

# UC San Diego

## UC San Diego Electronic Theses and Dissertations

### Title

Development of a short-term solar power forecasting capability using ground-based visible wavelength imagery

### Permalink

<https://escholarship.org/uc/item/6f8024vx>

### Author

Urquhart, Bryan Glenn

### Publication Date

2014

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Development of a short-term solar power forecasting capability  
using ground-based visible wavelength imagery

A dissertation submitted in partial satisfaction of the requirements  
for the degree of Doctor of Philosophy

in

Engineering Sciences (Mechanical Engineering)

by

Bryan Glenn Urquhart

Committee in charge:

Jan Kleissl, Chair  
Carlos Coimbra  
Joseph Ford  
Truong Nguyen  
Janet Shields  
George Tynan

2014



The Dissertation of Bryan Glenn Urquhart is approved and it is acceptable in quality and form for publication on microfilm and electronically:

---

---

---

---

---

---

---

Chair

University of California, San Diego

2014

## DEDICATION

This work is dedicated to my family who are always very supportive of my endeavors.

## TABLE OF CONTENTS

<b>Signature Page</b> .....	<b>iii</b>
<b>Dedication</b> .....	<b>iv</b>
<b>Table of Contents</b> .....	<b>v</b>
<b>List of Figures</b> .....	<b>xi</b>
<b>List of Tables</b> .....	<b>xvi</b>
<b>Acknowledgements</b> .....	<b>xvii</b>
<b>Abstract of the Dissertation</b> .....	<b>xix</b>
<b>1 Introduction</b> .....	<b>1</b>
<b>2 Sky Camera Development</b> .....	<b>7</b>
2.1 Summary .....	7
2.2 Introduction .....	7
2.2.1 Imaging for Solar Power Forecasting Applications .....	9
2.2.2 Existing sky imaging hardware .....	12
2.3 Hardware design and selection methods .....	15
2.3.1 Optical design .....	15
2.3.2 Camera and image sensor .....	21
2.3.3 Enclosure and balance of system design .....	22
2.4 System operation .....	24
2.4.1 Image capture and storage .....	24
2.4.2 System monitoring and control .....	25
2.5 Imaging performance characterization .....	26
2.5.1 Noise sources and photoresponse .....	26
2.5.2 Dark response.....	28

2.5.3 Sensor photoresponse uniformity correction .....	30
2.5.4 Photoresponse linearity.....	32
2.5.5 High dynamic range imaging.....	34
2.5.6 Brightness measurement uncertainty in HDR imagery .....	39
2.5.7 Stray light.....	41
2.5.8 Color balancing .....	45
2.6 Deployment experience .....	46
2.7 Conclusions .....	47
Acknowledgements.....	47
<b>3 Camera Calibration .....</b>	<b>49</b>
3.1 Summary .....	49
3.2 Introduction .....	49
3.3 Camera Model .....	52
3.3.1 Forward Camera Model .....	52
3.3.1.1 Projective Transformation Camera Model .....	52
3.3.1.2 Distortion Model .....	55
3.3.1.3 Forward Camera Model Overview .....	60
3.3.2 Backward Projection .....	60
3.4 Solar Position Input from Sky Imager Data .....	62
3.4.1 Imaging Equipment and Setup .....	62
3.4.2 Solar Position Modeling .....	63
3.4.3 Solar Position Calibration Input.....	64
3.5 Calibration Procedure .....	66
3.5.1 Model Initialization .....	67

3.5.1.4 Intrinsic Parameter Estimation.....	67
3.5.1.5 Pose Estimation .....	68
3.5.2 Nonlinear optimization of model parameters .....	69
3.5.3 Backward projection.....	71
3.6 Calibration Results.....	72
3.6.1 Calibration Performance Metrics .....	72
3.6.2 Calibration using the Solar Position.....	72
3.6.3 Camera Model Parameter Uncertainty .....	74
3.7 Conclusions .....	77
Acknowledgements.....	78
Appendix 3.1 Sun Position Detection .....	78
<b>4 Remote Sensing of Clouds: Detection and Velocity Estimation .....</b>	<b>81</b>
4.1 Introduction .....	81
4.2 Cloud Detection and Opacity Classification .....	83
4.2.1 Review of Cloud Detection Methodologies .....	83
4.2.2 Clear Sky Library .....	87
4.2.3 Cloud Detection with the CSL .....	90
4.2.4 The Dynamic Clear Sky Library for the USI.....	91
4.2.5 Classification of Cloud Type.....	93
4.3 Cloud Velocity Estimation .....	96
Acknowledgements.....	100
<b>5 Cloud Stereography .....</b>	<b>101</b>
5.1 Introduction .....	101
5.1.1 Statistical Whole Image Matching Analysis for Single Cloud Layer .....	102



5.1.2 Cloud depth estimation using epipolar geometry .....	103
5.2 Experimental Setup.....	106
5.2.1 Camera System .....	106
5.2.2 Ceilometer .....	106
5.2.3 Cloud Conditions.....	114
5.3 Methods .....	114
5.3.1 Coordinate System .....	115
5.3.2 Scene Plane Definition.....	116
5.3.3 Affine Homography between Calibrated and Uncalibrated Pseudo-image Coordinates .....	117
5.3.4 Pseudo-image Rectification.....	118
5.3.5 Pseudo-image Construction .....	121
5.3.6 Correspondence in Cloud Images Using Dynamic Programming .....	123
5.4 Results and Discussion.....	126
5.4.1 Discussion of individual daily results .....	128
5.5 Conclusions .....	131
Acknowledgements.....	142
<b>6 Solar Power Forecasting .....</b>	<b>143</b>
6.1 An Informal History of UCSD Solar Forecasting .....	143
6.2 Forecasting Methodology.....	144
6.3 Geometric Forecasting Approach .....	147
6.3.1 Application to UCSD DEMROES Network.....	150
6.4 Cloud Optical Depth Estimation .....	156
6.5 Case Study: 48 MW of Photovoltaics at Copper Mountain Solar 1 .....	159

6.5.1 Experimental Setup.....	159
6.5.2 Forecast Results .....	160
6.5.2.1 Forecast Error Metrics.....	160
6.5.2.2 Forecast Performance.....	161
6.5.3 Conclusions .....	164
Acknowledgements.....	165
<b>7 Comparison of the USI and TSI.....</b>	<b>166</b>
7.1 Summary .....	166
7.2 Introduction .....	167
7.3 Experimental Methods .....	168
7.3.1 TSI Site: Copper Moutnain Solar 1, 48MW.....	168
7.3.2 USI Site: UCSD solar energy testbed.....	169
7.3.3 Forecast methodology.....	170
7.3.4 Persistence Forecast .....	171
7.3.5 Comparison Metrics .....	172
7.3.6 Pearson correlation coefficient .....	172
7.3.7 Spearman correlation coefficient.....	172
7.4 Results.....	174
7.4.1 Forecast performance on a single day .....	174
7.4.2 Forecast performance for one month .....	179
7.5 Conclusions .....	182
Acknowledgements.....	183
<b>8 Conclusions and Future Work .....</b>	<b>185</b>
<b>Appendix A. Optimal Geographic Allotment of Distributed PV .....</b>	<b>188</b>

A.1 Summary .....	188
A.2 Introduction.....	189
A.3 Methodology .....	192
A.3.1 Experimental Setup .....	192
A.3.2 Power output simulation .....	193
A.3.3 Multiobjective optimization: maximizing energy generation while minimizing ramp rate .....	195
A.4 Results .....	198
A.5 Conclusions.....	206
Acknowledgements.....	209
<b>References .....</b>	<b>210</b>

## LIST OF FIGURES

<b>Figure 2.1.</b> (a) USI enclosure, (b) top view of enclosure with lid removed, (c) components of USI taken out of the enclosure.....	17
<b>Figure 2.2.</b> (a) Perspective (rectilinear), equidistant, and equisolid angle projection distances as a function of incidence angle, along with the projection for USI 1.2 determined from geometric calibration. The projection distance is normalized by the focal length. (b) Angular resolution of projections in (a).....	17
<b>Figure 2.3.</b> (a) Diffraction pattern measured with a 1000 $\mu\text{m}$ aperture on USI 1.8, with color red, green, and blue color components shown in (b), (c), and (d), respectively. (e) Diffraction of the hexagonal iris blades in the stock lens.....	20
<b>Figure 2.4.</b> Component layout of UCSD Sky Imager camera system. ....	24
<b>Figure 2.5.</b> (a) Occurrence frequency of signal measured in a dark room for 25 different integration times, ranging from 1ms (black) to 2s (lightest gray).....	29
<b>Figure 2.6.</b> (a) An example dark frame for a 100ms exposure and (b) the corrected dark frame. Typical pixel values in (a) range from 32 to 47 with a mean around 40 counts (of $2^{12}$ ). ....	30
<b>Figure 2.7.</b> (a) Raw red image of smooth source; (b) average of red ten frames, including (a); (c),(d) uniformity correction applied to (a),(b), respectively.....	32
<b>Figure 2.8.</b> Evaluation of sensor linearity using sky images under thin overcast conditions .....	34
<b>Figure 2.9.</b> USI 1.2 high dynamic range (HDR) exposure sequence for (a) May 23, 2013, 3:22 pm PDT and (b) May 23, 2013, 2:48 pm for integration times of (i) 30 , (ii) 120, and (iii) 480 ms. (a-iv) and (b-iv) show the final HDR composites.....	37
<b>Figure 2.10.</b> HDR images from USI 1.8 in April and May, 2013, showing a variety of sky conditions. Color correction has been applied. Images required intensity rescaling for display purposes.....	38
<b>Figure 2.11.</b> Photon transfer curve for a USI system for (a) a 12-bit image, and (b) and HDR image .....	40
<b>Figure 2.12.</b> Stray light from the dome (top), lens and neutral density filter (middle), and whole system (bottom).....	43
<b>Figure 2.13.</b> Stray light comparison between two designs of the USI; (a) design with filter, and (b) design with modified aperture.....	44

<b>Figure 2.14.</b> Red-blue-ratio (RBR) comparison of (a) a sky image (Figure 2.10) for (b) color corrected image and (c) original image. The RBR scale is given in the colorbar.....	46
<b>Figure 3.1.</b> USI 1.8 in the instrument field at the Department of Energy, Atmospheric Radiation Measurement Program, Southern Great Plains Climate Research Facility. ....	63
<b>Figure 3.2.</b> Solar position measurements.....	66
<b>Figure 3.3.</b> Synthetic dataset point distribution.....	75
<b>Figure 3.4.</b> Root mean square calibration error distribution.....	76
<b>Figure 3.5.</b> Distributions of (a) $x$ -focal length $\alpha_x$ ; (b) $y$ -focal length $\alpha_y$ ; and (c), (d) principal point $x_0, y_0$ are shown as a function of measurement error .....	77
<b>Figure 4.1.</b> (a) raw USI 1.7 HDR image 20130503T193330 rescaled for display; (b) RBR image from original HDR image (not rescaled image); (c) clustering with 100 clusters ; (d) clusters whose RBR mean is less than 0.62. ....	85
<b>Figure 4.2.</b> Characteristics of clusters. ....	86
<b>Figure 4.3.</b> Histograms of clear sky, thin cloud, and thick cloud generated from a set of 60 manually annotated images captured by the TSI.....	87
<b>Figure 4.4.</b> (a) Average RBR (colorbar) as a function of zenith angle and sun-pixel-angle based on TSI imagery on a clear day (September 24, 2009); (b) Standard deviation (colorbar) of clear sky RBR; and (c) number of occurrence (colorbar) for each zenith angle and sun-pixel-angle pair. ....	89
<b>Figure 4.5.</b> Processing chain of a sky image on October 4, 2009 15:45:30 PST (a) to obtain the cloud decision image.....	90
<b>Figure 4.6.</b> (a) Raw USI image captured on November 19, 2011; (b) RBR of the image. ....	92
<b>Figure 4.7.</b> Clear-Sky Library (CSL) lookup table as a function of pixel-zenith angle and scattering angle (Sun-pixel angle) for the USI over an entire day.....	93
<b>Figure 4.8.</b> Illustration of the normalized cross correlation method used to compute inter-image cloud motions.....	97
<b>Figure 4.9.</b> Cloud speed measured near Boulder City, Nevada, on November 12, 2011, using a TSI.....	99
<b>Figure 5.1.</b> (a) Satellite image showing the camera and ceilometer locations, along with geographic coordinates.....	107

<b>Figure 5.2.</b> Image pairs for which a dense estimation of cloud position was performed .....	108
<b>Figure 5.3.</b> Matching score (normalized cross correlation) for scanline 400 in testcase altocumulus-01 .....	125
<b>Figure 5.4.</b> Stereography output for case <b>altocumulus-01</b> .....	132
<b>Figure 5.5.</b> Stereography output for case <b>broken-02</b> .....	133
<b>Figure 5.6.</b> Stereography output for case <b>cirrocumulus-03</b> .....	134
<b>Figure 5.7.</b> Stereography output for case <b>cirrus-04</b> .....	135
<b>Figure 5.8.</b> Stereography output for case <b>cumulus-05</b> .....	136
<b>Figure 5.9.</b> Stereography output for case <b>cumulus-06</b> .....	137
<b>Figure 5.10.</b> Stereography output for case <b>lowspare-07</b> .....	138
<b>Figure 5.11.</b> Stereography output for case <b>lowspare-08</b> .....	139
<b>Figure 5.12.</b> Stereography output for case <b>overcast-09</b> .....	140
<b>Figure 5.13.</b> Stereography output for case <b>twolayers-10</b> .....	141
<b>Figure 6.1.</b> Flow chart showing basic operations for constructing the power forecast in theCopper Mountain case study.....	144
<b>Figure 6.2.</b> Sequential cloud advections for a single forecast issue .....	146
<b>Figure 6.3.</b> Side view of ray tracing a vector from ground station $x_g$ to intersection point $x_s$ .at cloud height $H$ for ground station height $h$ and solar zenith angle $\theta_s$ .....	148
<b>Figure 6.4.</b> Ray tracing to construct a georeferenced mapping of shadows .....	150
<b>Figure 6.5.</b> Map of UCSD showing sky imager coverage, weather stations, and PV arrays.....	151
<b>Figure 6.6.</b> Time series (a) of MOCC station for October 4, 2009 showing measured GHI (blue) versus the TSI nowcast .....	152
<b>Figure 6.7.</b> (a) TSI nowcast and measured GHI (blue) at HUBB station for March 10, 2010.....	153
<b>Figure 6.8.</b> Combined TSI forecasts results for all stations .....	156
<b>Figure 6.9.</b> Normalized power distribution for November 3, 2011 at the Copper Mountain Solar 1 power plant .....	159

<b>Figure 6.10.</b> Outline of the 48 MW section of Copper Mountain, with sky-imager locations indicated .....	160
<b>Figure 6.11.</b> Forecast performance as a function of horizon for a sky-imager forecast (solid) and persistence forecast (dashed) .....	162
<b>Figure 6.12.</b> Midday 10 min forecast performance on November 12, 2011 .....	164
<b>Figure 7.1.</b> (a) Footprint of the 48MW section of the Sempra Generation Copper Mountain Solar power plant. TSI locations are indicated with circle icons. The spatial resolution of the power data is indicated by the gray level shading. (b) University of California, San Diego campus .....	169
<b>Figure 7.2.</b> Scatter plots of measured vs. forecast clear sky index.....	175
<b>Figure 7.3.</b> Scatter plots of measured clear sky index vs. clear fraction .....	176
<b>Figure 7.4.</b> Pearson (PCC) and Spearman (SCC) correlation coefficients for forecasts with (a, b) the USI on November 8, 2012 and (c, d) TSI on July 12, 2012 .....	177
<b>Figure 7.5.</b> Euclidean dispersion of (a, c) measured vs. forecast clear sky index and (b, d) clear fraction vs. measured clear sky index.....	178
<b>Figure 7.6.</b> Scatter plots of measured vs. forecast clear sky index during the one month study period .....	179
<b>Figure 7.7.</b> Scatter plots of measured clear sky index vs. clear fraction during the one month study period .....	180
<b>Figure 7.8.</b> Pearson (PCC) and Spearman (SCC) correlation coefficients during the one month study period for (a, b) the USI and (c, d) the TSI forecasts .....	181
<b>Figure A.1.</b> Map of the island of Lanai, Hawaii, with four sensor locations indicated ..	191
<b>Figure A.2.</b> Global horizontal irradiance on 20 June 2011, showing typical large, cumulus-induced fluctuations on the island of Lanai, Hawaii.....	193
<b>Figure A.3.</b> The Pareto-optimal set for maximizing energy while minimizing the 99 <sup>th</sup> percentile of ramp rates .....	196
<b>Figure A.4.</b> (a) Probability distribution (PDF) and (b) cumulative probability distribution (CDF) of ramp rates for several time steps at the La Ola site .....	197
<b>Figure A.5.</b> Pareto-optimal sets for the (a) 97 <sup>th</sup> , 98 <sup>th</sup> , 99 <sup>th</sup> , and 99.5 <sup>th</sup> percentiles and (b) 99 <sup>th</sup> , 99.5 <sup>th</sup> , 99.9 <sup>th</sup> , 99.99 <sup>th</sup> , and 99.999 <sup>th</sup> percentiles.....	200

**Figure A.6.** Monthly ramp ratios for (a) the equal allocation scenario, (b) the 8-MWh design point scenario optimum allocation, and (c) a Lanai, Hawaii, case study reflecting currently installed capacity at La Ola. ....206

**Figure A.7.** Probability distribution of ramp rates above the 99<sup>th</sup> percentile .....207



## LIST OF TABLES

<b>Table 2.1.</b> Research systems for sky atmospheric observations .....	15
<b>Table 2.2.</b> Intrinsic parameters and lens focal length selection parameters measured for 7 USI units. The principal point $u_0, v_0$ and focal length $f$ are measured for each USI.....	18
<b>Table 2.3.</b> USI Locations in the United States and deployment time ranges.....	47
<b>Table 3.1.</b> Calibration test cases. ....	65
<b>Table 3.2.</b> Camera model parameters (excluding distortion terms) determined from the five test cases of solar input data. ....	73
<b>Table 3.3.</b> Calibration error metrics for each case .....	74
<b>Table 4.1.</b> Quality Control (QC) of Sky Imager Cloud Motion Vectors.....	99
<b>Table 5.1.</b> Cloud heights measured by the ceilometer and the pair of sky imagers ....	128
<b>Table 6.1.</b> Location and status information for weather stations used to validate TSI cloud shadow maps. ....	151
<b>Table 6.2.</b> Distribution of measured versus TSI nowcast clear (CLR) and cloudy (CLD) conditions for skies with cloud fraction $0.2 \leq cf \leq 0.8$ .....	154
<b>Table 6.3.</b> Sky Imager (SI) and persistence (P) forecast error at selected time horizons of 5, 10 and 15 minutes.....	163
<b>Table 7.1.</b> Locations of the pyranometers in the UCSD network and their distance to the USI.....	170
<b>Table A.1.</b> Measurement site geographic position and designation.....	193
<b>Table A.2.</b> Site allocation weight values on the Pareto-optimality curves for minimization of the 99 <sup>th</sup> , 99.9 <sup>th</sup> and 99.99 <sup>th</sup> percentile (pct) of ramp rates, shown as a function of aggregate daily energy produced. Aggregate system ramp rates are also shown. ....	199
<b>Table A.3.</b> Site geographic allocation percentage for 7 difference scenarios.....	202
<b>Table A.4.</b> Energy production and ramp rates at the 99 <sup>th</sup> , 99.9 <sup>th</sup> , 99.99 <sup>th</sup> , and 99.999 <sup>th</sup> percentiles (pct) for 7 different weighting scenarios.....	203

## ACKNOWLEDGEMENTS

I would like to acknowledge Professor Jan Kleissl for his support during the past several years of my research. Without his help, none of this would have been possible.

In the course of preparing this dissertation, a number of publications were used, the co-authors of which I am indebted. The works and coauthors are listed below. The specific chapter for which these publications have been used can be found at the end of each chapter.

**Publication:** Intra-hour forecasting with a total sky imager at the UC San Diego solar energy testbed. *Journal of Solar Energy*, volume 85, issue 11, pp. 2881–2893, 2011. doi: 10.1016/j.solener.2011.08.025

**Co-Authors:** Chi Wai Chow, Matthew Lave, Anthony Dominguez, Jan Kleissl, Janet Shields, and Byron Washom

**Publication:** Optimizing geographic allotment of photovoltaic capacity in a distributed generation setting. *Progress in Photovoltaics: Research and Applications*, volume 21, issue 6, pp. 1276–1285, 2012. doi: 10.1002/pip.2334

**Co-Authors:** Manajit Sengupta and James Keller

**Publication:** Sky imaging systems for short-term solar forecasting, chapter 9 in: Kleissl J, editor: *Solar Energy Forecasting and Resource Assessment*. Elsevier; 2013. ISBN: 978-0-12-397177-7, doi: 10.1016/B978-0-12-397177-7.00009-7

**Co-Authors:** Mohamed Ghonima, Dung Nguyen, Ben Kurtz, Chi Wai Chow and Jan Kleissl

Publication: Comparison of Solar Power Output Forecasting Performance of the Total Sky Imager and the University of California, San Diego Sky Imager. Energy Procedia, volume 49, pp. 2340–2350, 2014. doi: 10.1016/j.egypro.2014.03.248

Co-Authors: Seyyed Mohammad Iman Gohari, Handa Yang, Ben Kurtz, Dung Nguyen, Chi Wai Chow, Mohamed Ghonima and Jan Kleissl

Publication: Development of a sky imaging system for short-term solar power forecasting. Atmospheric Measurement Techniques Discussion, volume 7, pp. 4859–4907, 2014. doi:10.5194/amtd-7-4859-2014

Co-Authors: Ben Kurtz, Elliot Dahlin, Mohamed Ghonima, Janet E Shields and Jan Kleissl

Publication: Sky camera geometric calibration using solar observations. IEEE Transactions on Pattern Analysis and Machine Intelligence, in submission.

Co-Authors: Ben Kurtz and Jan Kleissl

Publication: Dense cloud stereo photogrammetry using dynamic programming. In preparation.

Co-Authors: Jan Kleissl

## ABSTRACT OF THE DISSERTATION

Development of a short-term solar power forecasting capability using ground-based visible wavelength imagery

by

Bryan Glenn Urquhart

Doctor of Philosophy in Engineering Sciences (Mechanical Engineering)

University of California, San Diego, 2014

Professor Jan Kleissl, Chair

A very short term solar power forecasting technology which uses ground-based visible wavelength imagery is presented. A sky camera system suitable for use as a solar power forecasting tool is described. Relevant imaging considerations are discussed, including the need for high dynamic range imaging of the daytime sky and an associated stray light assessment. To photogrammetrically calibrate this sky camera system, a general camera model applicable to a fixed focal length photo objective lens

with significant radially symmetric distortion is developed, and an accurate calibration technique for a stationary, skyward pointing daytime camera using the sun's position is given. Remote sensing algorithms used in the solar forecasting process are detailed, including clear sky characterization, cloud detection, cloud velocity estimation, and cloud height estimation using stereography. A cloud stereo photogrammetry method which provides dense 3D cloud position is presented. Correspondence is automatically determined using intra-scanline dynamic programming applied to a normalized cross correlation matching metric; an ordering constraint is implicit in the approach used.

Using the described remote sensing tools and methods, a complete solar power forecasting framework is detailed. The method is based on the estimation of cloud shadow position via ray tracing, and the forecast position of the cloud shadows relative to solar collectors. A ray tracing procedure that works with a planar mapping of cloud position is used to compute shadow position. Cloud transmissivity is characterized using past observations and applied to forecast cloud positions. The application of the procedure to two case studies: the UCSD DEMROES weather station network, and a 48MW solar photovoltaic power plant is presented. A comparison of the forecasting performance using a common Total Sky Imager is compared to the UCSD Sky Imager, where it is shown that the UCSD Sky Imager performs better overall.

# 1 Introduction

The use of cameras for atmospheric observations has a long history. Many instruments were devised to capture images of the clouds for varying purposes (e.g. Whipple 1890, Strachey and Whipple 1890, Acres 1893, Elkholtm 1893, Clayden 1898). Cloud height and velocities were measured with photographs by the end of the 19<sup>th</sup> century using a pair of cameras (Strachey and Whipple 1890). The basic triangulation procedure is described in the latter work, but several advances have been made in understanding imaging geometry since that time which facilitate the photogrammetry process (e.g. Longuet-Higgins 1981, Faugeras 1993, Hartley and Zisserman 2004).

The optical hardware used for whole sky imaging was developed during the course of the 20<sup>th</sup> century. The use of refraction at an air-water interface to describe the view from within a pond has been attributed to RW Wood in his text *Physical Optics* (1905) where he described how the entire 180° field of view could be seen within a 97° cone under water. One year later, WN Bond (1906) coined the term "fish-eye" view in a paper describing vision underwater and experimentation with an apparatus made from a water-filled lard bucket with photographic film placed at the bottom. He took the first nearly 180° photograph of the sky using this apparatus. Wood is often incorrectly given the credit of coining the term fish-eye because he used it, following Bond, in the second edition of *Physical Optics* (1911) where he described a more compact water-filled camera. Photos from Bond's 1906 paper appeared five years later in *Physical Optics* (1911). In 1922, Bond used a hemispherical glass lens with a pinhole to create wide angle images and suggested their use for sky photography and presented the first published whole sky imagery results on the subject. The primary drawback to the lens is that it could only be focused on a band of the sky (constant zenith). The first design of a true modern fish-eye lens is attributed to Robin Hill (1924) where he used a large

negative meniscus front element in the lens which allowed the full sky to be in focus. A caveat was that bandpass filters had to be used to restrict the wavelength range because of blurring due to chromatic aberration at high zenith angles caused by dispersion of colors (variation in refractive index with wavelength). A doublet was introduced (Miyamoto 1964) that corrected the dispersion and thus reduced issues due to chromatic aberration within the visible wavelength range. It is with this final development that sharp full color images of the sky could be obtained.

The development of a refractive lens for capturing full sky images opened up many research fields, including canopy research (e.g. The Canopy Camera developed by Harry E. Brown, 1962). The development of digital camera systems utilizing computers and semiconductor sensors began in the 1970s and 1980s. Some work was done by forestry community for canopy research (Chazdon and Field 1987), and in parallel the Marine Physical Laboratory (MPL) at the Scripps Institute of Oceanography (SIO) was developing a system designed to image clouds (Johnson et al. 1988, 1989). An excellent history of the Whole Sky Imager (WSI) can be found in Shields et al. (2013).

The application of sky imagery for solar power forecasting, however, is a recent development. Various parameterizations of irradiance at the surface have been devised which take into account cloud fraction (e.g. Kimball 1928, Lumb 1964, Pfister et al. 2003, Kalische and Macke 2008), which is a quantity easily measured with a sky imager. Crispim et al. (2008) appear to be the first to use cloud information derived from a sky imager for solar forecasting (specifically a TSI, section 2.2.2), where the cloud fraction and global horizontal irradiance (presumably from a collocated sensor) was input into an artificial neural network and tested for six days of data. Results seem reasonable,

although the work makes it difficult to discern the full details of the experiment and performance evaluation.

The first use of a sky imager to estimate sky condition at sites up to 2km away was performed by the author and reported in Chow et al. (2011). Five minute forecasts of sky condition for the same distant sites were reported in Urquhart et al. 2011. The first use of a sky imager to generate 15-minute ahead power forecasts for a large solar photovoltaic power plant was reported at the World Renewable Energy Forum in Denver, Colorado in May 2012 (Urquhart et al. 2012) and later in Urquhart et al. 2013. The deterministic forecasting work at UCSD is continuously evolving and being improved (Gohari et al. 2014, Yang et al. 2014). Because of the importance of solar power forecasting for economical grid integration, much contemporary work on solar power forecasting with a sky imager has also been performed (e.g. Marquez et al. 2012, Chu et al. 2013, Marquez and Coimbra 2013, Fu et al. 2013). The solar forecasting procedure developed by the author and his colleagues is discussed in Chapter 6, and the related remote sensing capabilities are presented in Chapter 4. Imaging issues with the hardware uncovered while forecasting for the solar power plant (Urquhart et al. 2012, 2013) were addressed by the development of a new camera system, presented in Chapter 2. A comparison of this new instrument, the UCSD Sky Imager (the "USI") and the previously used (and very common) Total Sky Imager is presented in Chapter 7.

The ray tracing method used to determine the obscuration of the sun at a particular ground site requires accurate cloud position information. To determine cloud position remotely, both active and passive sensors can be used. Active ground-based sensing methods use electromagnetic energy at different parts of the spectrum, e.g. lidar uses near infrared and cloud radar uses radio waves (see also Kazantzidis et al. 2012 about many cloud remote sensing technologies). Cloud radar is prohibitively expensive



for solar forecasting applications. The lidar instruments, in general, are used in a vertical pointing configuration only. They could be made to scan the sky, and some interesting results could be obtained in this manner, however, this avenue was not pursued here. Instead, passive sensing of clouds was the preferred method. To determine the position of a cloud element from passive observational systems, the cloud element must be triangulated by observing it from two different positions. Using cameras to determine positional information is known as stereography, and it has a rich history in the photogrammetric and computer vision literature (Barnard and Fischler 1982, Dhond and Aggarwal 1989, Brown et al. 2003). To triangulate the position of an object using a pair of cameras, the same object must be located in each image, and the pixel position within each image must have a known pointing angle in 3-space. The former problem of locating the object in two (or more) images is known as correspondence, and the latter problem of determining the pointing angle of pixel or set of pixels is known as calibration. Both must be performed accurately if the position of the desired object is to be located accurately.

With the forecast procedure developed as presented in Chapter 6 and Urquhart et al. 2013, neither an accurate calibration nor a cloud stereographic procedure was used. As a means to make targeted improvements in the forecast, these two topics were further pursued. A method to geometrically calibrate a whole sky camera with a 180 deg field of view was developed along with a general camera model in Urquhart et al. 2014b. This geometric calibration method is presented here in Chapter 3. A method for ground-based cloud stereography to produce dense cloud position maps was developed in Urquhart et al. 2014c. A majority of previous methods to determine cloud position stereographically either used analog imagery (Kassander and Sims 1957, Orville and Kassander 1961, Bradbury and Fujita 1968, Warner et al. 1973) and thus could not be

automated, or used digital imagery, but did not provide an automated means to determine correspondence (Zehnder et al. 2007, Öktem et al. 2014). Allmen and Kegelmeyer (1996) and Nguyen and Kleissl (2014) apply similar automated stereo approaches to whole sky imagery (i.e. 180 deg field of view) to generate dense cloud height maps, where matching was performed forward and backward to improve the confidence in the correspondence. Allmen and Kegelmeyer (1996) rectified image patches prior to matching whereas Nguyen and Kleissl (2014) did not. Seiz et al. 2007 developed a sparse matching routine where features were selected automatically and matched using a least squares matching procedure that incorporated intensity and geometric constraints. Results presented for this robust method are generally very good, although the primary drawback is the computational complexity and sparsity of the matching. The method presented here provides dense depth maps with automated determination of correspondence using dynamic programming. The triangulation method applied is optimal in the projective sense (Hartley and Sturm 1997). The automated cloud photogrammetric method is described in Chapter 5.

This dissertation is organized to follow a logical flow, starting from the hardware development aspects of a camera system used for solar power forecasting (Chapter 2). The camera must certainly be geometrically calibrated (Chapter 3) so that it can provide photogrammetric measurements of the sky scene. The fundamentals of cloud remote sensing from a ground-based visible wavelength camera are discussed in Chapter 4, along with many specific improvements and modifications which were developed in the course of the solar power forecasting work performed at UCSD. The remote sensing of cloud position (Chapter 5) was isolated in its own chapter because it relies on additional mathematical foundations, and it was (importantly) performed in a much later phase of the research work described here. Using the foundation established in the proceeding

chapters, Chapter 6 provides an overview of the deterministic geometric solar power forecasting developed at UCSD. Chapter 7 gives a brief comparison of the forecasting performance when applying the forecasting algorithms in Chapter 6 UCSD Sky Imager and the Total Sky Imager. Future work for each of the topics is discussed in Chapter 8. An appendix, containing bonus material, develops a means to optimally site photovoltaic generation in a distributed generation setting. This latter work was performed while the author was visiting the National Renewable Energy Laboratory in Golden, Colorado in the first half of 2012.

## **2 Sky Camera Development**

Contributing Authors:

B. Urquhart, B. Kurtz, E. Dahlin, M. Ghonima, J.E. Shields and J. Kleissl

### **2.1 Summary**

To facilitate the development of solar power forecasting algorithms based on ground-based visible wavelength remote sensing, we have developed a high dynamic range (HDR) camera system capable of providing hemispherical sky imagery from the circumsolar region to the horizon at a high spatial, temporal, and radiometric resolution. The University of California, San Diego Sky Imager (USI) captures multispectral, 16-bit, HDR images as fast as every 1.3 seconds. This article discusses the system design and operation in detail, provides a characterization of the system dark response and photoresponse linearity, and presents a method to evaluate noise in high dynamic range imagery. The system is shown to have radiometrically linear response to within 5% in a designated operating region of the sensor. Noise for HDR imagery is shown to be very close to the fundamental shot noise limit. The complication of directly imaging the sun and the impact on solar power forecasting is also discussed. The USI has performed reliably in a hot, dry environment, a tropical coastal location, several temperate coastal locations, and in the great plains of the United States.

### **2.2 Introduction**

Solar power output of an individual generator, or even a fleet of generators, will have some level of variability due to the nature of the input source of light from the sun. The source of short-term variability of irradiance at the earth's surface is clouds and atmospheric particulates, which are generally not controllable. To reliably integrate

increasing amounts solar power into the electric grid, forecasting, storage, additional transmission, and ancillary power generation services will constitute a portfolio of solutions to counteract variability.

From a planning and operations perspective, grid operators, require consumption and generation estimates from years to minutes ahead. On the scale of days to minutes, solar power output forecasts are provided by some combination of weather models and remote sensing of clouds, along with stochastic learning methods. Forecasting of solar radiation in the 0-30 minute ahead time frame poses unique challenges. High resolution models reported for both satellite and numerical weather prediction can issue forecasts that have 5 minute time steps for a one kilometer grid (Mathiesen et al. 2011, Perez et al. 2013), while the best operational models often have coarser resolutions in both space and time (Dupree et al. 2009, Rogers et al. 2012, Mathiesen et al. 2013). However, in numerical models, timing and/or positioning errors of clouds are inevitable, and for satellites, infrequent image capture and parallax effects can result in inaccurate georeferencing of clouds. These errors make it difficult to achieve an accurate, high resolution short term solar forecast. This motivates a need for other forecasting tools and observational methods.

One short-term forecasting technology that has emerged recently has been based on remote sensing of clouds from ground-based imaging systems (Chow et al. 2011, Marquez and Coimbra 2013, Urquhart et al. 2013, Yang et al. 2014, Fu and Cheng 2013). Urquhart et al. (2013) applied the forecasting method to 48MW of photovoltaic generation. One of the key conclusions from that work was that the Total Sky Imager (TSI), while providing the ability to monitor sky conditions, had shortcomings that limit its effectiveness for solar power forecasting.

This work describes the development of sky imaging hardware for short-term solar power forecasting at UCSD. The UCSD Sky Imager (USI) provides unique capabilities needed for forecasting research and applications. The goal of this article is to provide a details on the USI system, and report on its imaging performance so that other workers in this field may make more informed purchasing and design decisions. The remainder of section 2.2 discusses hardware requirements for sky imaging, and reviews relevant sky imaging hardware presented in the literature. Specifics of the USI hardware development and system operations are covered in sections 2.3 and 2.4, respectively. The imaging performance of the USI is characterized in section 2.5, and section 2.6 presents deployments of the USI to date.

### **2.2.1 Imaging for Solar Power Forecasting Applications**

Sky imagers were historically built for recording meteorological conditions such as sky cover. For this purpose it is not critical to image the area directly around the sun, so many systems have sun occluding devices to prevent direct sunlight from entering the optics. When the sun is unobstructed, more than 90% of the photons entering the optics can come from the direct solar beam. For most camera systems, the handful of pixels encompassing the sun saturate and thus direct-beam signal intensity is only known to exceed the saturation threshold. Immediately outside the direct beam is a region of intense forward scattering. Aerosols and dust scatter the direct beam predominantly in the forward direction, increasing the size of region around the sun that will potentially saturate in a sky image. Cloud droplets and ice crystals, when present, also predominantly scatter in the forward direction and, depending on the scattered intensity reaching the camera, can further extend the size of the region that will saturate. Obtaining on scale image information about the region around the sun requires

appropriate imaging hardware and methods, especially when the region around the sun has a very high intensity. Further, this high intensity region can cause image quality degradation through internal reflections, diffraction caused by the aperture, sensor saturation, smear and blooming, and, potentially, sensor damage (see section 2.5.7 for further discussion relating to the USI). This significance of each of these potential issues is imaging system dependent.

The use of occluding devices eliminates many of these potential issues, which is why they are often adopted. Blocking the sun and the surrounding area, however, eliminates important sky condition information needed to provide reliable forecasts in the first few minutes (< 5 min) of the forecast period. If the occluding device obstructed a minimal amount of the image along with precision-positioning mechanisms, it could then be used without adversely affecting immediate-term forecast accuracy. For cost and reliability reasons, these requirements are difficult to achieve in practice. The TSI, for example, has a shadowband that occludes 14% of the sky hemisphere, always in the region near the sun. Even for a 1.3 km<sup>2</sup> solar power plant, the shadowband on the TSI has been demonstrated to obscure sky condition information for over half of the plant (Urquhart et al. 2013). With proper design, and appropriate image capture and correction algorithms, sky imaging systems can acquire atmospheric information from an appreciable amount of the region around the sun (e.g. Stumpf et al. 2004). To this end, the high dynamic range imaging method described in section 2.5.5 provides a simple and robust approach.

The spectral content of the sky scene provides important information for the remote sensing of clouds and water vapor. Most camera systems capture visible wavelength imagery that spans between 350 and 800 nm. This allows the measurement of shortwave solar radiation that is scattered by the clouds, atmospheric gases and

aerosols. Silicon-based image sensors used in visible wavelength cameras are also sensitive up to 1.1  $\mu\text{m}$  in the near infrared. The sixteen bit (or higher) versions of these, with a set of selectable bandpass and neutral density filters, can be used for enhanced day and nighttime cloud detection (Shields et al. 2013). Long wave infrared (LWIR) imaging in the 8-12  $\mu\text{m}$  range measures cloud brightness temperatures which can be used to segment different cloud layers, estimate cloud heights, and potentially determine optical depth. LWIR imaging hardware costs significantly more than common visible wavelength imaging hardware, and it may not be practical for widespread deployment.

The image formation process in a sky imaging camera redirects radiant energy from the sky hemisphere onto the two dimensional image plane. Geometric and radiometric calibrations turn the brightness information at a given pixel position into a measurement of sky radiance at a given look angle. Geometric calibration relates a pixel position on the sensor array to a set of angles (azimuth and zenith) in a defined world coordinate system. This is a necessary step to accurately geolocate clouds. Solar forecasting methods such as Chow et al. (2011) and Yang et al. (2014) require geometric calibration of the imager because cloud positions are explicitly computed. Time of arrival methods such as Marquez and Coimbra (2013) or Wood-Bradley et al., (2012) do not require metric calibration because only a forecast of when a cloud will occlude the sun for the location of the imager is needed. A treatment of geometric camera calibration is beyond the scope of this article, but the interested reader is referred to Faugeras (1993) or Hartley and Zisserman (2004) for an introduction to the calibration process, and Yang et al. (2014) for more on the geometric calibration of the USI. Radiometric calibration makes it possible to determine the radiance of the scattered light coming from a portion of the sky (Shields et al. 1998b, Feister and Shields 2005, Roman et al. 2012), and can be used as input to retrieval algorithms for a number of



optical properties of atmospheric aerosols that impact solar energy generation (Nakajima et al. 1996).

### **2.2.2 Existing sky imaging hardware**

There are three fields where a majority of sky imaging work has been performed: atmospheric sciences, forestry and ecology, and astronomy. Camera developers in astronomy are typically concerned with having a high sensitivity and low noise so that a high percentage of incoming photons from stars, asteroids and other faint objects are converted into charge carriers on the image sensor. The sensors used are often full frame charge-coupled devices (CCDs) because of the high quantum efficiency and fill factor, but require mechanical shuttering which limits the frame rate of the system. One system similar to the USI from the astronomy field is the All Sky Infrared Visible Analyzer (ASIVA, Klebe et al. 2014, Sebag et al. 2008) with a dual camera system that captures both visible and LWIR images. It is one of the few LWIR dioptric (refraction-based) whole-sky designs (catadioptric [reflection and refraction] designs similar to the TSI are more common). It uses a  $640 \times 512$  uncooled microbolometer array sensitive in the 8–13  $\mu\text{m}$  range with a germanium fisheye lens. The system has an 8-slot filter wheel allowing for multiband LWIR measurements. The ASIVA also has a high-resolution visible camera with an 8-slot filter wheel (specific camera model has varied by ASIVA unit).

The area of forestry has extensively used hemispherical photography (Brown 1962, Anderson 1964). The high-dynamic-range all-sky-imaging system (HDR-ASIS) is a CMOS-based camera that leverages multiple exposures to create a high-dynamic-range (HDR) composite sky image for ecosystem and canopy research (Dye 2012).

Researchers in atmospheric science have very actively developed their own instruments over the years. In fact, Robin Hill (1924) mentions cloud photography as a

motivation in developing the first the first true fisheye lens design. Digital sky photography began in the 1980s with the development of personal computers, and one of the leading groups developing imaging systems for atmospheric observation was the Atmospheric Optics Group at the Scripps Institute of Oceanography's (SIO's) Marine Physical Laboratory (Johnson et al. 1988, 1989). Their well known Whole Sky Imager (WSI) is still to this day one of the highest quality, if not the highest quality, sky imaging systems ever developed (Shields et al. 2013). It was developed primarily for U.S. military applications in the 1980s and early 1990s. More recent designs of the system had a 512 × 512-pixel temperature-controlled, 16-bit low-noise monochrome CCD camera. It used a Nikon Nikkor 8 f/2.8 (8mm) fisheye lens (equidistant projection, section 2.3.1) and two filter wheels holding neutral density and spectral filters at multiple wavelengths. The image plane was the surface of a tapered fiber-optic bundle that interfaced directly to the CCD. Multiple corrections were made to the instrument to improve measurement quality: dark field correction; flat field correction (among other things, this corrected imaging issues caused by fiber optic imperfections); exposure corrections; linearity corrections; rolloff corrections; geometric calibration; and in some cases absolute radiometric calibration. By adjusting the neutral density and spectral filter selections, and/or the exposure time, the system achieved a wide dynamic range and could capture both daytime and nighttime imagery with high accuracy. The cloud detection algorithms developed over several decades were sophisticated, with accurate detection of haze, thin cloud, and opaque cloud (Shields et al. 1993a, 1993b, 1998a; Feister and Shields 2005).

The most widely used outdoor hemispheric camera system, first described by Long and DeLuisi (1998) as the Hemispheric Sky Imager (HSI), is the Total Sky Imager. It has been commercially available by Yankee Environmental Systems (YES) for over a

decade, and has a proven track record of reliably recording sky conditions. The catadioptric optical design uses a spherical mirror to reflect the sky hemisphere into a downward-pointing camera. The system has relatively low spatial and radiometric resolution (640×480 pixels, 48 dB max.), and there is little control of the camera capture settings. An antireflective black rubber strip ("shadowband") affixed to the mirror prevents direct sunlight from reflecting into the camera optics which improves image quality and avoids damage to the sensor. The shadowband covers approximately 0.70 steradians of the hemisphere, which is about 14% of the image region used for forecasting ( $< 80^\circ$  zenith angle). A comparison of the solar forecasting performance between the TSI and USI was performed by Gohari et al. (2014).

Beyond the WSI and TSI, a number of other imaging systems have been developed for atmospheric studies. A description of several of these can be found in Urquhart et al. (2013) and in Table 1. Outside of systems developed by research groups, there are alternatives to the TSI. The SONA (Sistema Automático de Observación de Nubes, Gonzales et al. 2012) uses a 1/3", 640 ×480 CCD, has integrated coolers, heaters and temperature sensors and is ruggedized for outdoor deployment. It has an integrated shadowband with azimuth control that shades part of the lens, but not the full optical system (i.e. it does not shade the entire dome). The Eko sky camera, built by Schreder, is reported to have 2M pixels, and like the SONA and TSI, has cloud detection software and a user interface. The Santa Barbara Instrument Group (SBIG) sells the Allsky-340C camera system based on a Truesense KAI-0340, 640 × 480 CCD with a specified dynamic range (defined section 2.3.2) of up 69dB, and uses a 1.4mm focal length Fujinon FE185C046HA-1 lens. The SBIG camera was used for solar forecasting research by Fu and Cheng (2013). The list of systems noted here is far from

comprehensive, and with the potential of sky imagery for solar energy applications, new systems are continuously being developed.

**Table 2.1.** Research systems for sky atmospheric observations

System	Camera	Sensor	Resolution	Lens	Reference
WSI	Photometrics S300	CCD	512 × 512	Nikon, equidistant	Shields et al. 2013
WSC	-	CCD, 1/3"	752 × 582	1.6mm	Long et al. 2006
ASI	QImaging RETIGA 1300C	Sony ICX085AK 1300 × 1030 CCD, 2/3", 12-bit		Fujinon FE185C057HA	Cazorla et al. 2008
IFM-GEOMAR	-	10-bit	3,648 × 2,736	-	Kalisch and Macke 2008
ASI	-	1/3" CCD	2,272 × 1,704	equidistant	Huo and Lu 2012

## 2.3 Hardware design and selection methods

### 2.3.1 Optical design

The University of California, San Diego has developed its own sky imager (the USI, Fig. 2.1) to address the instrument needs for short term forecasting. The USI uses a Sigma 4.5 mm focal length fisheye lens which allows the entire image circle to fit on the sensor. This can easily be verified from the focal length, the lens projection, and sensor size. A conventional camera lens has the rectilinear projection function


$$r_s = f \tan(\theta),$$

where  $f$  is the focal length,  $\theta$  is the angle from the optical axis, and  $r_s$  is the distance from the principal point in the image plane. It is evident that this pinhole camera model cannot image points at  $90^\circ$  from the optical axis with a sensor of finite size. In order to form the image of points that are  $90^\circ$  from the optical axis within a finite image plane, distortion is required, and the type of distortion can be selected by the optical designer. The two most common projections used in fisheye lenses are the equidistant and equisolid angle projections,  $r_{ed}$  and  $r_{es}$ , respectively:

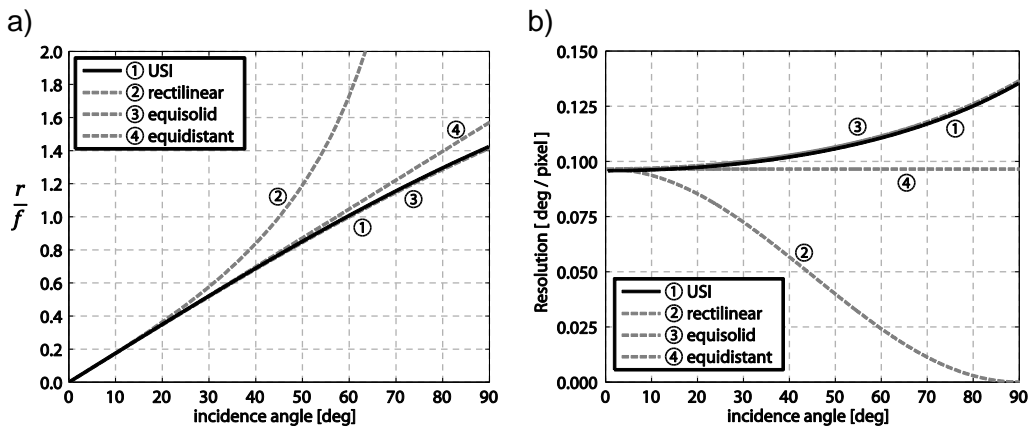
$$r_{ed} = f\theta,$$

$$r_{es} = 2f \sin\left(\frac{\theta}{2}\right).$$

Each of these projection models provides different performance characteristics. The equidistant model provides a linear relation between incidence angle and distance from the principle point, and it has slightly less distortion at large angles from the optical axis than the equisolid angle projection. The equisolid angle projection is so-named because the solid angle subtended by a unit area on the image plane is constant, regardless of incidence angle (Miyamoto 1964). A comparison of the different lens projections is shown in Fig. 2.2a, along with that measured for the USI system. The angular resolution per pixel is shown in Fig 2.2b. Figure 2.2b assumes the sensor is 15.15mm across containing 2048 pixels, and uses the specifications for USI 1.2 in Table 2.2. Even though the angular resolution at the horizon is coarser for an equisolid versus an equidistant projection at the same focal length, the former was selected for the USI because at large zenith angles, the horizontal configuration of clouds is difficult to determine because of self occlusion and perspective effects. Using more of the sensor area for the sky region overhead and near the sun (during midday) was preferred because these sky areas contain the clouds causing the current and near future solar power generation impacts when power output is highest.



**Figure 2.1.** (a) USI enclosure, (b) top view of enclosure with lid removed, (c) components of USI taken out of the enclosure.



**Figure 2.2.** (a) Perspective (rectilinear), equidistant, and equisolid angle projection distances as a function of incidence angle, along with the projection for USI 1.2 determined from geometric calibration. The projection distance is normalized by the focal length. (b) Angular resolution of projections in (a).

For a given sensor size, the selected projection places a limit on the maximum allowable focal length of a lens while still being able to capture the complete sky dome (or conversely, the minimum sensor size given a focal length). The maximum allowable focal length for the equidistant projection  $f_{ed,max}$  is  $2 r_{min}/\pi$  and for the equisolid angle projection  $f_{es,max}$  is  $r_{min}/\sqrt{2}$ , where  $r_{min}$  is the shortest distance from the principle point to the edge of the sensor. For the USI, with a sensor size of 15.15 mm,  $r_{min}$  is 7.575 mm (assuming the principle point is in the center of the image sensor), and a focal length of

less than 5.36 mm for the equisolid angle projection is required. Because the principal point will in general vary depending on machining and assembly tolerances of the components used, the value of  $r_{min}$  will vary. Table 2 shows the principal point location,  $r_{min}$ , and  $f_{max}$ , for several USI systems obtained from a nonlinear geometric calibration of extrinsic and intrinsic parameters that minimized the squared pixel error between actual sun position measurements and modeled sun position. The NREL solar position algorithm (Reda and Andreas, 2004) was used for modeled sun position input. The principal point shows significant variation because mounting location of the lens fluctuates by as much as 0.31 mm. As a result, the radial distance to the edge of the detector fluctuates, and thus the maximum allowable focal length.

