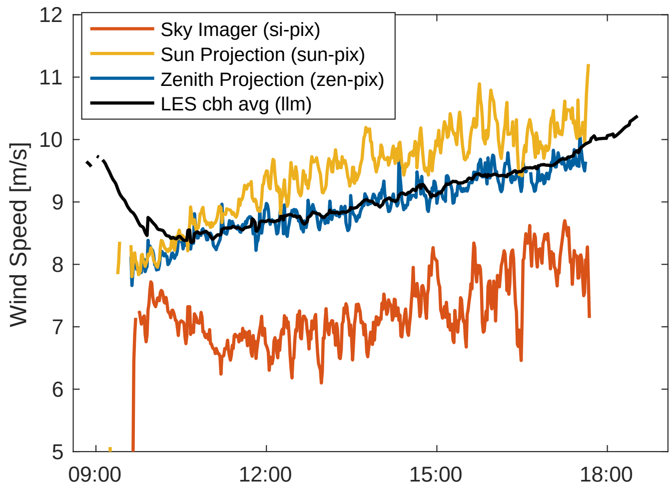

Figure 10: Cloud speed estimates based on pixel motion for the different projections in comparison to the LES reference speed. The black (llm) and red (si-pix) lines are also shown in Fig. 6.

detected speeds in the sun and zenith projections match the LES results much more closely. Furthermore, experiments with non-physical clouds occupying only a single grid layer showed no issues with motion estimation, suggesting that cloud depth or wind shear is involved. At present, a complete explanation for this under-prediction of velocity is lacking; it will be investigated in more detail in future work. As noted above, the more accurate projections also yield more accurate motion estimates without additional work, so this investigation is primarily of interest until it becomes possible to generate 3D cloud maps from sky imagery.

4. Discussion and Conclusions

The virtual sky imager testbed is a valuable and versatile tool, allowing us to validate the quality of outputs from many steps of the sky imager forecast algorithm, and to assess the source of remaining errors. Here, the testbed demonstrated that for a simple cloud scene with scattered cumulus clouds, nowcast errors already negated most of the utility of sky imager forecasting. Nowcast errors primarily originated in the converging-rays projection of 3D clouds into a 2D plane, while cloud detection contributed relatively minor errors.

Sky imager forecast errors further increase from the nowcast errors, never managing to outperform a persistence forecast. The virtual sky imager testbed allowed cloud motion estimation errors to be examined separately and these errors were found to be small except for converging-ray projections, and of minor consequence there. Further, the virtual sky imager testbed demonstrated that even with projection errors in the nowcast corrected, the frozen-cloud-advection assumption for forecasting future cloud positions increasingly deteriorates forecast accuracy at longer time horizons.

However, the virtual testbed suffers from a number of limitations as well. LES is mostly limited to boundary layer clouds over flat and homogeneous or at least idealized (periodic) ground surfaces. The current LES setup is therefore limited in its ability to produce high clouds, including cumulonimbus and cirrus, as well as multiple cloud layers and topographic clouds. In principle, use of a larger domain, non-idealized measured inputs, and advances in numerical codes can enable simulations of these other cloud types (e.g. as in [22, 10]), but with considerable computational and human resource investments. Varying types of clouds and topography would likely influence the measured errors quantitatively, but qualitative conclusions would likely be similar to those for cumulus clouds. For example, clouds with smaller vertical extent such as stratocumuli would likely reduce projection errors, but sun or zenith projection would still be expected to outperform converging-ray projection. Therefore, while not necessarily sufficient to validate forecasts under the variety of conditions seen in the real world, for development of generic forecast algorithms it is preferable to utilize simpler-to-implement, well-studied cases. Multiple cloud layers, on the other hand, considerably complicate cloud detection (shadows of upper layers on lower layers), cloud mapping (single-cloud-plane model is no longer accurate), and motion estimation (distinguish multiple layers moving independently), and are therefore more likely to reveal qualitatively different results. In a future iteration of the virtual testbed, multiple cloud layers might be approximated by running multiple separate LES simulations and stacking the results, though this is obviously not physically realistic. Finally, the process of producing virtual sky images currently omits both stray light and sensor noise. Noise, and in particular stray light tend to cause issues with cloud detection, so cloud detection in the virtual testbed is likely more accurate than for real images. Models for noise and stray light could be added in a future version of the virtual sky imager testbed.