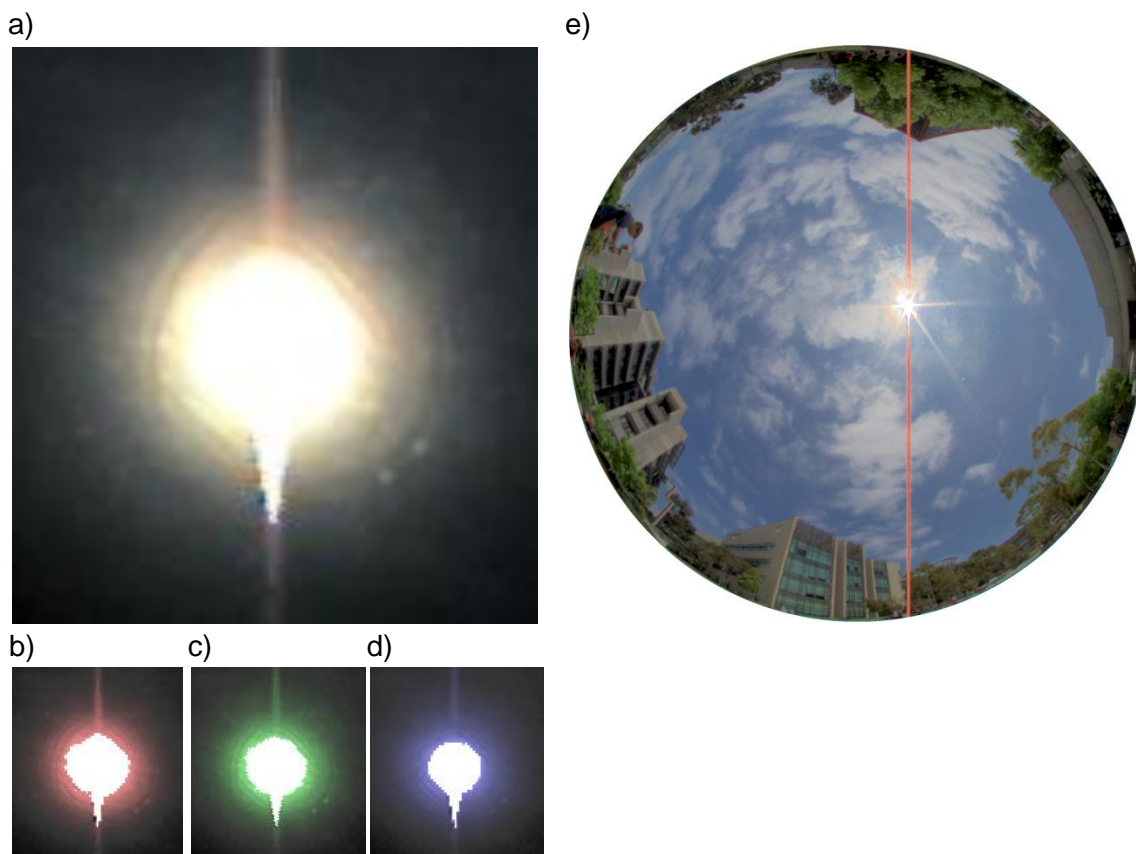
**Table 2.2.** Intrinsic parameters and lens focal length selection parameters measured for 7 USI units. The principal point  $(u_o, v_o)$  and focal length  $f$  are measured for each USI. The minimum distance to the sensor edge  $r_{min}$  from  $(u_o, v_o)$  yields the maximum allowable focal lengths  $f_{ed,max}$  and  $f_{es,max}$  for the equidistant and equisolid angle projections, respectively. Units are in mm, except for  $u_o$  and  $v_o$  which are given in pixels.

USI No.	$u_o$	$v_o$	$f$	$r_{min}$	$f_{ed,max}$	$f_{es,max}$
1.1	1032	965	4.437	7.139	4.545	5.048
1.2	1040	970	4.386	7.176	4.568	5.074
1.5	1033	963	4.429	7.124	4.535	5.037
1.6	1028	991	4.377	7.331	4.667	5.184
1.8	1023	1043	4.448	7.434	4.733	5.257
1.9	1045	976	4.474	7.220	4.596	5.105
mean	1033.5	984.7	4.425	7.237	4.607	5.118
std	7.3	27.7	0.034	0.111	0.071	0.079

Proper selection of the aperture diameter is important to ensure an appropriate flux of radiant energy impinges on the sensor plane. If the flux is high, very short exposure times are required to obtain quality sky images. Because there is no mechanical shutter in the USI, the sensor is always exposed and limiting the incoming radiant flux is a way to extend sensor life. If the aperture diameter is small, exposure

time must be increased, and motion blur of the clouds is possible. The Sigma lens comes with an iris diaphragm which was not used to avoid diffraction caused by the iris blades (e.g. Fig. 2.3e). To reduce the amount the incoming radiant flux without the iris diaphragm, two methods were tested: 1) a rear neutral density (ND) gelatin filter with a transmissivity of 0.1%, and 2) a fixed circular aperture for which several diameters were tested. Stray light and spectral effects of each approach are discussed in section 2.5.7. Undesirable diffraction patterns were observed on the USI for circular apertures of diameter 300  $\mu\text{m}$ , 700  $\mu\text{m}$ , and 1,000  $\mu\text{m}$  (Fig. 2.3). Because diffraction caused by a circular aperture generates a known Airy disk pattern, it is possible to partially correct the image with deconvolution processing, however this was not done in this work. To minimize the incoming flux while also minimizing diffraction, an aperture of 1,250  $\mu\text{m}$  was selected. In comparison, the aperture diameter with the ND filter is 9,520  $\mu\text{m}$ . This large diameter noticeably reduces the depth of focus of the camera compared to the 1,250  $\mu\text{m}$  aperture (depth of field is unaffected because a fisheye lens is used). The radiant flux is higher using an aperture of 1,250  $\mu\text{m}$  compared to the ND filter configuration by a factor of 18. This allows shorter exposures with less motion blur caused by longer integration times, but may also lead to increased sensor degradation in the long term due to the increased radiation on the sensor.





**Figure 2.3.** (a) Diffraction pattern measured with a 1000  $\mu\text{m}$  aperture on USI 1.8, with color red, green, and blue color components shown in (b), (c), and (d), respectively. (e) Diffraction of the hexagonal iris blades in the stock lens.

In order to develop a ruggedized system, it is necessary to protect the lens and properly seal the enclosure from the environment. For the lens to have full 180° access to the sky with this requirement, a 1/16<sup>th</sup> in. thick, neutral density acrylic dome was used on the USI. The dome has a UV hard-coat applied to minimize transmission of high energy solar radiation which helps reduce component degradation. Amorphous silicate glass has superior transmissivity and scratch resistance than acrylic, but is more difficult to machine and handle, and designing proper sealing for a glass dome is more complicated (and thus more expensive). Polycarbonate, while having similar transparency and machining characteristics to acrylic, becomes opaque due to oxidation, making it a poor choice as a dome material. The use of a neutral density

acrylic dome with a higher neutral density and an anti-reflective coating on the inner surface is being considered to improve image quality further.

### 2.3.2 Camera and image sensor

The USI uses an Allied Vision GE-2040C camera which contains a 15.15 × 15.15 mm, 2048 × 2048 pixel Truesense KAI-04022 interline transfer CCD sensor. The camera is connected to the computer with a gigabit ethernet interface, and customized control is achieved by using the PvAPI for Linux provided by Allied Vision. For solar forecasting research, we have found that the ability to adjust exposure integration times, frame rates, regions-of-interest, and other parameters is a necessary capability that systems such as the TSI do not have.

The USI imaging system was designed to generate images suitable for cloud detection and motion processing. Cloud detection requires spectral measurements, and thus a spectral filtering method must be employed in some capacity. Coupled with a high quality sensor, camera, and lens, a mechanical shutter and color filter wheel can provide very high quality still spectral measurements. These moving components, however, complicate system design and HDR capture, and limit frame rates, therefore no mechanical shutter or color filter wheel were used. Spectral measurements were instead obtained by using a Bayer color filter array (CFA, Bayer 1975).

The intensity range of the sky necessitates a sensor with a large dynamic range. Large dynamic range and global electronic shuttering is available from interline transfer CCDs, which is why this technology was selected for the USI. Dynamic range  $DR$  is defined by the ratio of maximum measurable signal to the noise floor:

$$DR = 20 \log_{10} \frac{c_{sat}}{c_{rd}},$$

where  $c_{sat}$  is the count value at saturation, which is 4095 for a single USI exposure, and  $c_{rd}$  is the read noise. Read noise is introduced by the camera readout electronics, including output amplifiers and analog-to-digital converters. For a single USI exposure, the dynamic range was measured to be 61 dB over the entire sensor. The sensor manufacturer specifies the dynamic range to be 72 dB. This indicates the readout electronics in the GE-2040C and the operating temperature of the camera increased the noise floor, which lowers the dynamic range. The large dynamic range for the KAI-04022 is achieved because it has large  $7.4\mu\text{m}$  pixels which have a charge capacity (also called full-well depth) of 23,600 e<sup>-</sup>.


The use of an interline transfer CCD is not without tradeoffs. Smear is very apparent in images with direct sun exposure. Smear has two sources: 1) stray light entering the VCCD (vertical transfer CCD) during readout; 2) charge generation occurring deeper in the silicon photodiode layer that diffuses to any of the charge collection or transfer electronics. The VCCD is the interline column near the exposed photodiode column, and is where the vertical readout step is performed. Longer wavelengths penetrate further into the silicon before being absorbed and can generate a hole-electron pair in undesirable locations. This is why the smear is noticeably worse in the red channel of Fig. 2.3b. Blooming, which is apparent as a saturated border of bright objects, is another problem for CCDs, and is noticeable in USI imagery near the sun. It is not significant however, because each KAI-04022 pixel has a vertical overflow drain to prevent large amounts of charge from diffusing to nearby collection sites.

### **2.3.3 Enclosure and balance of system design**

For solar forecasting, tough environmental conditions such as hot and dusty deserts will be encountered. The USI is designed to survive 60°C ambient air

temperature and direct sunlight conditions. It has a light colored exterior to reduce shortwave absorption and has two 80W thermoelectric coolers with a NEMA 4X rating. To monitor the system's environmental health, a suite of temperature and relative humidity sensors was added to measure camera, power supply, internal and external ambient, and dome conditions. The internal enclosure walls are all insulated to reduce thermal conductivity of the enclosure, which with the use of active thermal control, keeps it cooler on hot days and warmer on cold days. Internal water condensation was initially found to be an issue. Improved system sealing and thorough water testing was found to be necessary. Three 20 W resistive heating strips were installed on the base of the dome to reduce condensation on the exterior dome surface.

The USI camera is controlled by a 1.8Ghz dual core (Atom D525) embedded computer running Linux Ubuntu 12.04. The images can be stored locally on a set of internal and USB hard drives, or it can be transferred across a network connection. Using an embedded computer gives the system flexibility for customizing the configuration on a per deployment basis, and the capture software can easily be reconfigured, reprogrammed, or debugged remotely. A labeled CAD model of the USI is shown in Fig. 2.4.



**Figure 2.4.** Component layout of UCSD Sky Imager camera system.

## 2.4 System operation

### 2.4.1 Image capture and storage

Images are received from the camera as uncompressed single-channel 12-bit images with per-pixel color determined by the CFA. After three exposures are composited in the HDR process (section 2.5.5), the combined image is still a single channel, but with 16-bits per pixel. Images are compressed and stored in a lossless 16-bit PNG format as a single channel image. A single pixel contains information about only *one* color of red, green or blue light. To produce a full color image from the pixel array suitable for processing, linear demosaicing is applied prior to use. Current image sizes are around 3 MB per image, or between 3 and 6 GB/day depending on the time of year.

The maximum frame rate of the USI system in *single exposure* mode is 15 fps, which is relatively low. Future dynamic computer vision approaches to solar forecasting (e.g. optical flow) may require higher frame rates, and for these future methods, the

camera used on the USI may not be suitable. In *HDR mode*, which is the standard USI operational mode, three images are captured sequentially in 160ms, which is a frame rate of 18.8 fps (or HDR frame rate of 6.3 fps). This increase in frame rate is possible because a smaller 1748 × 1748 region of interest, extracted from the center of the 2048 × 2048 pixel array, is transferred off the camera. After subsequent HDR compositing and PNG image compression, the effective frame rate drops to 0.77 fps (i.e. 1.3s per HDR image).

#### **2.4.2 System monitoring and control**

The raw images generated by the camera are inconvenient for qualitative inspection on a user's screen because they are not in color (raw Bayer format), the file sizes are relatively large so loading is slow, and a majority of the sky resides within the lower end of the 16-bit dynamic range which means the image appears very dark except for the sun. Preview images are therefore generated, which are full color, but lower resolution, compressed, and tonemapped to 8 bits per color channel. These previews are small enough to be uploaded to the operator from all sites — including remote ones using a cellular modem — and serve as a system heartbeat so that image availability and quality can be inspected at a glance.

In addition, the data acquisition system reports temperature and humidity every 30 seconds. The internal temperature and dome temperature are used to control the heaters and coolers in the USI to ensure that the critical electronics are always within their operating temperature bounds, and to avoid conditions that might lead to condensation. A live plot of temperature and humidity is uploaded to the operator. An important feature of the microprocessor controlled data acquisition system is its ability to automatically power-cycle the USI if it fails to respond. This has proven to be a valuable

backup, particularly on remote systems that are hard to access and crash more often than the others due to bugs in the cellular modem driver.

## **2.5 Imaging performance characterization**

### **2.5.1 Noise sources and photoresponse**

Each pixel in a camera is an independent radiometric sensor, and has small response variations from its neighbors due to small manufacturing differences. After charge is collected on a pixel, it is converted to a voltage and then to a digital value, and at each step in the process noise is introduced. Common sources of noise include dark current generated by the semiconductor in the bulk and at the surface, reset noise from charge to voltage conversion (which is typically minimized by correlated double sampling), read noise from the camera's readout electronics, and photoresponse nonuniformity (PRNU) arising from small manufacturing differences of each pixel. Because there is a consistent spatial variation of many of these noise sources, it often forms a pattern called fixed pattern noise (FPN). Shot noise, arising from the quantum nature of the photons generating the signal, occurs in all imaging systems and acts as a lower bound to measurement uncertainty. It adds a random element to each image that is Poissonian in nature and it can only be reduced by averaging frames, which is not feasible for fast moving clouds or when high frame rates are desired.

Each pixel's response can be characterized and corrected so that under the same illumination, the corrected output is the same when averaged over several frames. A comparison of the average of several frames is required because shot noise will always be present in an individual frame. A polynomial can be used to model a pixel's response to light:

$$c_{ij}(I, t) = \sum_{n=0}^N \hat{a}_{ij,n} t^n + \sum_{m=1}^M d_{ij,m} (It)^m, \quad 2.1$$

where  $c_{ij}(I, t)$  is the camera measurement in counts at the  $i, j$  pixel location,  $I$  is the irradiance incident on the pixel,  $t$  is the integration time of the exposure, and  $\hat{a}_{ij,n}$  are coefficients that characterize the individual pixels' dark response, and  $d_{ij,m}$  characterize the pixels' photoresponse. Sensor noise and response characteristics are temperature dependent, so coefficients  $\hat{a}_{ij,n}$ , and  $d_{ij,m}$  will also vary with temperature. Here it has been assumed that dark response and photoresponse can be separated.

To determine the coefficients in Eq. 2.1, the irradiance  $I$  on the sensor plane must be known, which when using a lens implies the scene radiance must be known over the entire field of view. This can be achieved with a calibrated flat-field source. Many of the components of solar forecasting algorithms (e.g. Chow et al. 2011, Yang et al. 2014) have a training step where either relative brightness or brightness ratios are used to determine thresholds, or texture information is used and therefore calibrated radiance is not needed. Instead, these algorithms require spatially consistent measurements (i.e. consistent between pixels), for which a simpler radiometric uniformity correction (section 2.5.3) can be used. This has the advantage that it can also be employed in the field after the instrument has been deployed. The camera output signal  $s_{ij}$  after radiometric uniformity correction can be written as:

$$s_{ij}(I, t) = \sum_{m=0}^M b_{ij,m} \left( c_{ij}(I, t) - a_{ij}(t) \right)^m, \quad 2.2$$

where  $a_{ij}(t)$  provides the dark-field correction, and coefficients  $b_{ij,m}$  provide the flat-field (i.e. uniform illumination) correction.

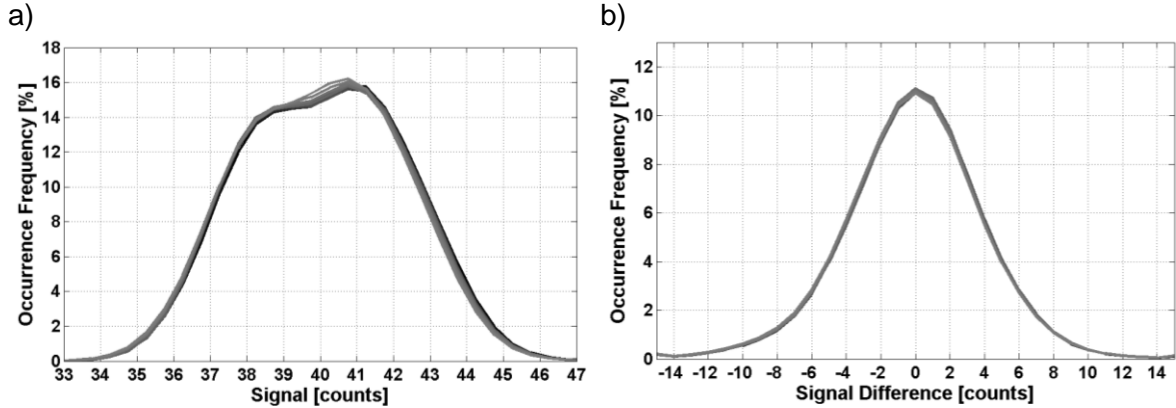


The parameters of the radiometric uniformity correction  $a_{ij}$  and  $b_{ij,m}$  have temperature dependencies that are not treated in the formulation of Eq. 2.2 or developments to follow. Sky imaging systems expecting large changes in sensor and camera temperature should perform the testing described in sections Dark response 2.5.2 to 2.5.4 at different temperatures to better understand the impacts. For the USI, the dark current of KAI-04022 roughly doubles for every 9°C increase in temperature in the system operating range. The USI camera temperature, measured with an LM335 thermal probe attached to the camera body, has been observed to change by over 20°C between day and night.

### 2.5.2 Dark response

The dark response of the sensor was measured by recording images in complete dark (unlit room, USI enclosure closed, lens cap on and covered with a thick, opaque cloth) at several integration times. Raw 12-bit images were taken at 25 different integration times ranging from 1 ms to 2 sec, and the sequence was repeated ten times for a total of 250 images. If the thermally generated dark current is low in comparison with the bias (defined below), there should be little increase in measured dark response signal as a function of integration time, i.e.  $a_{ij}(t)$  should *not* vary with time. The low dark current of the USI is illustrated in Fig. 2.5a as a set of histograms showing the occurrence frequency of each measured dark response count value. The average of ten frames at each integration time is used to reduce random noise present in a single measurement. Histograms of the difference between a 1ms average image frame (averaged from nine, 1ms exposures) and a single frame at each exposure time are shown in Fig 2.5b. Both sets of histograms show no strong change as a function of exposure time which confirms the thermally-generated dark current for the USI is low.

The A/D converters that convert voltage of each pixel to digital counts are calibrated to provide on-scale measurements throughout the range of the sensor. This sets the lower dark bound (or bias) to always be above zero, which for the USI camera this centered at approximately 40 counts (or ~1% of full scale, see dark bias distribution Fig. 2.5a).




**Figure 2.5.** (a) Occurrence frequency of signal measured in a dark room for 25 different integration times, ranging from 1ms (black) to 2s (lightest gray). Ten exposures at each integration time were averaged to construct each histogram. (b) Occurrence frequency of the signal in a single frame with an average 1ms frame subtracted. Individual labels for each integration time were not added because curves are not discernable.

The temporal component of the dark response for the exposure times used on the USI (< 1s) is small, but there is still a spatial component of the dark response called fixed pattern noise (FPN). The FPN is shown in Fig. 2.6a. There is relatively little variation within each column. Two distinct image halves are noticeable, an artifact caused by the use of two A/D converters, each serving half the sensor. Columns near the center of each half have lower readouts than columns near the edges. The dark FPN can be removed by subtracting the measured dark response to obtain the dark field corrected signal  $s_{ij}^d(I, t)$

$$s_{ij}^d(I, t) = c_{ij}(I, t) - a_{ij}(t), \quad 2.3$$

which is the same as the term in parenthesis in Eq. 2.2. The dark image term  $a_{ij}(t)$  is obtained by averaging several frames at integration time  $t$ . An image appears much

more uniform after dark correction (Fig 2.6b) which indicates the FPN has been eliminated. For over 99.9% of pixels,  $a_{ij}(t)$  does not show significant variation with time, however a small number of "hot" pixels have higher than average dark current and/or a nonlinear temporal dark response, and thus the time dependence of  $a_{ij}$  is retained.



**Figure 2.6.** (a) An example dark frame for a 100ms exposure and (b) the corrected dark frame. Typical pixel values in (a) range from 32 to 47 with a mean around 40 counts (of  $2^{12}$ ).

### 2.5.3 Sensor photoresponse uniformity correction


Photoresponse nonuniformity is caused by differing gains on each photodetector in the focal plane array; i.e.  $b_{ij,m}$  in Eq. 2.2 differs slightly for each pixel. The most direct approach to PRNU correction uses flat-field measurements (uniform lighting over the entire field of view) in order to adjust each pixel so that its response is uniform under uniform illumination. An alternative method is to use an illumination source that produces a smooth image without large brightness gradients. The resulting image can then be fit with a surface, and deviations of a given pixel from this surface can be considered the non-uniformity of that pixel. At each integration time, 10 exposures are used to obtain an average of the dark corrected signal  $s_{ij}^d(I, t)$  so that the effects of shot noise are reduced (the 10 frame average denoted by  $\overline{s_{ij}^d(I, t)}$ ). The same integration times used for the

characterizing the dark response in section 2.5.2 are used. At each integration time, a 5th order surface (denoted  $\langle s_{ij}^d(I, t) \rangle$ ) is then fit to the average dark corrected signal  $\overline{s_{ij}^d(I, t)}$  as a function of pixel location ( $i$  and  $j$ ). The resulting set of surfaces  $\langle s_{ij}^d(I, t) \rangle$  is used to determine the coefficients  $b_{ij,m}$  as a function of exposure time  $t$ :

$$\langle s_{ij}^d(I, t) \rangle = \sum_{m=0}^M b_{ij,m} \left( \overline{s_{ij}^d(I, t)} \right)^m, \quad 2.4$$

where for each pixel  $i, j$ , both  $\overline{s_{ij}^d(I, t)}$  and  $\langle s_{ij}^d(I, t) \rangle$  are a function of position ( $i, j$ ) and exposure time (here we assume the scene brightness  $I$  is not changing). The surface fit also assumes that if a CFA sensor is used, separate fits are used for each color channel. Before fitting a surface to  $\overline{s_{ij}^d(I, t)}$  (Fig. 2.7c) at each integration time, a row-by-row adjustment was applied to remove the imbalance in output from the A/D converters. A low-order fit of the row-by-row ratio of two columns on either side of the border between image halves was used to adjust the left side of the image.

An example of the results of the uniformity correction for the red channel is shown in Fig. 2.7. For this figure, the terms  $b_{ij,m}$  in Eq. 2.4 are obtained by using a training set of images, setting  $M = 2$ . The correction is then applied to a validation set using Eq. 2.2. The method corrects hot pixels that have not reached saturation, and corrects small-scale FPN, but it fails to correct large-scale nonuniformity. This occurs because the surface used for correction is fit to non-uniformities that occur across the whole image. It is therefore not as robust as the uniform illumination approach but is a useful substitute in field operations.



**Figure 2.7.** (a) Raw red image of smooth source; (b) average of red ten frames, including (a); (c),(d) uniformity correction applied to (a),(b), respectively.

### 2.5.4 Photoresponse linearity

Knowledge of the camera's response as a function of both intensity and exposure time is a prerequisite for the HDR process. The simplest model for a pixel's photoresponse is linear in the product of irradiance  $I$  on the sensor plane and exposure time  $t$

$$c_{ij}(I, t) = \hat{a}_{ij,o} + d_{ij,1}It, \quad 2.5$$

where  $M$  and  $N$  from Eq. 2.1 have been taken as zero and one, respectively. Assuming a constant intensity during the exposure sequence, we convert the value measured in an

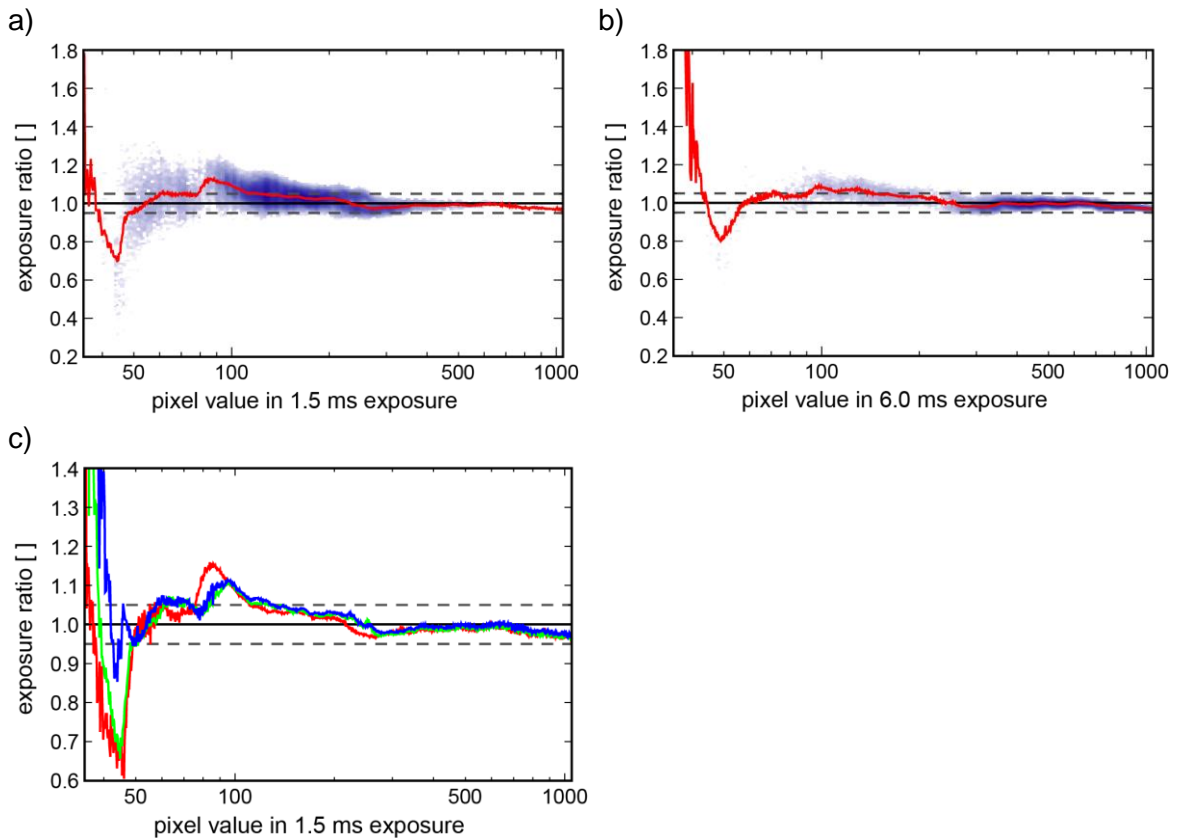
exposure of integration time  $t$  to the expected value had it been captured at integration time  $t_{ref}$ :

$$c_{ij}(I, t_{ref}) = (c_{ij}(I, t) - \hat{a}_{ij,o}) \frac{t_{ref}}{t} + \hat{a}_{ij,o}. \quad 2.6$$

This linear model predicts that the measurement values of the same scene should be scaled by the ratio of the exposure times from one image to the next. For example, we would expect that all the values in a 6 ms exposure would be 4 times as large as the values of the corresponding pixels in a 1.5 ms exposure. Figure 2.8 shows the ratio of modeled values based on a longer exposure to the measured values in a shorter exposure (i.e.  $t_{ref}/t = 0.25$ ). An average of five frames was used at each exposure time in making the comparison. To avoid negatively biasing the results, pixels that saturate in the longer image were removed, which corresponds to pixel values of over 1024 in the shorter exposure.

The observed deviation from unity is a measure of the error we introduce by scaling up a given value from the short exposure to place it in a composite with the longer exposure. Below 100 counts ( $\sim 2.5\%$  full scale), there appear to be significant non-linearity effects, and we do not recommend using signals below this level. Between about 400 and 800 counts, the median deviation is nearly zero. Deviations are small ( $< 5\%$ ) from around 150 counts to the end of the overlap range just above 1000 counts. Over the majority of the range, neither exposure time nor color has a significant effect on the result. The overlap of this "sufficiently linear" region on the abscissa of Fig. 2.8 extends from 409 counts (the lower limit in the short exposure) to 921 counts (the upper limit in the long exposure after multiplying by the integration time ratio, i.e.  $3684 \times 0.25 = 921$ ). We have therefore elected, for the purposes of this work, to consider pixel

response to be sufficiently linear if the value is between 10% and 90% of full scale, i.e. 409 to 3686 counts.



**Figure 2.8.** Evaluation of sensor linearity using sky images under thin overcast conditions. In (a), a point cloud (and median in red) showing the distribution of the ratio between a 6 ms exposure and a modeled 6 ms exposure generated from a 1.5 ms exposure, as a function of measured value in the 1.5 ms image. In (b), the same as (a), but with 6 ms and 24 ms exposures. In (c) the median line for each color is shown. To reduce random noise, each of the compared images is the average of five exposures captured over the course of approximately 3 seconds.

### 2.5.5 High dynamic range imaging

In order to image the daytime sky it is important that the camera have an extremely large dynamic range, since we wish to obtain images of both very bright objects (such as the sun and sunlit clouds) as well as very dark objects, such as the undersides of thick clouds. Unfortunately, 12 bit (or fewer) image sensors generally do not have sufficient dynamic range for this task in a single exposure. Instead, we capture

multiple exposures with different integration times in quick succession and combine those exposures into a single *high dynamic range* image (Debevec and Malik, 1997). Three 12-bit exposures are composited together to produce a single 16-bit image.

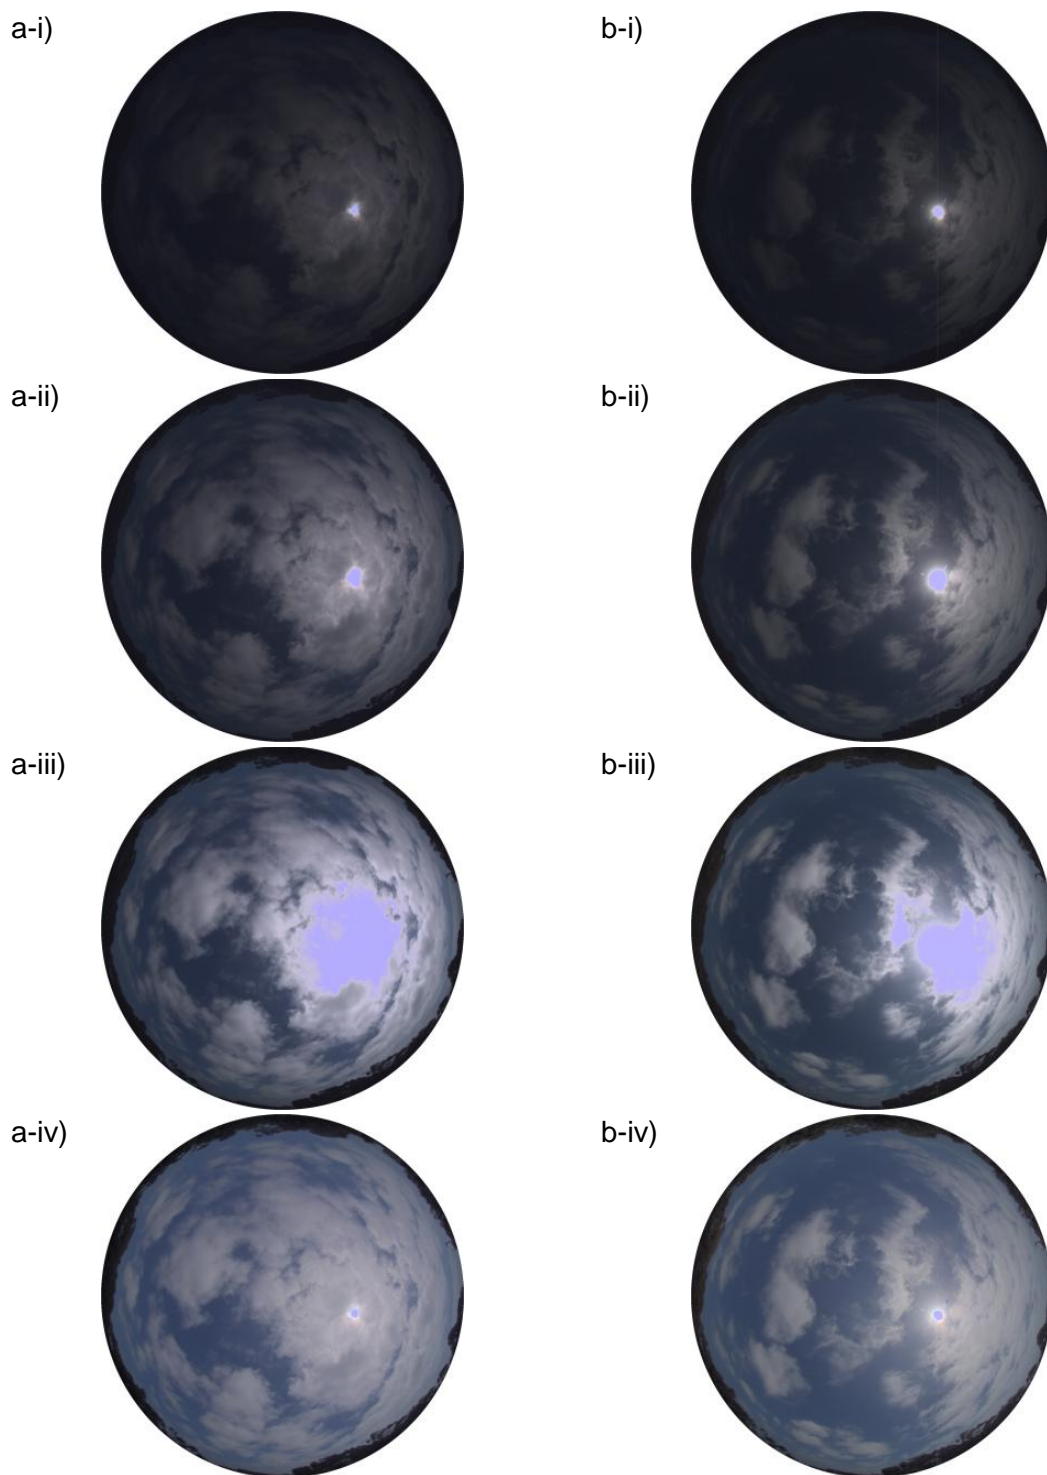
Although methods exist that would allow us to use a more sophisticated photoresponse model than Eq. 2.5 (e.g. Mann and Picard, 1994), by only using pixels in the linear region of the sensor photoresponse (section 2.5.4), we can apply the simple linear response model without significant error. For purposes of the HDR composite, this means that for a single exposure the pixels with values below 409 or above 3,686 counts are excluded. The integration times on the USI are separated by factors of four (i.e.  $t$ ,  $4t$ , and  $16t$ , where  $t$  is system dependent). This ensures that the region between 409 and 921.5 counts in a shorter exposure will overlap with the region between 1,636 and 3,686 counts in a longer exposure. Based on the results shown in Fig. 2.8, these settings ensure the linear approximation in Eq. 2.5 is applicable for the subset of overlapping pixels in the HDR image.

The HDR process is straightforward. First we select the pixels in each of the three exposures that are properly exposed, eliminating areas that are below 10% or above 90% of full scale. Next, using Eq. 2.6, we map the values for each pixel to what they would have been in the frame with the longest exposure time. This assumes that for short duration of an HDR exposure sequence, scene intensity is constant. Finally, we combine the exposures, using the average of all valid values for each pixel. This method is simple and effective, as demonstrated in Figs. 2.9 and 2.10. It is, however subject to small composition artifacts if the sensor response linearity is not properly characterized. If an image patch contains values for which sensor response is nonlinear *and* the HDR algorithm transitions from using a different subset of the three available exposures within




this patch, a small 1-2 pixel intensity step will occur, which after demosaicing into a color image, appears as a color fringe.

Figure 2.9 and 2.10 demonstrate the HDR method applied to two systems, USI 1.2 and USI 1.8 respectively (see Table 2.3). USI 1.2 used a 9,520  $\mu\text{m}$  diameter aperture and neutral density filter, whereas USI 1.8 used a modified aperture of diameter 1,000  $\mu\text{m}$  (note the spectral variation between instruments). Figure 2.9a and 2.9b highlights the differences between the HDR capture sequence in cloudy conditions for an obstructed and unobstructed sun. Figure 2.10 provides an overview of imaging performance in a variety of sky conditions, with both obstructed and unobstructed sun. Figure 2.10d shows a thin cloud in low lighting conditions, and in Fig. 2.10g a halo caused by the thin clouds can be seen.



**Figure 2.9.** USI 1.2 high dynamic range (HDR) exposure sequence for (a) May 23, 2013, 3:22 pm PDT and (b) May 23, 2013, 2:48 pm for integration times of (i) 30 , (ii) 120, and (iii) 480 ms. (a-iv) and (b-iv) show the final HDR composites.



**Figure 2.10.** HDR images from USI 1.8 in April and May, 2013, showing a variety of sky conditions. Color correction has been applied. Images required intensity rescaling for display purposes.

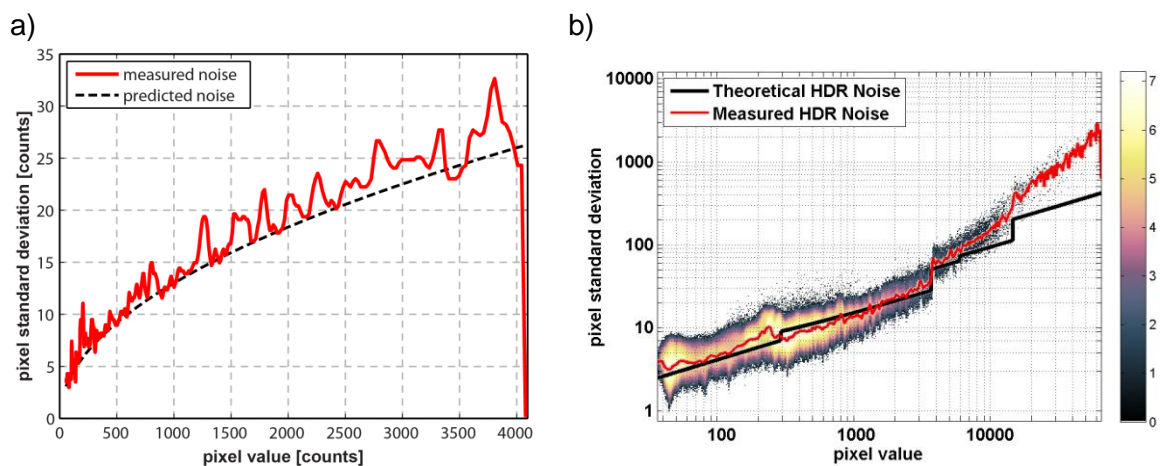
### 2.5.6 Brightness measurement uncertainty in HDR imagery

Two images of the exact same scene will not be identical due to the random shot noise present in the measurements. Electron generation in the sensor follows a Poisson distribution, so the root mean square (RMS) of the shot noise is expected to be  $e_{ij,shot} = \sqrt{e_{ij}}$ , where  $e_{ij}$  is the quantum unit being measured at pixel  $i, j$ . The quanta considered here is electrons. Assuming shot noise is the dominant noise source, this square root increase in RMS shot noise with stored electric charge  $e_{ij}$  implies the signal-to-noise ratio also increases as  $\sqrt{e_{ij}}$ . Shot noise places a fundamental limit on the lower bound of measurement uncertainty for an image sensor. The predicted RMS noise as a function of count value for a 12-bit image is shown in Fig. 2.11a. For this calculation, the manufacturer specified gain  $g$  of 0.174 counts per electron was used. Measured system noise as a function of pixel value (in counts) was quantified by computing the pixel-by-pixel standard deviation  $\sigma_{ij}$  for ten frames of a stationary scene, binning  $\sigma_{ij}$  by the pixel-by-pixel mean  $\mu_{ij}$  into bins 0 to 4095, and finally by taking the median  $\sigma$  of each bin. The drop-off near the maximum occurs because the upper bound that the saturation limit imposes causes the standard deviation of measured values to reduce.

When combining exposures in an HDR composite, the shot noise present in an individual pixel will depend on which exposures were compiled for that particular pixel, and the scaling factor  $t_{ref}/t$  for each pixel in the composition. For sufficiently large number of electrons, the Poisson distribution is approximately normal by the central limit theorem, and thus the RMS noise from each frame can be summed in quadrature to obtain the RMS shot noise  $c_{ij,shot}$  in an HDR exposure, i.e.

$$c_{ij,shot} = g \sqrt{\sum_k^P \frac{t_{ref}}{t_k} (e_{ij,shot})_k^2}, \quad 2.7$$

where  $k$  is the individual frame index,  $P$  is the number of frames, which ranges from one to three in this work. The actual RMS noise present in an HDR image was computed using the method described for Fig. 2.11a, and is shown in Fig. 2.11b. The noise is compared to the shot noise limit (Eq. 2.7, black line, Fig. 2.11b), where the number of frames in the HDR composition is determined using the algorithm described in the previous section. The use of different combinations of frames can be seen as sharp jumps in the theoretical minimum in Fig. 2.11b.



**Figure 2.11.** Photon transfer curve for a USI system for (a) a 12-bit image, and (b) and HDR image. The theoretical minimum shot noise limit is shown as a black line, and the median of the noise distribution at each count value is shown in red. In (b), the density of the pixel standard deviation distribution is shown behind the curves.

The curves presented in Fig. 2.11 are similar to photon transfer curves (PTCs) which characterize not only shot noise, but all random noise present in the image sensor. Noise sources such as dark current and read noise are subtracted out of a PTC. The closeness of the curves to the noise limit indicates that for the USI system, sources of noise other than shot noise are small in both a 12-bit image, and the HDR composition. The fluctuations in each curve, and the dips below the theoretical minimum occur because a limited number of samples were taken (10 frames). Above 15,000

counts, very few samples were present in the HDR images, so noise in this region is not well characterized here.

### 2.5.7 Stray light

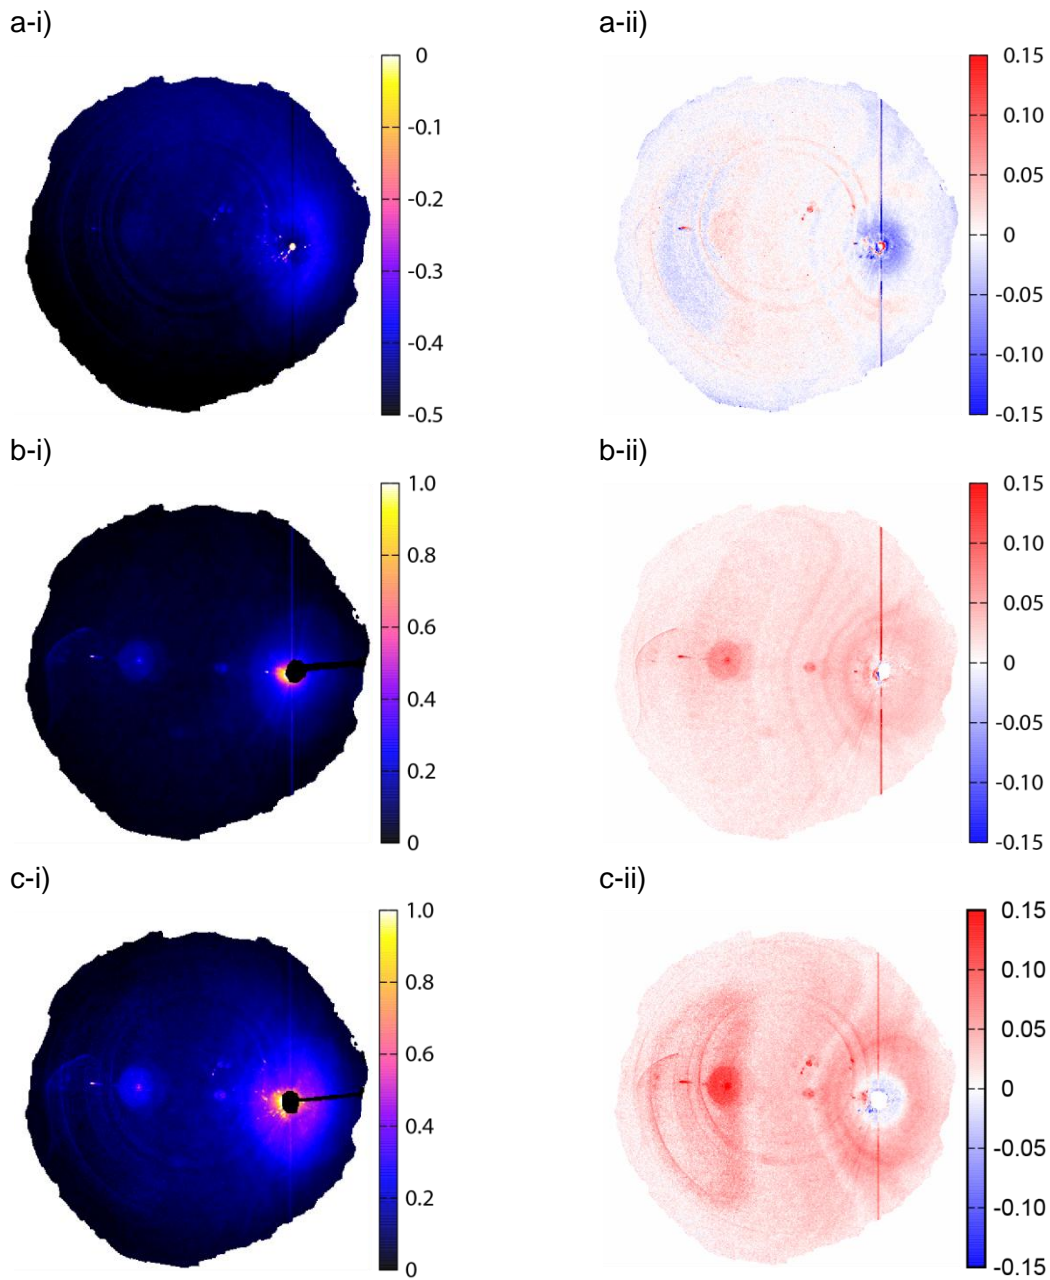
The red-blue-ratio image (RBR), defined as the ratio of the red channel to the blue channel, is the most common feature used for cloud detection. Clear sky has a relatively low RBR and clouds have a higher RBR. RBRs typically span between 0.4 and 1.2 for the USI, and the threshold for cloud is about 0.5. Stray light, due to exposure of the optical assembly to the direct beam, results in spots and artifacts in the image that are brighter generally whiter (i.e. more spectrally neutral) than they should be, resulting in either false positive cloud detections when the stray light pushes a hazy sky above the cloud threshold, or missing clouds due to contamination of the clear sky library (see Chow et al. 2011 or Yang et al. 2014 for details).

In order to characterize the stray light present in our system, we used a simple, hand-held shade to block the sunlight. Measurements were conducted on a clear day (May 13, 2013) and shaded and un-shaded images were taken 30 seconds apart. By comparing images captured with and without the shading device, we can observe the effect of stray light on the resulting images. Three different pairs of images are compared in Fig. 2.12. First, a normal image is compared to one taken with the dome removed. Second, with the dome removed, images taken with and without the shade are compared. The third and final comparison considers shaded and un-shaded images with the dome on. The latter comparison gives the best estimate of the total effect of stray light on the images produced by the USI, while the first two allow us to qualitatively separate effects due to the dome and the lens. To quantify the effects of stray light, the residual fractional intensity  $(I_2 - I_1)/I_1$  is computed and shown in the left column of Fig.

2.12, where  $I_1$  is the image with the shade (or without dome, pair 1), and  $I_2$  is the image without the shade (or with dome, pair 1).


Increases or decreases in residual fractional intensity affect the radiometric analysis of sky imagery, but for solar forecasting primarily the RBR is of interest. Therefore, it is primarily spectral variations in stray light that are of interest. Increases in RBR due to stray light are expected since the majority of stray light originates from the direct solar beam, which is whiter than most of the sky. To quantify the impacts of stray light on cloud detection, the difference  $RBR_2 - RBR_1$  of each of the three described pairs is shown in Fig. 2.12ii.

The following stray light effects were identified: 1) an overall increase in measured intensity averaging 12% across the image (Fig. 2.12c-i); 2) concentric ring-like reflections off the front face of the camera lens that reflect off the inner-surface of the dome (Fig. 2.12a versus b); 3) particularly strong (and bluish) forward scattering off the dome (bright circle in Fig. 2.12a-ii); 4) sharp reflections off of elements in the optical assembly, visible as spots along the intersection of the solar principal plane and the image plane (all); 5) a "swoopy" shape resulting from reflection of sunlight off the rear gelatin neutral density (ND) filter at the back of the lens (all); and 6) vertical smear that results near the sun from signal overflow during sensor readout (all); 7) at higher solar elevations (Fig. 2.13), a reflection of the sun off the surface of the image sensor. Here, the solar principal plane is defined by camera zenith vector and the vector to the sun. The dome decreases the stray light intensity by about 46% because of the ND acrylic used (Fig. 2.12a versus b and c). While the dome surface was clean during testing, in normal operations dirt or scratches on the dome will result in additional scattering with a specific pattern that changes not just with the position of the sun, but also as a function of time since last cleaning.



**Figure 2.12.** Stray light from the dome (top), lens and neutral density filter (middle), and whole system (bottom). The left column shows the fractional change in intensity due to stray light, while the right column shows the shift in the red-blue ratio from the shaded to unshaded image. Images were recorded against a clear (blue) sky, so stray (white) light shifts toward the red. Note the scale change between (a) and (b),(c) in the left column.





**Figure 2.13.** Stray light comparison between two designs of the USI; (a) design with filter, and (b) design with modified aperture.

Stray light impacts of the modified aperture versus the ND filter were qualitatively evaluated by visually inspecting a clear sky images such as those in Figure 2.13. The following differences between the modified aperture and the wide-open, filtered configuration are noted: i) reflection from the ND filter surface is, naturally, missing in the model without a filter; ii) the wide open aperture in the filtered configuration exhibits a pair of reflections of the sun striking the image sensor that become visible at high solar elevations (when the direct-beam is nearly orthogonal to the image plane); this has not been observed using the modified aperture; iii) the modified aperture shows a larger number of circles along the diameter containing the sun (i.e. intersection of the solar principal plane and the image plane); iv) a "feathery" radial pattern is sometimes observed near the sun with the modified aperture, arising from imperfections in the circularity of the aperture; v) the modified aperture has a more prominent smear stripe because the selected aperture diameter allows more light into the camera; and vi) prototypes with extremely small apertures exhibited diffraction rings around the sun (Fig. 2.30). Effect iii occurred because the antireflective black-oxide coating applied to the steel was mistakenly polished by the machinist, which increased its reflectivity.

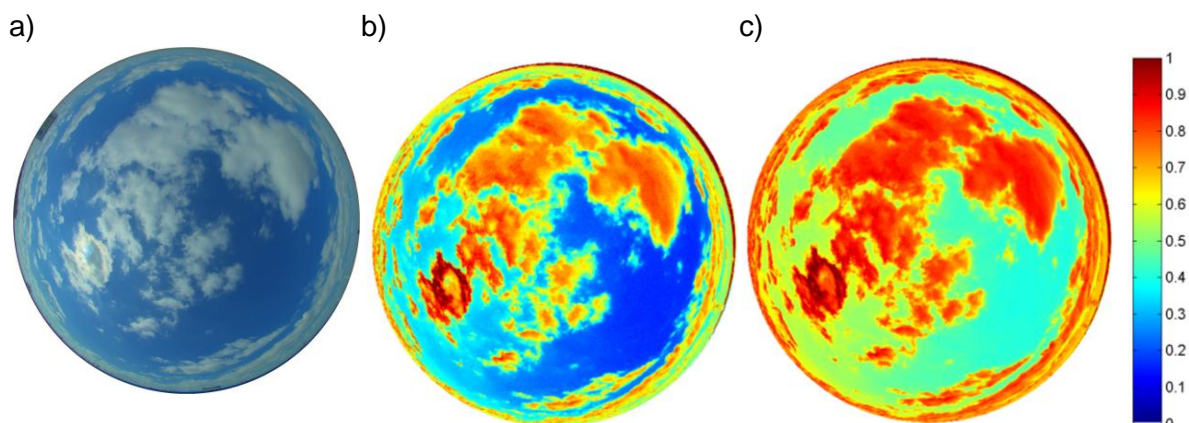
To correct these issues, we have performed experimentation with a stray light ratio lookup table as a function of solar zenith angle, sun-pixel angle, and image zenith angle (similar to the clear sky library, Chow et al. 2011). However, the results, while promising, were inconsistent and thus are not reported here. From our experience using the USI for forecasting, the stray light features discussed here negatively affects image quality and results in identifiable forecast performance degradation. Yang et al. (2014) have implemented adjustments to the cloud detection methods of Chow et al. (2011) to specifically address solar power forecast errors due to stray light. In future work we hope to develop corrections for the USI imagery so that stray light levels in imagery is reduced prior to being input into the cloud detection algorithms.

### 2.5.8 Color balancing

The neutral density filters currently used in the USI (Kodak Wratten 2, No 96 ND3.0) introduce a color cast to the image. Basic color correction is performed by applying the following color correction matrix provided by Truesense Imaging (pers. comm.):

$$D_{75} = \begin{bmatrix} 1.51018 & -0.24732 & -0.26286 \\ -0.06149 & 1.18892 & -0.12742 \\ 0.08708 & -0.53623 & 1.44915 \end{bmatrix}.$$

This color correction tends to increase the range over which RBR varies and decreases the lower bound (Figure 2.14). In the future we may use a color reference chart (e.g. the IT8.7/2-1993 calibration target) in order to improve the color balance of USI images in a way that might impact forecasting performance more.



**Figure 2.14.** Red-blue-ratio (RBR) comparison of (a) a sky image (Figure 2.10) for (b) color corrected image and (c) original image. The RBR scale is given in the colorbar.

## 2.6 Deployment experience

The UCSD USI system has been deployed across the United States (Table 2.3). The predominant cloud types in coastal California (USIs 1.1, 1.2, 1.9) are marine stratocumulus. In Kahului, Hawaii there are persistent orographic clouds over the West Maui Mountains to the west-northwest of USI 1.10 which makes it an interesting place to study non-advective solar forecast schemes. Redlands, California is hot and dry, and usually clear, but often sees higher ice clouds and larger synoptic systems. In Billings, Oklahoma there is a wide diversity of cloud conditions that occur from high ice clouds, to lower cumulus clouds. Solar forecasting algorithms may have location dependent performance, and testing components of an algorithm in multiple locations can help to identify shortcomings and areas for improvement.

The data gathered from the two instruments in Billings Oklahoma are of particular interest because they were fielded at a United States Department of Energy Atmospheric Radiation Measurement Program field site (the Southern Great Plains site). The site is outfit with a diverse suite of measurement equipment, including cloud radar covering a number of bands, several lidar systems, shortwave and longwave

radiometers, aerosol measurements, and a doppler wind profiler. These collocated measurements will be used to assess the performance of a number of remote sensing algorithms developed for the USI.

**Table 2.3.** USI Locations in the United States and deployment time ranges

<b>USI No.</b>	<b>Longitude [deg]</b>	<b>Latitude [deg]</b>	<b>Altitude [m]</b>	<b>State</b>	<b>City</b>	<b>Start Date</b>	<b>Stop Date</b>
1.1	-117.233088	32.881090	120	California	La Jolla	2012-04-21	-
1.2	-117.240987	32.872136	135	California	La Jolla	2012-06-06	-
1.5	-117.243111	34.076355	347	California	Redlands	2012-10-18	2014-03
1.6	-117.209333	34.079822	384	California	Redlands	2012-05-45	2014-03
1.7	-97.478766	36.618377	304	Oklahoma	Billings	2013-03-11	2013-11-04
1.8	-97.484871	36.604094	318	Oklahoma	Billings	2013-03-11	2013-11-04
1.9	-117.238378	32.707122	15	California	San Diego	2013-04-19	2014-03
1.10	-156.479136	20.890549	20	Hawaii	Kahului	2013-08-21	-

## 2.7 Conclusions

Clouds have a high degree of spatial complexity and the intensity range within a single scene can be over five orders of magnitude (including the sun). For solar forecasting applications, it is important to capture this information at a high spatial and radiometric resolution to facilitate the development of advanced algorithms and techniques. The UCSD Sky Imager system is a step in this direction. Ten instruments have been built and can be made available to other researchers. The units come with a camera and system control software and an extensive library of processing tools is available. The developers are also open to commercializing the instrument and extensive design documentation is available.

## Acknowledgements

The work presented in this chapter would not have been possible without extensive student support in the design, fabrication, and assembly of the USI. We would

like to thank, in no particular order, Caspar Hanselaar, Edmundo Godinez, William Gui, Felipe Mejia, Prithvi Sundar, Dan Erez, Scott Kato, Amy Chiang, Jessica Traynor, Kristen Ostosh, Tyler Capps, Sebastian Schwarzfischer, Sebastian Pangratz, Christian Faltermeier, Nick Truong, Salil Kektar, Jeff Yeh, Max Twogood, Alex Turchik, Danielle Donnelly, Emily Davis, and last, but not least the Victors Fung (1) and Piovano (2). We also appreciate funding from the Panasonic Corporation , the Department of Energy High Solar PV Penetration Award Number EE-0004680, and the California Energy Commission contract 500-10-060.

In the course of preparing this chapter, the following publications were used:

Publication: Development of a sky imaging system for short-term solar power forecasting. Atmospheric Measurement Techniques Discussion, volume 7, pp. 4859-4907, 2014. doi:10.5194/amtd-7-4859-2014

Co-Authors: Ben Kurtz, Elliot Dahlin, Mohamed Ghonima, Janet E Shields and Jan Kleissl

## **3 Camera Calibration**

Contributing Authors:

B. Urquhart, B. Kurtz and J. Kleissl

### **3.1 Summary**

A camera model and associated automated calibration procedure for stationary daytime sky imaging cameras is presented. The specific modeling and calibration needs are motivated by remotely deployed cameras used to forecast solar power production where cameras point skyward and use 180 degree fisheye lenses. Sun position in the sky and on the image plane provides a simple and automated approach to calibration; special equipment or calibration patterns are not required. Sun position in the sky is modeled using a solar position algorithm (requiring latitude, longitude, altitude and time as inputs). Sun position on the image plane is detected using a simple image processing algorithm. The performance evaluation focuses on the calibration of a camera employing a fisheye lens with an equisolid angle projection, but the camera model is general enough to treat most fixed focal length, central, dioptric camera systems with a photo objective lens. Calibration errors scale with the noise level of the sun position measurement in the image plane, but the calibration is robust across a large range of noise in the sun position. Calibration performance on clear days ranged from 0.94 to 1.24 pixel root mean square error.

### **3.2 Introduction**

The power output variability of renewable energy sources poses challenges to its integration into the electricity grid. Forecasting of renewable power generation (e.g. Monteiro et al. 2009, Perez et al. 2010, Kleissl 2013) enables more economical and

reliable scheduling and dispatch of all generation resources, including renewables, which in turn accommodates a larger amount of variable supply on the electricity grid. Specifically for solar power forecasting, a number of technologies are being applied: numerical weather prediction (e.g. Lorenz et al. 2009, Mathiesen and Kleissl 2011, Perez et al. 2013); satellite image-based forecasting (e.g. Hammer et al. 1999, Perez and Hoff 2013); and stochastic learning methods (e.g. Bacher et al. 2009, Marquez and Coimbra 2011, Pedro and Coimbra 2012). For very short term (15 minutes ahead) solar power forecasting on the kilometer scale, sky imaging from ground stations has demonstrated utility (Chow et al. 2011, Urquhart et al. 2013, Marquez and Coimbra 2013, Yang et al. 2014).

Some of these sky imaging methods require the camera to be geometrically calibrated, i.e., each pixel must be associated with a corresponding view direction. Together with cloud height estimates, the view direction allows geolocation of clouds and their shadow projections such that their position is known relative to solar power plants. Geometric calibration is a common task in photogrammetry and computer vision, and calibration methods have been developed for a variety of applications. Some methods for calibrating a stationary camera require the use of calibration equipment or setups (Tsai 1987, Weng et al. 1992, Heikkilä and Silvén 1996, Shah and Aggarwal 1996) or planar targets (Wei and Ma 1993, Sturm and Maybank 1999, Zhang 2000). Geometric scene information can be used to calibrate the camera's internal parameters (Liebowitz and Zisserman 1998) or estimate lens distortion (Brown 1971, Devernay and Faugeras 2001, Tardif et al. 2006). Scenes with parallel or perpendicular lines or primitive shapes are not generally available for skyward pointing cameras and thus there are no structures from the built environment around which to base a generic and automated calibration procedure.

Cameras used for solar power forecasting often employ fisheye lenses, which require appropriate camera modeling and associated model parameter estimation methods due to the large distortion required to achieve the approximately  $180^\circ$  field of view. Many models which include lens distortion cannot account for distortion present in lenses which have a field of view equal to or exceeding  $180^\circ$  because they rely on converting 'distorted' image coordinates (which are finite measurements on the image plane) to 'undistorted' image coordinates which are infinite at angles  $90^\circ$  from the optical axis (e.g. Tsai 1987). Gennery (2006) and Kannala and Brandt (2006) propose generic camera models suitable for fisheye lenses, and the form of the camera model presented here has features of both. The goal of the current work is to develop 1) a general camera model and 2) a calibration method for a wide angle dioptric sky camera with a photo objective lens that can be automated with little user input.

The calibration approach taken here is sometimes referred to as stellar calibration, where the 3D position of an object or set of objects is treated as known. In particular the sun position in the sky is treated as a known input which is used along with the corresponding measured sun position to calibrate a stationary camera of fixed focal length. Sun position has been used previously for camera calibration. Lalonde et al. (2010) have used manual image annotation to select the sun position in a few images, and with this estimated the focal length, principle point, and two of the three rotational degrees of freedom (the camera horizontal axis was assumed parallel to the ground). The work presented here builds on this idea and extends it using a more generalized camera model and automated sun detection. The camera model here allows any pose, non-square pixels, and both radially symmetric and decentering distortion components.

The layout of this paper is as follows. Section 3.3 discusses the forward and backward camera model. Section 3.4 discusses the imaging equipment and solar



position input used for the calibration process. Section 3.5 provides details of the calibration procedure: initialization, linear estimation, and nonlinear estimation. Section 3.6 provides results for both measured solar position input and synthetic data. Synthetic data is used to assess the uncertainty in calibration performance and parameter estimation as a function of measurement uncertainty.

### 3.3 Camera Model

The forward camera model projects points from a 3D scene onto the image plane. The backward camera model described in section 0. projects points on the image plane to rays in 3-space. Both models are developed assuming that the camera-lens system is central, i.e. all refracted rays within the lens pass through a single point. This, while not physically accurate, yields a close approximation (Ramalingam et al. 2005).

#### 3.3.1 Forward Camera Model

##### 3.3.1.1 Projective Transformation Camera Model

The standard model for a camera without distortion is a 3D to 2D projective transformation, mapping points  $\mathbf{X} = (X, Y, Z, T)^T$  in  $\mathbb{P}^3$  to  $\mathbf{x} = (x, y, w)^T$  in  $\mathbb{P}^2$ :

$$\mathbf{x} = P \mathbf{X}, \quad 3.1$$

where  $P$  is a  $3 \times 4$  perspective projection transformation with 11 degrees of freedom (it is defined up to scale), and  $\mathbb{P}^n$  is the  $n^{\text{th}}$  dimension of projective space. The points  $\mathbf{X} \in \mathbb{P}^3$  and  $\mathbf{x} \in \mathbb{P}^2$  are homogeneous quantities and thus are also defined only up to scale. The corresponding inhomogeneous points in Euclidean space are  $\tilde{\mathbf{X}} = (X/T, Y/T, Z/T)^T = (\tilde{X}, \tilde{Y}, \tilde{Z})^T$ ,  $\tilde{\mathbf{X}} \in \mathbb{R}^3$ , and  $\tilde{\mathbf{x}} = (x/w, y/w)^T = (\tilde{x}, \tilde{y})^T$ ,  $\tilde{\mathbf{x}} \in \mathbb{R}^2$ . The tilde overbar indicates inhomogeneous coordinates throughout this work. When scale factors  $T$  or  $w$  are zero, the corresponding Euclidean point is infinite. For points not lying on the plane or line at

infinity, we can write  $\mathbf{X} = (\tilde{\mathbf{X}}^\top, 1)^\top$  and  $\mathbf{x} = (\tilde{\mathbf{x}}^\top, 1)^\top$ , respectively. The point imaging transformation  $P$  is given by a composition of Euclidean, affine and perspective transformations

$$P = \underbrace{\begin{bmatrix} \alpha_x & s & x_o \\ 0 & \alpha_y & y_o \\ 0 & 0 & 1 \end{bmatrix}}_{\text{affine}} \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{\text{perspective}} \underbrace{\begin{bmatrix} r_{11} & r_{12} & r_{13} & t_x \\ r_{21} & r_{22} & r_{23} & t_y \\ r_{31} & r_{32} & r_{33} & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\text{Euclidean}}, \quad 3.2$$

where the affine transformation is known as the camera calibration matrix (denoted by  $K$ , parameters defined later), the perspective transformation  $P_n = [I \mid \mathbf{0}]$  projects 3D space points to 2D image points, and the Euclidean transformation gives the rotation and displacement of the camera center relative to the world coordinate system. The Euclidean transformation can be written in block matrix notation as

$$\begin{bmatrix} R & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{bmatrix},$$

where the upper left block  $R$  with components  $r_{ij}$  is a rotation from world coordinates into the camera coordinate system, and the upper right block  $\mathbf{t}$  with components  $t_i$  is a translation giving the displacement *from* the origin of the camera coordinate system (i.e. the camera center) *to* the origin of the world coordinate system. The rotation matrix has only three degrees of freedom and can be represented by the angle-axis three vector  $\mathbf{w}$ , where  $R = \text{expm}([\mathbf{w}]_\times)$ ;  $\text{expm}$  is the matrix exponential and the notation  $[\cdot]_\times$  indicates the  $3 \times 3$  skew symmetric matrix corresponding to the vector argument. The three rotation plus three translation parameters are known as the camera's *extrinsic* parameters. In inhomogeneous coordinates, the rigid body (Euclidean) transformation from world coordinates  $\tilde{\mathbf{X}}$  to camera coordinates  $\tilde{\mathbf{X}}_{cam}$  is

$$\tilde{\mathbf{X}}_{cam} = R\tilde{\mathbf{X}} + \mathbf{t},$$

or equivalently using homogeneous coordinates

$$\mathbf{X}_{cam} = \begin{bmatrix} R & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{bmatrix} \mathbf{X}. \quad 3.3$$

The calibrated points  $\hat{\mathbf{x}} = (\hat{x}, \hat{y}, \hat{w})^\top$  on the image plane given by  $\hat{\mathbf{x}} = [I | \mathbf{0}] \mathbf{X}_{cam}$ , can be converted to pixel coordinates  $\mathbf{x} = (x, y, w)^\top$  by the affine transformation  $K$

$$\mathbf{x} = \begin{bmatrix} \alpha_x & s & x_o \\ 0 & \alpha_y & y_o \\ 0 & 0 & 1 \end{bmatrix} \hat{\mathbf{x}} = K \hat{\mathbf{x}}, \quad 3.4$$

where  $\alpha_x$  and  $\alpha_y$  are the effective focal lengths (in pixels) in the  $x$  and  $y$  directions, respectively,  $(x_o, y_o)^\top = \mathbf{x}_o$  is the principal point (i.e. the point of intersection of the optical axis with the image plane), and  $s$  is the skewness of the pixel coordinate axes. The five parameters in matrix  $K$  are known as the camera's *intrinsic* parameters. The effective focal lengths  $\alpha_x = f k_x$  and  $\alpha_y = f k_y \csc \psi$  account for the actual focal length  $f$  (in meters) and the potential for pixel sizes  $1/k_x$  and  $1/k_y$  (in meters per pixel) to vary in the  $x$  and  $y$  directions, respectively. The angle  $\psi$  is the angle between the  $x$  and  $y$  axes, which is close to  $\pi/2$  for our camera, thus  $\alpha_y \approx f k_y$ . The skewness  $s = \alpha_x \cot \psi$  is the degree to which the rows and columns of the image sensor are not orthogonal.

In summary, the model of a camera given by equation 3.1 is known as the pinhole camera model. It contains six extrinsic (external) and five intrinsic (internal) camera parameter's and thus has 11 degrees of freedom. While the pinhole camera model has been widely used, it does not account for lens distortion and assumes that the camera is a central projection camera. Since we seek to develop a model for use with a fisheye lens exhibiting a significant amount of distortion, the above model must be modified appropriately.

### 3.3.1.2 Distortion Model

An equivalence class  $\mathbf{X}_{cam} \in \mathbb{P}^3$  in projective space (i.e. a point or vector in  $\mathbb{P}^3$ ) defines a ray  $\Phi = (\theta, \phi)^\top$  in Euclidean space ( $\mathbb{R}^3$ ):

$$\Phi = \begin{bmatrix} \theta \\ \phi \end{bmatrix} = \begin{bmatrix} \text{atan}\left(\frac{\sqrt{X_{cam}^2 + Y_{cam}^2}}{Z_{cam}}\right) \\ \text{atan}\left(\frac{Y_{cam}}{X_{cam}}\right) \end{bmatrix}, \quad 3.5$$

where  $\theta$  is the angle between the ray and the optical axis (i.e. the camera zenith angle), and  $\phi$  is the angle from the positive  $X_{cam}$ -axis to the projection of the ray onto the  $X_{cam} - Y_{cam}$  plane. The angle  $\phi$  is positive in the counterclockwise direction. The incoming ray  $\Phi$  is mapped onto the image plane by a mapping  $\mathcal{D}$  as

$$\begin{bmatrix} \hat{x} \\ \hat{y} \end{bmatrix} = \mathcal{D}(\Phi), \quad 3.6$$

where  $(\hat{x}, \hat{y})^\top$  are calibrated inhomogeneous coordinates in the image plane (the hat  $\hat{\phantom{x}}$  denotes a calibrated point, and the tilde  $\sim$  denotes an inhomogeneous coordinate, defined previously). The mapping  $\mathcal{D}$  is, in general, nonlinear and includes the distortion produced by the lens-camera system. Here we model  $\mathcal{D}$  following Brown (1971) as

$$\mathcal{D}(\Phi) = \hat{r}(\theta) \begin{bmatrix} \cos \phi \\ \sin \phi \end{bmatrix} + \begin{bmatrix} \delta_{cx}(\Phi) \\ \delta_{cy}(\Phi) \end{bmatrix}, \quad 3.7$$

where  $\hat{r}(\theta)$  is the normalized radius on the image plane, and  $\delta_{cx}$  and  $\delta_{cy}$  account for decentering distortion in the  $x$  and  $y$  directions, respectively. The normalized radial distance  $\hat{r}$  is obtained by dividing the actual radial distance in the image plane by the focal length  $f$ . These terms will be further discussed in the following subsections.

### Radially Symmetric Distortion

The most common form of distortion in dioptric imaging systems with a photo objective lens is radially symmetric distortion. Several adjustments to the pinhole model

to account for radially symmetric distortion have been proposed for small field of view lenses exhibiting moderate amounts of pincushion or barrel distortion (e.g. Slama 1980). In order to generate a one-to-one mapping of hemispherical radiance ( $180^\circ$  field of view) to the image plane, fisheye lenses must introduce extreme radial distortion. For a centered lens system,  $\delta_{cx}$  and  $\delta_{cy}$  can be taken as zero and  $\hat{r}(\theta)$  can be set to one of the following projection functions (Miyamoto 1964):

$$g_o(\theta) = \tan(\theta) \quad \text{perspective projection (not fisheye),} \quad 3.8a$$

$$g_o(\theta) = \theta \quad \text{equidistant projection,} \quad 3.8b$$

$$g_o(\theta) = 2 \sin(\theta/2) \quad \text{equisolid angle projection,} \quad 3.8c$$

$$g_o(\theta) = 2 \tan(\theta/2) \quad \text{stereographic projection,} \quad 3.8d$$

$$g_o(\theta) = \sin(\theta) \quad \text{orthographic projection,} \quad 3.8e$$

Equations 3.8b to 3.8e correspond to fisheye lens projections. Equation 3.8a is the undistorted perspective projection (i.e. same projection model as equation 3.1), but can still be used in the camera model and calibration as described here.

Fisheye lens designers generally strive to meet one of the above projections, but due to manufacturing and assembly tolerances, the standard projections (equations 3.8b, c, d, e) only approximate a particular lens-camera system. In order to model wide angle and fisheye lenses more accurately, a number of models have been proposed (e.g. Kannala et al. 2006, and Shah and Aggarwal 1996). Here, instead of modeling the radially symmetric distortion using a polynomial in  $\theta$  (e.g. Kannala et al. 2006), we follow a suggestion by Gennery (2006) and use one of the standard models  $g_o(\theta)$  in Eq. 3.8, and then fit a polynomial to the residual radial distortion as

$$\hat{r}(\theta) = g_o(\theta) + \sum_2^N k_n \theta^n, \quad 3.9$$

where  $\hat{r}(\theta)$  is the normalized radius on the image plane, and the polynomial in  $k_n$  models deviations from  $g_o(\theta)$ . In this work,  $N$  was set to nine. In section 3.5 the coefficients  $k_n$  are denoted as a vector  $\mathbf{k}$ , where  $\mathbf{k} \in \mathbb{R}^8$ .

### **Decentering Distortion**

In addition to radially symmetric distortion, lenses exhibit tangential distortion which is a deviation from the radial alignment constraint (Tsai 1987) and causes the measured azimuth of a point  $\varphi$  to differ from its true azimuth  $\phi$ . Tangential distortion is due in part to a decentering of lens elements (Conrady 1919, Brown 1966). Based on the paraxial optics assumption, Conrady (1919) developed the following radial  $\delta_{cr}^C$  and tangential  $\delta_{ct}^C$  distortion terms arising from decentering for a point located at  $(r, \chi)$  on the image plane:

$$\delta_{cr}^C = 3p_1 r^2 \cos(\chi - \chi_1) + p_2 r (2 + \cos(2(\chi - \chi_2))), \quad 3.10a$$

$$\delta_{ct}^C = p_1 r^2 \sin(\chi - \chi_1) + p_2 r \sin(2(\chi - \chi_2)), \quad 3.10b$$

where  $p_1$  and  $p_2$  are constants that determine the magnitude of each centering defect,  $r$  is the radius in the image plane taken from the principal point,  $\chi_1$  and  $\chi_2$  are reference axes for the distortion effects. The constants are proportional to the lens decentering magnitude  $\Delta$  as  $p_1 \propto \Delta$  and  $p_2 \propto \Delta^2$ . Conrady (1919) did not develop terms of higher than first order in  $\Delta$ , i.e. only terms containing  $p_1$  were developed. The terms containing  $p_2$  (investigated for this work) are the only higher order terms that are not constant or symmetric over the image.

The Brown-Conrady distortion model (Brown 1971) formulates the radial  $\delta_{cr}$  and tangential  $\delta_{ct}$  decentering distortion components with reference axis  $\phi_o$  to be that of maximum tangential distortion with  $\phi$  is positive counter clockwise from the  $x$ -axis ( $\chi$  is positive clockwise):

$$\begin{aligned}\delta_{cr} &= 3\mathcal{P} \sin(\phi - \phi_o), \\ \delta_{ct} &= \mathcal{P} \cos(\phi - \phi_o),\end{aligned}$$

where the terms containing  $p_2$  have been neglected, and the profile function  $\mathcal{P} = p_1 r^2$ . Because Conrady did not develop terms in  $\Delta$  of higher order than one, Brown speculated that  $\mathcal{P}$  could be extended as a polynomial in even powers of  $r$  (written here as a normalized radial distance):

$$\mathcal{P}(\hat{r}) = \sum_1^M J_m \hat{r}^{2m}.$$

Expanding the aberrations due to decentering as developed by Conrady (1919), one finds that the second and third order terms in  $\Delta$  produce only *lower* order terms in  $r$  (i.e.  $r^1$  and  $r^0$ ). The zeroth order term in  $\Delta$  (which is thus present in centered lens systems) produces a shift in the image proportional to  $r^3$  and is commonly known as pincushion or barrel distortion. Because decentering effects in most lenses are small ( $p_1 r^2 \sim 10^{-4}$  pixels for our lense), it is reasonable to neglect  $p_2$ .

The use of the Brown-Conrady decentering distortion model for a fisheye lens should only be considered as an expedient for model fitting, and not as a physical description of optical distortion. Conrady derived the decentering formulae following a paraxial method he devised to analytically obtain the five classical Seidel aberrations (Conrady 1918). Expressions 3.10 are therefore only valid under the small angle approximation  $\sin \phi \approx \phi$ , and are thus not valid for the large incidence angles in a fisheye lens. Additionally, there is no physical justification for Brown's extrapolation of  $\mathcal{P}$  as an even ordered polynomial in  $r$  (recall that Conrady's original model had no higher order terms than  $r^3$ ). The retention in this work of the Brown-Conrady decentering distortion model is for model fitting only.

The radial and tangential distortion can be converted to the corresponding Cartesian components as

$$\begin{bmatrix} \delta_{cx} \\ \delta_{cy} \end{bmatrix} = \begin{bmatrix} \cos \phi & -\sin \phi \\ \sin \phi & \cos \phi \end{bmatrix} \begin{bmatrix} \delta_{cr} \\ \delta_{ct} \end{bmatrix},$$

which upon expanding gives

$$\begin{bmatrix} \delta_{cx} \\ \delta_{cy} \end{bmatrix} = \mathcal{P} \begin{bmatrix} -(2 \cos^2 \phi + 1) \sin \phi_o + 2 \sin \phi \cos \phi \cos \phi_o \\ -2 \sin \phi \cos \phi \sin \phi_o + (2 \sin^2 \phi + 1) \cos \phi_o \end{bmatrix}. \quad 3.11$$

Following Brown by taking

$$\begin{aligned} p_1 &= -J_1 \sin \phi_o, \\ p_2 &= J_1 \cos \phi_o, \\ p_m &= \frac{J_{m-1}}{J_1}, \quad m > 2 \end{aligned}$$

it can be easily shown that

$$\delta_{cx}(\Phi) = [p_1(1 + 2 \cos^2 \phi) + 2p_2 \sin \phi \cos \phi](\hat{r}^2 + p_3 \hat{r}^4 + p_4 \hat{r}^6 + \dots), \quad 3.12a$$

$$\delta_{cy}(\Phi) = [2p_1 \sin \phi \cos \phi + p_2(1 + 2 \sin^2 \phi)](\hat{r}^2 + p_3 \hat{r}^4 + p_4 \hat{r}^6 + \dots). \quad 3.12b$$

Here only  $p_1$  through  $p_4$  are used. In section 3.5, the coefficients  $p_n$  are denoted as a vector  $\mathbf{p}$ , where  $\mathbf{p} \in \mathbb{R}^4$ .



### 3.3.1.3 Forward Camera Model Overview

Summarizing the results of this section, the forward projection of a 3D space point to 2D pixel coordinates consists of the following four steps:

1. Euclidean transformation

$$\mathbf{X}_{cam} = \begin{bmatrix} R & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \mathbf{X}$$

2. Cartesian to spherical coordinates

$$\Phi = \begin{bmatrix} \theta \\ \phi \end{bmatrix} = \begin{bmatrix} \text{atan} \left( \sqrt{X_{cam}^2 + Y_{cam}^2} / Z_{cam} \right) \\ \text{atan}(Y_{cam} / X_{cam}) \end{bmatrix}$$

3. Lens-camera projection with distortion

$$\begin{bmatrix} \hat{x} \\ \hat{y} \end{bmatrix} = \left( g_o(\theta) + \sum_2^9 k_n \theta^n \right) \begin{bmatrix} \cos \phi \\ \sin \phi \end{bmatrix} + (\hat{r}^2 + p_3 \hat{r}^4 + p_4 \hat{r}^6) \begin{bmatrix} p_1(1 + 2 \cos^2 \phi) + 2p_2 \sin \phi \cos \phi \\ 2p_1 \sin \phi \cos \phi + p_2(1 + 2 \sin^2 \phi) \end{bmatrix}$$

4. Affine transformation

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} \alpha_x & s & x_o \\ 0 & \alpha_y & y_o \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \hat{x} \\ \hat{y} \\ 1 \end{bmatrix}$$

### 3.3.2 Backward Projection

In many cases, one is given points  $\mathbf{x}$  in image coordinates and what is needed is the back projection of those points into world coordinates. This is true for the application of solar forecasting where many quantities derived from images are assigned a space angle  $\Phi$  according to their image coordinates. For example, Chow et al. (2011) back project cloud positions detected within an image to a 3D world plane to generate a mapping of the clouds, and subsequently used this cloud map to ray trace cloud shadows. Note that obtaining the distance from the camera to an object in the scene is not possible from a single image because the projective nature of the imaging process. Therefore, in general, converting from a back projected ray to a point in the world reference frame is not possible from a single image.

The inversion of mapping  $\mathcal{D}$  (equations 3.6 and 3.7) is the most difficult part of developing a back projection model from a forward projection model, and Kannala et al. (2006) suggest a function inversion approach to this end. An alternative is to formulate a separate back projection model and fit it using synthetic data generated from the forward projection.

After converting to calibrated inhomogeneous image coordinates using  $(\hat{x}, \hat{y}, 1)^\top = K^{-1}(x, y, 1)^\top$ , the decentering distortion is formulated as a function of the polar coordinate  $(r, \varphi)$  in the image plane

$$\begin{bmatrix} \delta_{cx}(r, \varphi) \\ \delta_{cy}(r, \varphi) \end{bmatrix} = (r^2 + q_3 r^4 + q_4 r^6 + q_5 r^8) \begin{bmatrix} q_1(1 + 2 \cos^2 \varphi) + 2q_2 \sin \varphi \cos \varphi \\ 2q_1 \sin \varphi \cos \varphi + q_2(1 + 2 \sin^2 \varphi) \end{bmatrix}, \quad 3.13$$

where

$$\begin{bmatrix} r \\ \varphi \end{bmatrix} = \begin{bmatrix} \sqrt{\hat{x}^2 + \hat{y}^2} \\ \text{atan}(\hat{y}/\hat{x}) \end{bmatrix}. \quad 3.14$$

The residual radially symmetric distortion polynomial (equation 3.9) is reformulated as a function of  $r$ :

$$\hat{r} - g_o(\theta) = \sum_2^N b_n r^n, \quad 3.15$$

where  $\hat{r}$ , equivalent to its definition in the forward projection, is the radial coordinate after adjustment for decentering:

$$\hat{r} = \sqrt{(\hat{x} - \delta_{cx}(r, \varphi))^2 + (\hat{y} - \delta_{cy}(r, \varphi))^2}. \quad 3.16$$

$N$  for the back-projection is set to nine. An image point can then be back-projected using

$$\Phi = \begin{bmatrix} \theta \\ \phi \end{bmatrix} = \begin{bmatrix} g_o^{-1} \left( \hat{r} - \sum_2^N b_n r^n \right) \\ \text{atan} \left( \frac{(\hat{y} - \delta_{cy}(r, \varphi))}{(\hat{x} - \delta_{cx}(r, \varphi))} \right) \end{bmatrix}. \quad 3.17$$

where inversion of  $g_o$  from any of the options listed in equation 3.8 is straightforward.

The ray  $\Phi$  can be parameterized in the world reference frame as

$$\begin{bmatrix} X(\lambda) \\ Y(\lambda) \\ Z(\lambda) \\ 1 \end{bmatrix} = \begin{bmatrix} \lambda R^T & -R^T \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} \cos \phi \sin \theta \\ \sin \phi \sin \theta \\ \cos \theta \\ 1 \end{bmatrix}, \quad 3.18$$

where  $\lambda$  is a scalar. For sky imaging, the camera center is often considered the origin of the world coordinate system and thus  $\mathbf{t} = \mathbf{0}$ .

### 3.4 Solar Position Input from Sky Imager Data

#### 3.4.1 Imaging Equipment and Setup

The University of California, San Diego (UCSD) sky imager (USI) camera system was developed for the purpose of solar power forecasting (Urquhart et al. 2013). The camera is an Allied Vision GE-2040C camera which contains a 15.15 mm x 15.15 mm, 2048 x 2048 pixel Truesense KAI-04022 interline transfer charge coupled device (CCD). The lens is a Sigma circular fisheye lens with a 4.5 mm nominal focal length and equisolid angle projection (equation 3.8c). Images are captured every 30 seconds during daylight hours, which for this experiment yielded over 1,400 images per day. The USI uses 3 exposures at integration times of 3, 12, and 48 ms to generate a composite HDR image. The system clock is regularly updated using the network time protocol, so image capture times are accurate to within a second. Extensive details of the USI can be found in Urquhart et al. 2014.

The USI used in this work was deployed at the Department of Energy, Atmospheric Radiation Measurement (ARM) Program, Southern Great Plains (SGP) Climate Research Facility from March 11th, 2013 to November 4th, 2013 at a longitude, latitude, altitude of -97.484856°E, 36.604043°N, 318 m. The horizon around the SGP

site is free of mountainous terrain, thus the USI has a nearly  $180^\circ$  field of view of the sky. The camera nominally points straight up, but has a slight angular offset due to the ground not being perfectly level. No leveling of the equipment was performed. Figure 3.1 shows the USI on its portable mounting stand.



**Figure 3.1.** USI 1.8 in the instrument field at the Department of Energy, Atmospheric Radiation Measurement Program, Southern Great Plains Climate Research Facility.

### 3.4.2 Solar Position Modeling

The input used in calibrating the camera model (i.e. fitting the camera model parameters) is the angular position of the sun  $\Phi_s$  and the corresponding position of the sun in a sky image  $(x_s, y_s)^\top$ . The ray  $\Phi_s = (\theta_s, \phi_s)^\top$  is defined following equation 3.5, but is a function of  $\mathbf{X}_s$  instead of  $\mathbf{X}_{cam}$ , where  $\tilde{\mathbf{X}}_s$  is the vector pointing to the sun in the world coordinate system. The angular solar position  $\Phi_s$  is estimated using the NREL solar position algorithm (Reda and Andreas 2004), which adopts the procedure from Meeus (1998). The algorithm takes observer position (latitude, longitude, altitude) and time as inputs, and outputs the topocentric solar zenith angle  $\theta_s$  and topocentric solar azimuth angle  $\phi_s$ . The refractive index of the air is a function of its density (hence a function of

temperature and pressure) along the optical path, and because the atmospheric density gradient is predominantly vertical, the apparent solar zenith angle must be corrected accordingly (Brown 1964). A correction using annual averages of surface air pressure and temperature is included in the algorithm, and the default value for refraction magnitude at sunrise/sunset is used. The uncertainty on solar zenith angle reported by Reda and Andreas is  $\pm 0.0003^\circ$ , however assuming the image capture time may be one second off, the error in solar hour angle may be  $\pm 0.004^\circ$ . In comparison, for our lens a one pixel measurement uncertainty in sun position measurements corresponds to approximately  $\Delta\theta = 0.14^\circ$  at the horizon. The sun detection process, therefore, introduces significantly more error than the solar position model. For the 'full' calibration dataset (case 3, below), 82% of measurements were within the one pixel measurement uncertainty bounds.

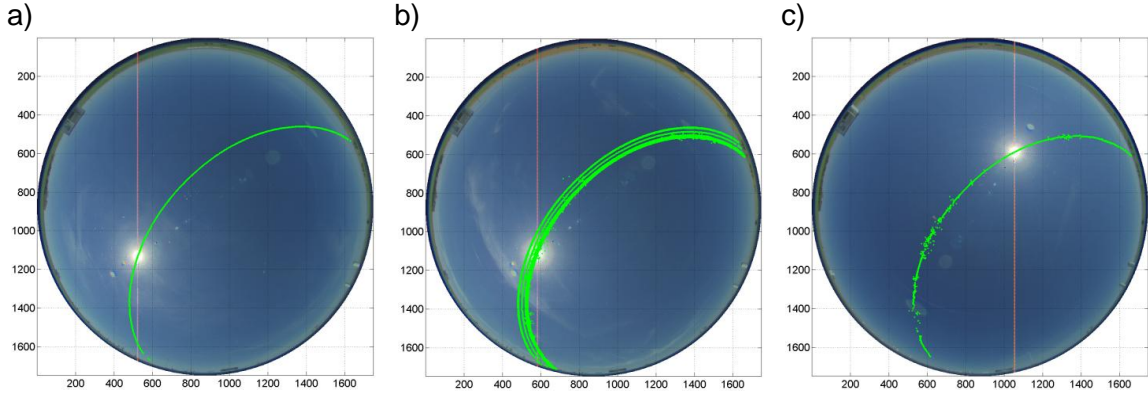
### 3.4.3 Solar Position Calibration Input

Measurement data consists of automated detection of the sun's position  $(x_s, y_s)^\top$  in an image using a set of methods described in Appendix 3.1. From the resulting set of sun coordinates  $\{(x_s, y_s)^\top\}_i$ , the median of each  $\{x_s\}_i$  and  $\{y_s\}_i$  were taken as the final sun coordinate  $(x_{s,i}, y_{s,i})^\top$  to be used for calibration. Here,  $i$  is the image index. The detection methods leverage the fact that the sun is the brightest object in a daytime sky image. For the days chosen, the sun could be seen at solar zenith angles near  $90^\circ$ , indicating that the horizon is at a similar altitude as the instrument. The sun could be detected reliably for images with  $\theta_s < 89.625^\circ$ . The sun detection algorithm described here was tested on predominantly clear days which simplifies detection because clouds cause occlusion of the sun or saturation of cloudy pixels near the sun.

The sun position is detected in a series of images collected from sunrise to sunset, yielding over 1,400 calibration points per day. The set of points collected throughout a single (clear) day nominally forms a smooth arc. To evaluate the camera model and calibration performance under different solar arc input possibilities, five input cases were tested: 1) a single solar arc, 2) two solar arcs on consecutive days, 3) four solar arcs, 4) ten solar arcs with measurement noise due to occasional clouds, 5) a single solar arc with noise due to clouds (Table 3.1). The solar arcs for cases 1, 4 and 5 are shown in Figure 3.2. Case 1 would be preferred in practice as it requires only an - admittedly perfectly clear - day of data. However, limitations in sun position availability during one day may not provide sufficient constraints for the optimization. The improvement associated with adding more days is evaluated in cases 2 and 3. Cases 4 and 5 were designed to provide more realistic and noisy data that would be found in climates without completely clear days.

**Table 3.1.** Calibration test cases. The days included in each test case are given along with the number of sun position points and an estimate of measurement standard deviation  $STD_m$  which represents the extent to which the data deviate from a smooth arc. See also Fig. 3.2 and Eq. 3.19.

	<b>Day(s)</b>	<b>points</b>	<b><math>STD_m</math> [pixels]</b>
case 1	May 13	1,582	0.4346
case 2	May 13,14	3,195	0.4313
case 3	May 13,14 June 2,11	6,543	0.6237
case 4	May 13,14,22 June 1,2,3,7,9,10,11	15,966	2.4883
case 5	June 7	1,482	5.2107



**Figure 3.2.** Solar position measurements on (a) May 13, 2013 (case 1); (b) three days in May and seven days in June 2013 with an image on June 11, 2013 (case 4); (c) June 7, 2013 (case 5). Measurements are overlaid on example images.

The sequence of sun position detections forms an arc that should be a smooth curve. The detection process, however, is associated with errors, especially when clouds are present. The deviation of the measured data from a smooth arc can be used to quantify the error. Separately for each day, a 9th order polynomial is fit to the  $x$  and  $y$  pixel coordinates as a function of solar hour angle  $H$  (obtained from the NREL solar position algorithm). A separate polynomial is obtained for  $x$  and  $y$ , which after obtaining the polynomial coefficients  $a_n$  and  $b_n$  can be written as

$$\tilde{x}_f = \sum_0^N a_n H^n, \quad \tilde{y}_f = \sum_0^N b_n H^n \quad 3.19a,b$$

where  $(\tilde{x}_f, \tilde{y}_f)$  is the pixel coordinate. The polynomials are (quite apparently) valid only for the single day being investigated. The standard deviation of the pixel-by-pixel distance between the measurements and the polynomial fit (equation 3.25c) is given in Table 3.1 as  $STD_m$  and is a useful estimate of the detection error.

### 3.5 Calibration Procedure

The calibration procedure is a three step process: 1) generate a rough estimate of the intrinsic parameters; 2) estimate the camera pose (rotation and translation) assuming one of the projections in equation 3.8; 3) perform a three stage nonlinear

parameter estimation using the Levenberg-Marquardt algorithm to obtain the final intrinsic and extrinsic parameters. Steps 1 and 2 will be described in section 3.5.1 and step 3 will be discussed in section 3.5.2. Calibration results are given in section 3.6.

### 3.5.1 Model Initialization

In order to apply the Levenberg-Marquardt (LM) algorithm to estimate the model parameters, the parameter vector  $\boldsymbol{\beta} = (\alpha_x, \alpha_y, s, \mathbf{x}_o^\top, \mathbf{w}^\top, \mathbf{t}^\top, \mathbf{k}^\top, \mathbf{p}^\top)^\top$  must be initialized with a reasonably close estimate (see section 3.3 for the definitions of the components of  $\boldsymbol{\beta}$ ).

#### 3.5.1.4 Intrinsic Parameter Estimation

In whole sky imagery, the entire sky hemisphere is visible and forms an ellipse on the image plane with eccentricity near unity (e.g. Figure 3.2 or 3.3). A Hough circle transform is used to obtain the approximate center  $\mathbf{x}_{img}^\top$  of and radius  $r_{img}$  of this near circular ellipse. The principal point  $\mathbf{x}_o^\top$  is initialized to  $\mathbf{x}_{img}^\top$ . The  $x$  and  $y$  focal lengths are assumed to be equal, i.e.  $\alpha_x = \alpha_y = \alpha$ , and are determined using an unnormalized version of equation 3.9 where  $r = \alpha g_o(\theta)$ :

$$\alpha = r_{img}/g_o(\theta_{max}), \quad 3.20$$

where the radius  $r_{img}$  from the Hough circle detection process corresponds to the maximum field of view. For the USI,  $\theta_{max}$  is taken to be  $\pi/2$  and  $g_o$  is given by equation 3.8c, thus  $\alpha = r_{img}/\sqrt{2}$ . Initially it is assumed that  $s$ ,  $\mathbf{k}$ , and  $\mathbf{p}$  are all zero, i.e. there is negligible distortion beyond the nominal equisolid angle projection, and the camera pixel axes are orthogonal. The initial estimate of the camera calibration matrix  $K_o$  is then

$$K_o = \begin{bmatrix} \alpha & 0 & x_{img} \\ 0 & \alpha & y_{img} \\ 0 & 0 & 1 \end{bmatrix}. \quad 3.21$$



### 3.5.1.5 Pose Estimation

The camera pose is estimated by computing the linear transformation between the inhomogeneous camera coordinates and the homogeneous world coordinates:

$$\tilde{\mathbf{X}}_{cam,i} = [R | \mathbf{t}] \mathbf{X}_{s,i}, \quad 3.22$$

where  $i$  is the data point index for the set of points to be used in calibration. The camera coordinates  $\tilde{\mathbf{X}}_{cam,i}$  are obtained by first computing the calibrated image coordinates  $\hat{\mathbf{x}}_i = K_o^{-1} \mathbf{x}_i$ , and then by using

$$\begin{bmatrix} \theta_{cam,i} \\ \varphi_{cam,i} \end{bmatrix} = \begin{bmatrix} g_o^{-1}(\hat{r}_i) \\ \text{atan}(\hat{y}_i/\hat{x}_i) \end{bmatrix}, \quad 3.23$$

where  $\hat{r}_i$  is given by equation 3.16 with decentering distortion set to zero, and finally by projecting onto the unit sphere:

$$\tilde{\mathbf{X}}_{cam,i} = \begin{bmatrix} X_{cam,i} \\ Y_{cam,i} \\ Z_{cam,i} \end{bmatrix} = \begin{bmatrix} \sin(\theta_{cam,i}) \cos(\varphi_{cam,i}) \\ \sin(\theta_{cam,i}) \sin(\varphi_{cam,i}) \\ \cos(\theta_{cam,i}) \end{bmatrix}. \quad 3.24$$

The calibrated perspective projection matrix  $\hat{P} = [R | \mathbf{t}]$  from equation 3.22 can be obtained using the Direct Linear Transform (DLT) algorithm. Premultiplying equation 3.22 by the left nullspace of  $\tilde{\mathbf{X}}_{cam,i}$ :

$$[\tilde{\mathbf{X}}_{cam,i}^T]^{\perp T} \hat{P} \mathbf{X}_{s,i} = \mathbf{0}$$

where  $[\cdot]^{\perp}$  denotes the nullspace of the argument. Applying the *vec* operator yields

$$\left( \mathbf{X}_{s,i}^T \otimes [\tilde{\mathbf{X}}_{cam,i}^T]^{\perp T} \right) \text{vec}(\hat{P}) = \mathbf{0}.$$

where  $\otimes$  is the Kronecker product. Our design matrix then contains subblocks  $A_i = \mathbf{X}_{s,i}^T \otimes [\tilde{\mathbf{X}}_{cam,i}^T]^{\perp T}$ . Stacking rows  $A_i$  to form a matrix  $A$  gives the homogeneous linear equation

$$A\hat{p} = \mathbf{0},$$

where  $\hat{\mathbf{p}} = \text{vec}(\hat{P})$ . Due to measurement noise, the right hand side is not identically zero. A least squares solution is obtained by computing the singular value decomposition of  $A$  and taking  $\hat{\mathbf{p}}$  as the right singular vector corresponding to the smallest singular value (Hartley and Zisserman 2004). As always with the DLT algorithm, appropriate data normalization is required (Hartley 1997).

Due to imperfect data, the left  $3 \times 3$  subblock of  $\hat{P}$  is likely not an orthogonal matrix in  $SO(3)$  as is required for rotation matrices. To obtain  $R$  and  $\mathbf{t}$  from the DLT estimate of  $\hat{P}$  we take  $\hat{P} = [M | \mathbf{v}]$  where  $M$  is the  $3 \times 3$  left subblock of  $\hat{P}$  and  $\mathbf{v}$  is the rightmost column vector. We then use singular value decomposition to write  $M = UDV^T$  where  $U$  comprises the left singular vectors and  $V$  comprises the right singular vectors of  $M$  (both  $U$  and  $V$  are orthogonal matrices), while  $D$  contains the singular values of  $M$ . An orthogonal in  $SO(3)$  is obtained by taking  $R = \mu UV^T$  where  $\mu = \text{sign}(\det(UV^T))$ . This gives the closest matrix  $R$  to  $M$  in the sense of the Frobenius norm. The translation vector  $\mathbf{t}$ , which is nominally zero here, is given by  $\mathbf{t} = -R[\tilde{\hat{P}}]^\perp$ , where the tilde indicates that after computing the nullspace, the resulting homogeneous 4-vector is converted to an inhomogeneous 3-vector before multiplication by  $R$ .