Despite these limitations, the virtual testbed is expected to be a valuable tool for validating and improving sky imager forecast algorithms. The authors would be happy to share the virtual sky images and ancillary data with other researchers.

5. Acknowledgements

We acknowledge funding from the California Energy Commission EPIC program under EPC-14-005. Furthermore, we are greatful to the creators of UCLA-LES for making their code publicly available.

6. References

[1] H. Yang, B. Kurtz, D. Nguyen, B. Urquhart, C. W. Chow, M. Ghonima, J. Kleissl, Solar irradiance forecasting using a ground-based sky imager developed at UC San Diego, Solar Energy 103 (2014) 502–524. doi:10.1016/j.solener.2014.02.044.
URL http://dx.doi.org/10.1016/j.solener.2014.02.044

[2] C. Gauchet, P. Blanc, B. Espinar, B. Charbonnier, D. Demengel, Surface solar irradiance estimation with low-cost fish-eye camera, in: Workshop on "Remote Sensing Measurements for Renewable Energy", Risoe, Denmark, 2012.
URL https://hal-mines-paristech.archives-ouvertes.fr/hal-00741620

[3] T. Schmidt, J. Kalisch, E. Lorenz, D. Heinemann, Evaluating the spatio-temporal performance of sky-imager-based solar irradiance analysis and forecasts, Atmospheric Chemistry and Physics 16 (5) (2016) 3399–3412. doi:10.5194/acp-16-3399-2016.
URL http://www.atmos-chem-phys.net/16/3399/2016/

[4] A. Cazorla, C. Husillos, M. Antón, L. Alados-Arboledas, Multi-exposure adaptive threshold technique for cloud detection with sky imagers, Solar Energy 114 (2015) 268 – 277. doi:10.1016/j.solener.2015.02.006.
URL http://www.sciencedirect.com/science/article/pii/S0038092X1500064X

[5] Z. Peng, D. Yu, D. Huang, J. Heiser, S. Yoo, P. Kalb, 3D cloud detection and tracking system for solar forecast using multiple sky imagers, Solar Energy 118 (2015) 496 – 519. doi:10.1016/j.solener.2015.05.037.
URL http://www.sciencedirect.com/science/article/pii/S0038092X15002972

[6] C. Meneveau, J. Katz, Scale-invariance and turbulence models for large-eddy simulation, Annual Review of Fluid Mechanics 32 (1) (2000) 1–32. arXiv:http://dx.doi.org/10.1146/annurev.fluid.32.1.1, doi:10.1146/annurev.fluid.32.1.1.
URL http://dx.doi.org/10.1146/annurev.fluid.32.1.1

[7] C.-H. Moeng, A large-eddy-simulation model for the study of planetary boundary-layer turbulence, Journal of the Atmospheric Sciences 41 (13) (1984) 2052–2062. arXiv:http://dx.doi.org/

10.1175/1520-0469(1984)041<2052:ALESMF>2.0.CO;2, doi:10.1175/1520-0469(1984)041<2052:ALESMF>2.0.CO;2.
URL http://dx.doi.org/10.1175/1520-0469(1984)041<2052:ALESMF>2.0.CO;2

[8] B. Stevens, A. Seifert, Understanding macrophysical outcomes of microphysical choices in simulations of shallow cumulus convection, Journal of the Meteorological Society of Japan. Ser. II 86A (2008) 143–162. doi:10.2151/jmsj.86A.143.

[9] J. Schalkwijk, E. J. Griffith, F. H. Post, H. J. J. Jonker, High-performance simulations of turbulent clouds on a desktop pc: Exploiting the gpu, Bulletin of the American Meteorological Society 93 (3) (2012) 307–314. doi:10.1175/BAMS-D-11-00059.1.
URL http://dx.doi.org/10.1175/BAMS-D-11-00059.1

[10] W. I. Gustafson, A. M. Vogelmann, X. Cheng, S. Endo, B. Krishna, Z. Li, T. Toto, H. Xiao, Large-eddy simulation (LES) ARM symbiotic simulation and observation (2016). doi:10.5439/1256454.
URL http://www.arm.gov/science/themes/lasso/releases/

[11] R. A. J. Neggers, A. P. Siebesma, T. Heus, Continuous single-column model evaluation at a permanent meteorological supersite, Bulletin of the American Meteorological Society 93 (9) (2012) 1389–1400. doi:10.1175/BAMS-D-11-00162.1.