### 3.5.2 Nonlinear optimization of model parameters

The nonlinear calibration of the forward model is performed in three successive stages: 1) take  $\mathbf{k} = \mathbf{0}$  and  $\mathbf{p} = \mathbf{0}$ , i.e. do not include residual radial distortion and decentering distortion; 2) include residual radial distortion terms  $\mathbf{k}$ , but take  $\mathbf{p} = \mathbf{0}$ ; 3) include both residual radial distortion  $\mathbf{k}$  and decentering distortion  $\mathbf{p}$ . Three stages of nonlinear optimization were used because it was found that this approach was more consistent across the different test cases. The multi-stage optimization process first fits the most basic model parameters, and additional degrees of complexity are sequentially

added (i.e. radial followed by decentering distortion). The motivation in doing so is to avoid local minima that would result in errors in the estimation of the basic parameters  $(\alpha_x, \alpha_y, s, \mathbf{x}_o^\top, \mathbf{w}^\top, \mathbf{t}^\top)$ .

The model fitting for each stage is accomplished by using the Levenberg-Marquardt (LM) algorithm, for which an excellent introduction is given in Hartley and Zisserman (2004). Initial intrinsic and extrinsic parameter estimates for the first calibration stage are those determined in section 3.5.1, and the subsequent stages are initialized with the results of the previous stage. The model is fit by minimizing the sum of squared distances of the measured and modeled solar inhomogeneous pixel coordinates:

$$\sum_i^N d(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_{s,i})^2,$$

where  $d$  is the Euclidean distance function. In the case of the synthetic data of section 3.6.3, the sum of squared distances is taken for the ground truth data with noise added  $\tilde{\tilde{\mathbf{x}}}_i + \tilde{\delta}_i$  and the modeled data  $\tilde{\mathbf{x}}_i$  (i.e.  $\sum_i^N d(\tilde{\mathbf{x}}_i, \tilde{\tilde{\mathbf{x}}}_i + \tilde{\delta}_i)^2$ ).

### **Calibration Constraints**

It was found necessary to enforce additional constraints in the model fitting process to ensure consistent and physically significant results. The first constraint

$$\int_0^{\theta_{max}} g_o(\theta) \sum_2^N k_n \theta^n d\theta = \sum_2^N k_n \int_0^{\theta_{max}} g_o(\theta) \theta^n d\theta = 0,$$

ensures that the residual and nominal radially symmetric distortion are orthogonal functions over the field of view. Without this constraint, the LM algorithm tended to decrease  $\alpha_x, \alpha_y$  and increase the  $k_n$  to compensate, leaving the focal lengths at values that were obviously incorrect based on the nominal lens and sensor specifications. This

constraint is very important if the formulation in equation 3.9 is to be used for the radially symmetric distortion.

The specific shape of the solar arc used to calibrate the camera, particularly when only a single day was used, resulted in a falsely large skewness  $s$ . This was corrected by applying a penalty on deviations from circularity of the ellipse formed at  $\theta_{max}$  parameterized by varying the azimuth angle. A simple metric such as the standard deviation of the radius of the ellipse at different azimuth angles, taken from the center  $\mathbf{x}_o^T$  is simple and effective for this purpose. A similar and simpler approach would be to place a penalty that is proportional to  $|s|$ , however this was not tested in this work. The last constraint applied was that  $\hat{r}(\theta)$  was forced to be monotonically increasing with  $\theta$  (the lens mapping would not be one-to-one if it was not!) by applying a penalty if  $d\hat{r}(\theta)/d\theta < 0$ .

### 3.5.3 Backward projection

The calibration of backwards projection model parameters was performed with a single stage. The parameter vector used in the LM algorithm consisted only of  $b_n$  and  $q_n$ . The constraints were found to be unnecessary because the process involves fitting only the residual radially symmetric and decentering distortion. The focal lengths  $\alpha_x$  and  $\alpha_y$  are already set, thus the orthogonality constraint is not required. The other two constraints treat the specific shape of the solar arc, and the back projection parameters are fit using synthetic data points generated from the forward projection which cover the whole image, and therefore are not required. To initialize LM for back-projection fitting, coefficients  $b_n$  can be set to the  $k_n$  obtained in the forward projection. It was found empirically that  $k_n$  are very close to  $b_n$  for the equisolid angle lens used on the USI, and should be even closer if an equidistant lens is used.

### 3.6 Calibration Results

#### 3.6.1 Calibration Performance Metrics

The root mean square error (*RMSE*), mean absolute error (*MAE*) and standard deviation (*STD*) are computed as

$$RMSE = \left[ \frac{1}{M} \sum_{i=1}^M \mathbf{dx}_i^\top \mathbf{dx}_i \right]^{1/2}, \quad 3.25a$$

$$MAE = \frac{1}{M} \sum_{i=1}^M |\mathbf{dx}_i|, \quad 3.25b$$

$$STD = \left[ \frac{1}{M} \sum_{i=1}^M (\mathbf{dx}_i - \overline{\mathbf{dx}})^2 \right]^{1/2}, \quad 3.25c$$

where the total number of measurements is  $M$ , and  $\mathbf{dx}_i = \tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_{s,i}$  is the distance vector from the measured solar position  $\mathbf{x}_{s,i}$  to the modeled point  $\mathbf{x}_i$ . The vertical bars  $|\cdot|$  denote the 2-norm of the argument, and  $\overline{\mathbf{dx}}$  is the mean displacement vector for all points  $i$ . These definitions hold for the evaluation of measurement error as well, where instead  $\mathbf{dx}_i = \tilde{\mathbf{x}}_{s,i} - \tilde{\mathbf{x}}_{f,i}$  (see section 3.4.3 for description of the polynomial fit  $\tilde{\mathbf{x}}_{f,i}$ ).

#### 3.6.2 Calibration using the Solar Position

The results of calibrating the USI for the five different solar arc cases is shown in Table 3.2. In the cases using more than one solar arc (cases 2-4), the principle point is consistent to within 0.60 pixels (4.4  $\mu\text{m}$ ). The  $x$  and  $y$  focal lengths are consistent to within 0.77 pixels (5.7  $\mu\text{m}$ ) for all cases and consistent to within 0.29 pixels (2.1  $\mu\text{m}$ ) for cases 2-4. The camera pose results presented here are represented by three angles in Table 3.2:  $\phi_{xz}$  is the angle of rotation of the camera  $X_{cam}$ -axis from the world  $X$ -axis about the world  $Z$ -axis (effectively the instrument's rotation from a northern alignment);  $\theta_{zz}$  is the angle between the camera  $Z_{cam}$  and world  $Z$ -axis (i.e. the degree to which the system is tilted); and  $\psi_z$  is the azimuthal direction towards which the  $Z_{cam}$ -axis is tilted.

The pose determined in all cases was very consistent, with a maximum difference in  $\theta_{zz}$  and  $\phi_{xz}$  of 0.2 degrees. Because the tilt angle  $\theta_{zz}$  was very small, there was increased variability in the tilt direction  $\psi_z$  which is to be expected. Anecdotal observations of the USI 1.8 system as deployed at the SGP site indicate that it was tilted slightly southwest, which is supported by the estimated pose.

**Table 3.2.** Camera model parameters (excluding distortion terms) determined from the five test cases of solar input data. The mean and standard deviation are also given. The units denoted [pixels/f] is pixels per focal length.

	$\alpha_x$ [pixels/f]	$\alpha_y$ [pixels/f]	$s$ [pixels/f]	$x_o$ [pixels]	$y_o$ [pixels]	$\phi_{xz}$ [deg.]	$\theta_{zz}$ [deg.]	$\psi_z$ [deg.]
case 1	601.44	601.44	$-1.94 \cdot 10^{-3}$	873.24	881.62	47.89	2.58	136.51
case 2	601.40	601.40	$1.72 \cdot 10^{-3}$	871.48	882.97	47.90	2.47	141.48
case 3	601.32	601.32	$2.57 \cdot 10^{-3}$	871.96	883.00	47.91	2.46	140.16
case 4	601.12	601.11	$4.14 \cdot 10^{-3}$	871.97	883.31	47.93	2.43	140.06
case 5	601.88	601.88	$2.53 \cdot 10^{-3}$	874.20	880.84	47.90	2.63	134.02
mean	601.43	601.43	$1.80 \cdot 10^{-3}$	872.57	882.35	47.91	2.51	138.45
std	0.25	0.25	$2.03 \cdot 10^{-3}$	1.00	0.95	0.01	0.06	2.76

The performance of camera calibration using solar position is given in Table 3.3 along with the estimated measurement error of the sun position (subscript 'm' in Table 3.3). While not a true lower bound on calibration accuracy, the measurement errors given here can be used to assess the calibration accuracy relative to the estimated accuracy of the input data. The polynomial fit to the measurement data (equation 3.19) does not have the same constraints as fitting the camera model parameters to the measurement data, thus the measurement standard deviation ( $STD_m$ ) and root mean square difference ( $RMSD_m$ ) are lower than the  $STD$  and  $RMSE$  obtained for camera calibration, with  $RMSE/RMSD_m$  lying between 1.8% to 24%. This proportionality of calibration and measurement error, along with the consistency of the parameter estimation indicates that the camera model presented here reasonably approximates the

imaging process for the camera tested. It also indicates that the calibration procedure is consistently obtaining reasonable parameter estimates for the model used. Additionally, the robustness of the model and calibration to larger measurement errors and outliers is demonstrated in case 5.


**Table 3.3.** Calibration error metrics for each case: root mean square error (RMSE); mean absolute error (MAE); standard deviation (STD); and measurement root mean square difference (RMSD<sub>m</sub>) and standard deviation (STD<sub>m</sub>).

	RMSE		MAE		STD		RMSD <sub>m</sub>		STD <sub>m</sub>	
	[pixels]	[ $\mu\text{m}$ ]	[pixels]	[ $\mu\text{m}$ ]	[pixels]	[ $\mu\text{m}$ ]	[pixels]	[ $\mu\text{m}$ ]	[pixels]	[ $\mu\text{m}$ ]
case 1	0.9370	6.931	0.7775	5.752	0.5229	3.868	0.8099	5.991	0.4346	3.215
case 2	0.9635	7.128	0.8089	5.984	0.5235	3.873	0.7802	5.772	0.4313	3.191
case 3	1.2381	9.159	1.0241	7.576	0.6956	5.146	1.0538	7.795	0.6237	4.614
case 4	2.9351	21.712	1.4852	10.987	2.5316	18.727	2.8080	20.772	2.4883	18.407
case 5	6.2994	46.510	3.4777	25.726	5.2525	38.855	6.1871	45.769	5.2107	38.546

### 3.6.3 Camera Model Parameter Uncertainty

As with any image detection algorithm, there are errors in the position of the sun obtained from the detection algorithm (Tables 3.1 and 3.3). Depending on the content of each image, such as the possibility of thin clouds veiling a still visible sun, or more opaque clouds passing near or occluding the sun, the magnitude of the detection error will vary. A Monte Carlo method was used to assess the uncertainty in model performance and parameter estimation as a function of measurement error. A ground truth synthetic calibration dataset was constructed with  $Q = 1673$  data points by simulating a single solar arc on May 13, 2013 (Figure 3.3). The points in the world coordinate system  $\{\mathbf{X}_i = (X_i, Y_i, Z_i, 1)^T \in \mathbb{P}^3, i = 1 \dots Q\}$  were obtained by computing solar position every 30 seconds from sunrise to sunset with the NREL solar position algorithm (Reda and Andreas 2004), and then by projecting the solar position onto the unit sphere centered at the camera center. The ground truth pixel coordinates  $\{\bar{\mathbf{x}}_i =$

$(\bar{x}_i, \bar{y}_i, 1)^\top \in \mathbb{P}^2, i = 1 \dots Q$  were obtained by applying the forward camera model (section 3.3.1) to points  $\mathbf{X}_i$ . The ground truth camera model parameters were set according to the results from solar calibration case 3. The points  $\mathbf{X}_i$  were treated as known points in space corresponding to synthetic measurements  $\bar{\mathbf{x}}_i + \boldsymbol{\delta}_i$ , where  $\boldsymbol{\delta}_i$  is the measurement error generated as follows. The points  $\bar{\mathbf{x}}_i$  were taken as the mean measurement values for  $Q \times S$  independent normal probability distributions  $P_{ij}(\bar{\mathbf{x}}_i)$  with standard deviations  $\sigma_j$  where  $j = 1 \dots S$ . Standard deviation  $\sigma_j$  was varied from 0 to 10 pixels in steps of 0.25 pixels (thus  $S = 41$ ). The point measurements  $(\bar{\mathbf{x}}_i + \boldsymbol{\delta}_i)_{kj}$  used in calibration trial  $k, j$  were obtained by sampling  $P_{ij}(\bar{\mathbf{x}}_i)$ . A number of trials  $N_t = 1000$  was performed for each  $j$ , yielding  $N_t \times S$  calibration trials. For the  $k^{\text{th}}$  trial at error level  $\sigma_j$ , a set of model parameters  $\{\boldsymbol{\beta}_{kj} = (\alpha_x, \alpha_y, s, \mathbf{x}_o^\top, \mathbf{w}^\top, \mathbf{t}^\top, \mathbf{k}^\top, \mathbf{p}^\top)_{kj}^\top, k = 1 \dots N_t, j = 1 \dots S\}$  is obtained. The distribution of  $\boldsymbol{\beta}_{kj}$  at each  $j$  (i.e. along dimension  $k$ ) is a measure of the uncertainty in the model parameters at error level  $\sigma_j$ .



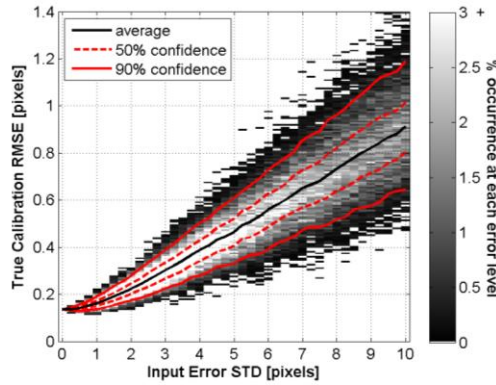
**Figure 3.3.** Synthetic dataset point distribution. The 1673 points are generated from taking the solar position every 30 seconds from sunrise to sunset on May 13, 2013, and projecting onto the image plane using a set of ground truth camera model parameters. Background image (for visual reference only) is from May 3, 2013.



The distribution of true root mean square calibration error ( $RMSE$ ) for  $\sigma_j \in [0,10]$  pixels is shown in Figure 3.4.  $RMSE$  is computed as

$$RMSE_{kj} = \frac{1}{Q} \sum_{i=1}^Q (\mathbf{x}_{i,kj} - \bar{\mathbf{x}}_i)^2, \quad 3.26$$

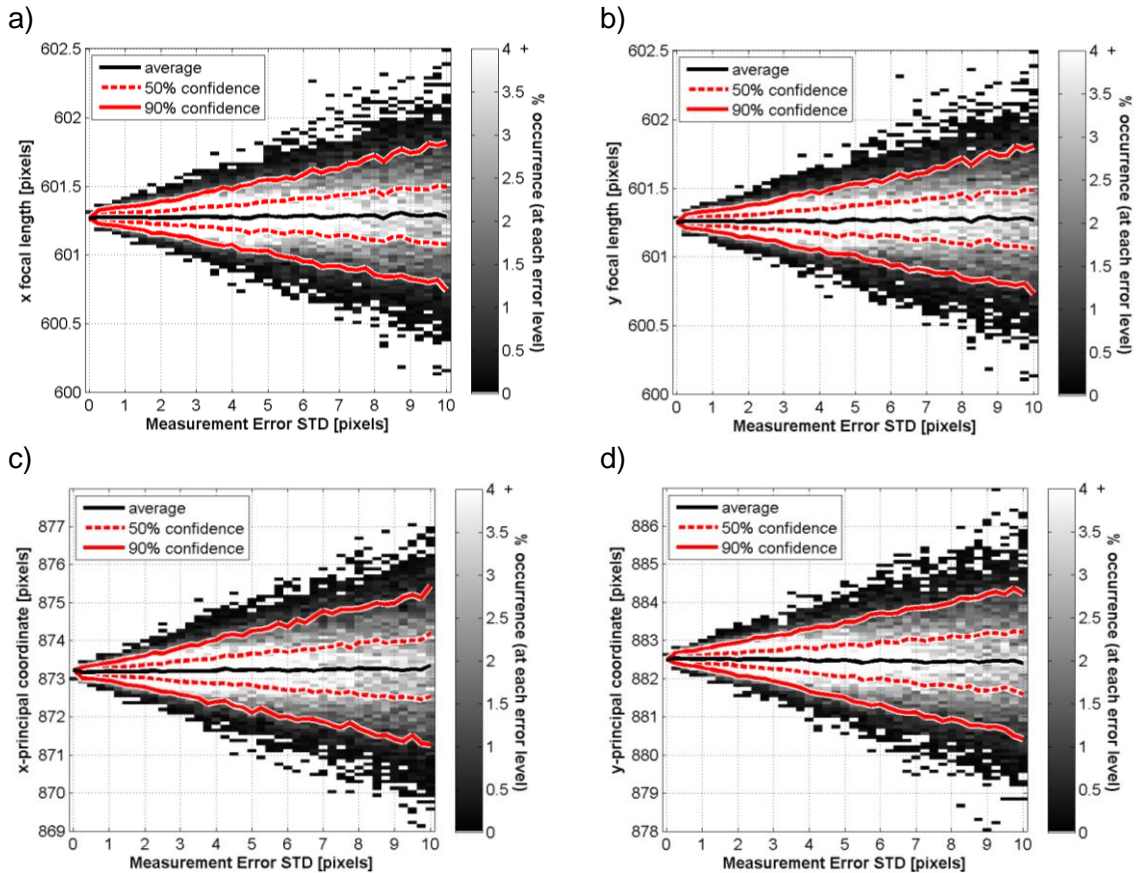
where  $\bar{\mathbf{x}}_i$  is the ground truth pixel position of the  $i^{\text{th}}$  point (as described above), and  $\mathbf{x}_{i,kj}$  is the modeled pixel position of the  $i^{\text{th}}$  point (i.e. the projected pixel position of  $\mathbf{X}_i$ ) for calibration trial  $k$  at error level  $\sigma_j$ . Even for  $\sigma_j = 10$  pixels, the average  $RMSE$  is below 1 pixel. For reference, the measurement error standard deviation in Table 3.1 for clear days is less than 0.75 pixels, and the worst case tested here (Case 5) has a measurement standard deviation of 5.21 pixels. Based on Fig 3.4, these measurement errors correspond to true errors of  $0.14 \pm 0.03$  pixels and  $0.37 \pm 0.11$  pixels (mean  $\pm$  90% confidence interval), respectively. This assumes measurement errors are normally distributed.



**Figure 3.4.** Root mean square calibration error distribution (Eq. 3.25a) as a function of simulated measurement error standard deviation  $\sigma_j$ . The mean, 50% and 90% confidence intervals are shown as curves.

Distributions of parameter estimation for four of the intrinsic parameters are shown in Figure 3.5. For both  $\alpha_x$  and  $\alpha_y$ , the 90% uncertainty bounds are nearly linear and approximately follow  $\alpha_{p90} = \bar{\alpha} \pm 0.05 STD_m$ , which is about 0.09% error at  $STD_m = 10$  pixels. The overbar indicates the mean value. Similar results hold for the 90%

uncertainty bounds of  $(x_o, y_o)$ , which follow  $\mathbf{x}_{o,p90} = \bar{\mathbf{x}}_o \pm 0.25 \text{ STD}_m$  and gives 0.29% error at  $\text{STD}_m = 10$  pixels. The latter error percentage is computed using  $0.25 \text{ STD}_m / N_p \cdot 100\%$ , where  $N_p = 874$  pixels is the radius of the usable sky image circle (Fig. 3.3) For the application of solar forecasting using sky imagery, these error levels are satisfactory.



**Figure 3.5.** Distributions of (a)  $x$ -focal length  $\alpha_x$ ; (b)  $y$ -focal length  $\alpha_y$ ; and (c), (d) principal point  $(x_o, y_o)$  are shown as a function of measurement error standard deviation  $\sigma_j$ . The mean, 50% and 90% confidence intervals are shown as curves.

### 3.7 Conclusions

The increasing use of stationary daytime sky imagery instruments for solar forecasting applications has motivated the need to develop automatic geometric camera calibration methods and an associated general camera model. The camera model presented is not specific to fisheye lenses, and is generally applicable to most dioptic

camera systems with a photo objective lens. We have proposed a method to automatically detect and use the sun position over a sequence of images to calibrate the proposed camera model. Calibration performance on clear days ranged from 0.94 to 1.24 pixel root mean square error (RMSE). An uncertainty analysis indicated that if measurement errors are normally distributed, this corresponds to a true calibration error of  $0.14 \pm 0.03$  to  $0.16 \pm 0.03$  pixels RMSE ( $0.07 \pm 0.02$  to  $0.08 \pm 0.02$  pixels STD), respectively. A back-projection model, which may be more useful for many applications, is proposed as a straightforward extension of the forward projection model. The uncertainty in the forward model parameters was analyzed and is provided graphically as a function of solar position measurement error.

### **Acknowledgements**

In the course of preparing this chapter, the following publications were used:

Publication: Sky camera geometric calibration using solar observations. IEEE Transactions on Pattern Analysis and Machine Intelligence, in submission.

Co-Authors: Ben Kurtz and Jan Kleissl

### **Appendix 3.1 Sun Position Detection**

The sun is only detected for images with solar zenith angles  $\theta_s < 89.625^\circ$ . For approximately  $\theta_s < 87^\circ$ , the pixels surrounding the sun's location saturate for the USI camera. For the purposes of image detection, this saturated region, which is larger than the sun itself, will be referred to as the 'sun' when discussing the image of the daytime sky. The USI uses 3 exposures at different integration times to generate a composite HDR image which reduces the number of saturated pixels encompassing the sun. When

the sky is clear, the sun is the only saturated object in the sky which simplifies its detection. With the lens used on the USI, the sun appears as a nearly circular ellipse. The high intensity and near circularity of the sun along with the vertical smear stripe (occurring in columns containing the sun) are the primary image features used in the sun position detection process.

Both the red-green-blue (RGB) and hue-saturation-value (HSV) color spaces were used for detection, and each color image matrix will be referred to as an X-image, e.g. the R-image (the red image). The approximate diameter of the sun is detected by constructing a binary image by thresholding the V-image at the 99.99th percentile, and then performing an erosion and dilation to remove noise. The diameter of the largest connected binary entity is taken. The apparent sun diameter changes with solar zenith angle, and this size metric is used in constructing detection filters. Three filters are then constructed and subsequently convolved with the V-image: 1) a binary circular kernel; 2) a Gaussian kernel; and 3) a modified Gaussian kernel which has a flattened top. The standard deviation  $\sigma$  (in units of pixels) used for constructing the Gaussian kernels is

$$\sigma = \min(3 \ln(91 - \theta_s)^2 + 1, 24),$$

which was obtained empirically for our camera. Kernel 3 was 'flattened' such that the circular flat top of the Gaussian was the diameter of the sun. For each kernel, the row and column of the maximum value of the convolution image was taken to be the solar position.

The columns containing the vertical smear (Fig. 3.2) are detected by extracting the first row of the V-image and the sum of the first row for the R,G and B-images (i.e. three times the first row of the equal weight grayscale image). A measure of the local mean is subtracted from each row separately using a 100 pixel moving average filter. The corresponding pixels from the modified V and gray scale rows are multiplied which

gives a very strong peak at the smear column which is taken as the column of the sun. A sub image extracted from the original image consisting of the set of columns surrounding the sun column ( $\sim 10\sigma$  columns) is used for further sun position detection. A Förstner circle detector (Förstner and Gülch 1987) is applied to this sub-image with a window size of  $7.5\sigma$  columns and the resulting maximum minor eigenvalue is taken as the sun location.

The detection processes described yield four row-column pairs (three from the circular kernel convolutions and a fourth from the Förstner operator), and the detection of maximum smear gives a fifth column estimate for a total of 4 detected rows and 5 detected columns. The median of the row and column position is taken as the true position. Generally these methods are consistent to within 3 pixels. The detection process would be simpler and more accurate if the camera had been set up to take very short (microsecond) exposures because the saturated sun region would be only a few pixels in diameter instead of 10s of pixels. This was not available for the calibration set used here, but is strongly recommended for operational autocalibration when using the solar position as calibration input.

It should be noted that the sun detection method is purely empirical and was not designed to have the fastest performance. In practice, any reasonable algorithm can be used for the sun position detection. If the position errors are zero mean and normally distributed, then the uncertainty analysis in section 3.6.3 can be used as a guide for expectations of calibration accuracy. The detection method described here is one of many that can be used, and the authors expect that other superior algorithms could be constructed. Since small calibration errors were obtained, the present algorithm is sufficient to demonstrate the calibration methodology.

## 4 Remote Sensing of Clouds: Detection and Velocity Estimation

Contributing Authors:

B. Urquhart, C.W. Chow, M. Ghonima, J Shields and J. Kleissl

### 4.1 Introduction

The hardware used for sky imaging was developed during the course of the 20<sup>th</sup> century. The use of refraction at an air-water interface to describe the view from within a pond has been attributed to RW Wood in his text *Physical Optics* (1905) where he described how the entire 180° field of view could be seen within a 97° cone under water. One year later, WN Bond (1906) coined the term "fish-eye" view in a paper describing vision underwater and experimentation with an apparatus made from a water-filled lard bucket with photographic film placed at the bottom. He took the first nearly 180° photograph of the sky using this apparatus. Wood is often incorrectly given the credit of coining the term fish-eye because he used it, following Bond, in the second edition of *Physical Optics* (1911) where he described a more compact water-filled camera. Photos from Bond's 1906 paper appeared five years later in *Physical Optics* (1911). In 1922, Bond used a hemispherical glass lens with a pinhole to create wide angle images and suggested their use for sky photography and presented the first published imagery results on the subject. The primary drawback to the lens is that it could only be focused on a band of the sky (constant zenith). The first design of a true modern fish-eye lens is attributed to Robin Hill (1924) where he used a large negative meniscus front element in the lens which allowed the full sky to be in focus. A caveat was that bandpass filters had to be used to restrict the wavelength range because of blurring due to chromatic aberration at high zenith angles caused by dispersion of colors (variation in refractive index with wavelength). In 1964, Kenro Miyamoto introduced a doublet that corrected the

dispersion and thus reduced issues due to chromatic aberration within the visible wavelength range. It is with this final development that sharp full color images of the sky could be obtained.

The light gathering element of an optical system is called the objective, which for sky photography is a camera lens. The lens accepts light from the scene of interest, which in this case is the sky, and through refraction it is redirected onto the image plane of the sensor. A standard camera lens is typically a grouping of multiple elements (e.g. 8-12 elements) to form a lens system. To image the entire sky within a single frame a wide angle lens with a 180° field of view (FOV) must be used. Lenses that meet this FOV requirement, conventionally called fisheye lenses, use different projection systems to redirect light onto the sensor (compared to the rectilinear projection of smaller FOV lenses). The most common projection is the equidistant projection where the incident zenith angle is linearly proportional to radial distance of the pixel measuring the light. This projection provides a higher resolution at the horizon because each pixel will subtend a smaller solid angle as radial location on the sensor increases. The equisolid angle projection is the second most common approach which attempts to maintain consistency of the subtended solid angle of each pixel. Mathematical descriptions of each lens are given in section 3.3.1.2.

The development of a refractive lens for capturing full sky imaging opened up many research fields, including canopy research (e.g. The Canopy Camera developed by Harry E. Brown, 1962), and daylighting research (Shahriar et al. 2009). The development of digital systems utilizing computers and semiconductor sensors began in the 1980s. Some work was done by forestry community for canopy research (Chazdon and Field 1987), and in parallel the Marine Physical Laboratory (MPL) at the Scripps Institute of Oceanography (SIO) was developing a system designed to image clouds

(Johnson et al. 1988, 1989). This system and several other notable developments following it are described in chapter 2.

At the heart of using a sky imager to as a means to generate power forecasts is a retrieval of the cloud field configuration from ground-based imaging devices. Once the position and motion are determined, the future positions can be estimated. This chapter focuses on visible wavelength passive remote sensing from the ground to detect clouds (section 4.2), and determine their velocity (4.3). A discussion of cloud height estimation is deferred until Chapter 5.

## **4.2 Cloud Detection and Opacity Classification**

Sky imagers provide a visual measurement of a whole sky dome with high temporal and spatial resolution. Spatial resolution is determined by the camera's ability to resolve a single picture element (pixel), the solid angle subtended by this pixel, and the distance of the cloud from the imager. Temporal resolution is dictated primarily by the operational ability of the image processing algorithm. In order to predict cloud positions in the 0-20 minute forecast horizon, current cloud locations, and specifically cloudy pixels, in the image must first be detected accurately. A review of image processing techniques in the literature is presented in section 4.2.1 along with basic concepts of detection. In section 4.2.4 a brief overview of the cloud detection and opacity classification method (Ghonima, Urquhart et al. 2012) developed at UCSD is provided. An overview of cloud type classification is given in section 4.2.5.

### **4.2.1 Review of Cloud Detection Methodologies**

After the development of digital sky photography and its application for atmospheric research, image processing techniques began to develop for better detection of clouds. In order to detect clouds in digital images, the ratio of the red channel to the blue




channel of an image, better known as the red-blue ratio (RBR), is used (Johnson et al. 1989, 1991; Shields et al. 1993, 2007, 2009). Use of the RBR takes advantage of a fundamental difference in scattering by clouds versus a clear sky: molecular scattering in the clear sky has a strong wavelength dependence where shorter wavelengths are scattered more heavily resulting in an observable blue color; scattering by clouds whose particles are much larger is nearly uniform across the visible spectrum and results in a gray color. Taking a spectral ratio at the extremes of the visible provides an image with high contrast between clear sky and cloud. The RBR of cloudy pixels in a digital image are close to one whereas for clear pixels the ratio is close to 0.5. By characterizing the typical RBR of clear and cloudy pixels, it is possible to segment images into two states, clear or cloudy, by thresholding the RBR image. The result is a binary mapping of sky condition.

The RBR method utilizes the RGB color space, but techniques for pixel classification involving other color spaces have been developed as well. With the hue, saturation, value (HSV) color space, the saturation channel of the image can be used for cloud detection. Value is a measure of total brightness, hue is a measurement of the spectral content (i.e. the color), and saturation ( $S$ ) describes color “purity”, and is expressed as:

$$S = 1 - \frac{3}{(R + G + B)} [\min(R, G, B)] \quad 4.1$$

Clear skies, which scatter much less red light than blue, have a low  $\min(R, G, B)$  which causes the saturation value to be higher which indicates that the clear sky color is “pure”. Clouds have similar red, green, and blue content and thus saturation is low and the color is not as “pure”. Martins et al. (2003) and Souza-Echer et al. (2006) used the saturation value of digital images for cloud detection. Image pixels were classified into clear or cloudy if they fell within three standard deviations of the mean for each class.

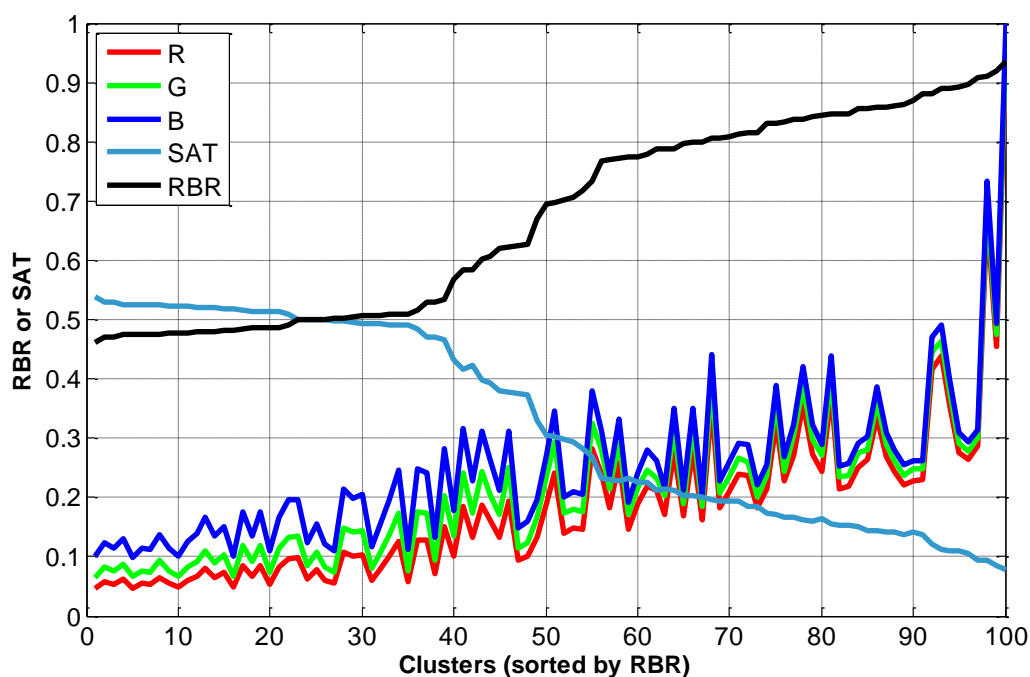
Using a k-means clustering technique, it can be shown that RBR and saturation have an approximately inverse monotonic relationship. Figure 4.1 shows the results for a single image where the clustering process where the number of clusters was set to 100 and the 10 iterations were performed. Figure 4.1d shows clusters which have a mean RBR exceeding 0.62 which correspond well with cloudy regions of the image.



**Figure 4.1.** (a) raw USI 1.7 HDR image 20130503T193330 rescaled for display; (b) RBR image from original HDR image (not rescaled image); (c) clustering with 100 clusters ; (d) clusters whose RBR mean is less than 0.62.

Figure 4.2 shows the mean RBR, saturation, and normalized intensity of each cluster (intensity was normalized by the max observed for plotting). A distinct step in

RBR (and saturation) can be attributed primarily to the clear/cloudy bimodality of RBR which is leveraged for cloud detection. There are three noticeable regions, a flat initial region of low RBR in the clear sky parts of the image, a region of rapid change, which is between clear, thin and thick clouds, and finally 'flatter' region of higher RBR, attributable thick clouds. It should be noted that the image selected does not have thick 'dark' clouds with low intensity (a case which is often observed in other images). In cases of dark thick clouds, RBR tends to be much lower and approaches that of clear sky.



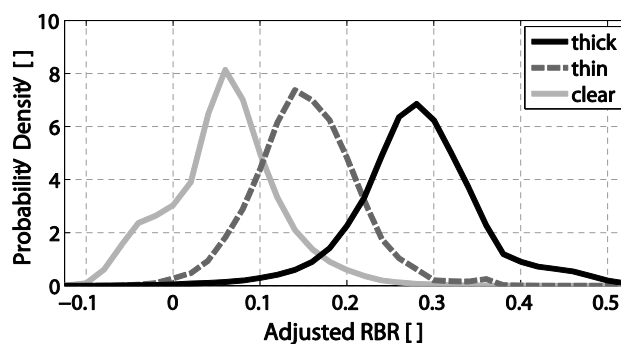
**Figure 4.2.** Characteristics of clusters. Clusters have been sorted in ascending RBR order, nominally placing the clear sky on the left and thicker clouds on the right of the plot.

Neural networks have been used by Cazorla et al. (2008a, 2008b, 2009) for cloud detection in all-sky-images. Multiple image parameters such as the red channel magnitude, blue-green ratio, RBR, etc. were input into a neural network to classify pixels into clear skies, thin and thick clouds. Another image processing technique utilizes the Euclidean geometric distance and pattern statistical analysis to classify cloudy and clear pixels (Neto et al. 2010). A combination of both fixed and adaptive thresholding methods for

cloud detection was presented by Li and Yang (2011). In this method the RBR of all-sky digital images is first obtained, next the image is classified as either unimodal, a sky image that is predominately clear or cloudy, or bimodal, a sky images that contains a mixture of clear and cloudy pixels. Unimodal images are then classified according to the fixed thresholding technique. Bimodal images, on the other hand, are classified based on the minimum cross entropy method. A comprehensive review of cloud detection methodologies can be found in Tapakis and Charalambides (2013).

#### 4.2.2 Clear Sky Library

One major drawback of a fixed threshold RBR method is that it is frequently unable to distinguish between thin clouds and clear sky. Figure 4.3 shows histograms of the RBR of clear, thick cloud, and thin cloud states from a set of 60 manually annotated images captured using the TSI system (section 2.2.2). Significant overlap of the thin cloud histogram can be seen with both clear and thick cloud histograms indicating a single threshold is problematic.




**Figure 4.3.** Histograms of clear sky, thin cloud, and thick cloud generated from a set of 60 manually annotated images captured by the TSI. These histograms are assessed to select red-blue ratio detection thresholds. Adjusted RBR is  $RBR - CSL RBR$ .

Cloud droplets, ice particles and aerosols particles, such as dust or sea salt have nearly *spectrally* uniform scattered intensities within the visible spectrum. This is due to the size distribution of these particles relative to the wavelength of incident visible light -

the interaction of light with these particular particles shows a weaker wavelength dependence than scattering by molecules (McCartney, 1976). This acts to increase the relative red content of a pixel over that observed in clean, clear skies. The increase in RBR is attributable to the number concentration of the particles: up to an optical depth of about 5, the RBR increases with increasing particle number concentration. Thin clouds, like aerosols, have lower number concentrations and the RBR lies somewhere between that of a clean, clear sky and thick clouds. Additionally, the RBR in the whole sky hemisphere in clear conditions is not *spatially* uniform due to aerosol and airmass effects. The spatial pattern of scattering by aerosols is apparent near the horizon where relative optical airmass increases and more total aerosols occupy the optical path through the atmosphere. In the circumsolar region the forward scattering effects of aerosols are also apparent as a whiter region around the sun. A final complication is that, in general, different instruments will vary in measured RBR distribution throughout an image due to specific measurement hardware characteristics and the level and type of instrument calibration. Specifically, this may include impacts by such factors as scattering off the front optical surface. With some instruments a single RBR threshold can be used effectively for opaque clouds, but not for thin clouds. These phenomenon render a single RBR threshold for the image prone to significant detection error. Specifically, for the TSI440 and the USI we found that a single RBR threshold for the entire image is not appropriate. To address the issue of spatial uniformity in a clear sky, a Clear Sky Library (CSL based on Shields et al., 2009) was developed which serves to provide a background of expected clear sky conditions (Chow, Urquhart et al. 2011). Our implementation of the CSL was compiled as a sun-pixel-angle by zenith angle lookup table from a full day of images. In Figure 4.4a, sky imager zenith angles larger than  $75^\circ$  correspond to the horizon, and small sun-pixel-angles (approximately  $<35^\circ$ ) correspond


to the circumsolar region. It can be seen in Figure 4.4a that clear pixels away from the sun and horizon have rather uniform RBR. Decreasing the distance from the sun increases RBR because the forward scattering of sunlight turns the circumsolar region whiter and brighter. Near the horizon, the higher aerosol concentration and optical depth (due to the airmass effect) causes increased Mie scattering and thus also acts to increase RBR. Figure 4.4b presents the standard deviation of the clear sky RBR to quantify the RBR variation during the (clear) day in Figure 4.4a. Larger variation of RBR leads to less accurate cloud decisions since it makes it more likely that a clear sky RBR would be larger than the threshold and classified as cloudy. For sun-pixel-angles greater than  $50^\circ$  and sky imager zenith angle less than  $60^\circ$  (i.e. in the region with the highest occurrence, Figure 4.4c) the variation in the clear sky RBR is small and our method performs well. However, the standard deviation of the RBR in the circumsolar region is found to be largest when the sun is near the horizon (low sun-pixel-angle and high zenith angle). Anecdotal observations confirm that clear conditions in the circumsolar region at low sun elevation are often misclassified as cloudy while the horizon and circumsolar region appear white. In addition, at larger solar zenith angles less sunlight reaches the TSI due to the increased optical path length through the atmosphere.



**Figure 4.4.** (a) Average RBR (colorbar) as a function of zenith angle and sun-pixel-angle based on TSI imagery on a clear day (September 24, 2009); (b) Standard deviation (colorbar) of clear sky RBR; and (c) number of occurrence (colorbar) for each zenith angle and sun-pixel-angle pair.

### 4.2.3 Cloud Detection with the CSL

Nominally, a pixel is considered cloudy if its RBR is above a threshold value determined from instrument characterization (e.g. Figure 4.3). To use the CSL for improving cloud detection accuracy, a clear sky background image is generated for each sky image based on the current solar zenith angle (Figure 4.5c) by extracting the RBR for each pixel from the lookup table shown in Figure 4.4. The CSL threshold is defined to be the RBR in the clear sky background image plus an additional constant threshold value. A pixel is classified as cloudy if its RBR (Fig. 4.5b) is larger than the CSL threshold. In general, the method using the CSL is able to detect bright opaque clouds accurately. However, within the circumsolar region, thick dark clouds cannot be identified since they have a lower RBR than the CSL threshold. This is still an issue at the time of writing and active research is engaged at improving detection in this troublesome region.



**Figure 4.5.** Processing chain of a sky image on October 4, 2009 15:45:30 PST (a) to obtain the cloud decision image (d). The dotted black lines show the borders of the circumsolar region defined here as solar azimuth  $\pm 35^\circ$ . (b) RBR image and (c) clear sky RBR background image plus the threshold. Pixels in (b) with  $RBR > CSL + \tau$  are assumed to be cloudy (d), where  $\tau$  is a constant threshold.

Particularly challenging are cases with high concentrations of aerosols or haze. When the aerosol concentration is high, more of this spectrally uniform scattering occurs, which results in cloud free pixels having larger than average RBR. This can result in haze and aerosols being inaccurately classified as thin clouds, to address this, a Dynamic Clear Sky Library was developed.

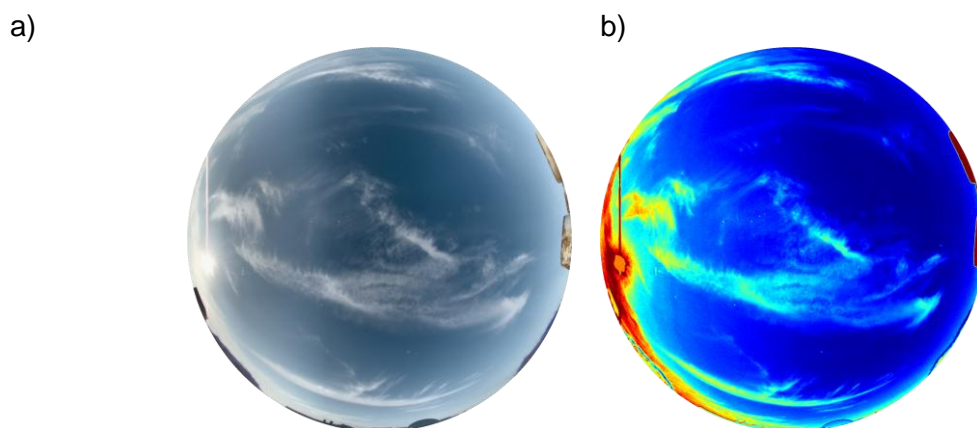
#### **4.2.4 The Dynamic Clear Sky Library for the USI**

To address detection problems of thin clouds in conditions with increased levels of aerosol and haze, Shields et al. (2010) developed a technique to characterize the RBR of clear skies and store it as a function of solar zenith angle and look angle (a pixel's zenith-azimuth coordinate pair). Thick clouds are detected and classified based on a fixed threshold. Next, a clear sky background image is constructed using the stored RBR. A perturbation ratio is then computed, which is the ratio of RBR for the remaining unclassified (non-thick) pixels to the background RBR of the generated clear sky. Finally, a fixed threshold is used to classify the remaining pixels into either clear sky pixels or thin cloud pixels (Shields et al. 2010).

To be able to detect cloudy pixels and differentiate between thick and thin clouds in various atmospheric conditions, an algorithm was developed by Ghonima et al. (2012). The RBR of pixels inside the solar region is close to 1 as the inputs to all three channels of the CCD camera are saturated. In the circumsolar region, the RBR of pixels is also always close to 1 because of forward scattering by aerosols. Outside the solar region, clear-sky and cloudy pixels can be classified based on their RBR. However, for a given image, the RBR of clear pixels varies as a function of the pixels' angular distance from the sun, as well the atmospheric aerosol optical depth (AOD) for that particular day (Shields et al. 2010, Ghonima et al. 2012; see Figure 4.6). A three-dimensional CSL was




constructed in which the RBR of clear-sky pixels was stored as a function of pixel zenith angle, sun-pixel-angle, and solar zenith angle (refer to Figure 4.6b).



**Figure 4.6.** (a) Raw USI image captured on November 19, 2011; (b) RBR of the image.

For each captured sky image, the algorithm constructs a clear-sky background image by looking up each pixel's clear-sky RBR for a given SPA and SZA from the library. The difference is then computed between the RBR of the sky image and the constructed CSL RBR image. Next, pixels with a difference value greater than a certain thick-cloud threshold value are classified as thick cloud. To account for the variations in the CSL RBR image caused by varying AOD, a haze correction factor is applied to the CSL RBR image (Ghonima et al. 2012). Finally, utilizing the haze-corrected difference and a fixed clear-sky threshold value, any pixels not already classified as thick are classified as thin-cloud and clear.



**Figure 4.7.** Clear-Sky Library (CSL) lookup table as a function of pixel-zenith angle and scattering angle (Sun-pixel angle) for the USI over an entire day (a) and for the USI at selected solar-zenith angles (b). Near the Sun and the horizon, the scattered intensity measured on the red channel increases and thus the RBR is greater.

#### 4.2.5 Classification of Cloud Type

The methods described thus far have treated the problem of cloud detection - determining if a cloud exists in a given pixel or not - primarily through RBR or saturation thresholding. Determination of cloud type extends the detection capability of the sky imager to identify phenomenological characteristics of clouds. While microphysical parameters and macroscopic optical properties of clouds can have large variations within a genera of cloud, the structural features of clouds, which are closely associated

with type, do impact variability characteristics. This is exhibited as an identifiable cloud type 'fingerprint' of oscillations within the irradiance signal measured at the ground. The cloud type can potentially inform uncertainty assigned to a forecast, or to augment the forecast with a variability or confidence metric.

The identification of cloud type falls under the general discipline of classification. The field of machine learning defines classification as a form of supervised learning because information about the desired categories, to be determined by a detection algorithm, are available for training. Supervision comes in the form of a user generated set of labeled or annotated observations that form a training set. For clouds, these labels or annotations could be user identifications of different cloud types, such as cirrocumulus, altostratus, stratocumulus, etc. From the set of annotations, features (also called descriptors) are computed. In computer vision, a feature is the formal term for a quantity like RBR or saturation that denotes a single piece of information. Multiple relevant features can be constructed into a feature vector, which for a set of observations will span a feature space. A single feature for a single cloud category will not have a unique value, but rather will have a distribution of values. Figure 4.3 is an example of a scalar feature space (where only one feature, RBR, is considered), and the histogram provides information about how the three selected categories occupy this feature space. With only a single feature, reliable distinction between each of the three categories is not possible, but consideration of additional independent features will allow for further separation of the categories within the expanded feature space and accuracy of classification will improve.

Much work has been done on the classification of clouds from satellite images, but surprisingly little research has been conducted on the classification of clouds using sky imager systems. One of the first studies on cloud classification using a sky imager was

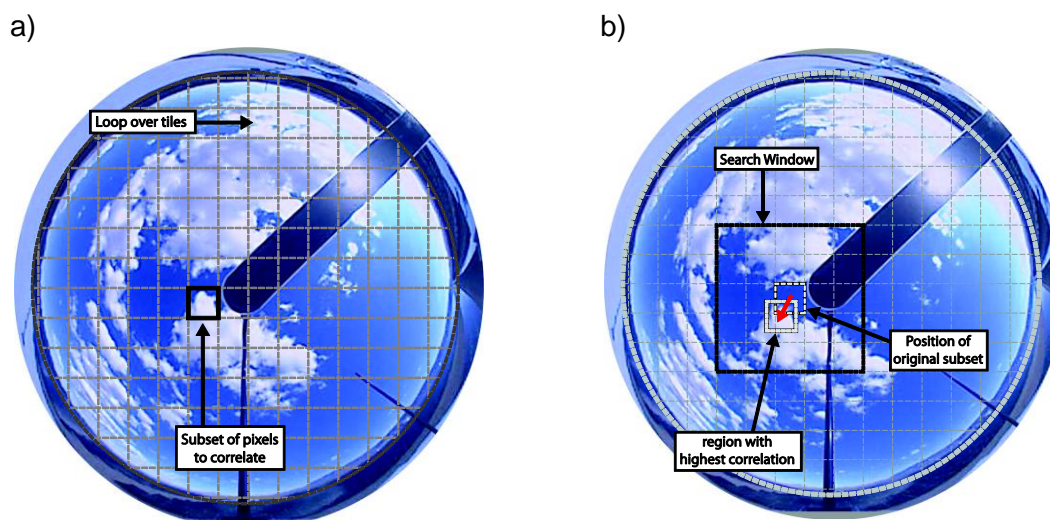
conducted by Buch et al. (1995) using a WSI instrument (section 2.2.2). The authors utilize a binary decision tree and nine features for classification, and report moderate success. An excellent review of cloud classification methodologies applied in remote sensing is given by Singh and Glennen (2005). The focus of this work is automatic identification of convective clouds for air traffic control, but the work documents well the differences between satellite and ground-based detection, and illuminates challenges specific to ground-based cloud classification. The authors apply several pattern recognition techniques to distill images into a form suitable for classification, and utilized both k-nearest neighbor and neural network classifiers. Due to the similarity of many convective cloudforms, results were moderate. Calbó and Sabburg (2008) classified images from both a TSI and the WSC (section 2.2.2) located in Australia and Spain, respectively. This work clearly explains the set of textural and spectral features used, and in addition the authors attempt a novel Fourier power spectrum approach to incorporate information about larger scale cloud structure not contained within the other selected features. Using the distribution of each feature to construct a parallelepiped classifier, the authors report a 76% success rate for five cloud classes. A very successful reporting of cloud classification is that of Heinle et al. (2010). The authors use several of the same features employed by Calbó and Sabburg (2008), but were using a higher resolution camera (IFM-GEOMAR, section 2.2.2) along with a k-nearest neighbor classifier. While the reported success rate was still at the 75% benchmark, this was for seven cloud classes. The selected set of cloud classes is relevant because in the literature this is the set that most closely resembles the 10 classes of the International Cloud Atlas (WMO, 1987). Kazantzidis et al. (2012) follow closely the method of Heinle et al. (2010), using the nearly identical features, the same classifier, and the same 7 cloud classes, but with the distinction of separating processing into regimes based upon

cloud cover percent, solar zenith angle, and solar disk visibility fraction. Both sensor and cloud illumination change with these three quantities which will in turn broaden feature distributions and reduce classification performance. This insightful approach, which the authors term 'subclassing,' yielded an average classification success rate of 88%.

### **4.3 Cloud Velocity Estimation**

Cloud speed and direction are determined by analyzing the change in cloud position between consecutive images. Change in cloud position is detected using a normalized cross correlation (NCC) procedure. The process begins by first breaking one image into small tiles and then cross correlated each tile with the second image. The set of displacements between each tile and its matching location yields a vector field quantifying how the clouds have transformed between images.

The size of the region in the second image that a tile is cross-correlated with is restricted to prevent matches from outside of a physically realistic area. The tiling procedure and the corresponding search area are illustrated in Figure 4.8 for two consecutive images. In Figure 4.8b the original tile position and the position of maximum cross-correlation within the search window are shown. The figure illustrates the process using raw images in the coordinates of the imaging system; in other words, each image as shown has not been projected into pseudo-Cartesian coordinates (Koehler and Shields 1990, Allmen and Kegelmeier 1996). This figure is for illustration only, and the actual NCC is computed using images transformed into pseudo-Cartesian coordinates.



**Figure 4.8.** Illustration of the normalized cross correlation method used to compute inter-image cloud motions. The image (a) at  $t_0 - 30$  sec is broken into small tiles, and each of these tiles cross correlated with the corresponding search window in (b) the image at  $t_0$ .

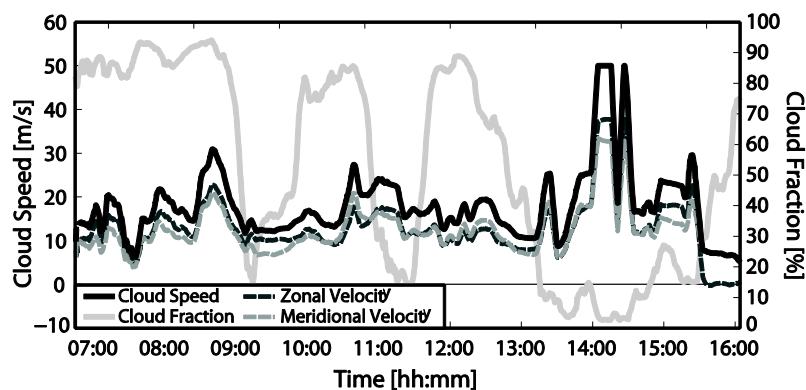
Cross correlation between two consecutive images yields a motion vector field at a single instant in time, consisting of one vector for each tile. This vector field may have erroneous vectors and thus requires quality control (QC). Furthermore, in Chow et al. (2011) the long term trend in the velocity field was shown to be stable, but there was considerable inter-image fluctuation. This indicates the procedure, on average, will provide a stable measurements, but there will be short-term variance in velocity estimation due to the particular clouds occupying the imager field of view and their evolution over a short time window. To address this issue a second level of QC was devised to reduce inter-image fluctuation. The first level (L1) of quality control applied to the raw vector field of a single image pair is described in Table 4.1 and generates a single representative velocity for the image pair being correlated. For the second level (L2) of QC, L1 output and raw velocity field is screened for criteria listed in Table 4.1, and if the velocity passes the screening it is incorporated into the cloud velocity estimate through a low pass filter:

$$y_n = \sum_j^n \alpha_j x_j + \sum_i^{n-1} \beta_i y_i, \quad 4.2a$$

$$\sum_i^{n-1} (\alpha_i + \beta_i) + \alpha_n = 1, \quad 4.2b$$

$$\sum_i^n \alpha_i = 0.2, \quad \sum_i^{n-1} \beta_i = 0.8, \quad 4.2c,d$$

where  $y_i$  are the filtered velocities,  $x_j$  are the raw cross correlation outputs, and  $\alpha_i$  and  $\beta_i$  are the corresponding weights. Both weights  $\alpha_i$  and  $\beta_i$  logarithmically decrease to zero at ten minutes in the past and are normalized such that the previous filtered estimates ( $y_i$ ) are weighted at 80% and previous unfiltered estimates ( $x_i$ ) at 20% (Eq. 4.2c,d). This filtering is independently applied to the east-west and north-south velocity. Figure 4.9 shows a velocity profile with much less high frequency noise than previous methods. A large spike in velocity was not filtered out by the method because false cloud detection near the sun in a period of low cloud fraction caused erroneously high correlations that passed both L1 and L2 QC. Another contribution to this event is that cloud velocity linearly scales with cloud height, which was determined to be 7 km during the spike versus 4 km otherwise. Accuracy in determination of cloud height is also impacted during low cloud fraction periods. In lieu of adding additional screening to address this, better cloud detection in the solar region has been developed (Ghonima et al., 2012), along with higher quality imaging equipment (USI, Chapter 2).



**Figure 4.9.** Cloud speed measured near Boulder City, Nevada, on November 12, 2011, using a TSI. The cloud fraction indicates how much data is available to detect motion.

**Table 4.1.** Quality Control (QC) of Sky Imager Cloud Motion Vectors

---

#### Level 1 (L1) QC

---

1. Vector must have inter-image correlation exceeding 0.8
  2. Vector removal process:
    - a. magnitude is greater than 1.5 standard deviations from the mean
    - b. azimuth is greater than a standard deviation from the mean
    - c. magnitude is greater than a standard deviation from the mean
    - d. azimuth is greater than 1.5 standard deviations from the mean
  3. Average vector field for most representative vector
- 

#### Level 2 (L2) QC

---

Screening for inclusion of L1 output in L2 filtering:

1. Require that raw vector field has enough vectors to consider it reliable (>3)
  2. Require that the step velocity change for L1 output be within a tolerance (computed using standard deviation of past L2 estimates)
  3. Require cloud fraction be above a threshold (>5%)
-



## Acknowledgements

In the course of preparing this chapter, the following publications were used:

**Publication:** Intra-hour forecasting with a total sky imager at the UC San Diego solar energy testbed. *Journal of Solar Energy*, volume 85, issue 11, pp. 2881–2893, 2011.  
doi: 10.1016/j.solener.2011.08.025

**Co-Authors:** Chi Wai Chow, Matthew Lave, Anthony Dominguez, Jan Kleissl, Janet Shields, and Byron Washom

**Publication:** Sky imaging systems for short-term solar forecasting, chapter 9 in: Kleissl J, editor: *Solar Energy Forecasting and Resource Assessment*. Elsevier; 2013.  
ISBN: 978-0-12-397177-7, doi: 10.1016/B978-0-12-397177-7.00009-7

**Co-Authors:** Mohamed Ghonima, Dung Nguyen, Ben Kurtz, Chi Wai Chow and Jan Kleissl

## 5 Cloud Stereography

Contributing Authors:

B. Urquhart, D. Nguyen and J. Kleissl

### 5.1 Introduction

Cloud height plays a vital role in intra-hour solar forecasting. The distance between a vertical projection of a cloud onto the ground versus the actual shadow location increases linearly with cloud height. For typical midlatitude solar zenith angles of  $45^\circ$ , a change of 1 km in cloud height causes a 1 km translation of the cloud shadow on the ground. Ceilometers, typically as part of automated airport weather stations, are the most common ground-based cloud height observational tools. Ceilometers provide a vertical profile of atmospheric backscatter and compute cloud base height (CBH) directly above a single ground location because retrievals are made using a vertically pointing high intensity near infrared laser beam. Satellite imaging is another popular approach for estimating cloud top height, but such measurements require atmospheric temperature profiles and the spatial and temporal resolution is coarse. Radiosondes can also provide accurate cloud height profiles, but with a 12 hour repeat time the temporal resolution is insufficient and long term operation is not feasible for short-term solar forecasting.

Stereographic methods applied to sky imagery can provide cloud base height at a high resolution. With a single sky imager, whole sky visualization is possible. Two imagers allow cloud position to be triangulated using viewing geometry and the distance between sky imagers. There have been a limited number of studies for which cloud height (or position) have been calculated. Approaches reviewed here are grouped into two categories: 1) a 2D framework with a single cloud layer at a constant height (Kassianov et al. 2005, Nguyen et al. 2013); 2) the height of each image segment is

determined separately through intra-image matching techniques and triangulation between instruments (Allmen and Kegelmeyer 1997, Seiz et al. 2007).

### **5.1.1 Statistical Whole Image Matching Analysis for Single Cloud Layer**

Kassianov et al. (2005) proposed a statistical approach of CBH retrieval with the assumption that there is only a single cloud layer and a single CBH for this layer. The procedure begins by collecting two images simultaneously from two closely located instruments. The images are cropped at 100° field of view to remove pixels with large zenith angles. Cloud detection is applied (section 4.2) to identify the cloudy and clear sky regions. A pseudo-cartesian transform is applied (Allmen and Kegelmeyer 1996, section 6.3) to remove distortion due to the projection used in fisheye lenses (section 2.3.1 and 3.3.1.2). The matching process begins by placing the two images on top of each other and computing the mean square error (MSE), i.e. compute the sum of the squared pixel matching error (difference in intensity values) divided by the number of overlapping pixels. The process is repeated, moving the images apart pixel by pixel, and the MSE is recorded as a function of the displacement  $R$  between image centers. The minimum MSE provides the displacement  $R^*$  that yields the best match, and the CBH is computed from  $R^*$  and the geometry of the system.

Nguyen and Kleissl (2014) developed a similar whole image matching approach which projected the saturation images (HSV colorspace, section 4.2.1) from two imagers to georeferenced planes at different heights (section 6.2), referred to here as the georeferenced projection method (GPM). The two images are projected to successive height levels, and the height with the minimum root mean square matching error  $\varepsilon$  is selected as the cloud height. The height search is separated into two steps: first, errors are computed at intervals of 100m; second a binary search at a resolution of 1m is used

to find the cloud height at a high resolution. This method is very efficient which lends itself well to operational implementation. On four days tested, the root mean square height error was 281 meters (14.9%). The results presented in that work are by far the most comprehensive cloud stereography results from a ground based imager presented to date, and the time series comparisons between a collocated ceilometer and the sky imager are highly correlated (by visual inspection). A method to determine optimal imager spacing for the 2D method is also presented.

### **5.1.2 Cloud depth estimation using epipolar geometry**

Allmen and Kegelmeyer (1996) developed a matching method for cloud images captured with a fisheye camera that determined correspondence using the intensity and optical flow as matching metrics. While performing the matching along epipolar lines, a pseudo-Cartesian transform was applied to both the image intensity window (a small neighborhood about the point of interest) and the optical flow field neighborhood. Matching was performed asymmetrically (i.e. only one image was taken as reference) following Algorithm 1. The 'hysteresis' in a forward and reverse matching procedure was used to determine the confidence of the match. The method was tested on both real and simulated cloud data. The tests on the simulated data yielded decent results, with standard deviation of the error ranging from 0.3 to 1.3 km. The standard deviation of the error in the real results was 0.75 and 1.25 km for the two days tested with optimal parameter settings (many tests were reported at different settings to assess their impact). The use of optical flow to aid in determining correspondence actually reduced the accuracy of the results.

**Algorithm 1. Allmen and Kegelmeyer epipolar matching procedure**

*adapted from Allmen and Kegelmeyer (1996). Notation here is different from the remainder of the dissertation.*

1. For each point  $x_1$  in image 1
  - a) compute the epipolar curve for  $x_1$  in image 2.
  - b) for each point  $\tilde{x}_2$  along the epipolar curve, compute the correlation coefficient for the neighborhoods around  $x_1$  and  $\tilde{x}_2$
  - c) find the point  $x_2$  along the epipolar curve with maximum correlation coefficient
  
2. For each point  $x_2$  located within image 2
  - a) compute the epipolar curve for  $x_2$  in image 1.
  - b) for each point  $\tilde{x}_1$  along the epipolar curve, compute the correlation coefficient for the neighborhoods around  $x_2$  and  $\tilde{x}_1$
  - c) find the point  $x_3$  along the epipolar curve with maximum correlation coefficient
  
3. Compute the cloud height using  $x_1$  and  $x_2$ .
4. The difference between  $x_1$  and  $x_3$  is the confidence metric for match

Seiz (2003) and Seiz et al. (2007) report on stereographic determination of cloud height from ground based cameras. The camera system used in both works was an ETH Zurich cloud stereographic system. The ETH system does not have large radial distortion like the fisheye lens presented here, although three radial and two decentering distortion parameters were added to the camera model following Brown (1971). A Whole sky imager (Shields et al. 2013) was presented in Seiz et al. (2007), but was not used for stereographic determination. A discussion was included indicating such a system may not be well designed for stereography applications. Matching was performed using both a cross correlation method and what is termed the 'multi-photo geometrically constrained' matching (MPGC, Baltsavias 1991). MPGC is a least squares matching method that simultaneously enforces collinearity (i.e. geometric) constraints across the patch being matched, and minimizes the error in the area matching metric selected (e.g. grey level error, cross correlation coefficient, etc.). It also incorporates a rectification procedure to ensure that when the two images are compared, local distortion is removed. The method was tested for three days in a sparse matching configuration with

47, 62, and 64 points selected match points (a Förstner operator was used to select features, Förstner and Gulch 1987). The RMS difference in height between correspondences measured manually and determined through MPGC was 60.1, 8.7, and 60.1 meters, with a maximum height error of 413.6 meters. Heights for the remaining cloud pixels was interpolated from this sparse set. A time series of cloud heights for a single location above a ceilometer was generated for a 30 minute window at 5 minute steps (only a few pixels were used corresponding to the projected position of the ceilometer). Comparisons to the ceilometer give a mean height difference of 49 meters and a standard deviation of 74 meters. It was observed that the ceilometer had limitations with broken clouds whereby the ceilometer could not retrieve a cloud height because the cloud was not directly overhead, a limitation that the sky imager does not have. While MPGC algorithm is robust, it also requires solving a linear system for each match point which takes significant computational power for dense depth estimation.

Nguyen and Kleissl (2014) report on dense 3D cloud field determination for a single whole sky image. Matching was performed along epipolar curves (due to lens distortion, the *apparent* intersection of the epipolar plane and the image plane is a curve in the image plane, in general. Variable window cross correlation matching was applied to the saturation image, with a window size set such that the standard deviation of the pixel values within the window exceeded a threshold. Matching was performed following Algorithm 1, and  $L_{max} = |\mathbf{x}_1 - \mathbf{x}_3|$  was used confidence metric. The difference in median cloud height for the single image and ceilometer measurements in a  $\pm 10$  minute window of the image capture time was +8 meters, and 67% of the pixels were within 2 standard deviations (341 meters) of the reported ceilometer height of 1826 meters MSL.


## 5.2 Experimental Setup

### 5.2.1 Camera System

The University of California, San Diego (UCSD) sky imager (USI) camera system was developed for the purpose of solar power forecasting (Urquhart et al. 2013a). The camera is an Allied Vision GE-2040C camera which contains a 15.15 mm × 15.15 mm, 2048 × 2048 pixel Truesense KAI-04022 interline transfer charge coupled device. The lens is a Sigma circular fisheye lens with a 4.5 mm nominal focal length and equisolid angle projection. The USI uses 3 exposures at integration times of 3, 12, and 48 ms to generate a composite 16-bit HDR image. The final sky images have an effective spatial resolution of 1748 × 1748 and radiometric resolution of 65,536 levels per color channel. The two USIs used in this work were deployed at the Department of Energy, Atmospheric Radiation Measurement (ARM) Program, Southern Great Plains (SGP) Climate Research Facility from March 11th, 2013 to November 4th, 2013. The longitude, latitude, altitude of USI 1.7 was -97.478766°E, 36.6183435°N, 304.5 m, and for USI 1.8 was -97.484856°E, 36.604043°N, 318 m (Figure 5.1). Both cameras were geometrically calibrated using the method described in Urquhart et al. 2014b using solar position data on March 31, 2013. Extensive details of the USI can be found in Urquhart et al. 2014a.

### 5.2.2 Ceilometer


A Vaisala Ceilometer Model CL31 (hereinafter 'ceilometer') was used for cloud height validation. The ceilometer emits pulses of near infrared light, and detects the backscattered near infrared signal within an 18.7 deg field of view. It can detect up to three cloud layers and has a 7,600 m vertical range. The ceilometer was located at a longitude, latitude, altitude of -97.485516°E, 36.605128°N, 316 m, which is 1,590 m from USI 1.7 at 202 deg azimuth (SSW), and 134 m from USI 1.8 at 334 deg azimuth (NNW).




**Figure 5.1.** (a) Satellite image showing the camera and ceilometer locations, along with geographic coordinates. Close-ups of the locations for (b) USI 1.7 near the Bobbit roundtop and (c) USI 1.8 in the ARM SGP instrument field. Angled views showing (d) the height of the instrument field relative to the roundtop and (e) a north-south looking view.



**altocumulus-01, 2013-04-13 T 00:15:00**




**broken-02, 2013-04-09 T 00:04:00**



**Figure 5.2.** Image pairs for which a dense estimation of cloud position was performed. Images on the left are from USI 1.7 (note Farmer Bobbit's roundtop barn), and images on the right are from USI 1.8. All images are color corrected HDR with a logarithmic intensity rescaling (similar to gamma correction). (*continued*)

**cirrocumulus-03, 2013-04-12 T 16:52:00**




**cirrus-04, 2013-04-06 T 14:47:30**




**Figure 5.2.** Image pairs for which a dense estimation of cloud position was performed. Images on the left are from USI 1.7 (note Farmer Bobbit's roundtop barn), and images on the right are from USI 1.8. All images are color corrected HDR with a logarithmic intensity rescaling (similar to gamma correction). (*continued*)

**cirrocumulus-03, 2013-04-12 T 16:52:00**




**cirrus-04, 2013-04-06 T 14:47:30**




**Figure 5.2.** Image pairs for which a dense estimation of cloud position was performed. Images on the left are from USI 1.7 (note Farmer Bobbit's roundtop barn), and images on the right are from USI 1.8. All images are color corrected HDR with a logarithmic intensity rescaling (similar to gamma correction). (*continued*)

**cumulus-05**, 2013-05-03 T 19:33:30




**cumulus-06**, 2013-04-04 T 22:30:30



**Figure 5.2.** Image pairs for which a dense estimation of cloud position was performed. Images on the left are from USI 1.7 (note Farmer Bobbit's roundtop barn), and images on the right are from USI 1.8. All images are color corrected HDR with a logarithmic intensity rescaling (similar to gamma correction). (*continued*)

lowsparse-07, 2013-04-11 T 15:08:00



lowsparse-08, 2013-04-11 T 21:03:30




**Figure 5.2.** Image pairs for which a dense estimation of cloud position was performed. Images on the left are from USI 1.7 (note Farmer Bobbit's roundtop barn), and images on the right are from USI 1.8. All images are color corrected HDR with a logarithmic intensity rescaling (similar to gamma correction). (*continued*)

**overcast-09**, 2013-04-07 T 17:41:00



**twolayers-10**, 2013-04-14 T 16:08:00



**Figure 5.2.** Image pairs for which a dense estimation of cloud position was performed. Images on the left are from USI 1.7 (note Farmer Bobbit's roundtop barn), and images on the right are from USI 1.8. All images are color corrected HDR with a logarithmic intensity rescaling (similar to gamma correction).

### 5.2.3 Cloud Conditions

A variety of conditions were selected for validation (Figure 5.2). Due to the wide baseline and occlusion of higher layers when lower layers have sufficient areal coverage, cloud height detection in two layers is often not possible and only a single image with multiple layers was tested (twolayers-10). The single-layer images contain different texture length scales (a spatial feature) and intensity scales (altocumulus-01, cirrocumulus-03, cirrus-04, overcast-09). The image set also tests the stereo algorithm under extreme dissimilarity in perspective (e.g. low clouds, broken-02, cumulus-06, lowsparse-07, lowsparse-08). For good measure, an easy case was also included, where the main features are near the image centers (cumulus-05).

### 5.3 Methods

Stereo determination of scene depth from images captured with a fisheye lens has the additional complication that, in general, correspondence must be determined along curves instead of straight lines. While it is possible to determine correspondence along these curves, it is not as computationally efficient as matching along frame buffer axes (pixel axes) in the image. Additionally, depending on the feature location within the field of view, each camera will image the feature with potentially differing amounts of distortion, further complicating the matching process. To simplify the correspondence problem, a pseudo-image which follows an ideal perspective projection is constructed by sampling the original image. The points to sample are determined by defining a viewing direction and desired field of view of a 'pseudo-camera' from which to take a 'snapshot'. This section discusses the construction of a pseudo-image which can greatly simplify matching and convert the originally distorted scene information into a format that can be readily used with a host of computer vision stereography techniques. The generation of

a pseudo-image is a process which involves obtaining the perspective projection matrix of the pseudo-camera, and performing epipolar rectification such that matching can be performed along corresponding columns.

### 5.3.1 Coordinate System

To obtain georeferenced position information about the cloud field, it is important to clearly define the coordinate systems used. The topocentric 'world' coordinate system is defined such that the origin is at a point  $O$  which is the center of the reference camera (here, USI 1.7). The coordinate axes are aligned with the cardinal directions such that the  $X$ -axis is positive east, the  $Y$ -axis is positive north and the  $Z$ -axis is positive upwards aligned with the zenith. A point in this coordinate system is denoted  $\tilde{\mathbf{X}} = (\tilde{X}, \tilde{Y}, \tilde{Z})^\top$ . The tilde overbar indicates that the point is inhomogeneous, that is, it is in Euclidean space  $\tilde{\mathbf{X}} \in \mathbb{R}^3$ . This notation is to distinguish it from the corresponding set of homogeneous points in projective space  $\mathbf{X} \in \mathbb{P}^3$ , where  $\mathbf{X} = (X, Y, Z, T)^\top$  and  $\tilde{\mathbf{X}} = (X/T, Y/T, Z/T)^\top = (\tilde{X}, \tilde{Y}, \tilde{Z})^\top$ . Similar notation holds for a point in 2-dimensional projective space  $\mathbf{x} \in \mathbb{P}^2$ ,  $\mathbf{x} = (x, y, w)^\top$ , with corresponding Euclidean point  $\tilde{\mathbf{x}} = (x/w, y/w)^\top = (\tilde{x}, \tilde{y})^\top$ ,  $\tilde{\mathbf{x}} \in \mathbb{R}^2$ . When scale factors  $T$  or  $w$  are zero, the corresponding Euclidean point is infinite. For points not lying on the plane or line at infinity, we can write  $\mathbf{X} = (\tilde{\mathbf{X}}^\top, 1)^\top$  and  $\mathbf{x} = (\tilde{\mathbf{x}}^\top, 1)^\top$ , respectively.

The  $n^{\text{th}}$  camera has a center  $\tilde{\mathbf{C}}_n \in \mathbb{R}^3$  in the world coordinate system. A point  $\tilde{\mathbf{X}}_{cam}$  in the camera coordinate system (where the camera index  $n$  has been dropped) can be written using homogenous coordinates as

$$\mathbf{X}_{cam} = \begin{bmatrix} R & -R\tilde{\mathbf{C}} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \mathbf{X}, \quad 5.1$$



where  $R \in \mathbb{R}^{3 \times 3}$  is the rotation from the world to camera coordinate system (i.e. the rows of  $R$  form an orthonormal basis for the camera coordinate system). Equivalently, in inhomogeneous coordinates we can write  $\tilde{\mathbf{X}}_{cam} = R(\tilde{\mathbf{X}} - \tilde{\mathbf{C}})$ .

### 5.3.2 Scene Plane Definition

The desired viewing direction for a single camera can be defined by zenith angle  $\theta$  and azimuth angle  $\phi$ , where  $\theta$  is the angle from the vertical direction ( $Z$ -axis) and  $\phi$  is the angle positive east from north in the plane parallel to the ground (i.e. the angle from the positive  $Y$ -axis in the  $X - Y$  plane, positive clockwise). The vector  $\mathbf{n}$  pointing in direction  $\Phi = (\theta, \phi)^T$  can be defined as

$$\mathbf{n} = \begin{bmatrix} \sin \theta \sin \phi \\ \sin \theta \cos \phi \\ \cos \theta \end{bmatrix}.$$

Orthogonal axes to span the plane defined by normal  $\mathbf{n}$  can be defined:

$$\mathbf{e}_x = \begin{bmatrix} -\cos \phi \\ \sin \phi \\ 0 \end{bmatrix}, \quad \mathbf{e}_y = \begin{bmatrix} -\cos \theta \sin \phi \\ -\cos \theta \cos \phi \\ \sin \theta \end{bmatrix},$$

which gives an orthonormal basis for the pseudo-camera coordinate system:

$$R^T = [\mathbf{e}_x \quad \mathbf{e}_y \quad \mathbf{n}]. \quad 5.2$$

For  $\tilde{\mathbf{C}} = \mathbf{0}$ , we have

$$\tilde{\mathbf{X}} = R^T \tilde{\mathbf{X}}_{cam}. \quad 5.3$$

The field of view, which is a pyramidal in shape, can be defined by the horizontal and vertical field of view angles,  $\psi_h$  and  $\psi_v$ , respectively. The four corners of our plane in the camera coordinate system  $\tilde{\mathbf{X}}_{cam,i}$ ,  $i = 1,2,3,4$ , can be defined using these angles:

$$\tilde{\mathbf{X}}_{cam,1} = \left(-\tan \frac{1}{2}\psi_h, \tan \frac{1}{2}\psi_v, 1\right)^T, \quad \tilde{\mathbf{X}}_{cam,2} = \left(\tan \frac{1}{2}\psi_h, \tan \frac{1}{2}\psi_v, 1\right)^T, \quad 5.4a,b$$

$$\tilde{\mathbf{X}}_{cam,3} = \left(\tan \frac{1}{2}\psi_h, -\tan \frac{1}{2}\psi_v, 1\right)^T, \quad \tilde{\mathbf{X}}_{cam,4} = \left(-\tan \frac{1}{2}\psi_h, -\tan \frac{1}{2}\psi_v, 1\right)^T, \quad 5.4c,d$$

which are, sequentially, the top-left, top-right, bottom-right, and bottom-left corners. We can use Eq. 5.3 to compute the corresponding world coordinates of the scene plane corners.

### 5.3.3 Affine Homography between Calibrated and Uncalibrated Pseudo-image Coordinates

The projective point transformation that describes the perspective camera model can be written as

$$\mathbf{x} = P \mathbf{X}, \quad 5.5$$

where the perspective projection matrix  $P \in \mathbb{R}^{3 \times 4}$  is given by

$$P = K \begin{bmatrix} R & -R\tilde{\mathbf{C}} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix}. \quad 5.6$$

The matrix  $K \in \mathbb{R}^{3 \times 3}$  is an affine transformation known as the camera calibration matrix.

The calibrated points  $\hat{\mathbf{x}} = (\hat{x}, \hat{y}, \hat{w})^\top$  on the image plane given by  $\hat{\mathbf{x}} = [I \mid \mathbf{0}] \mathbf{X}_{cam}$ , can be converted to pixel coordinates  $\mathbf{x} = (x, y, w)^\top$  by the affine transformation  $K$ :  $\mathbf{x} = K \hat{\mathbf{x}}$ . The calibrated point  $\hat{\mathbf{x}}$  in homogeneous coordinates is equivalent to the inhomogeneous coordinate  $\tilde{\mathbf{X}}_{cam}$ , hence we also have

$$\mathbf{x} = K \tilde{\mathbf{X}}_{cam}. \quad 5.7$$

The corners in Eq. 5.4 are defined to correspond to

$$\mathbf{x}_1 = (1, 1, 1)^\top, \quad \mathbf{x}_2 = (N, 1, 1)^\top, \quad 5.8a,b$$

$$\mathbf{x}_3 = (N, M, 1)^\top, \quad \mathbf{x}_4 = (1, M, 1)^\top, \quad 5.8c,d$$

where  $\mathbf{x}_i = (x_i, y_i, 1)^\top$  is a homogeneous point in the image plane,  $N$  is the number of desired columns, and  $M$  is the number of desired rows in the pseudo-image. The values of  $N$  and  $M$  determine the sampling resolution for the pseudo-image. Multiplying by the left nullspace of  $\mathbf{x}_i$ , Eq. 5.7 (for the  $i^{\text{th}}$  corner) becomes

$$[\mathbf{x}_i^T]^{\perp T} K \tilde{\mathbf{X}}_{cam,i} = \mathbf{0} \quad 5.9$$

where  $[\cdot]^{\perp}$  denotes the nullspace of the argument (hence  $[\mathbf{x}_i^T]^{\perp T}$  is the left nullspace).

Applying the *vec* operator to Eq. 5.9 yields

$$\left(\tilde{\mathbf{X}}_{cam,i}^T \otimes [\mathbf{x}_i^T]^{\perp T}\right) \text{vec}(K) = \mathbf{0},$$

which is a homogeneous linear equation in  $\mathbf{k} = \text{vec}(K)$ . The Direct Linear Transform algorithm can be used to solve for  $K$  (see Hartley and Zisserman 2004 for an excellent introduction and Hartley 1997 for the importance of data normalization when using the DLT algorithm; see also section 3.5.1.5).

The perspective projection matrix  $P = KR[I_{3 \times 3} \mid -\tilde{\mathbf{C}}]$ , with  $R^T = [\mathbf{e}_x \ \mathbf{e}_y \ \mathbf{n}]$ , and  $K$  obtained as just described, defines our 'pseudo-camera'. For the reference camera,  $\tilde{\mathbf{C}} = \mathbf{0}$ , and for all other cameras  $\tilde{\mathbf{C}}$  is the baseline vector between cameras (i.e. the displacement vector *from* the secondary camera *to* the reference camera). Matrices  $R$  and  $K$  will vary for each camera based on the parameters  $\theta$ ,  $\phi$ ,  $\psi_h$ ,  $\psi_v$ ,  $M$ , and  $N$  selected for that camera. To 'point' two pseudo-cameras at the same volume of object space,  $\theta$ ,  $\phi$ ,  $\psi_h$ , and  $\psi_v$  must be set appropriately for each camera.

### 5.3.4 Pseudo-image Rectification

The geometry between two cameras is well known (e.g. Longuet-Higgins 1981, Faugeras 1993, Gu and Zhang 1993, Hartley and Zisserman 2004, etc.), so only a cursory introduction is provided here. Two cameras with centers  $\mathbf{C}_1$  and  $\mathbf{C}_2$ , record a point  $\mathbf{X}$  in object space at image points  $\mathbf{x}_1$  and  $\mathbf{x}_2$  (where the subscript corresponds to the camera number). Points  $\mathbf{C}_1$ ,  $\mathbf{C}_2$ , and  $\mathbf{X}$  define an epipolar plane in object space. For an object at  $\mathbf{X}$ , imaged at point  $\mathbf{x}_1$  in image 1, the corresponding point  $\mathbf{x}_2$  in image 2 must lie somewhere within this plane. An epipolar plane intersects the image plane in a line

known as an epipolar line, thus when searching in image 2 for point  $\mathbf{x}_2$  corresponding to the point  $\mathbf{x}_1$  in image 1, we must only search along this line. All epipolar planes pass through the 3-space baseline connecting the two cameras (i.e.  $\mathbf{b} = \tilde{\mathbf{C}}_1 - \tilde{\mathbf{C}}_2$ ,  $\mathbf{b} \in \mathbb{R}^3$ ). The point of intersection of  $\mathbf{b}$  and an image plane is known as an epipole  $\mathbf{e} \in \mathbb{P}^2$ , which is equivalently the position in the image of other camera, i.e. it is the image of the camera center of the opposing camera:

$$\mathbf{e}_1 = P_1 \mathbf{C}_2, \quad \mathbf{e}_2 = P_2 \mathbf{C}_1 \quad 5.10$$

where  $\mathbf{e}_1$  is epipole in image 1,  $\mathbf{e}_2$  is the epipole in image 2,  $\mathbf{C}_1 = [P_2]^\perp$  and  $\mathbf{C}_2 = [P_1]^\perp$ . All epipolar lines intersect the epipole.

If the two image planes are parallel (non-verged geometry), then the epipoles are at infinity (i.e.  $\mathbf{e} = (e_x, e_y, e_w)^\top$  where  $e_w = 0$ ). When this occurs the epipolar lines intersect at infinity, and thus they are parallel. The goal of image rectification in binocular stereo is to transform the pair of images such that the epipolar lines are parallel and aligned with the frame buffer axes (i.e. pixel axes). This simplifies the algorithmic implementation of determining correspondence and reduces computational complexity of the matching process. We seek to rectify the images such that corresponding *columns* are in correspondence, and matching can be performed along each column to determine which rows within that column are in correspondence between the two images. This breaks from the typical convention to rectify images such that *rows* are in correspondence. For convenience, Algorithm 2 describes how rectify an image column-wise (the column-wise version of the algorithm is uncommon and generally not included in texts). Once the matrix  $H$  (Algorithm 2) is obtained, the rectified camera matrix is

$$\hat{P} = HKR [ I \mid -\tilde{\mathbf{C}} ]. \quad 5.11$$

**Algorithm 2. Column-wise image rectification in binocular stereo**

Column-wise image rectification in binocular stereo. The steps here develop a projective transform  $H$  that shift the epipole  $\mathbf{e}$  to infinity along the positive  $y$ -axis, i.e.  $\mathbf{e} = (0,1,0)^\top$ . Epipolar lines are thus parallel to columns in the image.

*Note: algorithm update equations should not be confused with proper mathematical equations*

1. Shift a point of interest  $\mathbf{x}_o = (x_o, y_o, 1)$ , often the image center, to the origin

$$\mathbf{e}_t = \begin{bmatrix} 1 & 0 & -x_o \\ 0 & 1 & -y_o \\ 0 & 0 & 1 \end{bmatrix} \mathbf{e} = T\mathbf{e}$$

2. Rotate the image such that the epipole  $\mathbf{e}_t$  has zero  $x$ -component  $e_{tx}$

$$\varphi = \text{atan2}(e_{tx}, e_{ty})$$

$$\mathbf{e}_r = \begin{bmatrix} \cos \varphi & -\sin \varphi & 0 \\ \sin \varphi & \cos \varphi & 0 \\ 0 & 0 & 1 \end{bmatrix} \mathbf{e}_t = R\mathbf{e}_t$$

3. Shift the epipole to infinity such that it has zero  $w$ -component  $e_{rw}$

$$f = \mathbf{r}^{2\top} \mathbf{e}_r, \text{ where } \mathbf{r}^2 \text{ is the second row of } R$$

$$\mathbf{e}_g = \frac{1}{f} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -e_{rw}/e_{ry} & 1 \end{bmatrix} \mathbf{e}_r = G\mathbf{e}_r$$

4. Once the epipole is at infinity, reverse the rotation and translation

$$H = T^{-1}R^\top GRT$$

Update  $H$  such that the sign of the bottom right element is positive:

$$H = H \cdot \text{sign}(H_{33})$$

The Fundamental Matrix  $F$ , which is a point-to-line transform, defines the epipolar line in the second image, given a point in the first, i.e.

$$\mathbf{l}_2 = F\mathbf{x}_1 \quad 12$$

where  $\mathbf{l}_2 \in \mathbb{P}^2$  is the epipolar line in image 2 associated with a point  $\mathbf{x}_1$  in image 1. Similarly, we can define an epipolar line  $\mathbf{l}_1 \in \mathbb{P}^2$  in image 1 associated with a point  $\mathbf{x}_2$  as  $\mathbf{l}_1 = F^T\mathbf{x}_2$ , where the transpose of  $F$  is required. The fundamental matrix for the rectified pair image cameras  $\hat{P}_1 = H_1P_1$  and  $\hat{P}_2 = H_2P_2$  is

$$F = [\hat{\mathbf{e}}_2]_{\times} \hat{P}_2 \hat{P}_1^+ \quad 13$$

where the notation  $[\cdot]_{\times}$  indicates the  $3 \times 3$  skew symmetric matrix corresponding to the vector argument,  $\hat{\mathbf{e}}_2 = \hat{P}_2[\hat{P}_1^+]^{\perp}$ , and  $\hat{P}_1^+$  is the pseudoinverse of  $\hat{P}_1$ .

The epipolar lines for cameras  $\hat{P}_1$  and  $\hat{P}_2$  are parallel with the columns in each image, however, the columns corresponding to pairs of epipolar lines are not the same in each image in general, e.g. if the epipolar line in image 1 corresponds to column 1, the paired epipolar line in image 2 does not necessarily correspond to column 1 in image 2. We must apply an additional transformation to one of the images so that columns are in correspondence. Algorithm 3 provides the approach used.

### 5.3.5 Pseudo-image Construction

Once the pseudo cameras have been properly defined (obtaining matrices  $K$  and  $R$ ), and rectified (obtaining  $H$  from Algorithm 1 and  $H_s$  from Algorithm 3), a pseudo-image can be constructed (Algorithm 4).

**Algorithm 3. Rescaling of rectified images**

Rescaling the  $x$ -axis of an image so that paired epipolar lines correspond to the same column number in each image. The angles  $\theta_1$  and  $\theta_2$  are the zenith angles of the view direction for camera 1 and camera 2, respectively. The image coordinates  $\mathbf{x}_i$  are defined in Eq. 5.8.

*Note: algorithm update equations should not be confused with proper mathematical equations*

## 1. Compute epipolar lines

if ( $\theta_1 \leq \theta_2$ )

$\mathbf{l}_i = F\mathbf{x}_i$  % epipolar lines in image 2

else

$\mathbf{l}_i = F^T\mathbf{x}_i$  % epipolar lines in image 1

end

scale each  $\mathbf{l}_i = (a_i, b_i, c_i)^T$  such that  $(a_i^2 + b_i^2)^{1/2} = 1$

## 2. Determine mapping between corners

The corners in the opposite image correspond to:

$$\mathbf{y}_1 = (-\mathbf{l}_{1,w}, 1, 1)^T \quad \mathbf{y}_2 = (-\mathbf{l}_{2,w}, 1, 1)^T$$

$$\mathbf{y}_3 = (-\mathbf{l}_{3,w}, M, 1)^T \quad \mathbf{y}_4 = (-\mathbf{l}_{4,w}, M, 1)^T$$

We seek  $H_s$  such that corners  $\mathbf{y}_i$  are remapped to  $\mathbf{x}_i$ :  $\mathbf{x}_i = H_s\mathbf{y}_i$ .

Use the DLT algorithm to find  $H_s$ .

## 3. Update camera projection and fundamental matrices

if ( $\theta_1 \leq \theta_2$ )

Update  $\hat{P}_2$ :  $\hat{P}_2 = H_s \hat{P}_2$

else

Update  $\hat{P}_1$ :  $\hat{P}_1 = H_s \hat{P}_1$

end

The new fundamental matrix is then (all 3 terms must be updated):

$$F = [\hat{\mathbf{e}}_2]_{\times} \hat{P}_2 \hat{P}_1^+$$

**Algorithm 4: Construction of a pseudo-image**

Given camera calibration matrix  $K$ , camera rotation matrix  $R$ , epipolar rectification matrix  $H$  and scaling matrix  $H_s$ , construct a pseudo-image. (The notation  $\mathbf{x}_i$  here differs from main text.)

*Note: algorithm update equations should not be confused with proper mathematical equations*

## 1. Construct a discretized grid representing the coordinates of each pixel

$$S = \{\mathbf{x}_i^g = (x_i, y_i, 1), x_i = 1 \dots N, y_i = 1 \dots M\}$$

## 2. Backproject to world coordinates

$$\mathbf{X}_i = R^T K^{-1} H^{-1} H_s^{-1} \mathbf{x}_i^g$$

## 3. Compute forward projection using calibrated camera model (Urquhart et al. 2014b, Chapter 3)

$$\mathbf{x}_i = \mathcal{P}(\mathbf{X}_i)$$

4. Sample the distorted image at points  $\mathbf{x}_i$  using bilinear interpolation to construct the pseudo image.


### 5.3.6 Correspondence in Cloud Images Using Dynamic Programming

With the images rectified such that corresponding epipolar lines lie along corresponding columns, what is left is to match rows within each column. The correspondence problem can be decomposed into local and global methods, where the former seeks only to determine correspondence by applying constraints only to a small number of pixels, and the latter uses information from entire scanlines or the entire image. For both local and global methods, a matching metric must be selected which is typically derived from image primitives which are computed from neighborhoods of varying size and shape within the image. Common examples are pixel intensity, cross correlation, sum of squared differences, edge or corner elements, etc. An excellent overview of stereo methods can be found in Brown et al. (2003). The methods for full 3D reconstruction of the cloud base in both Allmen and Kegelmeyer (1996) and Nguyen and Kleissl (2014) use a local method. Seiz (2003) and Seiz et al. (2007) apply a least squares method that is nominally local, however it is set up in a hierarchical configuration that enables information from a subsampled version of the entire image to influence the final stereo matching results. The method applied here is dynamic programming which uses scanline constraints, and thus is a semi-global method, i.e. it is global within each scanline, but not global to the entire image.

Dynamic programming is the decomposition of a problem into simpler sub-problems. Ohta and Kanade (1985) developed a framework for intra-scanline search which integrated an implicit ordering constraint: "when we examine the correspondence of two edges, one on the right and one of the left scanline, the edges which are on the left of these edges on each scanline must already be processed." This method is implemented here by finding the maximum path through the 2D matching score matrix (an alternative formulation is the minimum path through the matching cost



matrix). Each scanline has a matching score  $D(x, y)$  associated with position  $x \in [1: N_x]$  along the left scanline and position  $y \in [1: N_y]$  along the right scanline (here  $N_x = N_y$ ). Figure 5.3 shows the matching score for (vertical) scanline 400 in the test case *altocumulus-01*. The maximum path, shown in black, determines which rows are in correspondence. The path is determined by using Algorithm 5, following a similar algorithm described in Ohta and Kanade (1985). The missing rows (in grey) are the rows in image 1 for which there was no cloud detected; similarly, the missing columns (in grey) are the rows in image 2 where no cloud was detected. An ordering constraint is implicit because the maximum path can only move down and to the right, starting from the top left (row 1 in both images). This means that if row  $x$  in image 1 is matched to row  $y$  in image 2, row  $x + 1$  can only match from row  $y$  to  $M$  in image 2, (i.e. cannot match to anything less than row  $y$ ). When a single row from one image corresponds to several in the other, the maximum path will have a vertical or horizontal segment. This often indicates an occlusion or discontinuity in one of the images. In Figure 5.3, a vertical segment indicates a stretch of pixels visible in image 1 is not visible in image 2 and thus are occluded in image 2. A horizontal segment of the maximum path indicates an object in image 2 is occluded in image 1.



**Figure 5.3.** Matching score (normalized cross correlation) for scanline 400 in testcase altocumulus-01. The maximum path, shown in black, indicates which rows are in correspondence. The column (scanline 400) is defined to be in correspondence because it is an epipolar line. The missing rows (in grey) are the rows in image 1 for which there was no cloud detected; similarly, the missing columns (in grey) are the rows in image 2 where no cloud was detected.

Dynamic programming works with any block matching metric (e.g. pixel intensity, cross correlation, sum of squared differences). The block matching metric used here is normalized cross correlation (Lewis 1995). This is one of the most commonly used matching metrics because of its accuracy and robustness; however, it is at the expense of computational speed. No other metrics were tested here, and it is left to future work to explore more computationally efficient methods for matching cloud features. For a single cloud layer, the assumption of ordered correspondence implicit to dynamic programming is applicable. When two or more layers are present, it is not common that an isolated low-level cloud is positioned between the two imagers while the upper cloud layers are clearly visible in both cameras. In this latter case, the ordering constraint breaks down and this method cannot be used.

**Algorithm 5. Dynamic programming for intra-scanline search**

The function  $D(x, y)$  represents the matching score for row  $x \in [1, N]$  in image 1 and row  $y \in [1, M]$  in image 2.

1. Compute maximum running sum table  $M$  with the following recursive formulae

```

M(1,1) = D(1,1) % Origin
M(1,i) = D(1,i) + M(1,i-1) % First row, image 1
M(i,1) = D(i,1) + M(i-1,1) % First row, image 2
M(i,j) = D(i,j) + max( M(i,j-1), M(i-1,j) ) % First row, image 2

```

2. Find the maximum path by traversing  $M$  from  $M(N, M)$  (in matlab pseudocode)

```

P = false( N , M ); % Allocate M x N matrix for path
i = N; j = M; % start at the bottom right
flag = false;
while( i > 1 )
    while( j > 1 )

        P(i,j) = true;

        if( M(i-1,j) > M(i,j-1) )
            i = i - 1;
        else
            j = j - 1;
        end

        % Exit conditions
        if( i == 1 )
            P( i , 1:j ) = true;
            flag = true;
            break;
        end

        if( j == 1 )
            P( 1:i , j ) = true;
            flag = true;
            break;
        end

    end
end
if( flag ) , break; end
end

```

**5.4 Results and Discussion**

The methods described in section 5.3 provide a dense mapping of cloud position over the field of view pseudo-image. Only clouds elements that are visible are assigned a position. The ceilometer, however, provides a backscatter intensity as a function of

height within the 18.7 deg field of view measurement cone. Because the ceilometer is an active sensing device, it can 'see through' lower layers of clouds and return a backscatter signal from clouds that would otherwise be occluded. Noise in the ceilometer data was filtered and a threshold was applied to eliminate erroneous returns. The resulting filtered time-height signals are given in Figures 5.4 through 5.13, subfigure (c). All height values computed from the ceilometer were obtained from a  $\pm 10$  minute window about the associated image capture time. The height map (dense  $Z$ -coordinate map of cloud position) from the sky imager was *not* filtered or corrected for occlusions prior to making comparisons, and results are expected to improve after further post-processing of the height maps. Height maps for each case are given in Figures 5.4 through 5.13, subfigure (b).

The two instruments do not provide a one-to-one comparison of cloud height. To assess the performance of the height retrieval from the USI, the median and the 20th and 80th percentiles of cloud height measurements are given in Table 5.1 (given as P50, P20, and P80, respectively). For the ceilometer, Table 5.1 contains the backscatter-weighted cloud height along with the 20th and 80th percentile of the heights which returned a backscatter signal (after filtering and noise elimination). Overall, the height estimated from the USI was positively biased, ranging from 31 to 332 meters when comparing the ceilometer mean and USI P50 (Table 5.1). The mean difference between the USI P50 and the ceilometer mean was 136 meters, and the root mean square difference (RMSD) was 215 meters. The consistency of the height estimation for both methods (ceilometer and USI stereo) can be observed when looking at the P20 and P80 levels which give mean differences of 285 and 208 meters respectively, and RMSDs of 369 and 125 meters, respectively. These statistics do not include case lowspare-07 which was significantly impacted by the sun's presence for much of the image. Proper

filtering of the USI stereo results should reduce the problems associated with the sun. Additionally, the cloud detection scheme used here to segment the images into clear sky and cloudy pixels was rudimentary and did not leverage the techniques described in Chapter 4.

**Table 5.1.** Cloud heights measured by the ceilometer and the pair of sky imagers. Ceilometer height statistics are derived from the time-height data in a  $\pm 10$  minute window about the image capture time. The mean ceilometer height is a backscatter weighted height of a filtered dataset (filtered to remove noise). The sky imager data is unfiltered. The 20<sup>th</sup>, 50<sup>th</sup> (median) and 80<sup>th</sup> percentiles are given as P20, P50 and P80, respectively. The bias between the P50 USI and mean ceilometer measurements is given.

Case	ceilometer			sky imager			P50 - mean [m]
	mean [m]	P20 [m]	P80 [m]	P50 [m]	P20 [m]	P80 [m]	
altocumulus-01	4,360	4,095	4,545	4,575	4,464	4,670	215
broken-02	1,295	615	1,335	1,408	1,369	1,438	113
cirrocumulus-03	3,805	3,465	4,035	3,996	3,945	4,073	191
cirrus-04	5,656	5,385	5,895	5,988	5,525	6,455	332
cumulus-05	2,353	2,235	2,655	2,486	2,318	2,680	133
cumulus-06	955	825	1,095	996	909	1,090	41
lowspare-07	-			808	521	11,096	
lowspare-08	1,380	1,275	1,485	1,411	1,326	2,112	31
overcast-09	558	225	645	636	575	826	78
twolayers-10	1,318	1,155	1,425	1,369	1,283	1,456	51
	5,014	4,695	5,145	5,190	5,106	5,539	176

#### 5.4.1 Discussion of individual daily results

Reviewing the cloud height maps for the individual days (Figures 5.4 through 5.13, subfigure b), it can be seen that, with the exception of overcast-09 and lowspare-07 and -08, the height results are consistent and comparable with the ceilometer measurements (Table 5.1 and Figures 5.4 through 5.13, subfigure c). The presence of the sun in the images significantly impacted the results in lowspare-07 and -08 because this is a saturated region of low texture and there was not a significant amount of cloud texture surrounding it to limit its impact on the circumsolar region. In cirrocumulus-03

and twolayers-10, the two other images with a visible sun, there was an impact in the immediate vicinity of the sun, but outside of this circumsolar region, results were consistent. This was due to the texture of the cirrocumulus clouds surrounding the sun in both cases.

As expected with the intra-scanline approach to dynamic programming, the height maps have many vertical streaks. These streaks occur when the maximum path through the matching score matrix has vertical or horizontal segments. These can either occur because of an occlusion or because a patch in one image is falsely matched with a several consecutive rows in the other image. Occlusions typically occur on the upper or lower boundary of the cloud edges as sharp height gradients, whereas the false matches occur in the core of the clouds as smaller height gradients which are distinct from the surrounding cloud (e.g. Figures 5.5, 5.9 and 5.12) If these horizontal or vertical segments of the maximum path are filtered out, any occlusions associated with these segments are properly filtered out. Falsely matched regions will be filtered out as well because looking at the maximum path of a single column alone does not distinguish between occlusions and false matches. To improve the results, interscanline dynamic programming (Ohta and Kanade 1983) could be implemented which will better identify occlusions and reduce the number of falsely matched segments. The number of falsely matched segments is inherently reduced in this work because the matching metric used is normalized cross correlation from a 41x41 pixel window. This window size incorporates support from 20 columns on either side of the column being matched, which reduces the 'streakiness' due to falsely matched cloud patches.