[12] B. Stevens, C.-H. Moeng, P. P. Sullivan, Large-eddy simulations of radiatively driven convection: Sensitivities to the representation of small scales, Journal of the Atmospheric Sciences 56 (23) (1999) 3963–3984. doi:10.1175/1520-0469(1999)056<3963:LESORD>2.0.CO;2.

[13] B. Stevens, C.-H. Moeng, A. S. Ackerman, C. S. Bretherton, A. Chlond, S. de Roode, J. Edwards, J.-C. Golaz, H. Jiang, M. Khairoutdinov, M. P. Kirkpatrick, D. C. Lewellen, A. Lock, F. Müller, D. E. Stevens,

E. Whelan, P. Zhu, Evaluation of large-eddy simulations via observations of nocturnal marine stratocumulus, Monthly Weather Review 133 (6) (2005) 1443–1462. arXiv:http://dx.doi.org/10.1175/MWR2930.1, doi:10.1175/MWR2930.1.
URL http://dx.doi.org/10.1175/MWR2930.1

[14] B. Stevens, T. Heus, B. van Stratum, A. Seifert, UCLA LES, https://github.com/uclales/uclales (2015).

[15] A. R. Brown, R. T. Cederwall, A. Chlond, P. G. Duynkerke, J.-C. Golaz, M. Khairoutdinov, D. C. Lewellen, A. P. Lock, M. K. MacVean, C.-H. Moeng, R. A. J. Neggers, A. P. Siebesma, B. Stevens, Large-eddy simulation of the diurnal cycle of shallow cumulus convection over land, Quarterly Journal of the Royal Meteorological Society 128 (582) (2002) 1075–1093. doi:10.1256/003590002320373210.
URL http://dx.doi.org/10.1256/003590002320373210

[16] R. Pincus, B. Stevens, Monte carlo spectral integration: a consistent approximation for radiative transfer in large eddy simulations, Journal of Advances in Modeling Earth Systems 1 (2), 1. doi:10.3894/JAMES.2009.1.1.

[17] K.-N. Liou, Q. Fu, T. P. Ackerman, A simple formulation of the delta-four-stream approximation for radiative transfer parameterizations, Journal of the Atmospheric Sciences 45 (13) (1988) 1940–1948. doi:10.1175/1520-0469(1988)045<1940:ASFOTD>2.0.CO;2.

[18] L. M. Hinkelman, B. Stevens, K. F. Evans, A large-eddy simulation study of anisotropy in fair-weather cumulus cloud fields, Journal of the Atmospheric Sciences 62 (7) (2005) 2155–2171. doi:10.1175/JAS3463.1.
URL http://dx.doi.org/10.1175/JAS3463.1

[19] K. F. Evans, The spherical harmonics discrete ordinate method for three-dimensional atmospheric radiative transfer, Journal of the Atmospheric Sciences 55 (3) (1998) 429–446. doi:10.1175/1520-0469(1998)055$<$0429:

TSHDOM$>$2.0.CO;2.
URL http://dx.doi.org/10.1175/1520-0469(1998)055<0429:TSHDOM>2.0.CO;2

[20] C. A. Gueymard, D. Thevenard, Monthly average clear-sky broadband irradiance database for worldwide solar heat gain and building cooling load calculations, Solar Energy 83 (11) (2009) 1998 – 2018. doi:10.1016/j.solener.2009.07.011.
URL http://www.sciencedirect.com/science/article/pii/S0038092X09001790

[21] E. Kassianov, C. N. Long, M. Ovtchinnikov, Cloud sky cover versus cloud fraction: Whole-sky simulations and observations, Journal of Applied Meteorology 44 (1) (2005) 86–98. doi:10.1175/JAM-2184.1.
URL http://dx.doi.org/10.1175/JAM-2184.1

[22] J. Schalkwijk, H. J. J. Jonker, A. P. Siebesma, F. C. Bosveld, A year-long large-eddy simulation of the weather over cabauw: An overview, Monthly Weather Review 143 (3) (2015) 828–844. doi:10.1175/MWR-D-14-00293.1.
URL http://dx.doi.org/10.1175/MWR-D-14-00293.1