The altocumulus-01, broken-02, cirrocumulus-03, cumulus-05 and twolayer-10 give the best results when evaluating the height map qualitatively. Occlusions are properly identified when they occur (by a horizontal or vertical segment in the maximum

path leading to a height gradient on the edge of a cloud) and the height map is consistent across the image with heights comparable to the ceilometer readings. In the twolayer-10 case, the method was able to correctly identify the height of each layer present. This was one of the simplest possible two layer cases to test because there was no occlusion of the higher layer due to the lower layer in one. The two layers were on differing sets of epipolar lines which meant that the row matching procedure for each column (i.e. determining the maximum path) did not have to distinguish between layers. The performance of the stereo algorithm for the cirrus-04 and cumulus-06 was qualitatively poorer than the five cases just mentioned, with cumulus-06 case having a frequent occurrence of false matches. There was a large degree of dissimilarity of perspective between the two cameras in cumulus-06, with the low clouds being directly overhead in USI 1.8 and down near the horizon in USI 1.7 (Figure 5.9). Matching in this case was expected to be difficult. Case cirrus-04 was expected to perform better, although the 'wispy smooth' texture of the cirrus clouds has a larger scale than the 41x41 matching window used. Visually the two rectified images have a high enough degree of similarity that quality matching should be possible if a larger window is used. The lowsparse-07 and -08 cases were attempted to test the ability of the stereo camera setup to determine correspondence between clouds that are visually difficult to identify as the same cloud in the raw images. While the sun plagued the overall results in each case, the clouds and image columns not adversely affected by the sun were matched reasonably well and the triangulated cloud heights are reasonable when compared to the ceilometer measurements. These very low cloud fraction cases would benefit greatly from a space carving method, because much of the volume over the test site would have been labeled as clear which can help in the post processing of the stereo matching results. The overcast-09 case was attempted because it is nearly impossible to

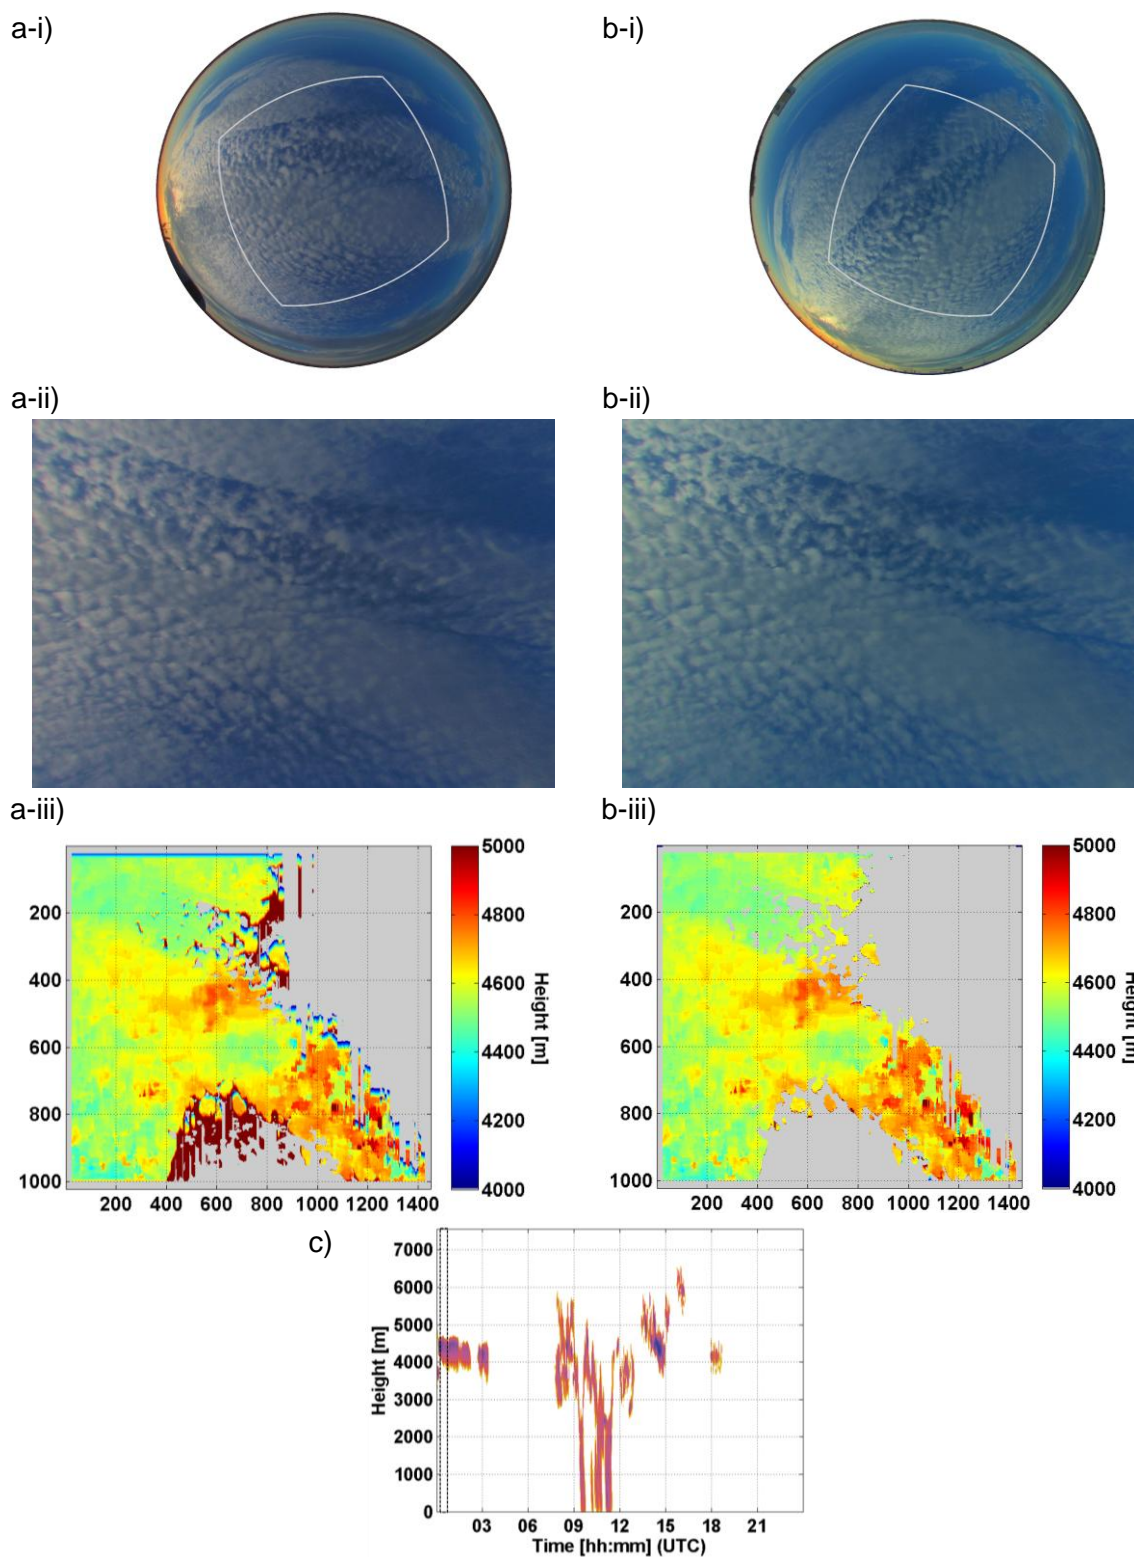
determine corresponding cloud features in this textureless full-sky cloud type. The quantitative results are reasonable, but this is mostly due to the pointing angles and field of view selected (section 5.3.2) and has less to do with correspondence being accurately determined. In other words, this case was designed to fail and the results should be completely discarded.

## **5.5 Conclusions**

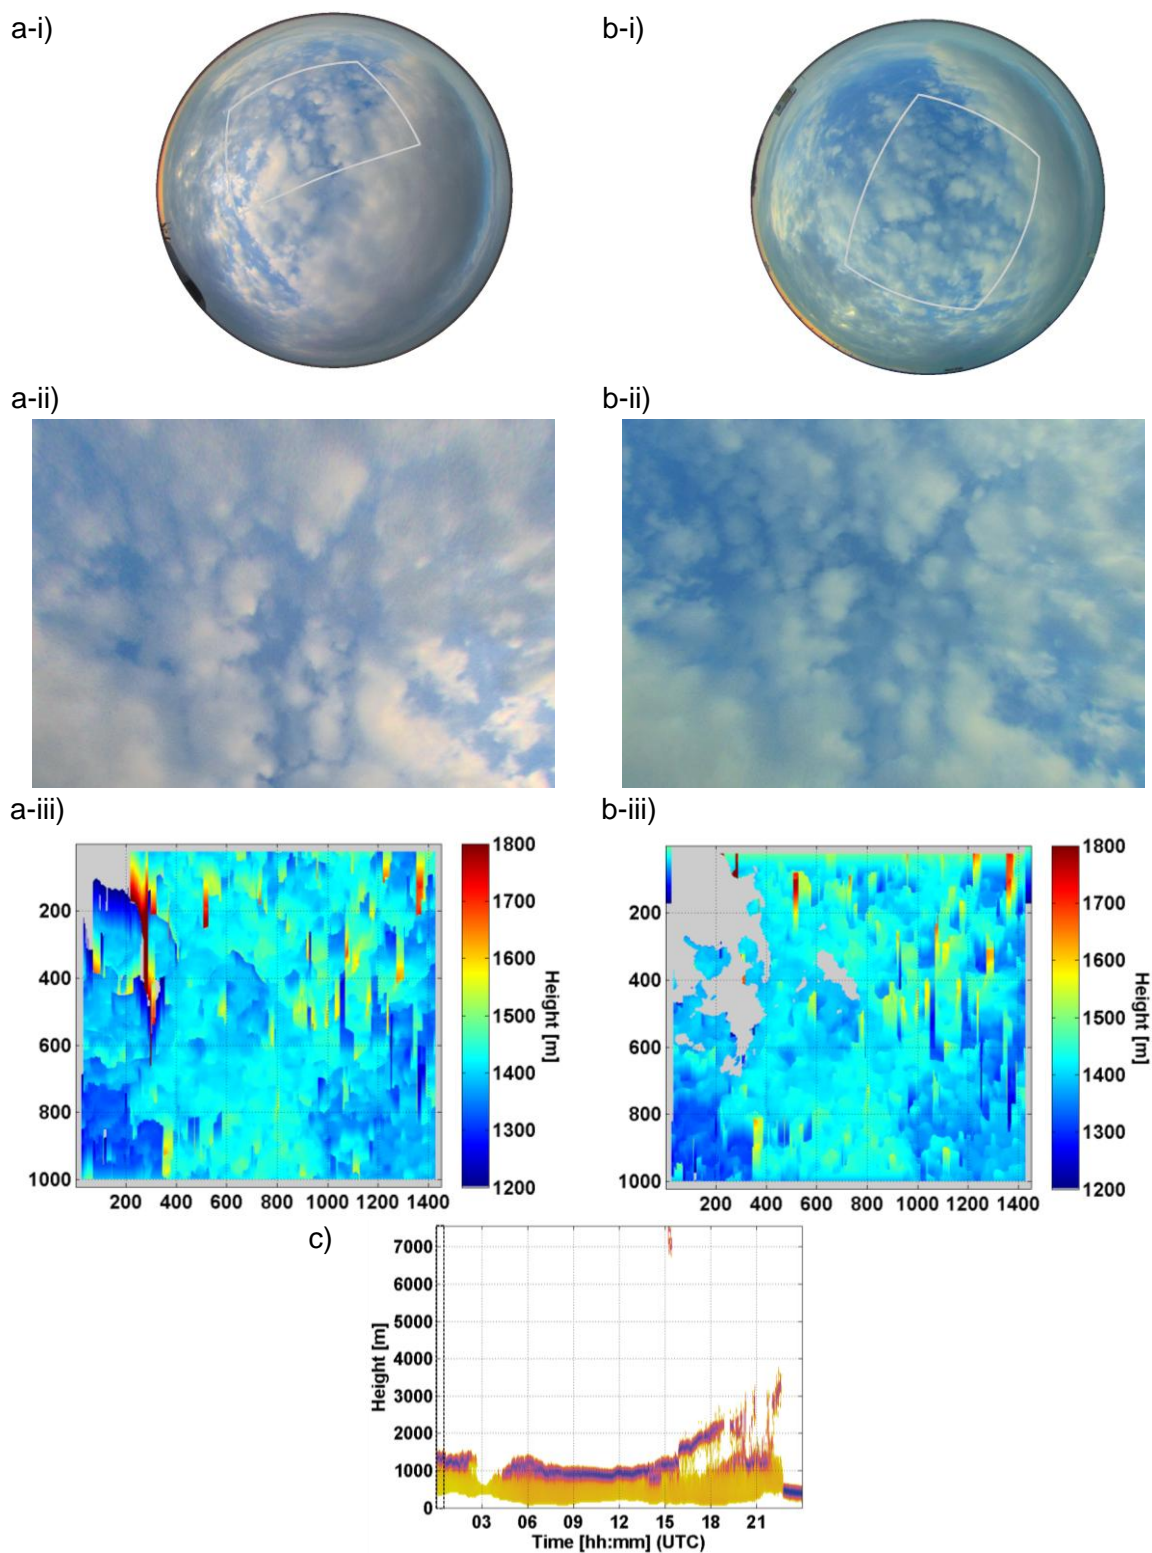
A method generate dense cloud position estimates from sky imagery with a large amount of radial distortion in a computationally efficient way was presented. The method performs well for higher clouds containing a reasonable amount of texture with which to determine correspondence. Due to the 1.7km baseline, the matching results for lower clouds yielded poorer results than the high clouds, in general. The use of dynamic programming for single cloud layer matching was shown to work very well and provide consistent and reasonable dense cloud height estimates. Multiple cloud layers were successfully distinguished and the heights determined for each layer was consistent with ceilometer measurements. This indicates the method may have promising outcomes for more complicated multi-layer cloud height cases assuming that the ordering constraint is not violated.

Improved cloud position information will benefit geometrically based deterministic solar power forecasting with sky imagers by allowing more accurate shadow position estimation. As future work, the algorithm developed here will be automated and parallelized on a GPU platform for operational forecasting use.

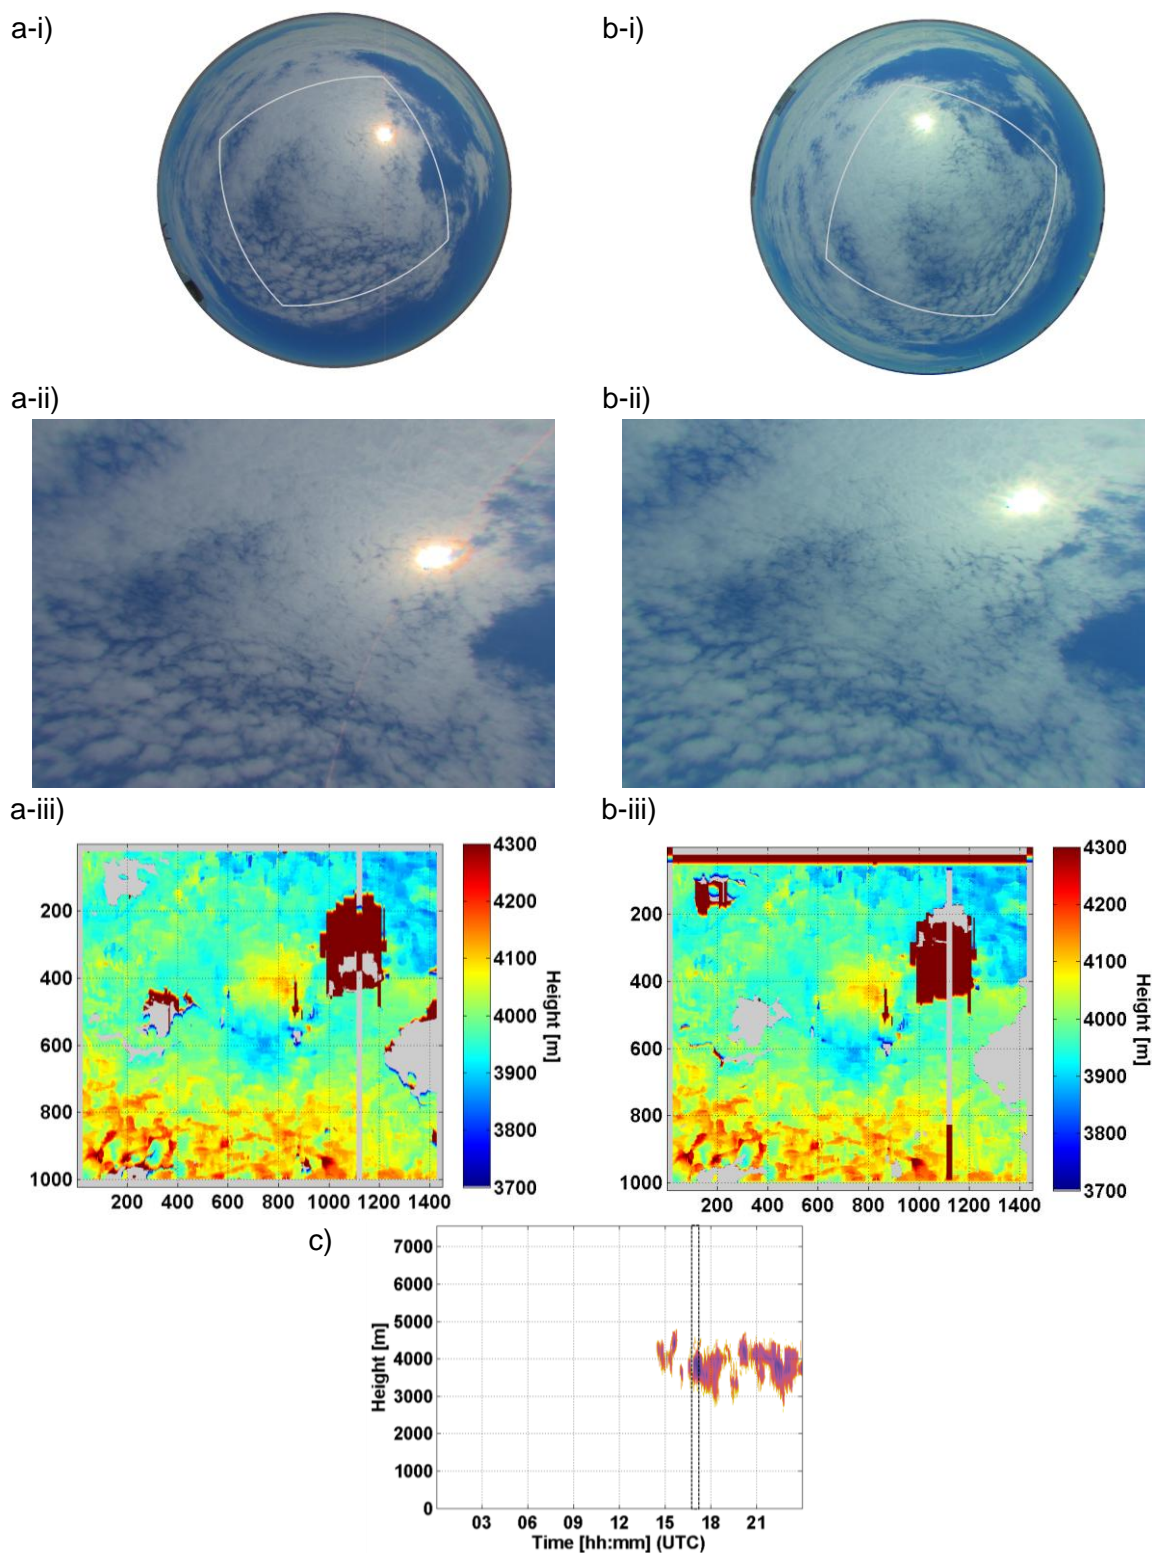





**Figure 5.4.** Stereography output for case **altocumulus-01** for (a) USI 1.7 and (b) USI 1.8. (i) Input images with field of view bounds; (ii) rectified images; (iii) cloud height map. Ceilometer returns are shown in (c) with image capture time indicated by dashed lines.



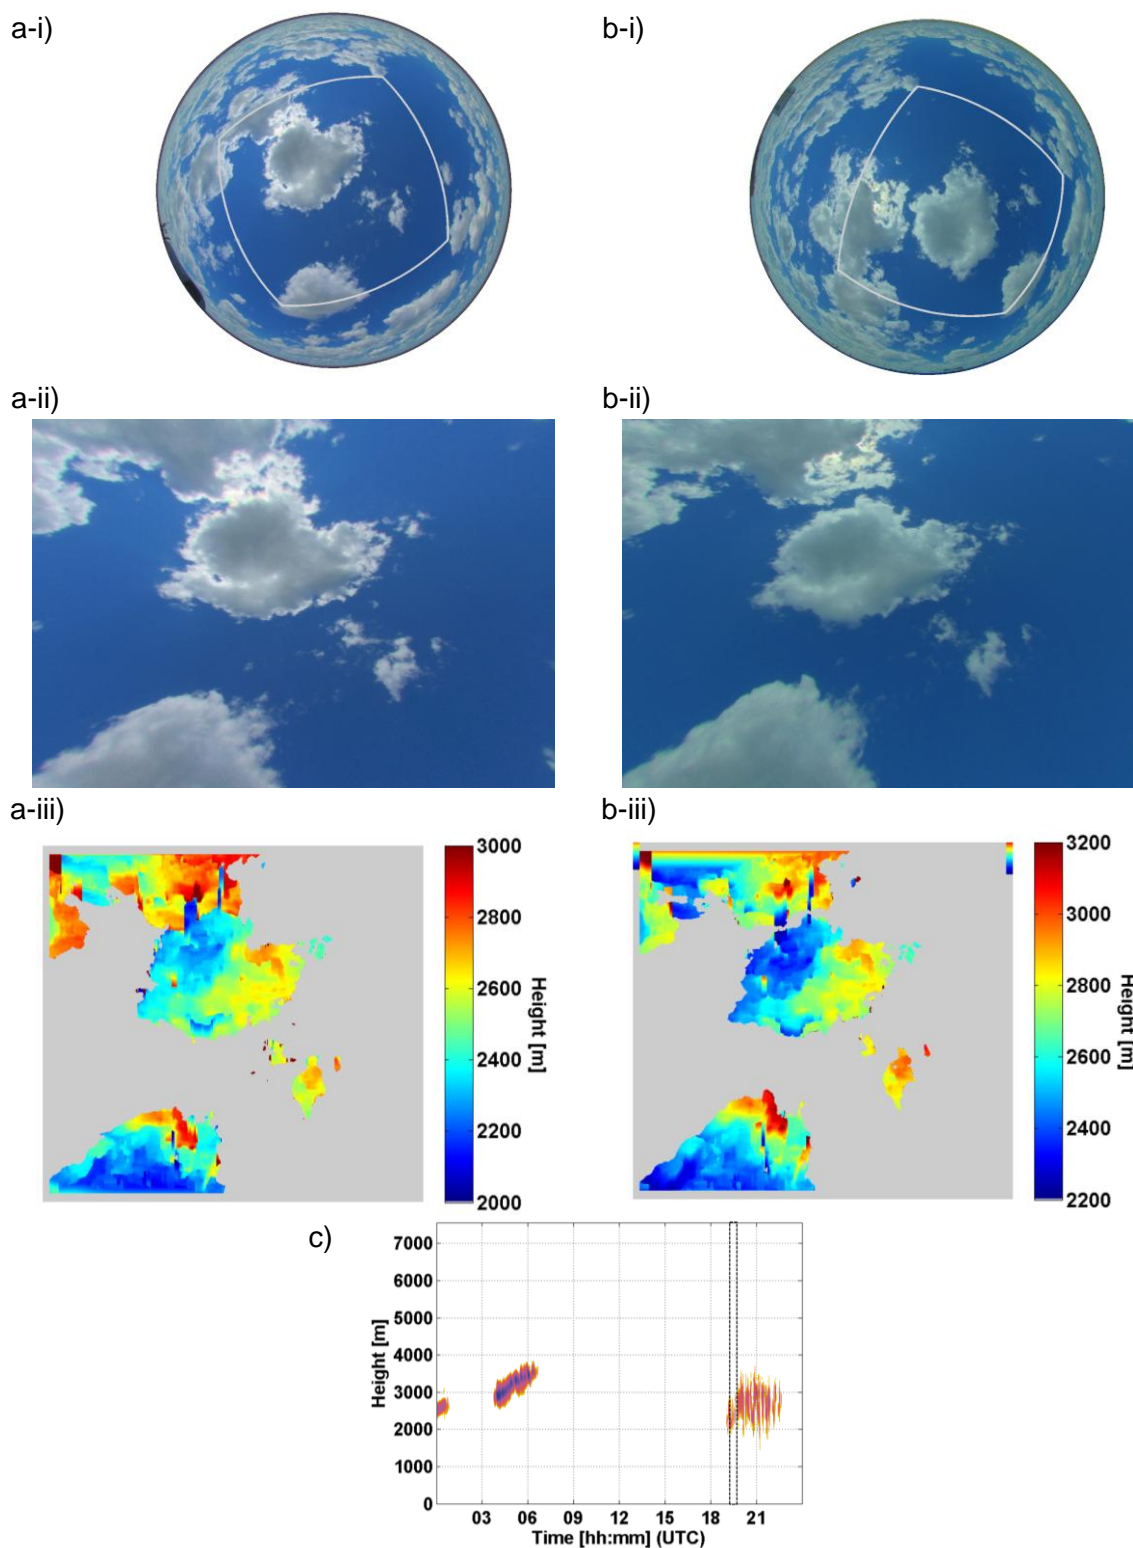
**Figure 5.5.** Stereography output for case **broken-02** for (a) USI 1.7 and (b) USI 1.8. (i) Input images with field of view bounds; (ii) rectified images; (iii) cloud height map. Ceilometer returns are shown in (c) with image capture time indicated by dashed lines.



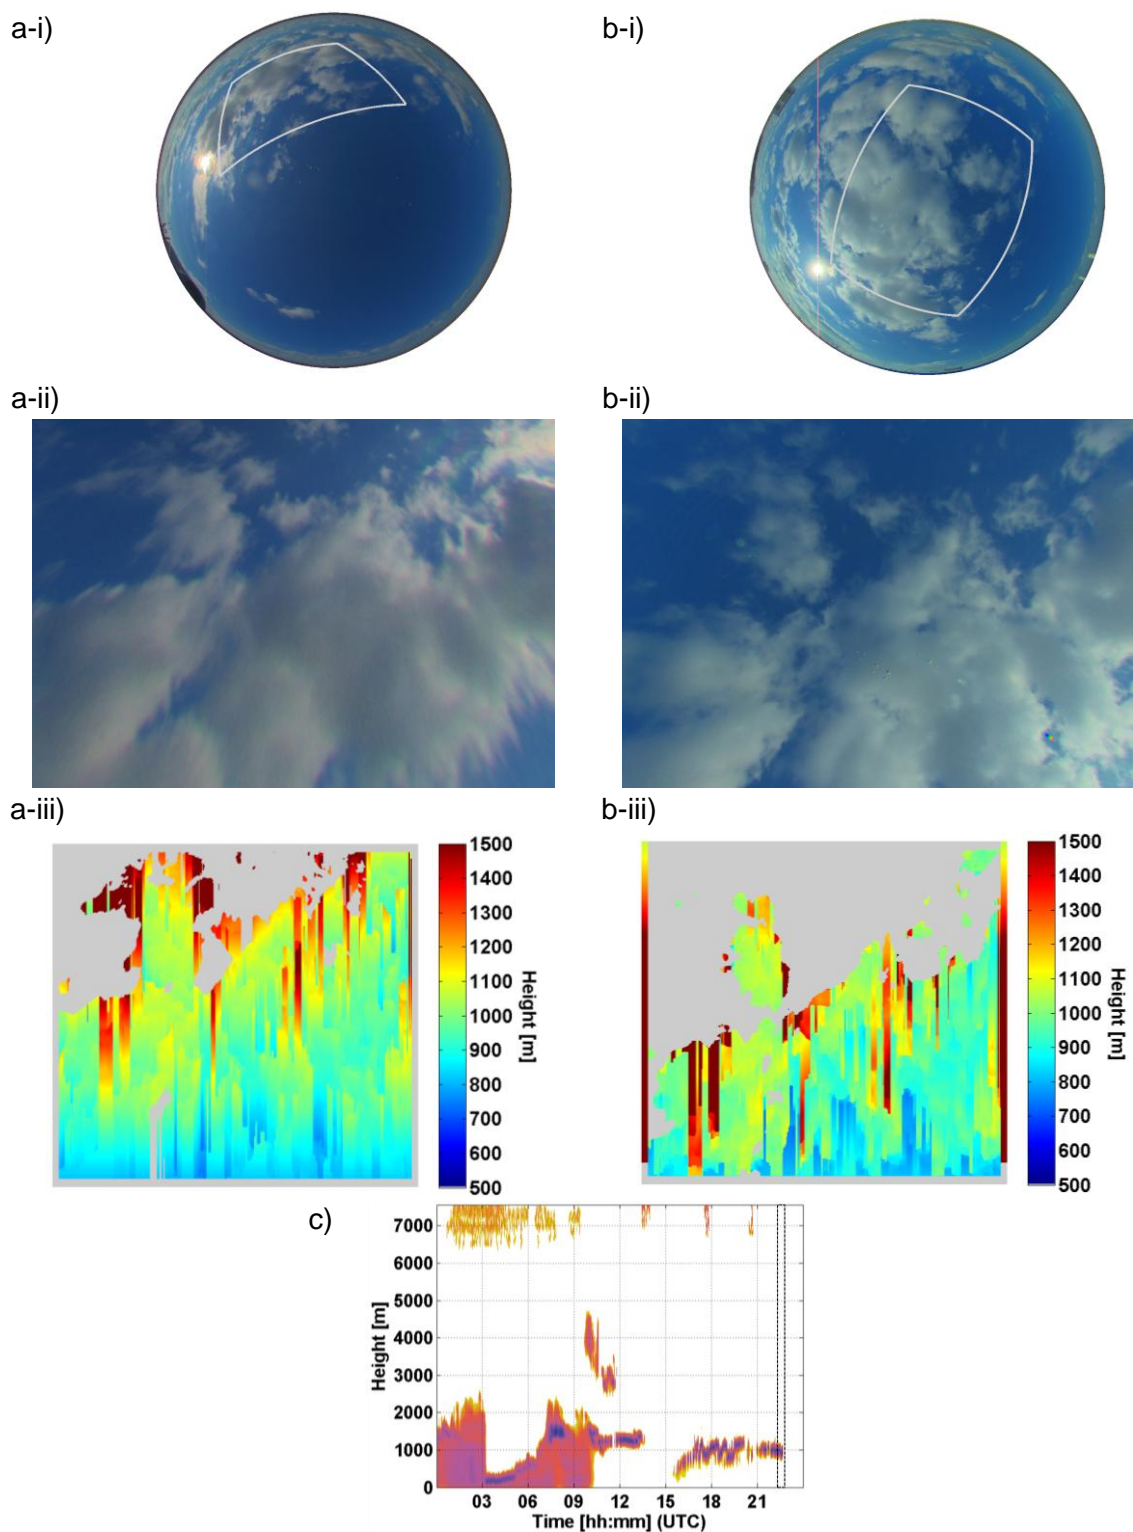
**Figure 5.6.** Stereography output for case **cirrocumulus-03** for (a) USI 1.7 and (b) USI 1.8. (i) Input images with field of view bounds; (ii) rectified images; (iii) cloud height map. Ceilometer returns are shown in (c) with image capture time indicated by dashed lines.



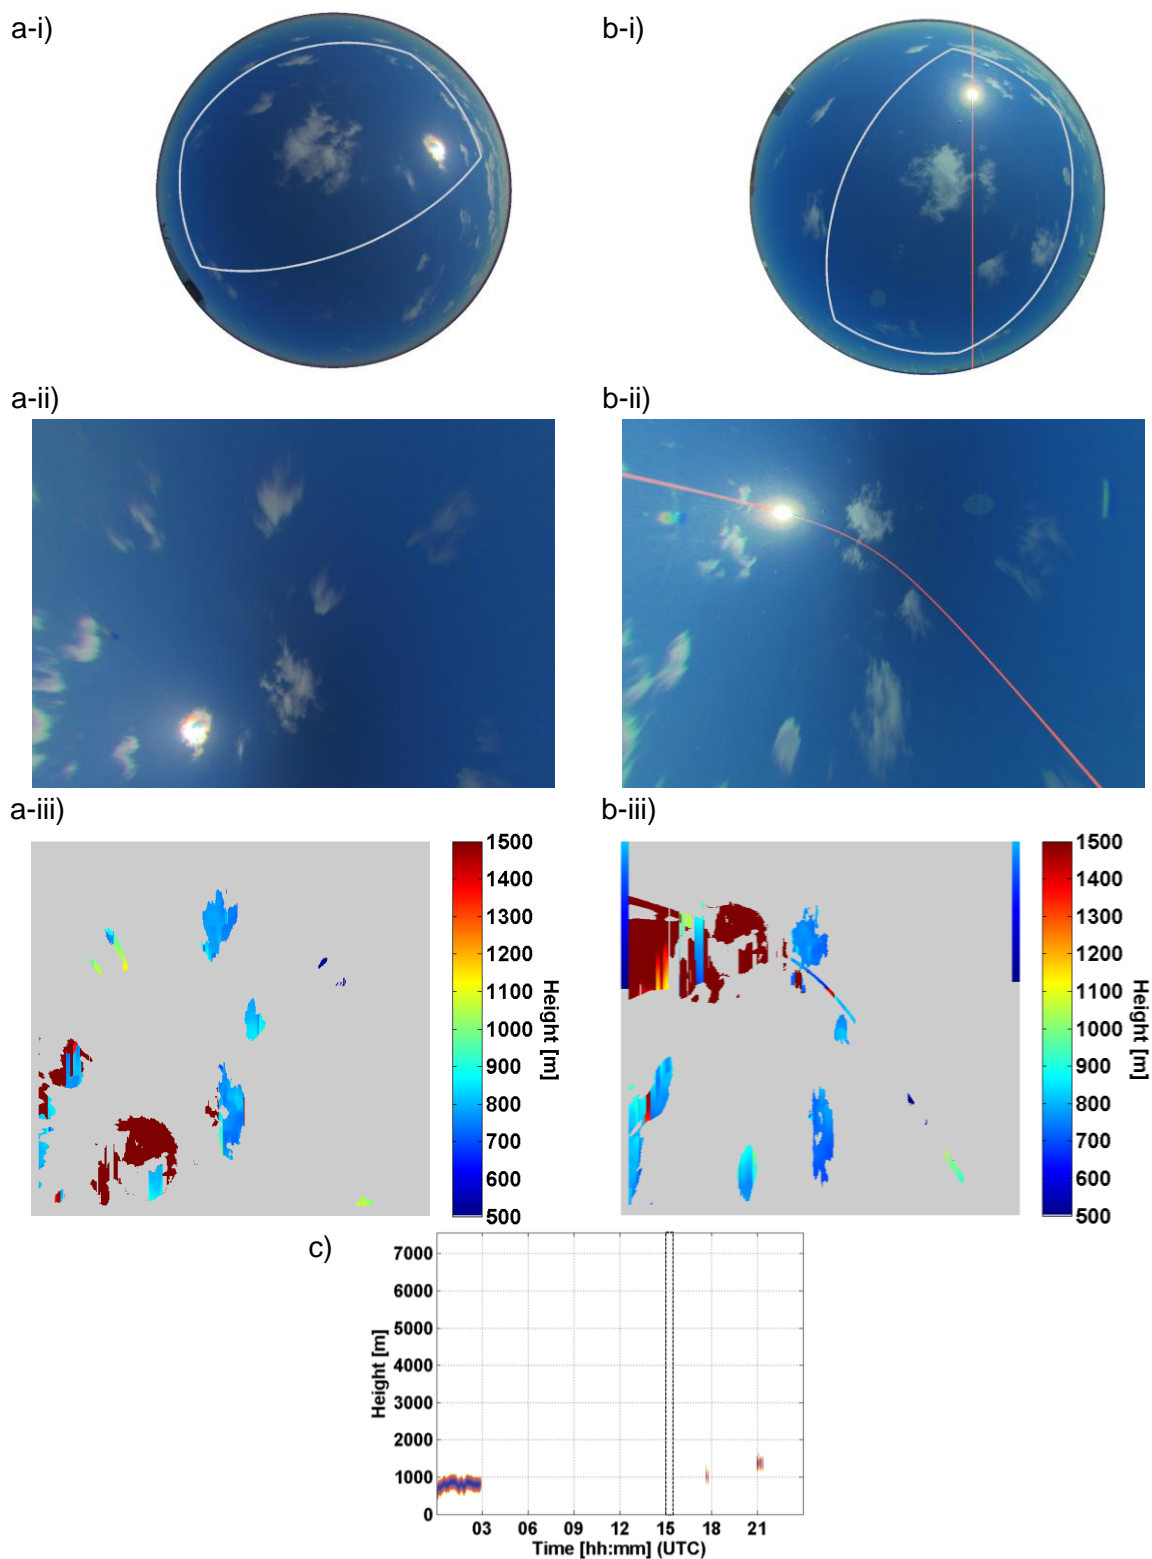
**Figure 5.7.** Stereography output for case **cirrus-04** for (a) USI 1.7 and (b) USI 1.8. (i) Input images with field of view bounds; (ii) rectified images; (iii) cloud height map. Ceilometer returns are shown in (c) with image capture time indicated by dashed lines.



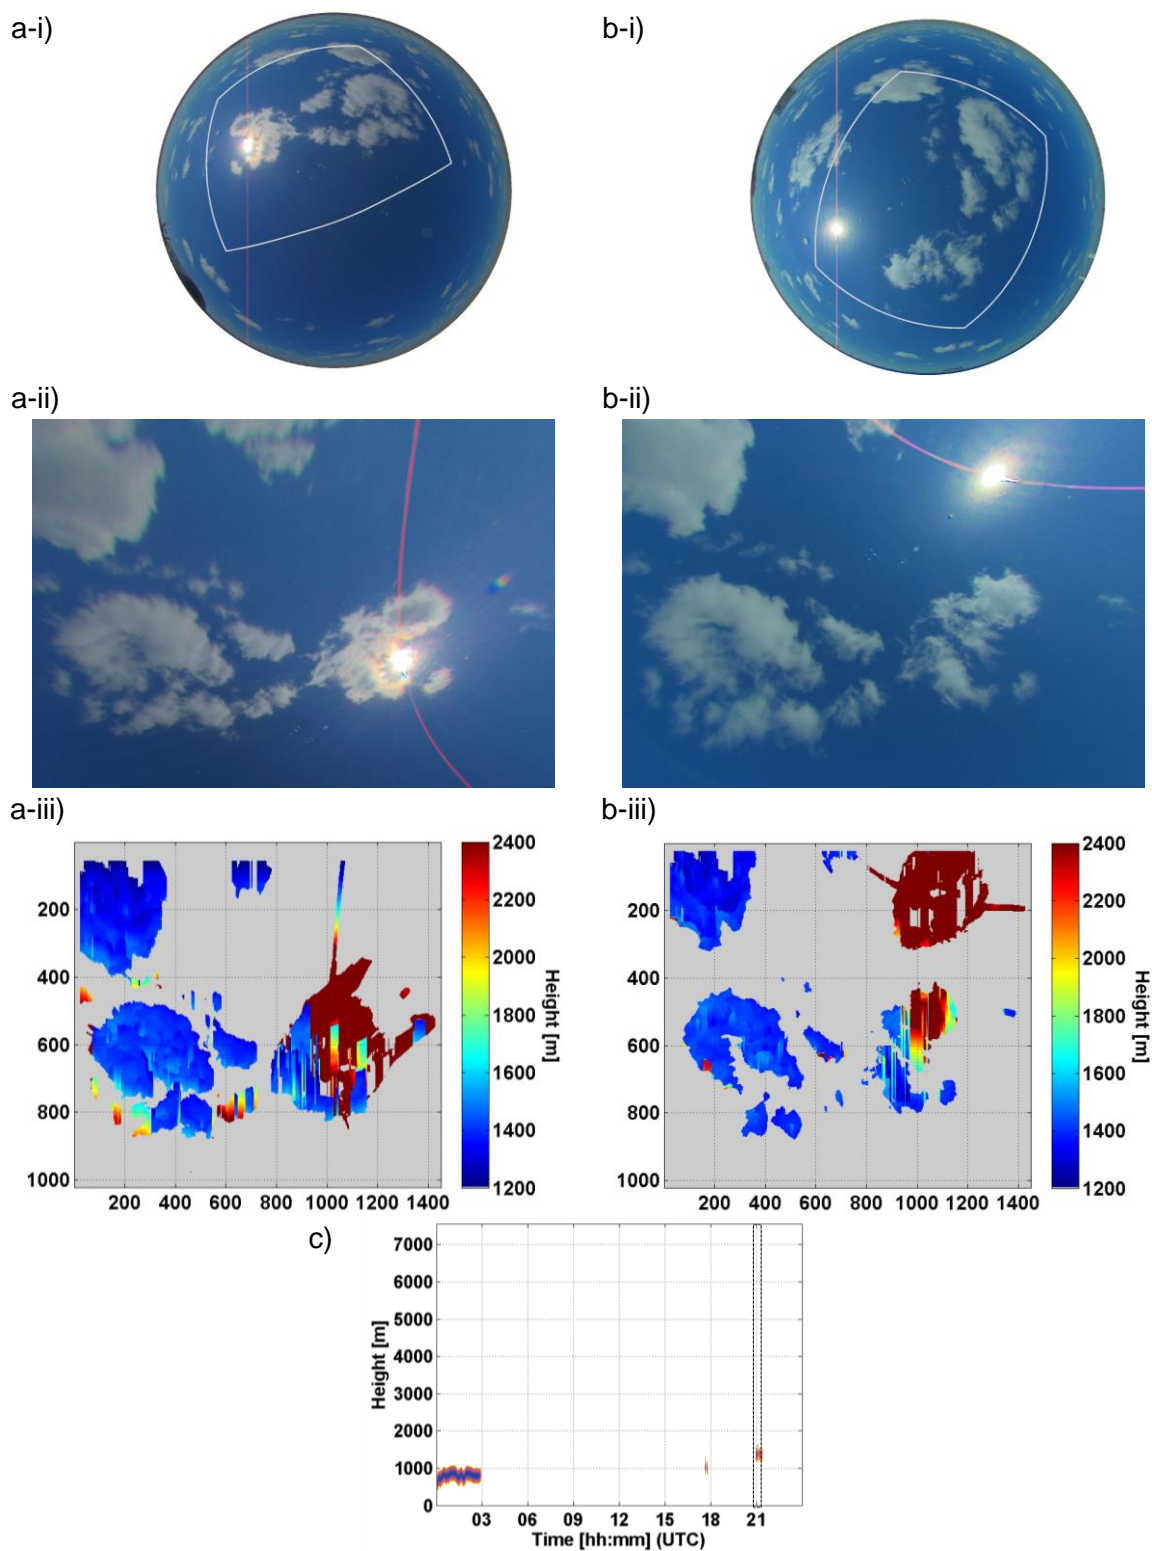
**Figure 5.8.** Stereography output for case **cumulus-05** for (a) USI 1.7 and (b) USI 1.8. (i) Input images with field of view bounds; (ii) rectified images; (iii) cloud height map. Ceilometer returns are shown in (c) with image capture time indicated by dashed lines.



**Figure 5.9.** Stereography output for case **cumulus-06** for (a) USI 1.7 and (b) USI 1.8. (i) Input images with field of view bounds; (ii) rectified images; (iii) cloud height map. Ceilometer returns are shown in (c) with image capture time indicated by dashed lines.

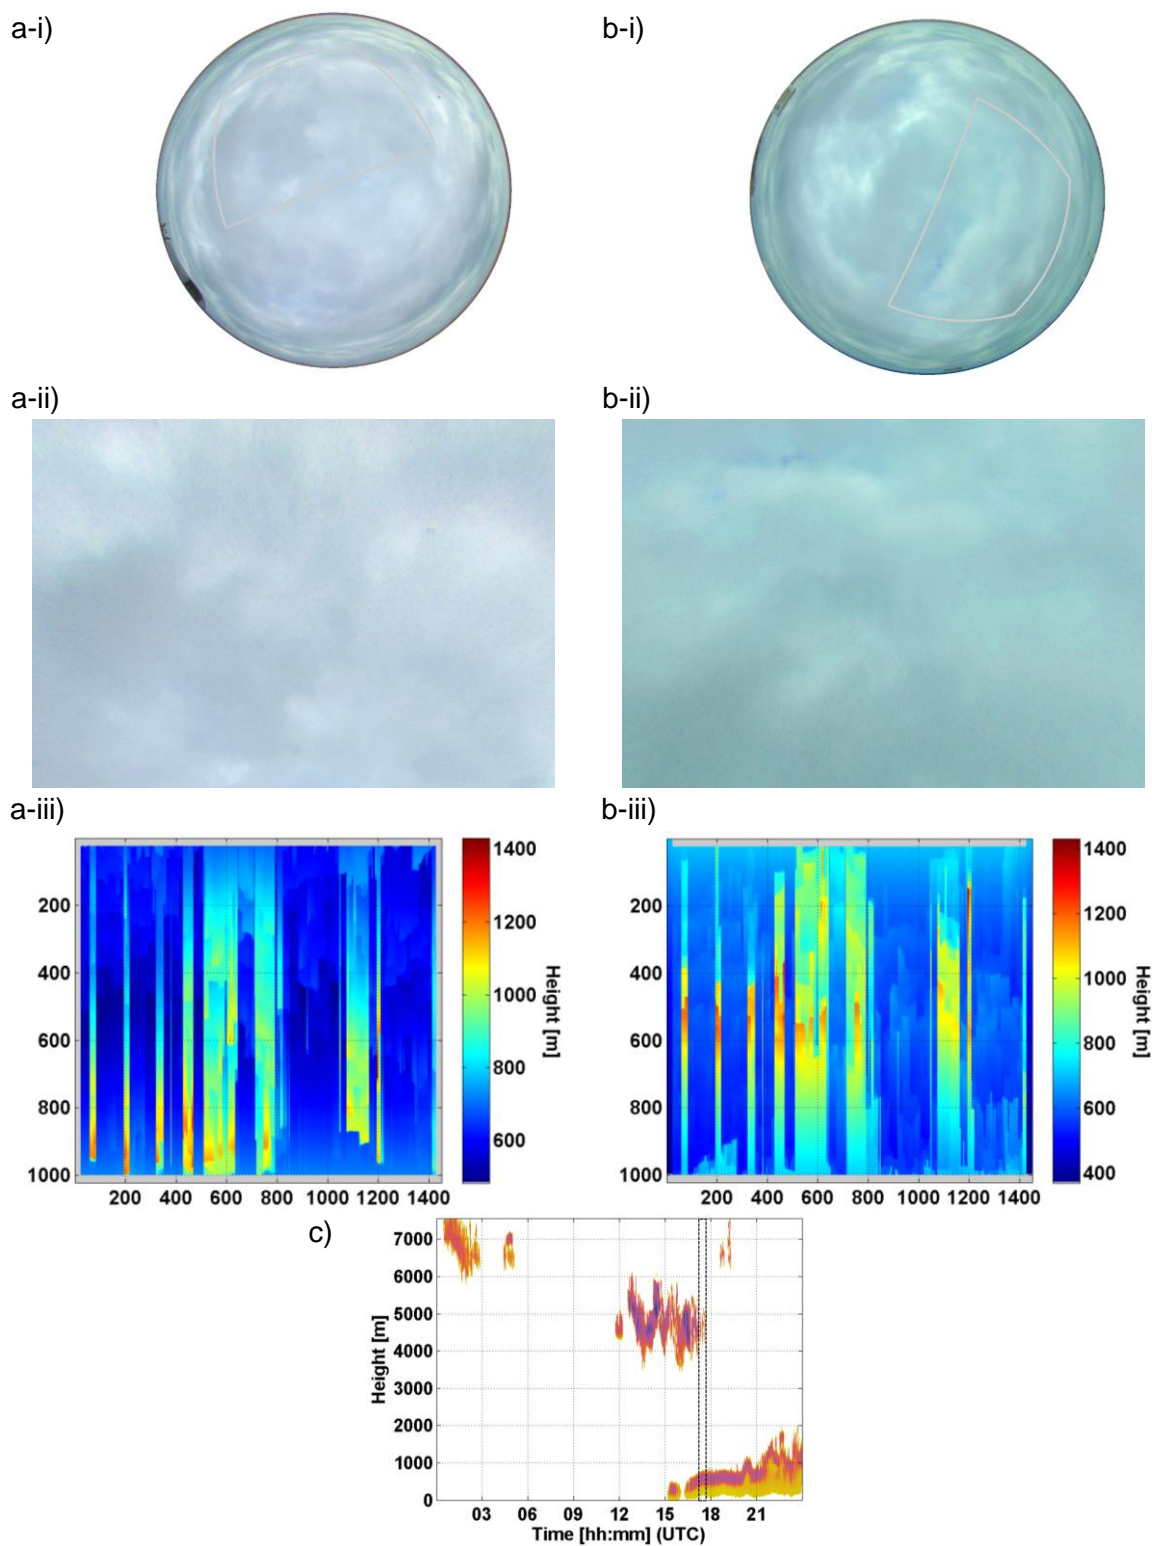


**Figure 5.10.** Stereography output for case **lowsparse-07** for (a) USI 1.7 and (b) USI 1.8. (i) Input images with field of view bounds; (ii) rectified images; (iii) cloud height map. Ceilometer returns are shown in (c) with image capture time indicated by dashed lines.

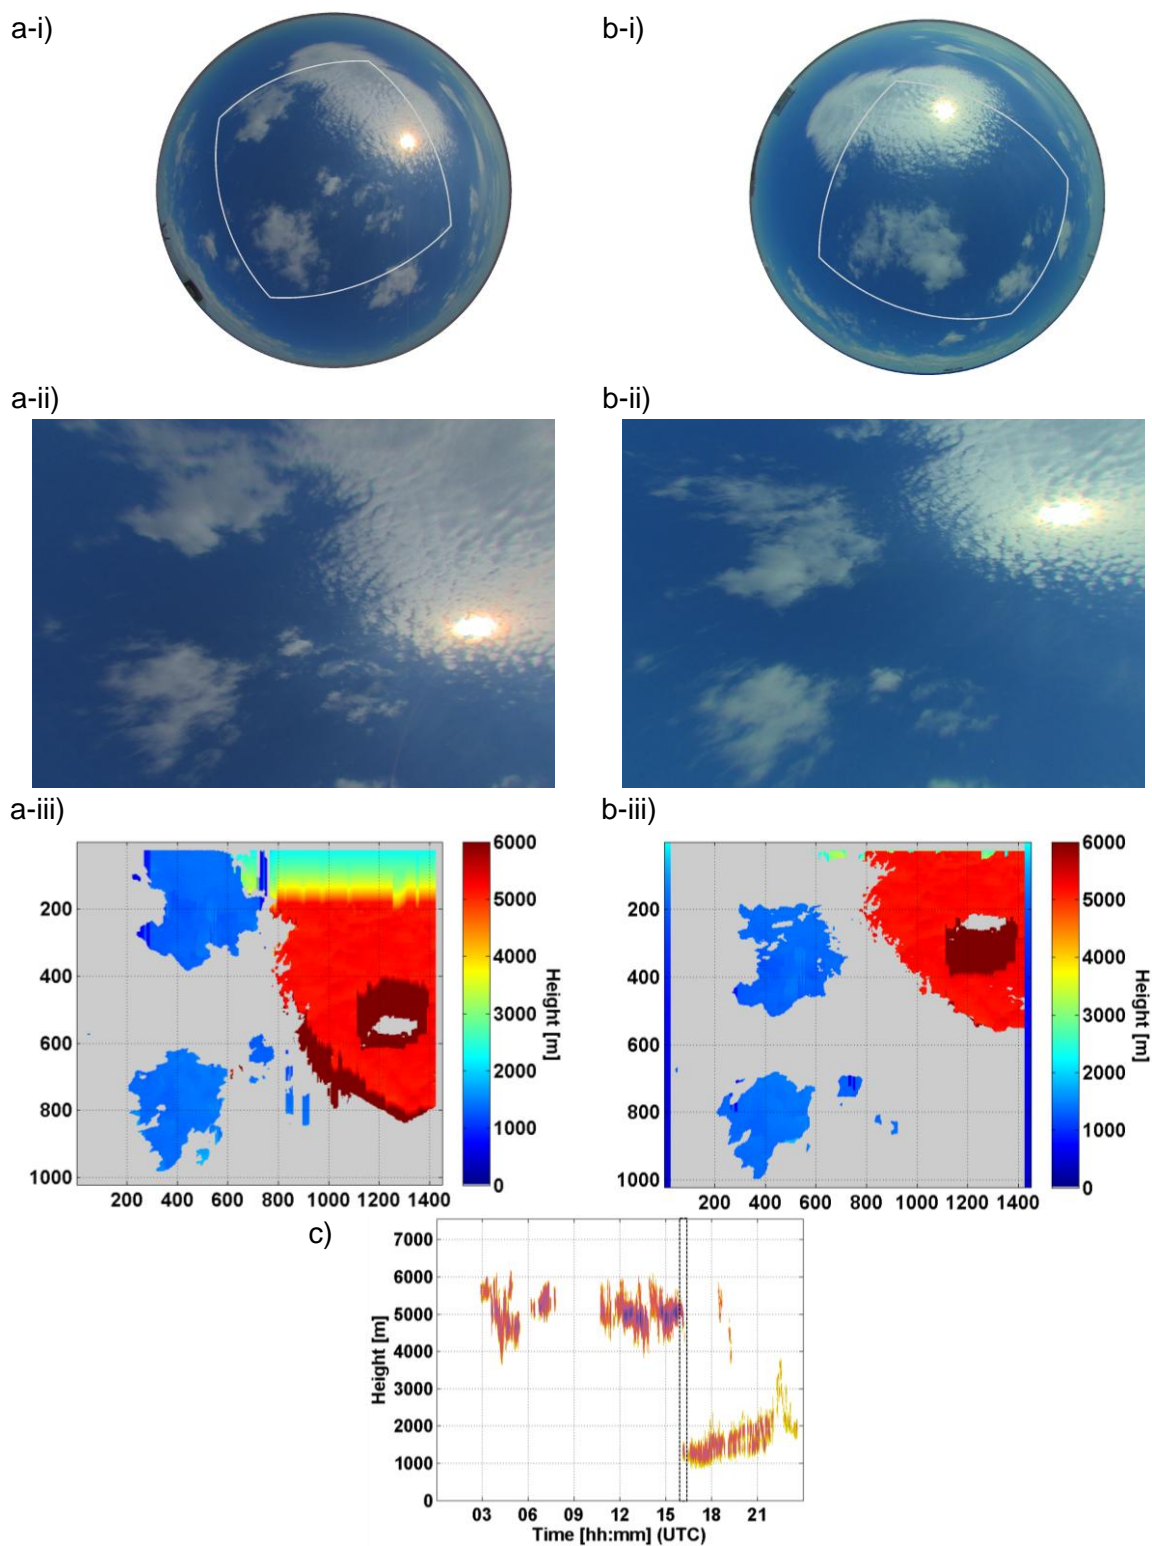


**Figure 5.11.** Stereography output for case **lowspare-08** for (a) USI 1.7 and (b) USI 1.8. (i) Input images with field of view bounds; (ii) rectified images; (iii) cloud height map. Ceilometer returns are shown in (c) with image capture time indicated by dashed lines.





**Figure 5.12.** Stereography output for case **overcast-09** for (a) USI 1.7 and (b) USI 1.8. (i) Input images with field of view bounds; (ii) rectified images; (iii) cloud height map. Ceilometer returns are shown in (c) with image capture time indicated by dashed lines.



**Figure 5.13.** Stereography output for case **twolayers-10** for (a) USI 1.7 and (b) USI 1.8. (i) Input images with field of view bounds; (ii) rectified images; (iii) cloud height map. Ceilometer returns are shown in (c) with image capture time indicated by dashed lines.

## **Acknowledgements**

In the course of preparing this chapter, the following publications were used:

Publication: Dense cloud stereo photogrammetry using dynamic programming. In preparation.

Co-Authors: Jan Kleissl

## 6 Solar Power Forecasting

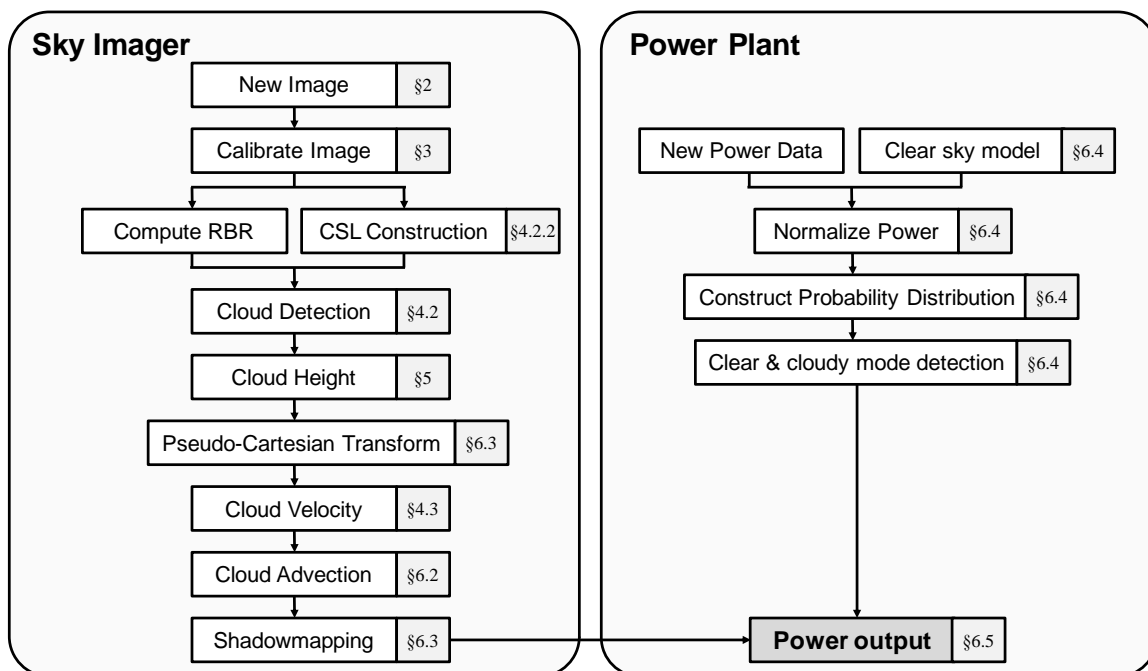
### 6.1 An Informal History of UCSD Solar Forecasting

While sky imagers have a long history of providing atmospheric observations, prior to 2010, they had seen little use for forecasting solar radiation. With the exception of Crispim et al. 2008, no other work with which the author is familiar had been published. Forecasting using cloud imagery had primarily been performed using data from satellites using techniques like frozen cloud advection. The University of California has had a long history of cloud observation using sky imaging systems at the Scripps Institute of Oceanography (Shields et al. 2013). In the fall of 2009, a Total Sky Imager (TSI) was installed on the campus gymnasium with the intent to monitor the campus microclimate. Based on past experience with satellite-based irradiance estimation (Urquhart 2010), the idea was conceived to use the sky imager for forecasting following similar approaches to those of satellites. The remote sensing techniques of Chapter 4 were subsequently developed from 2009 through 2010.

The question we asked of ourselves was how do we turn sky cover information into a forecast for an irradiance sensor? The answer to this question was ultimately what spawned the overall forecasting approach currently used. Instead of forecasting solely for the location of the sky imager (a procedure followed by many current forecast methods as of this writing), we attempted to generate forecasts for a large spatial extent surrounding the imaging system. To accomplish this, several techniques described in the remainder of this Chapter were devised.

## 6.2 Forecasting Methodology

The deterministic geometric projection sky imager solar forecasting method is constantly evolving as it is an active area of research, thus freezing the procedure in time to describe it in writing can only capture the current state of the art. The forecast procedure described here was developed by the author and used for the case study presented in section 6.5. Improvements and adaptations (e.g. Yang et al. 2014) are constantly building on this approach, but those developments will not be described here. The flowchart in Figure 6.1 gives a basic overview of the forecast procedure. The procedure is broken up into two steps: one that relies on sky-imager data and one that is designed for use with ancillary sensor data from the forecast target of interest (e.g. a pyranometer or a set of photovoltaic panels). This section provides an overview of the procedure, and details of each procedure are left to their respective sections.



**Figure 6.1.** Flow chart showing basic operations for constructing the power forecast in the Copper Mountain case study. Relevant sections of the dissertation are indicated where appropriate.

After a new image is collected, it is cropped, corrected for nonuniformity (section 2.5.2 and 2.5.3), and image-specific masks are generated (for the sun, shadowband, etc.). Due to potential instabilities in the internal configuration of the camera (i.e. the intrinsic geometric calibration parameters may vary if the system is moved, bumped, or maintenance is performed), the geometric calibration (section 3) should be kept current. Using the geometric calibration, a map of scattering angle based on the current solar zenith angle is constructed for the entire image (also referred to as the sun-pixel-angle). Radiance calibration similar to that for the WSI could be performed as well (Shields et al. 1998a).

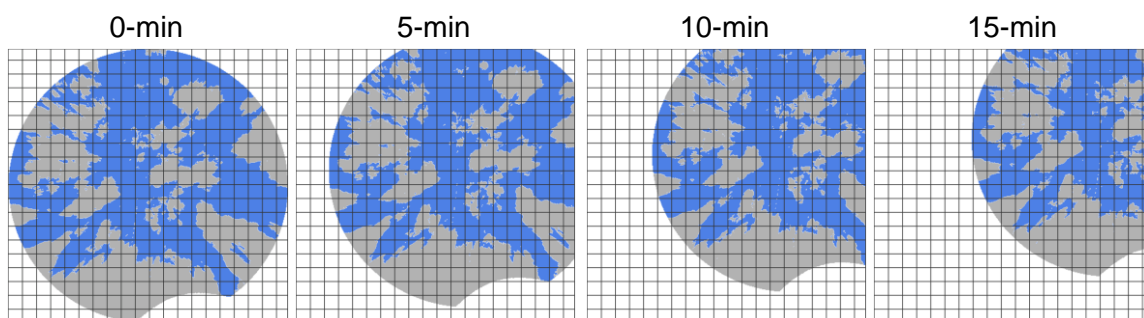
Following the above set of preprocessing steps, clouds are detected (section 4.2) and cloud altitude is computed (e.g. section 5, however nearby ceilometers located at airports are often used as a supplement for stereographically determined cloud height). The binary cloud/no-cloud (or trinary clear/thin cloud/thick cloud) information is still in the original distorted image coordinates  $(i, j)^T$ , but what is needed is georeferenced cloud mapping (see section 6.3 regarding georeferencing). To obtain this, the pseudo-Cartesian transform following Allmen and Kegelmeyer (1996) is applied:

$$\mathbf{x}_{i,j} = (x_{i,j}, y_{i,j}) = H \tan \theta_{ij} \begin{bmatrix} \sin \phi_{ij} \\ \cos \phi_{ij} \end{bmatrix}, \quad 6.1$$

where  $H$  is the cloud height above the camera,  $\Phi_{ij} = (\theta_{ij}, \phi_{ij})^T$  is view direction of pixel  $(i, j)$ , and  $\mathbf{x}$  parameterizes the plane at  $Z = H$ . This transform requires geometric calibration of the imaging system such that each pixel has a known look angle  $\Phi_{ij}$  in the world coordinate system (zenith-azimuth coordinate pair  $\Phi_{ij}$ ). Each transformed pixel  $\mathbf{x}_{i,j}$  is assigned a latitude and longitude based on the sky imager latitude and longitude, and a conversion from meters to either degrees latitude to or degrees longitude computed for the sky imager position (these conversions are functions of latitude).

Altitude is considered constant for this cloud mapping. The resulting georeferenced map of clouds is termed the “cloudmap,” which is a planar (constant altitude) mapping of cloud position above the forecast site.

Cloud velocity (section 4.3) is then used to advect the planar cloudmap to generate a cloud-position forecast for each forecast interval. Typically, the cloud position is advected at 30 second intervals out to a 15 min forecast horizon for every new image captured (Figure 6.2). The use of 30 seconds is a throwback to the capture frequency of the TSI used in the initial forecasting investigation (Chow et al. 2011). The forecast horizon of 15 minutes is limited by practical considerations such as significant changes in cloud morphology or viewing angles which render cloud images older than 15 minutes almost unrecognizable when compared to current conditions.



**Figure 6.2.** Sequential cloud advectations for a single forecast issue. The cloud positions are shown for the nowcast (a), along with the 5 min (b), 10 min (c), and 15 min (d), cloud-position forecasts.

The forecast domain can be defined as any set of latitude, longitude coordinates, either, for example, in a regular grid, a set of regular grids or polygons, or a list of latitude, longitude pairs. For each latitude, longitude coordinate, a ray is traced along the vector to the sun and the intersection with the cloudmap is determined (Figure 6.4). If the intersected point is clear, that ground location is deemed clear, whereas if the intersection is cloudy, the ground point is deemed shaded by cloud. Repeating the shadow-mapping process for each latitude, longitude coordinate point constructs a map

of cloud shadows (termed a shadowmap). When a regular grid of latitude, longitude coordinates is used, the resulting cloud shadows form a single map. This ray tracing process provides information about which locations on the ground are shaded. Shadowmaps are constructed for each advected cloudmap out to the 15 min forecast horizon.

The set of shadowmaps contained in a single forecast issue provide only a prediction of the progression of cloud shadows as they move over the ground. This information is binary (or trinary) and does not provide power output or irradiance estimates. To estimate the solar power striking the points of interest, past observations are used to characterize the transmissivity of the clouds and clear sky (section 6.4). The characteristics transmissivities are then weighted by the corresponding ground coverage fractions (from the progression of shadowmaps) for the subset of latitude, longitude coordinates of interest to determine an average sky transmissivity for the set of coordinates. The average sky transmissivity is then converted to either power or irradiance time series, depending on the application.

### 6.3 Geometric Forecasting Approach

The geometric forecasting approach seeks to determine cloud-caused sun-to-ground obstructions by using the position information of detected clouds. A single camera gives only the angular position of the cloud observed; i.e. the ray  $\Phi_{ij}$  to the cloud at pixel  $(i, j)$  (section 3.3.1.2). To determine the distance to the cloud along ray  $\Phi_{ij}$ , additional information from another instrument is needed. One option is to use a ceilometer which provides the cloud based height, from which the position of the cloud can be determined by computing the intersection of ray  $\Phi_{ij}$  with a plane at  $Z = H$ , where  $H$  is the cloud height. Another possibility explored in Chapter 5 is to use a second




camera to determine the position of clouds in the sky. The planar approximation method will be discussed here.

A ground station at  $\mathbf{X}_g = (x_g, y_g, h) = (\mathbf{x}_g, h)$  is shaded when a cloud is at position  $\mathbf{X}_s = (x_s, y_s, H) = (\mathbf{x}_s, H)$ , where  $\mathbf{X}_s - \mathbf{X}_g$  points in the direction of the sun (Figure 6.3). The use of the tilde overbar  $\tilde{\phantom{x}}$  will be dropped in this section and all coordinates given are Euclidean (i.e. no homogeneous coordinates will be used). The intersection of the ray to the sun  $\Phi_s = (\theta_s, \phi_s)^T$  extended from ground station  $\mathbf{X}_g$  with a plane at  $Z = H$  is

$$\mathbf{x}_s = \mathbf{x}_g + (H - h) \tan \theta_s \begin{bmatrix} \sin \phi_s \\ \cos \phi_s \end{bmatrix}, \quad 6.2$$

where the expression is written as a 2-vector with a lowercase notation. For convenience, Figure 6.3 is drawn with  $\mathbf{x}_s$  aligned with  $\mathbf{x}_g$ , however it should be noted that this is not true in general.



**Figure 6.3.** Side view of ray tracing a vector from ground station  $\mathbf{x}_g$  to intersection point  $\mathbf{x}_s$  at cloud height  $H$  for ground station height  $h$  and solar zenith angle  $\theta_s$ . The sky imager is located at  $O$ . Note that  $\mathbf{x}_g/|\mathbf{x}_g| = \mathbf{x}_s/|\mathbf{x}_s|$  only for convenience of drawing (i.e. they point the same direction), but this is not true in general.


By assuming that the spatial coverage of the sky imager is small enough, an analogous expression to Eq. 6.3 can be devised such that georeferenced coordinates  $\mathbf{X}_g$  (i.e. expressing position  $\mathbf{X}_g$  in latitude, longitude altitude, Chow et al. 2011):

$$\mathbf{x}_s = S_1 \left( \begin{bmatrix} lon_g - lon_o \\ lat_g - lat_o \end{bmatrix} + (H - h) \tan \theta_s S_2 \begin{bmatrix} \sin \phi_s \\ \cos \phi_s \end{bmatrix} \right), \quad 6.3$$

where  $\mathbf{x}_s$  is the pixel position in the cloud map,  $lat$  and  $lon$  are the latitude and longitude, respectively, of the ground station (subscript  $g$ ) and the cloudmap origin (subscript  $o$ ), and  $S_1$  and  $S_2$  are a scaling matrices

$$S_1 = \begin{bmatrix} \lambda_{lon} & 0 \\ 0 & \lambda_{lat} \end{bmatrix} \quad S_2 = \begin{bmatrix} \mu_{lon} & 0 \\ 0 & \mu_{lat} \end{bmatrix}$$

where  $\lambda$  is a conversion from degrees longitude or latitude to cloudmap pixel coordinates, and  $\mu$  converts from meters to degrees longitude or latitude (depending upon the subscript). The assumption of small spatial scale allows the latitude and longitude to locally be considered as a regular Cartesian grid (an assumption which becomes invalid as the poles are approached. The georeferenced geometric approach is illustrated for UCSD campus in Figure 6.4.



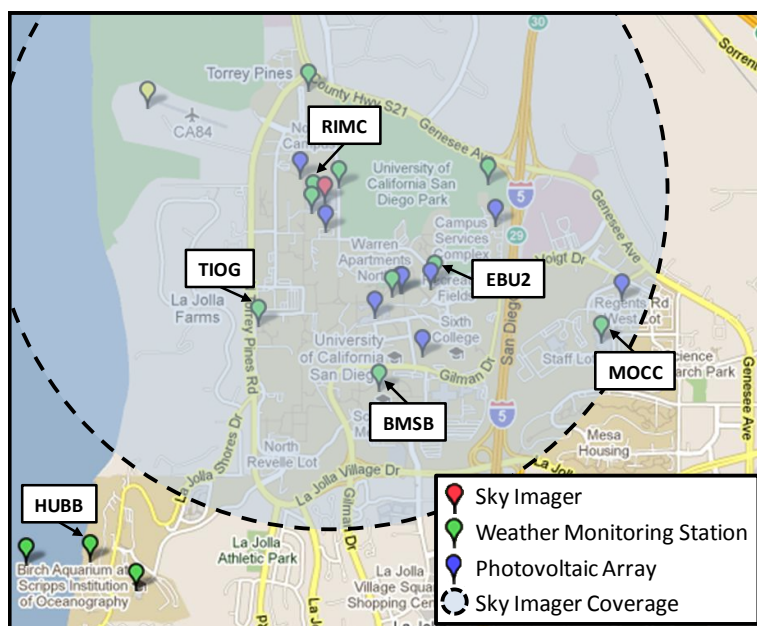
**Figure 6.4.** Ray tracing to construct a georeferenced mapping of shadows. The shadow value for a given point in the forecast domain grid is determined by tracing a ray along the solar vector and determining the cloud value at the intersection with the cloudmap.

The shadowmap introduced in section 6.2 is the result of computing intersections for a set of ground station locations  $\{\mathbf{X}_{g,k}\}$ . Typically, but not required, the  $\mathbf{X}_{g,k}$  form a regular grid such that actual maps can be produced. The value of the shadowmap associated with  $\mathbf{X}_{g,k}$  is set by the state of the cloudmap (clear, cloud, etc.) at the point intersection of the ray  $\Phi_s$  extending from  $\mathbf{X}_{g,k}$  and the plane  $Z = H$ .

### 6.3.1 Application to UCSD DEMROES Network

The UCSD testbed had six meteorological stations monitoring global horizontal irradiance (GHI) at 1 Hz with Li-COR 200SZ pyranometers (Figure 6.5, Table 6.1). The instruments are spatially close enough that they can be calibrated against each other on clear days to ensure consistency (see Lave et al., 2011 for a more detailed study of the irradiance fluctuations at these sites). The 1 Hz GHI data collection is faster than the image capture frequency, so validation GHI time series were constructed for each

station by using only the data points gathered at the time when the sky image was taken (no averaging is performed).




**Figure 6.5.** Map of UCSD showing sky imager coverage, weather stations, and PV arrays. The coverage area of the sky imager is a function of cloud base height. The region displayed here has a radius of approximately 2 km which would apply to a cloud base height of 1150 m at a maximum zenith angle of  $65^\circ$ .

**Table 6.1.** Location and status information for weather stations used to validate TSI cloud shadow maps.

Station Identifier	RIMC	BMSB	EBU2	TIOG	HUBB	MOCC
	Location					
latitude [deg]	32.8852	32.8758	32.8813	32.8790	32.8673	32.8782
longitude [deg]	-117.2401	-117.2362	-117.2330	-117.2434	-117.2535	-117.2231
altitude [m]	124	134	104	158	12	104
distance to TSI [m]	15	1,107	789	760	2,362	1,770
	Operational Status					
Sept 14, 2009	■	■	■	■	■	■
Oct 4, 2009	■	■	■	N/A	N/A	■
Mar 4, 2010	N/A	N/A	■	N/A	▨	N/A
Mar 10, 2010	■	N/A	■	N/A	■	N/A
	■ Operational	N/A Unavailable	▨ Cloud too low, out of range			


Binary irradiance time series were produced for ground stations in the UCSD DEMROES weather station network. When the sun is occluded the GHI is assumed to equal 40% of the clear sky value (i.e. clear sky index  $kt = 0.4$ , section 6.4). This value is chosen because the clouds that occurred on the days selected for validation were optically thick and reduced the irradiance to approximately this level. While the binary irradiance time series is not used in a quantitative evaluation, it is useful for a qualitative evaluation of the method.

Figure 6.6, showing MOCC on October 4, 2009 indicates that GHI variability can be estimated using a TSI. The morning period is dominated by false clouds due to MOCC's projected sky position being located in the shadowband area of the circumsolar region ( $\phi_{ij} < 35^\circ$ ). Once the projected sky position moved outside the circumsolar region (outer region) the TSI nowcasts the sky conditions between the sun and MOCC correctly 68.3% of the time in clear conditions and 80.4% in cloudy conditions, where clear is defined as  $kt > 0.7$ .



**Figure 6.6.** Time series (a) of MOCC station for October 4, 2009 showing measured GHI (blue) versus the TSI nowcast. The TSI shadow map is separated into circumsolar (red) and outer-region (green). Nowcasts of the TSI after 1100 PST (b),(c) and (d) match the trends in the measured data well.

Nowcast accuracy is expected to decrease with distance from the TSI due to the lower resolution at large sky imager zenith angles (in an extreme case on March 4, low clouds over the HUBB station 2.4 km to the southwest were out of sight of the TSI). However, the TSI was able to capture the trends in GHI at the HUBB station (Figure 6.7) on the other days. When HUBB was in the outer region, the TSI correctly predicted clear skies 60.0% of the time, and cloudy skies 90.6% of the time.



**Figure 6.7.** (a) TSI nowcast and measured GHI (blue) at HUBB station for March 10, 2010. Zooming in, (b–d) show the ability of the TSI to produce nowcasts for a distant site.

Since binary descriptive statistics are most appropriate to validate the binary TSI cloud estimate the clear-cloudy co-occurrence was tabulated (Table 6.2) for images taken when the cloud fraction  $cf$  was in the range  $0.2 < cf < 0.8$ . This restriction on cloud fraction is used to ensure that the results are not biased by nearly cloud free or overcast conditions which would result in a near perfect sky condition nowcast. For the four days chosen and for all available stations (Table 6.1), the TSI correctly estimated the condition of the sky 69.7% of the time in the outer region. This ratio was significantly different in measured clear (59.4%) and cloudy conditions (81.5%). It was more often predicted cloudy when it was clear than predicted clear when it was cloudy.

**Table 6.2.** Distribution of measured versus TSI nowcast clear (CLR) and cloudy (CLD) conditions for skies with cloud fraction  $0.2 \leq cf \leq 0.8$  (fraction of hemispheric solid angle). In line 1, the first term denotes measured clear sky index, and the second is the TSI estimate, thus CLD/CLR means the ground station measured cloudy conditions ( $kt \leq 0.7$ ) but the TSI estimated clear.  $CLR_m$  and  $CLD_m$  are the number of clear and cloudy measurements, respectively, in the outer region. The last column shows the overall statistics (not conditioned on cloudy or clear). Data is for all available stations.

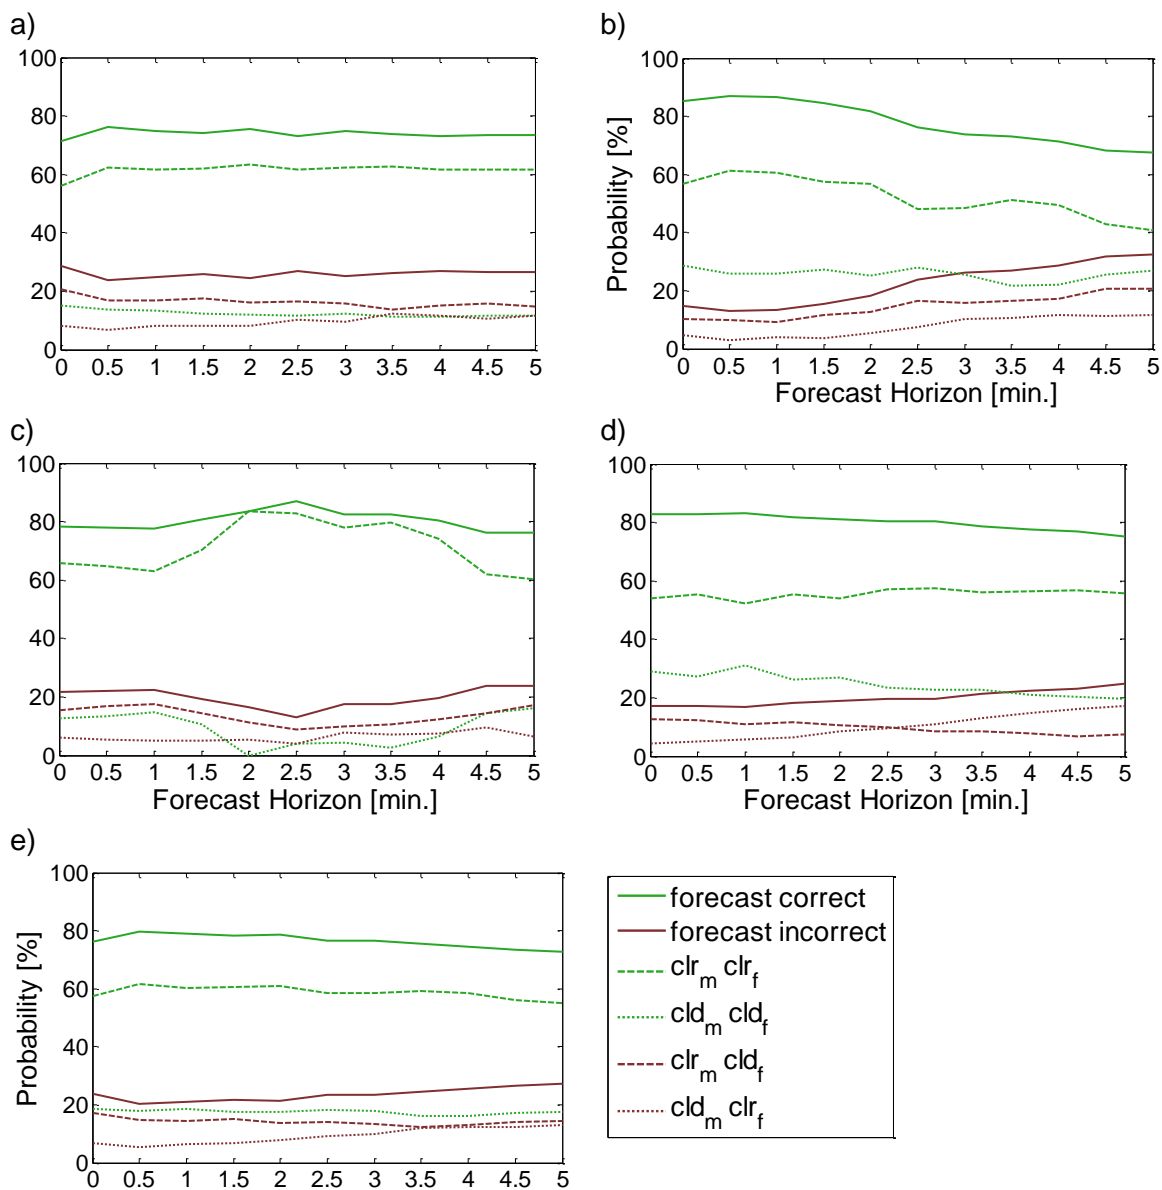
	CLR/CLR [%]	CLR/CLD [%]	CLR <sub>m</sub>	CLD/CLR [%]	CLD/CLD [%]	CLD <sub>m</sub>	% correct
Sept 14, 2009	47.1	52.9	594	28.3	71.7	368	56.6
Oct 4, 2009	76.2	23.8	441	10.3	89.7	419	82.8
Mar 4, 2010	63.1	36.9	157	33.3	66.7	144	64.8
Mar 10, 2010	52.7	47.2	110	7.0	93.0	201	78.8
Total	59.4	40.6	1302	18.5	81.5	1132	69.7

Forecasts for up to five minutes ahead are shown in Figure 6.8 for individual stations (Figure 6.8a-d) and for all four stations (Figure 6.8e). In general the sky condition was forecast correctly more often than incorrectly. However, this was primarily due to clear conditions being both measured and forecast. The most interesting comparison is when it is measured cloudy ( $cl_{d,m}$ , Figure 6.8): cloudy conditions are more frequently forecast than clear on average. On two of the days when the sensor condition was measured as cloudy, the forecast at certain horizons was more frequently forecast clear than cloudy, which was a significant error during this test. The specific source of these particular errors was not identified due to the forecast complexity, but the primary sources of error are considered to be cloud position inaccuracies due to calibration offsets, uncertainty in cloud height and velocity, and cloud detection errors (essentially each major component of the forecast is suspected to play a role). Cloud detection outside the circumsolar region is considered to be reasonable, however if the camera is not geometrically calibrated (which for this work it was a very rough calibration), the detected cloud is assumed to be at the incorrect position. The distance of many of these

sites poses the additional complication of projective distortion. Gaps in the clouds cannot be seen, and cloud sides are treated as cloud base, which leads to further errors for more distant stations. Cloud velocity is difficult to validate without the aid of a wind profiler or radiosonde, so the effect of cloud velocity is unclear. Finally, clouds are dynamic and evolving, which makes forecast horizons longer than 15 minutes nearly impossible without a dynamic and thermodynamic simulation (i.e. LES or NWP). Even at five minutes, clouds change significantly, further making the estimation of irradiance at a distant site challenging.

The root mean square error of the sky imager GHI estimates (Figures 6.6 and 6.7) indicate a negative bias (not shown). The negative bias in the circumsolar region is primarily due to false clouds decreasing the estimated GHI. In the outer region the negative bias is mostly caused by the cloudy bias noted above and cloud enhancement effects whereby the solar disk is not obscured but forward scattering of radiation by clouds located near the solar disk increases the diffuse radiation component and thus the overall GHI. Since the TSI estimate assumes a clear sky index of unity when the sun is not obscured it underestimates the clear sky GHI as well.





**Figure 6.8.** Combined TSI forecasts results for all stations on (a) September 14, 2009, (b) October 4, 2009, (c) March 4, 2010, and (d) March 10, 2010, (e) all four days. Correct forecasts are indicated in green and incorrect forecasts are indicated in maroon.

## 6.4 Cloud Optical Depth Estimation

With a sky imager and a ceilometer (or two sky imagers) we can obtain position and velocity information about clouds, but we do not know how much of the solar power is attenuated as it passes through the atmosphere. In order to obtain this information, we

use local observations of irradiance or power output to characterize the cloud transmissivity for use in future estimates.

Transmissivity, also known as clear sky index  $kt$ , is a ratio of the actual irradiance  $I$  to the irradiance observed in a cloudless sky (or clear sky, subscript  $clr$ ):

$$kt = I/I_{clr} ,$$

The clear sky index gives a normalized measurement wherein time-of-day effects are (nominally) removed, isolating the effects of clouds. The irradiance of interest in this work is the global irradiance  $GI$ , frequently defined with a *horizontal* reference plane (thus  $GHI$ ). The clear sky index for a horizontal surface is thus  $kt = GHI/GHI_{clr}$ . An analogous expression to  $kt$  can be devised for power output  $P$  by normalizing the power output by the expected clear sky power  $P_{clr}$ :

$$P_{norm} = P / P_{clr} .$$

Assuming standard testing conditions (STC) for rating the panels (ASTM E1036 - 08),  $P_{clr}$  can be estimated

$$P_{clr} = P_{cap} GI_{csk} / GI_{stc} ,$$

where  $P_{cap}$  is the nameplate capacity of the set of panels,  $GI_{stc}$  is the STC irradiance taken as 1000 [W/m<sup>2</sup>], and  $GI_{csk}$  is the clear sky plane-of-array irradiance.

The clear sky irradiance on a horizontal plane can be modeled using any clear sky model. Here, the Kasten clear sky model (modified by Ineichen and Perez, 2002) is used:


$$GHI_{clr} = a_1 I_o \cos(\theta_s) \exp\left(-a_2 m (f_{h1} + f_{h2}(T_L - 1))\right) \exp(0.01m^{1.8})$$

$$a_1 = 5.09 \times 10^{-5} alt + 0.868$$

$$a_2 = 3.92 \times 10^{-5} alt + 0.0387$$

where  $I_o = 1366 \text{ W/m}^2$  is the solar constant,  $alt$  is the altitude above mean sea level,  $m$  is relative optical airmass (Kasten and Young 1989),  $T_L$  is the Linke turbidity coefficient,  $f_{h1} = \exp(-alt/8000)$ , and  $f_{h2} = \exp(-alt/8000)$ . The Linke turbidity is a measure of the aerosol content of the atmosphere and is defined as the number of clean, dry atmospheres required to produce the extinction observed in a real atmosphere containing aerosols (Linke 1922). To convert from  $GHI$  to  $GI$  (i.e. the irradiance on a horizontal plane to that on a plane of arbitrary orientation), the Muneer transposition model is used (Page 2003). The diffuse fraction, which is a required input into the Muneer model, is obtained following Boland et al. (2008).

The cloudmap is binary (or trinary) with a clear and cloudy state (or clear/thin cloud/thick cloud states), so the surface projected shadowmap of the clouds is also binary (or trinary). Characteristic values of clear sky index  $kt$  or normalized power  $P_{norm}$  are determined using the probability density function (PDF) generated from recent observations (usually 2 hours). The characteristic values of  $kt$  or  $P_{norm}$  are extracted from the peaks of the respective PDF. The distribution of  $P_{norm}$  for an entire day at the Copper Mountain Solar 1 power plant is shown in Figure 6.9, where the peak near  $P_{norm} = 1.05$  corresponds to clear sky and the peak near  $P_{norm} = 0.45$  is the modal cloud transmissivity. The bimodal shape exhibited by  $P_{norm}$  (and also  $kt$ ) is common, but does not always occur. Clouds are very complex in nature which is reflected in daily distributions of normalized power and irradiance. It is interesting to note that the small peak near  $kt = 1.3$  corresponds to cloud enhancement effects Luoma et al. 2011.




**Figure 6.9.** Normalized power distribution for November 3, 2011 at the Copper Mountain Solar 1 power plant. The ordinate values were normalized such that the integral of the *unfiltered* probability density is one.

## 6.5 Case Study: 48 MW of Photovoltaics at Copper Mountain Solar 1

### 6.5.1 Experimental Setup

Two TSIs were installed at Sempra Generation's Copper Mountain Solar 1 power plant to validate the sky-imager forecast methodology in a utility-scale environment. The cadmium telluride thin-film panels for the 96 inverters covered approximately 1.3 km<sup>2</sup> and were tilted at 25° with a due south azimuth. The TSIs were spaced 1.8 km apart using the configuration shown in Figure 6.1. Fifteen calibrated reference cells provided plane-of-array (POA) global irradiance (*GI*) at 1 s, and five weather stations provided standard meteorological measurements at 1 s, including POA *GI* and *GHI* from Kipp & Zonen CMP11 broadband pyranometers. The forecast intervals selected matched the image-capture frequency; forecasts were issued every 30 s out to 15 min. The forecast domain was defined by a 4 × 4 km grid overlaying the plant with a resolution of 2.5 m per

forecast cell (1,600 × 1,600 cells), where each cell is resolved to a latitude, longitude, and altitude (altitude is obtained from a digital elevation model).



**Figure 6.10.** Outline of the 48 MW section of Copper Mountain, with sky-imager locations indicated. Each inverter's panel footprint is shaded with a different gray level.

## 6.5.2 Forecast Results

### 6.5.2.1 Forecast Error Metrics

To evaluate the forecast, mean bias error (MBE), mean absolute error (MAE), and root mean square error (RMSE) were computed over the given period during daylight hours ( $\theta_s < 80^\circ$ ):

$$MBE = \frac{1}{\langle P \rangle} \frac{1}{N} \sum_i^N (P_i^f - P_i^a),$$

$$MAE = \frac{1}{\langle P \rangle} \frac{1}{N} \sum_i^N |P_i^f - P_i^a|,$$

$$RMSE = \frac{1}{\langle P \rangle} \left[ \frac{1}{N} \sum_i^N (P_i^f - P_i^a)^2 \right]^{1/2}.$$

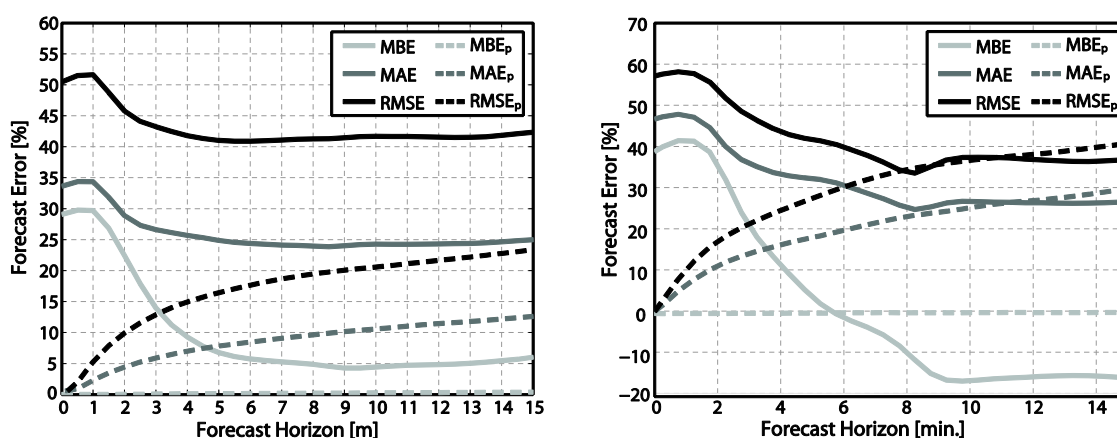
where the  $P_i^f$  is the  $i^{th}$  forecast power,  $P_i^a$  is the  $i^{th}$  power measurement, and  $\langle P \rangle$  is the average power production over the given period (i.e.  $\theta_s < 80^\circ$ ). Normalizing by  $\langle P \rangle$  produces higher relative errors than normalizing by capacity, and in general will give

lower errors on days with extended clear periods, and higher errors on mostly cloudy days. The sky imager generates a forecast every 30 s, whereas the plant reports power output every 1 s, so to compare the forecast to actual power production, a 30 s average of power output data centered on the image capture time was used. Error metrics were computed for each of the 31 forecast intervals out to a 15 min forecast horizon. While the metrics used provide a numerical evaluation of forecast accuracy, they are difficult to assess without a baseline comparison. The use of persistence as a baseline forecast is especially useful for short term forecasts. To generate a persistence forecast for comparison, the plant's aggregate normalized power was averaged for 1 min prior to forecast issue and was then applied to the remainder of the 15 min forecast window. Adjustments were made for changing solar geometry throughout the 15 min forecast window by computing the clear sky  $GI$  for each of the 30 s intervals.

#### **6.5.2.2 Forecast Performance**

The results presented here are for the week of November 9–15, 2011, which provided a variety of conditions with clear, partly cloudy, and overcast days. Forecast performance as a function of forecast horizon is shown in Figure 6.11. The forecast error of persistence steadily increases, whereas the forecast error of the sky imager starts off at a larger value because of shadowband issues and cloud decision errors near the Sun, and then levels off after about 3–4 min. The shadowband can block the entire sky region over the plant that contains the clouds, actually causing the irradiance impact, and as a result minimal or no data to generate an immediate-term forecast is available. As the shadowband (or circumsolar cloud decision error) is advected away, valid data from another part of the image moves to the region of sky over the plant in the path of the Sun, and thus a more accurate forecast can be generated.

Looking at individual days provides performance information for different cloud regimes (Table 6.3). On clear days the error is small but nonzero, largely because of the offset in absolute power predicted using the normalized power histogram to extract the modal clear value. Persistence uses a recent average of normalized power, which is more accurate than the most frequent daily value (i.e., the mode; see Figure 6.9) when the input solar signal is not affected by clouds. When there are clouds, the sky imager adds value because it can forecast when a ramp will occur and it can provide a reasonable approximation of the magnitude. Partly cloudy days with significant ramping occurred on the 10th, 12th, and 13th. The error on the 13th is shown as a function of forecast horizon in Figure 6.11b. Because of frequent ramping, the persistence forecast error increases significantly after a few minutes and the sky imager performs better.




**Figure 6.11.** Forecast performance as a function of horizon for a sky-imager forecast (solid) and persistence forecast (dashed) shown for (a) November 9–15, 2011, and (b) November 13, 2011.

**Table 6.3.** Sky Imager (SI) and persistence (P) forecast error at selected time horizons of 5, 10 and 15 minutes. Error is given as mean absolute error. Error is reported for individual days and the aggregate set of days as a percentage of average power generated during daylight hours. Superscript 'c' indicates the day was clear. The large errors on November 11<sup>th</sup> are due to cloud decision errors.

	5 min. [%]		10 min. [%]		15 min. [%]	
	SI	P	SI	P	SI	P
9 <sup>c</sup>	4.5	0.9	5	1.3	5.1	1.7
10	42.6	14.9	39	18.5	42	22.1
11	152.7	8	161.8	13.9	157.7	18.6
12	33.9	23.2	33.6	30.7	38.8	35.6
13	32	17.5	26.5	24.7	26.4	29.3
14 <sup>c</sup>	4.9	1.5	4.2	1.9	4.1	2.2
15 <sup>c</sup>	6.7	1.3	6.7	1.8	6.7	2
1 week	24.9	7.8	24.3	10.6	25	12.6

The ability of a sky imager to capture ramps is illustrated in Figure 9.21 for the 10 min forecast horizon slice. Constant values in the sky-imager forecast indicate periods when the plant is forecast to be entirely clear or entirely cloudy. Much temporal shifting of when a ramp is forecast to occur can be seen, in both the early and late temporal directions. Ramps are also missed and falsely predicted. The ramp forecast is directly related to how well the shadows predicted by the sky imager match plant observations. Errors in ramp timing are caused by any combination of inaccuracies in cloud decision, cloud height, camera resolution, geometric calibration, and cloud advection, as well as by differences in cloud morphology due to viewing angle. Because of the novelty of the system, each error source listed can be markedly improved, and overall ramp-forecast performance is expected to improve as well.





**Figure 6.12.** Midday 10 min forecast performance on November 12, 2011, showing how well (or, if your glass is half empty, how poorly) the sky imager captures ramps at 10 min in the future. A perfect forecast would have both curves matching exactly.

### 6.5.3 Conclusions

The results presented here represent the first attempt at a power forecast at a large solar plant using sky-imager forecasting. The MBE, MAE, and RMSE are larger than persistence in general. Error values reported (MAE, etc.) are gross metrics and do not focus on ramp-forecasting, which is a key value of the sky imager. In the future, stochastic-learning techniques (Coimbra et al. 2013) will be added and ramp-forecasting metrics will be developed to improve forecast accuracy and to better quantify skill.

Improvement in techniques is still needed for construction of short-term forecasts of high spatial and temporal resolution. Skill is demonstrated at longer horizons where the gross error metrics of MAE and so forth are comparable to persistence, and when the prediction of ramp timing is reviewed. This skill, however, is not currently sufficient for industry needs, and more work needs to be done. Many of these errors stem from inaccurate cloud detection or cloud-height determination because each impacts the

accuracy of determining whether or not a cloud is obstructing a given ground location. As described in Chapter 5, the next targeted improvement will be better geometric determination of clouds. Improvements here will reduce errors in ray traced shadow position.

## **Acknowledgements**

In the course of preparing this chapter, the following publications were used:

**Publication:** Intra-hour forecasting with a total sky imager at the UC San Diego solar energy testbed. *Journal of Solar Energy*, volume 85, issue 11, pp. 2881–2893, 2011.  
doi: 10.1016/j.solener.2011.08.025

**Co-Authors:** Chi Wai Chow, Matthew Lave, Anthony Dominguez, Jan Kleissl, Janet Shields, and Byron Washom

**Publication:** Sky imaging systems for short-term solar forecasting, chapter 9 in: Kleissl J, editor: *Solar Energy Forecasting and Resource Assessment*. Elsevier; 2013.  
ISBN: 978-0-12-397177-7, doi: 10.1016/B978-0-12-397177-7.00009-7

**Co-Authors:** Mohamed Ghonima, Dung Nguyen, Ben Kurtz, Chi Wai Chow and Jan Kleissl

## 7 Comparison of the USI and TSI

Contributing Authors:

S. M. I. Gohari, B. Urquhart, H. Yang, B. Kurtz, D. Nguyen, C.W. Chow, M. Ghonima and J. Kleissl

### 7.1 Summary

The performance of the Total Sky Imager (TSI) and University of California, San Diego Sky Imager (USI) instruments has been evaluated using correlation and dispersion metrics applied to clear sky index. The TSI was located at a 48MW PV plant in Nevada, USA, and the USI was located at the UCSD campus in San Diego, USA. Distributed pyranometer measurements over several square kilometers were available at both locations to provide spatial averages of irradiance. Clear sky index was forecast for a 15 minute horizon at 30 sec. intervals, and the correlation and Euclidean dispersion were evaluated. A persistence forecast was generated for each location to provide a reference for comparison. For the aggregate dataset, the forecast performance of the USI exceeded that of persistence beyond the 10 minute forecast horizon, whereas the TSI did not show skill over persistence. For the period of study, the USI gave correlations of clear sky index between 0.70 and 0.82, whereas the TSI provided correlations between 0.55 and 0.69. Overall, the mean value and the trend of the correlation results were more consistent for the USI than the TSI on the days evaluated. It is concluded that the USI is a more reliable solar resource assessment instrument for the wide range of atmospheric conditions evaluated.

## 7.2 Introduction

As nations around the world push for cleaner sources of energy, the need for technologies that facilitate the integration of renewable energy into the electricity grid is growing. Forecasting is a grid integration tool that can help the scheduling of power generation and dispatch. Rules and regulations for the energy industry differ by region, but all electric grids have multiple relevant time scales at which planners, operators, power generators, and utilities need advanced information about how much power will be produced. Solar power is one of the most scalable choices for renewable energy generation, and so forecasting of this resource is vital to reliably integrate it into the electricity grid.

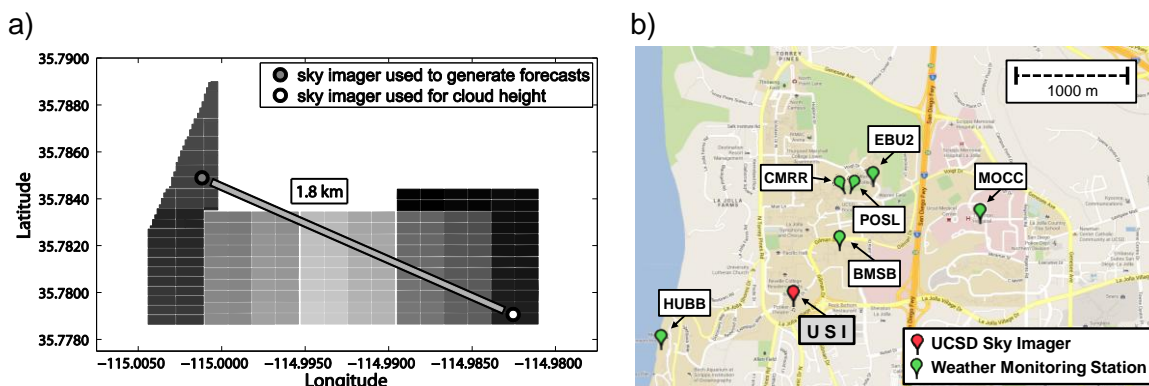
Short term power output forecasting, on the scale of 1 to 30 minutes, can help to enable the economic dispatch of solar power. Sky imaging systems have a history of providing atmospheric observation and monitoring, and have the potential to be effective monitoring systems for solar plant operations by providing both situational awareness and estimates of power output over a large spatial extent. Because of their ability to continuously monitor cloud conditions, the University of California, San Diego (UCSD) has developed a short term power output forecasting capability based on visible wavelength sky imagery (Chow et al. 2011). A related forecasting procedure was developed in Marquez and Coimbra (2013). Initial results on power output forecasting for 48MW of photovoltaic generation using the UCSD method were reported in Urquhart et al. (2013). Experience with their use at large solar power plants, however, has highlighted shortcomings in common imaging hardware Urquhart et al. (2012). The Total Sky Imager (TSI) is the most common device used for sky imaging applications, but no significant improvements have been made in the last 15 years. One of the key conclusions in Urquhart et al. (2012, 2013) is that the Total Sky Imager (TSI) instrument

had shortcomings for the purposes of solar forecasting. The University of California San Diego (UCSD), leveraging expertise in sky imaging technology (Shields et al. 2013), has designed a system specifically for solar resource assessment work (Urquhart et al. 2013, 2014). This system, called the USI, exceeds the TSI in both spatial and intensity resolution (i.e. number of pixels and bit depth, respectively). Here, a comparison of the two instruments is made to quantify the forecasting performance in both a relative and absolute sense.

## **7.3 Experimental Methods**

### **7.3.1 TSI Site: Copper Mountain Solar 1, 48MW**

To assess forecast performance, 1 sec real AC power data from a 48MW section of Sempra Generation's Copper Mountain Solar 1 was used (Figure 7.1a). This 1.3 km<sup>2</sup> plant section had 96 inverters receiving power from cadmium telluride thin film panels tilted at 25° with a due south azimuth. This section of the plant contained fifteen calibrated reference cells providing plane-of-array global irradiance (*GI*) at 1 sec and five weather stations providing standard meteorological measurements including plane-of-array *GI* and Global Horizontal Irradiance (*GHI*) from Kipp & Zonen CMP11 broadband pyranometers at 1 sec. Two TSIs were installed at 35.7849° N, 115.0012° W, 572 m MSL, and 35.7791° N, 114.9825° W, 580 m MSL, which are at the northwest and southeast corners, respectively. The imagery for the northwest unit was used to generate forecasts. The forecast intervals selected match the 30 sec image capture frequency; forecasts were issued every 30 seconds out to 15 minutes. Forecasts were generated for November 1<sup>st</sup> to the 21<sup>st</sup>, 2011, and July 12<sup>th</sup> to the 17<sup>th</sup>, 2012 (27 days total). Measured power output data is compared to the forecast by applying a 30 sec moving averaging to plant measurements.



**Figure 7.1.** (a) Footprint of the 48MW section of the Sempra Generation Copper Mountain Solar power plant. TSI locations are indicated with circle icons. The spatial resolution of the power data is indicated by the gray level shading. (b) University of California, San Diego campus with USI and weather stations indicated by teardrop icons.

### 7.3.2 USI Site: UCSD solar energy testbed

The forecast domain for UCSD was defined as a  $2.6 \times 3.6$  km grid at a resolution of 2.5 m per grid point. Each grid point was mapped to a latitude and longitude, as well as an altitude obtained from the SRTM1 dataset (Farr et al. 2007). Six Li-COR 200SZ pyranometers, sampling at 1 sec, are located within this forecast domain (Figure 7.1b), with positions given in Table 7.1. The USI was located on a rooftop at  $32.8722^\circ\text{N}$ ,  $117.2410^\circ\text{W}$ , 145 m MSL. The USI captures high dynamic range images (Urquhart et al. 2014) every 30 seconds during times when the sun is above an elevation of  $-3^\circ$ . Forecasts were generated from November 1<sup>st</sup> to the 29<sup>th</sup>, 2012, excluding November 12<sup>th</sup> (28 days total). A single domain average irradiance measurement is obtained by averaging irradiance measurements from the six stations. Measured irradiance data is compared to the forecast by applying a 30 sec moving averaging to the domain-averaged time series.

**Table 7.1.** Locations of the pyranometers in the UCSD network and their distance to the USI.

Station Identifier	Longitude [deg]	Latitude [deg]	Altitude [m]	Distance to USI [m]
BMSB	-117.2362	32.8758	111	603
CMRR	-117.2353	32.8806	111	1074
EBU2	-117.2330	32.8813	101	1257
HUBB	-117.2534	32.8672	24	1288
MOCC	-117.2225	32.8784	103	1857
POSL	-117.2350	32.8807	110	1103

### 7.3.3 Forecast methodology

For brevity, this work focuses on the comparison of TSI and USI forecast performance. Only an overview of the forecast procedure is provided here. Chapters 4 and 6 contain detailed information on the forecast procedure. The basic principle of the forecast is to locate clouds in the sky using a geometrically calibrated camera system, to track the motions over a few images, and then to advect the clouds forward in time. Ray tracing is performed over the forecast domain to estimate the positions of the shadows for every advection. A shadow fraction for relevant forecast domain gridpoints (i.e. those with solar collectors) is then computed. Power and/or irradiance is then parameterized using the shadow fraction. The compliment to the shadow fraction is the clear fraction, which is simply one minus shadow fraction.

The parameterization of power or irradiance requires simultaneous measurements from the PV array or an irradiance sensor, where the former will provide a power output forecast and the latter will provide an irradiance forecast. The power  $P$  or irradiance  $I$  parameterization is based on the construction of a histogram of clear sky normalized measurements (section 6.4). The clear sky normalized power  $k_{t,P}$  or irradiance  $k_{t,I}$  can be computed as

$$k_{t,P} = P/P_{clr} , \quad \text{or} \quad k_{t,I} = I/I_{clr} , \quad 7.1a, b$$

where the subscript  $clr$  indicates a clear sky modeled value appropriate for the time, location, and orientation of the solar collectors or irradiance sensor. The following text will omit the  $I$  or  $P$  from  $k_t$ , and refer to it as the clear sky index. In this work, the clear sky index for the TSI is computed using Eq. 7.1a and computed for the USI using Eq. 7.1b. It should be noted that in this work, only global irradiance is considered. The method of using the clear sky index histogram to physically constrain power output between lower and upper bounds can instead use direct normal irradiance (DNI) as input. The fractional cloud cover method developed in Urquhart et al. 2012, 2013 and Yang et al. 2014, which is similar to that applied in Marquez and Coimbra (2013) for DNI forecasting, can then be used to determine the forecast DNI between these two bounds.

There is an important difference between the deployment configurations of the TSI (section 7.3.1) and the USI (section 7.3.2): to generate forecast performance results, the USI uses only six 6.25 m<sup>2</sup> forecast gridpoints, whereas the TSI uses over 190,000 forecast grid points (also 6.25 m<sup>2</sup>). Spatial errors in cloud and shadow positions are averaged over the entire plant for the TSI whereas for the USI, it is averaged over the six point measurements. To provide a reference for comparing instrument performance, a persistence forecast is used.

#### **7.3.4 Persistence Forecast**

Any forecast technique must outperform persistence (i.e. using the current value as the future estimate) to be considered to have forecast skill. For comparison, a persistence forecast is generated for each forecast issue by averaging the previous 60 seconds of ground-measured clear sky index.



### 7.3.5 Comparison Metrics

This section provides definitions of the statistical metrics used to compare the forecast performance of the TSI and USI. The utility of each metric for assessing forecast performance will be discussed. The three different metrics used for comparison are the Pearson correlation coefficient, the Spearman correlation coefficient and Euclidean dispersion.

### 7.3.6 Pearson correlation coefficient

Pearson Correlation Coefficient (PCC;  $\rho_p$ ) is a measure of the linear dependence between a set of two variables  $X$  and  $Y$ , giving an value between  $-1$  and  $+1$ . PCC is formulated for  $n$  samples by using the variance and covariance of two vectors  $X$  and  $Y$  as

$$\rho_p = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2} \sqrt{\sum_{i=1}^n (Y_i - \bar{Y})^2}}, \quad 7.2$$

where the overbar indicates the mean of the set. PCC is a normalized assessment of the covariance of the forecast and measured  $k_t$ , and indicates how well the two signals match, irrespective of magnitude.

### 7.3.7 Spearman correlation coefficient

Spearman correlation coefficient (SCC;  $\rho_s$ ) assesses the monotonicity of the relationship between two variables. When each variable is a perfect monotone function of the other, SCC has a value of  $+1$  or  $-1$ . SCC, when applied to a set of two vectors is defined as the PCC between the ranked vectors. For  $n$  samples, the vectors  $X$  and  $Y$  are converted to ranked vectors  $x$  and  $y$ , respectively, and  $\rho_s$  is computed

$$\rho_s = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^n (y_i - \bar{y})^2}}, \quad 7.3$$

where the overbar indicates the mean of the set, which here is a set of ranks. In this work, fractional ranking is used because repeated values for  $k_t$  or shadow fraction are common. SCC evaluates the overall trend of the forecast  $k_t$  signal compared to the measured  $k_t$ . The SCC value is not affected by the linearity of the relationship between  $X$  and  $Y$  like the PCC value.

The equation of a line in two dimensions is  $aX + bY + c = 0$ , where the line is defined by the vector  $\ell = (a, b, c)$ . The Euclidean distance vector  $D_i$  assesses how far each ordered pair is located from the line defined by  $\ell$ :

$$D_i = \frac{|aX_i + bY_i + c|}{\sqrt{a^2 + b^2}}. \quad 7.4$$

The Euclidean distance vector can be condensed to a scalar metric, which we call Euclidean dispersion, by looking at the mean absolute distance (MAD) or root mean square distance (RMSD). For evaluating the measured versus forecast clear sky index, a one-to-one line is used for  $\ell$  because the relation between both clear sky index sets should have unity scaling and zero bias. For comparing forecast clear fraction (section 7.3.3) and measured clear sky index, a reference line defined by the points  $(0, P_5(k_t))$  and  $(1, P_{95}(k_t))$  is used, where  $P_5(k_t)$  and  $P_{95}(k_t)$  represent the 5<sup>th</sup> and 95<sup>th</sup> percentile of the measured clear sky index set. These points were selected because the working hypothesis of the sky imager forecast is that clear sky index is linearly proportional to clear fraction through the range measured clear sky values for a given day. Euclidean dispersion assesses the validity of this hypothesis.

## 7.4 Results

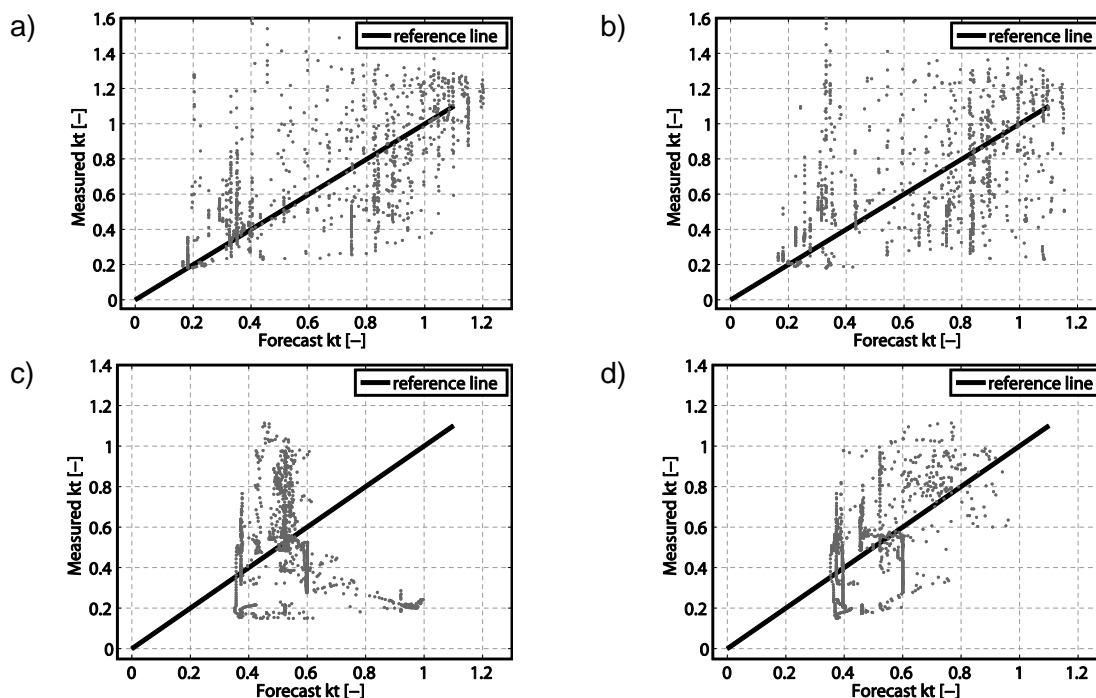
Forecast performance for each instrument was characterized for one month of data. The aggregate dataset is useful in assessing the overall forecast skill of the forecast method applied to both instruments, but a deeper understanding of performance is obtained only when analyzing individual days. This section is divided into two sections: one which presents results for a single day, and one which presents results for the aggregate dataset.

### 7.4.1 Forecast performance on a single day

To evaluate single day forecast performance, July 12, 2012 was selected for the TSI, and November 8, 2012 was selected for the USI. These days were chosen because they had high variability, which is the most interesting condition in which to evaluate forecast skill.

Scatter plots of measured versus forecast clear sky index for the 0-min (nowcast) and 10-min horizon for both instruments are shown in 7.2. The scatter for both instruments at each time scale which indicates that there is significant error in the forecast. In Figure 7.2c, the TSI nowcast is particularly poor. This is due to the shadowband on the TSI mirror that blocks the view of the sky dome near the sun. Because the TSI is collocated with the solar collectors, this region of sky is where clouds causing current shading exist (e.g. a cloud between an observer and the sun casts a shadow on the observer). Even for a plant of  $1.3 \text{ km}^2$  in area, the shadowband can eliminate all (or nearly all) sky image data needed for a 0 to 5 minute forecast. The size of the plant area impacted by the shadowband scales linearly with cloud height (in the positive sense). The average cloud height on July 12, 2012 was 4.6 km. When the ground projection of the shadowband advects away from the plant, valid image data can

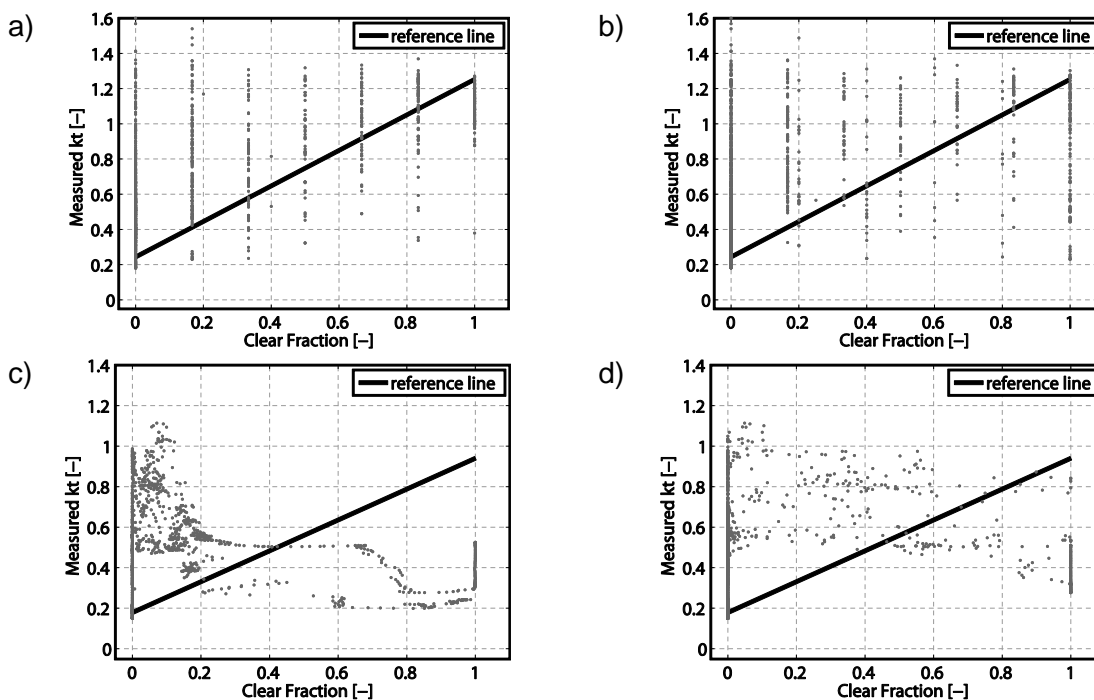
be used to generate a forecast (Chow et al. 2011, Urquhart et al. 2012), such as that in Figure 7.2d which shows a much better result than Figure 7.2c. The USI, with no occulter blocking the sun, does not have this problem as is evidenced in Figure 7.2a and 7.2b.



**Figure 7.2.** Scatter plots of measured vs. forecast clear sky index for (a, c) nowcast and (b, d) 10-min horizon for (a, b) the USI on November 8, 2012 and (c, d) the TSI on July 12, 2012.

Scatter plots of measured clear sky index versus clear fraction are shown in Figure 7.3. From a single day of measurements, it is not apparent that this relation is exhibited by either instrument. The results for the TSI (7.3c) are counter intuitive in that it appears that the more shading the plant receives, the higher the irradiance. This is not a real result, but rather an artifact of the image quality, the shadowband, and the processing algorithms applied. It should be noted that the vertical lines in Figure 7.3a and 7.3b (USI) are due to the six stations used for generating a clear fraction which nominally has seven possible clear fraction values unless data is missing or the clouds advect too far from a measurement station to determine if it is shaded. For the TSI

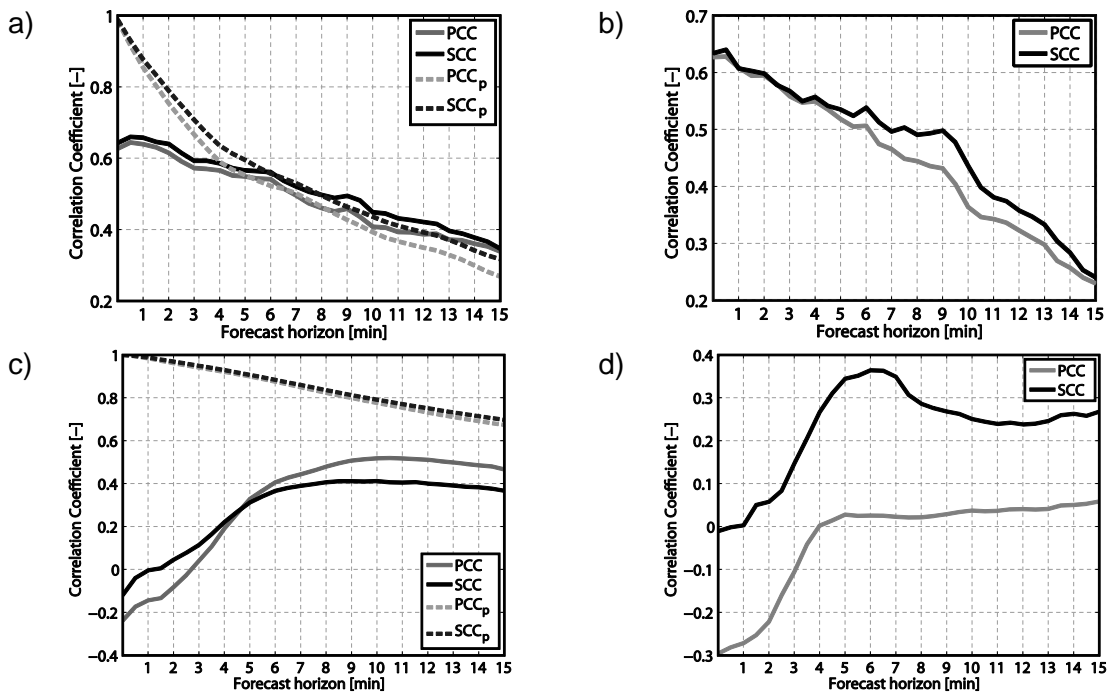
(Figure 7.3c and 7.3d), the grid points are  $6.25 \text{ m}^2$  squares covering a  $1.3 \text{ km}^2$  plant, giving many possible values for clear fraction.



**Figure 7.3.** Scatter plots of measured clear sky index vs. clear fraction for (a, c) nowcast and (b, d) 10-min forecast for (a, b) the USI on November 8, 2012 and (c, d) the TSI on July 12, 2012.

The correlation (PCC and SCC) between the forecast and measured clear sky index is shown in Figure 7.4a for the USI and Figure 7.4c for the TSI. The correlation of the persistence forecast with the measurements is also given for comparison. Persistence is better correlated with the measurements than the USI forecast until about 5 minutes. After this time, the USI and persistence are comparable until about 10 minutes where the USI begins outperforming persistence. The TSI forecast performance is poor (even anti-correlated) until the projection of the shadowband advects away from the plant, after which the forecast performance improves, but is still not as high as persistence. Both sky imager forecasts and persistence show the expected decrease in accuracy with forecast horizon. The correlation (PCC and SCC) of clear fraction and measured clear sky index is shown in Figure 7.4b for the USI and Figure 7.4d for the

TSI. For the USI, clear fraction is positively correlated with clear sky index, but it shows a marked reduction in correlation with forecast horizon. In Figure 7.4d the effect of the shadowband can again be seen until approximately the 6 minute forecast horizon.

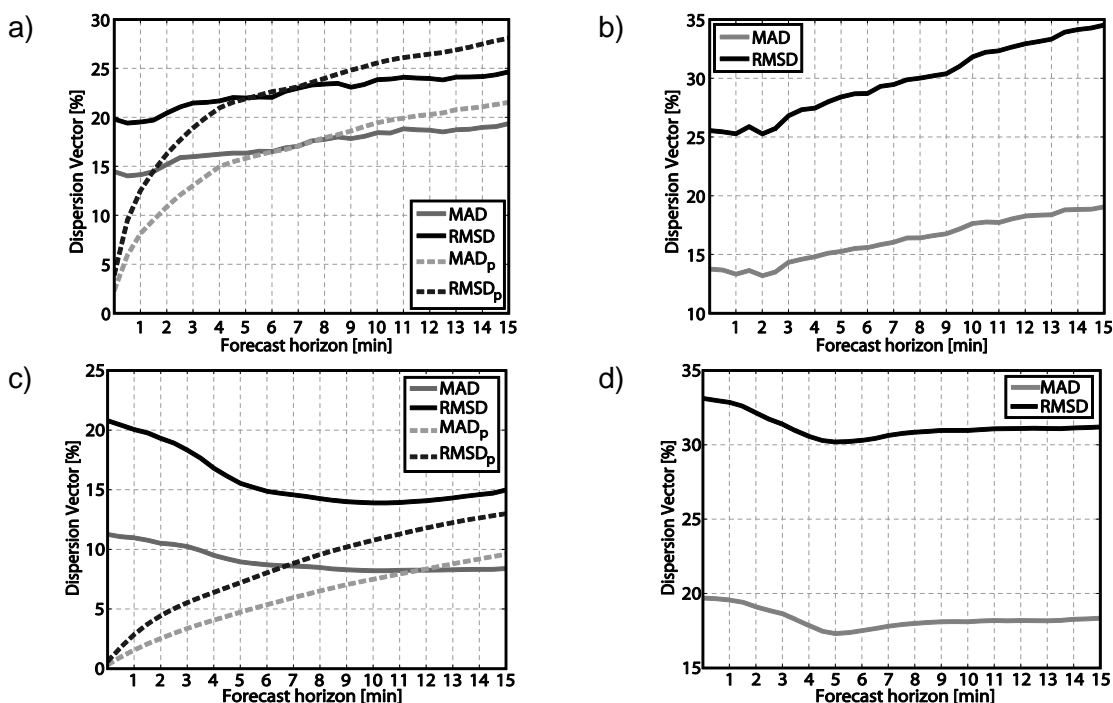


**Figure 7.4.** Pearson (PCC) and Spearman (SCC) correlation coefficients for forecasts with (a, b) the USI on November 8, 2012 and (c, d) TSI on July 12, 2012. Forecast vs. measured clear sky index is shown in (a, c), and clear fraction vs. measured clear sky index in (b, d). Performance of the persistence forecast is shown in (a, c) as a dashed line (subscript  $p$  in legend).

The persistence forecast for UCSD and for Copper Mountain (the USI and TSI, respectively) decay at much different rates, with UCSD experiencing much poorer performance when using persistence than for Copper Mountain. The conditions at Copper Mountain are no less variable than for UCSD on the day selected, but quite apparently, the measurements made within the last 15 minutes have a higher correlation with current measurements at Copper Mountain. There is a spatial smoothing factor at Copper Mountain that does not exist at UCSD; as the area increases, the timescale of correlation also increases (i.e. the time it takes for correlation to decay to a given value increases; Lave et al. 2012, 2013). This results in longer time scales having higher

correlations, the converse of which is that point sensors such as those at UCSD have the same level of correlation on shorter timescales. This is clearly exhibited by the persistence forecasts in Figure 7.4 and 7.8.

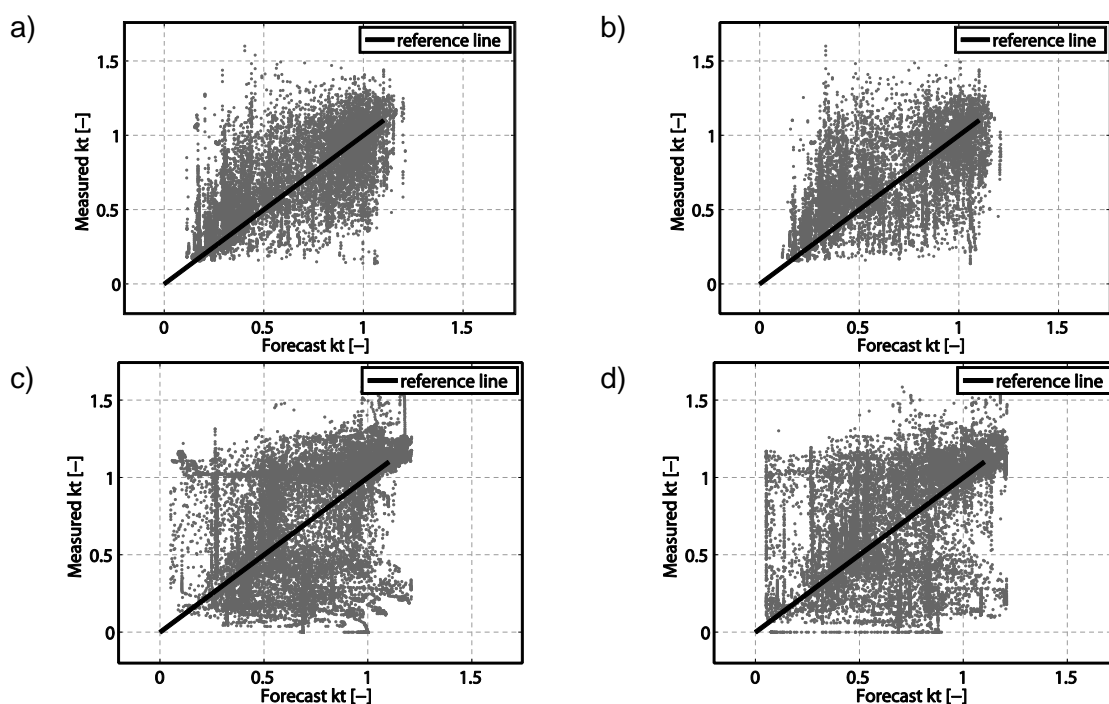
The Euclidean dispersion for the USI (Figure 7.5a) is consistent throughout the 15 minute forecast horizon, whereas for persistence it increases sharply, and after 7-min is higher than for the USI. The dispersion for the TSI (Figure 7.5c) is actually lower than for the USI, which occurs because of the smaller range of clear sky index for the TSI on the day selected (0.2 to 1.6 for the USI vs. 0.1 to 1.1 for the TSI). The Euclidean dispersion between clear fraction and measured clear sky index for the USI (Figure 7.5b) is lower than that of the TSI (Figure 7.5d), even with the range disparity of measured clear sky index, indicating it performs better when viewing this metric.



**Figure 7.5.** Euclidean dispersion of (a, c) measured vs. forecast clear sky index and (b, d) clear fraction vs. measured clear sky index for (a, b) the USI on November 8, 2012 and (c, d) the TSI on July 12, 2012. Performance of the persistence forecast is shown in (a, c) as a dashed line (subscript  $p$  in legend).

### 7.4.2 Forecast performance for one month

Reviewing forecast results for a single day highlights performance in particular conditions, but to view the overall performance assessment, a longer dataset is required. Scatter plots of measured versus forecast clear sky index during the one month study period is shown in Figure 7.6. For both the nowcast (left column) and the 10-min forecast horizon (right column), the USI (top row) shows less (but still considerable) scatter than the TSI (bottom row). For the USI, there appear to be less frequent occurrences of inverted forecasts (points in the top left or bottom right).

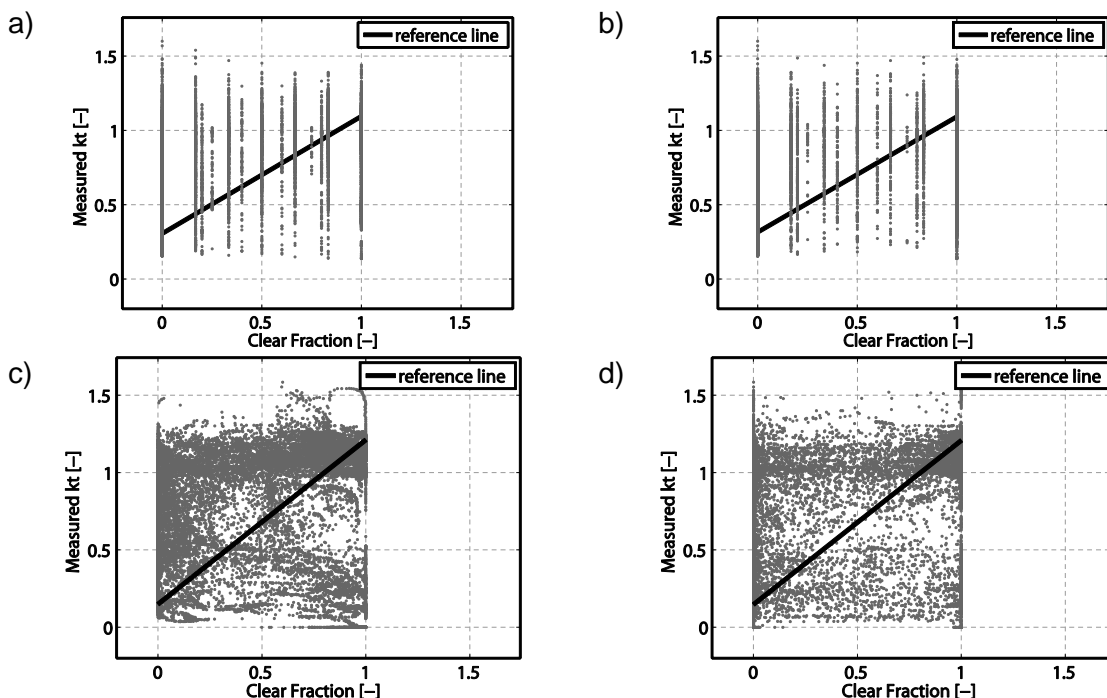


**Figure 7.6.** Scatter plots of measured vs. forecast clear sky index during the one month study period for (a, c) nowcast and (b, d) 10-min horizon for (a, b) the USI and (c, d) the TSI.

Scatter plots of measured clear sky index versus clear fraction for the one month study period are shown in Figure 7.7. There is again no clear relation between the clear fraction and the clear sky index for either the USI (top row) or the TSI (bottom row). It is also not apparent from the scatter plots that the nowcast clear fraction (left column)



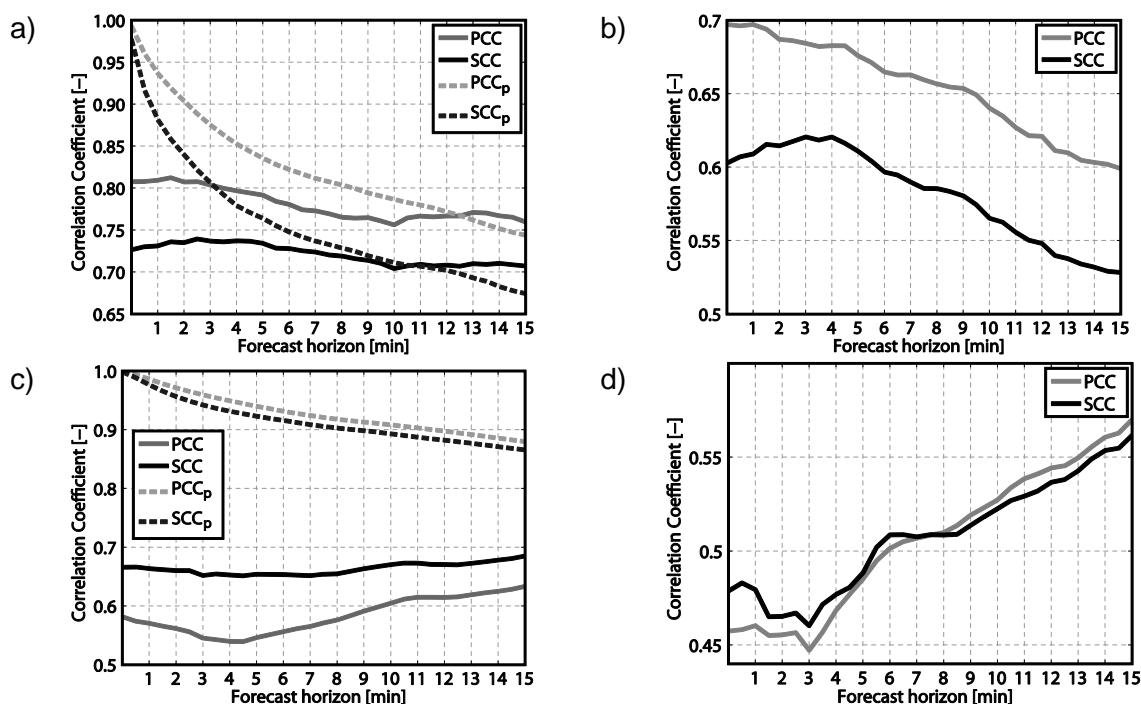
produces better results than the 10-min forecast clear fraction (right column). The TSI has the inherent advantage of the large spatial average over the power plant, but still does not show significant improvement. The higher density of points above the reference line in the bottom row of Figure 7.7 indicates that the TSI biases forecasts to contain cloud when the clear sky index indicates power output is near 100% of clear sky levels.



**Figure 7.7.** Scatter plots of measured clear sky index vs. clear fraction during the one month study period for (a, c) nowcast and (b, d) 10-min forecast for (a, b) the USI and (c, d) the TSI.

The correlation results for the one month data set are shown in Figure 7.8. Reviewing the measured versus forecast clear sky index, the average PCC was 0.78 and the SCC was 0.72 for the USI, whereas for the TSI, the respective values were 0.58 and 0.67. The SCC for the USI and TSI are similar, which indicates both systems are capturing the trends of the irradiance or power fluctuations. The significantly higher PCC value for the USI indicates that the relation between forecast and measured clear sky index is much more linear than for the TSI, which is positive result for the USI because

the relationship should be one-to-one. The effects of the shadowband on the TSI can clearly be seen in Figure 7.8d, where correlation increases nearly monotonically with time. The TSI actually exceeds performance of the USI with the SCC value for clear fraction versus measured clear sky index at the longest forecast horizons.



**Figure 7.8.** Pearson (PCC) and Spearman (SCC) correlation coefficients during the one month study period for (a, b) the USI and (c, d) the TSI forecasts. Forecast vs. measured clear sky index is shown in (a, c), and clear fraction vs. measured clear sky index in (b, d). Performance of the persistence forecast is shown in (a, c) as a dashed line (subscript  $p$  in legend).

Comparing the performance of both the USI and TSI to persistence at their respective locations shows that the USI performed significantly better for the period of study. As noted in section 7.4.1, forecasting at a large plant like Copper Mountain should produce higher correlations for the 15 minute forecast window because of the increased correlation time scale, yet the TSI shows correlation results that are consistently below persistence by 0.15 to 0.30, even for longer forecast horizons. The USI, however, outperforms persistence at longer time scales. The number of clear days at Copper

Mountain may be a suspect cause for higher correlations of persistence, but both the UCSD and Copper Mountain data sets contained 8 clear days.

## 7.5 Conclusions

The overall forecast results for both instruments indicates that there is much error in the deterministic forecasting process. The working hypothesis that clear fraction is proportional to clear sky index was not conclusively validated in this work. The correlation was positive between these two parameters, but the low correlation values and scatter in the data suggests this hypothesis is only partially valid. The clear fraction estimated for the solar collectors or irradiance sensors is dependent primarily upon cloud detection, cloud speed and advection assumptions, and the geometric calibration of the instrument. If cloud is not correctly detected, and if the positional information of each pixel is not accurate, then the ray traced clear fraction will be incorrect. Similarly, if the cloud advection is erroneous, due to either an error in cloud speed or in the assumption that clouds can be advected forward in time without accounting for condensation/evaporation, errors in clear fraction will result. These results also indicate that more investigation into the cloud coverage (i.e. clear fraction) to power output parameterization is required.

Comparing the USI performance to that of the TSI indicates that the USI outperforms the TSI overall, particularly for the first 10 minutes of the forecast window. This is an expected result because the USI does not have missing data due to the shadowband, it does not lose information to image compression, and has both a higher spatial and intensity resolution (i.e. it has more horizontal and vertical pixels, and provides 48-bit images as opposed to 24-bit images). In the experimental configuration used here, the TSI had the advantage of a significantly larger spatial average (35,500:1;

TSI:USI) because it was placed at a power plant. It is reasonable to expect some of the spatial errors in clear fraction to average out in this scenario, however this result was not realized and the USI still showed better performance.

Moving forward with this forecast work, emphasis will be placed on improving results with the USI. It has been shown here to be a superior instrument for short term forecasting, and will provide higher quality image data for this computer vision-based forecasting work.

### **Acknowledgements**

Hardware development for the USI was funded by Panasonic Corporation. Advanced forecast algorithms and demonstrations were funded by Panasonic Corporation, the Department of Energy High Solar PV Penetration Grant 10DE-EE002055, the California Energy Commission contracts 500-10-057 and 500-10-043, and the California Public Utilities Commission California Solar Initiative RD&D Program. The help of David Jeon, Leslie Padilla, Ron Clabaugh, Darrel Lopez and Shiva Bahuman from Sempra US Gas and Power, Jim Blatchford from the California Independent System Operator, David Zeglinski and Megan Fox from OSISoft, and Jo Frabetti from the San Diego Supercomputer Center are appreciated. This work would have not been possible without the assistance of Elliot Dahlin, Caspar Hanselaar, Edmundo Godinez, William Gui, Prithvi Sundar, Dan Erez, Scott Kato, Amy Chiang, Jessica Traynor, Kristen Otosh, Tyler Capps, Sebastian Schwarzfischer, Nick Truong, Jeff Yeh, Max Twogood, Alex Turchik, Danielle Donnelly, and Emily Davis, whom helped with hardware development, fabrication, assembly, and deployment of the USI systems. Anthony Dominguez, Patrick Mathiesen, Matt Lave and Felipe Mejia set up and maintained the UCSD pyranometer network.

In the course of preparing this chapter, the following publications were used:

- Publication:** Comparison of Solar Power Output Forecasting Performance of the Total Sky Imager and the University of California, San Diego Sky Imager. *Energy Procedia*, volume 49, pp. 2340–2350, 2014. doi: 10.1016/j.egypro.2014.03.248
- Co-Authors:** Seyyed Mohammad Iman Gohari, Handa Yang, Ben Kurtz, Dung Nguyen, Chi Wai Chow, Mohamed Ghonima and Jan Kleissl

## 8 Conclusions and Future Work

A novel solar power forecasting procedure using ground-based sky imagery to geometrically locate clouds and estimate the future positions of their shadows was presented. This use of sky imagery for solar power forecast was one of the first of its kind, with only one prior study reported (which did not see future development). The initial forecasts developed in this work using a TSI instrument were shown to have limited skill over persistence at a solar power plant, however using the new USI sky camera system described here, the forecast approach has consistently outperformed persistence after the first few minutes. This computer vision solar power forecast system is still being actively developed, and as the performance improves in accuracy and reliability, it will serve to facilitate the control and automation of power systems which rely on solar energy as a power source. Depending on the application, it can serve purely to save money, or it can provide critical functionality as a controls input to fast acting devices (e.g. battery charge/discharge optimization).

The USI, along with associated remote sensing applications provides a lower cost solution than the high quality WSI, while providing intermediate quality imaging performance with more features than are currently available from lower cost solutions. The remote sensing algorithms presented here are fast, automated and reasonably accurate. Taken as a package, the system can currently detect cloud, determine cloud velocity, and determine cloud position. The determination of cloud optical depth is a topic of current work within our research group. A 3D radiative transfer model is being used to construct synthetic USI images which can be used to create a lookup table or functional relation between the cloud optical depth and the observed brightness on the red, green and blue channels. The measurement of GHI and DNI directly from images is also being investigated. To accomplish this, the brightness of the charge overflow from smear is

being correlated with DNI measurements, and the DNI estimates along with the radiance distribution over the sky dome will be used to estimate GHI. Taken together, these algorithms will yield a potent cloud remote sensing device. Other potential future work is radiometrically calibrating the system so that it can measure sky radiance distribution. For clear sky analysis, the determination of aerosol optical depth has been shown possible, and could be added to the remote sensing suite of capabilities of the USI.

For the cloud position and velocity estimates, an accurate photogrammetric calibration is required. In this work a generic camera model and a fully automated solar-position-based calibration algorithm was developed for the USI. This method provides an accurate photogrammetric calibration with no user input (only "factory settings" must be input to indicate the type of lens). The calibration algorithm was shown to have sub pixel uncertainty, even when the calibration input has an RMS uncertainty of 10 pixels. Given the specific format of the calibration input, which is a single arc of solar position throughout a day that covers a very limited portion of the field, this low level of uncertainty is quite good. Future work on this topic will involve the use of a checkerboard or other patterns to calibrate the intrinsic and distortion parameters of the general model. Similar calibration work using a checkerboard has been reported in the literature, and the implementation for the USI is expected to be an extension of this work where possibly additional developments must be made to account for the specifics of the camera model.

The cloud stereo photogrammetry described here is the first dense cloud position estimation with automatic feature registration and correspondence. The work was also the first reported use of dynamic programming for determining correspondence in cloud images. Results for several cloud types, heights, and relative projective distortions was presented. The method does require manual input of zenith view direction and field of

view (azimuth is set automatically), and thus future work on this topic will pursue the full automation of this parameter selection. Preprocessing the image for select, high quality features is expected to yield a path forward for determining the appropriate zenith angle and field of view for optimal matching performance.



## **Appendix A. Optimal Geographic Allotment of Distributed PV**

Contributing Authors:

B.Urquhart, M. Sengupta, and J. Keller

### **A.1 Summary**

A multi-objective optimization was performed to allocate 2MW of photovoltaic (PV) among four candidate sites on the island of Lanai, Hawaii, such that energy was maximized and variability in the form of ramp rates was minimized. This resulted in the Pareto-optimal set, an optimal solution set that provides a range of geographic allotment alternatives for fixed PV capacity. Within the Pareto-optimal set, a trade-off was found between energy produced and variability experienced, whereby a decrease in variability always necessitates a simultaneous decrease in energy. With this development, system designers have a method to select the preferred combination of energy generation and variability within the set of optimal alternatives to meet their needs. A design point within the optimal set was selected for study that decreased extreme ramp rates by more than 50% while decreasing annual energy generation by only 3% above the maximum generation allocation. To quantify the allotment mix selected, a new metric called the “ramp ratio” was developed. It compares ramping magnitude when all capacity is allotted to a single location to the aggregate ramping magnitude in a distributed scenario. The ramp ratio quantifies simultaneously how much more smoothing a distributed scenario would experience than single-site allotment and how much a single site is being underutilized for its ability to reduce aggregate variability. This paper creates a framework for use by cities and municipal utilities to reduce variability impacts while planning for high penetration of PV on the distribution grid, thereby maximizing the value of investments.

## A.2 Introduction

Variability of solar power provides integration challenges as a primary power source on the transmission system and poses potential power quality issues for distribution networks. A known level of fluctuation in solar power output will always exist because of natural variations in the solar position. Superimposed on the known fluctuations will exist additional variability in production from installed photovoltaic (PV) that depends primarily on weather conditions influencing clouds and solar radiation at the location. It is well studied that aggregation of sites produces a smoother output of power on a per-capacity basis (Murata and Otani 1997, Otani et al. 1997, Wiemken et al. 2001, Roy 2006, Curtright and Apt 2008, Collins and Crowther 2011, Lave et al. 2011). These studies primarily address smoothing through geographic dispersion, and attempts have been made to mathematically model this phenomenon. On the other hand, there has been an apparent dearth of work on how to select the best geographic allocation of PV among candidate sites. This study looks at the trade-off between energy maximization and variability minimization while selecting and allocating PV generation among multiple candidate sites. In this paper, the geographic separation of the sites is small enough that we can assume that the problem we address is limited to the distribution grid. A related study looking at minimizing the levelized cost of energy of a set of 12 PV generators using daily energy as input was performed by Roy (2006). Collins and Crowther (2011) performed a multiobjective optimization that maximizes daily energy while minimizing generation shortfall using hourly input data and sites distributed across the state of Virginia. The latter is similar to this study in that it evaluated the trade-off between energy maximization and variability minimization; however, in the latter, the sites were separated far enough apart to be considered a transmission level problem. As previously mentioned, the scope of this study is restricted

to geographic dispersion that can be assumed to not have the constraints imposed by balancing area coordination, line congestion, and other transmission-related issues.

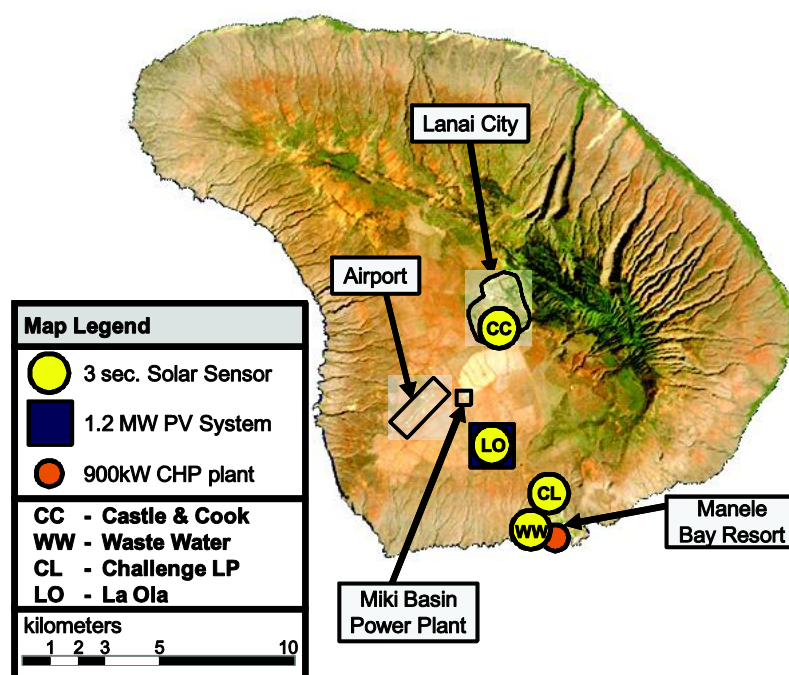
The Maui Electric Company (MECO) owns and operates the island of Lanai's electric power system. Compared with mainland grids, the electric power system is small. System energy is produced by a set of diesel generators located at the main power plant, one or two of which (depending on system load requirements) provide system frequency regulation using isochronous frequency control. Currently there are two large distributed generation systems in addition to the main power plant. La Ola (LO), a large central station (1.2MW) PV installation owned by Castle & Cooke, Inc., installed near the main diesel power plant, provides power to the grid through a power purchase agreement. Recently, the LO PV system began to generate power at full capacity because the power purchase agreement required that the system incorporate a battery energy storage system to help mitigate ramping issues. A combined heat and power generator is located near the end of the distribution circuit as well.

A team led by the US Department of Energy's National Renewable Energy Laboratory (NREL) and MECO are working with local developers on the island of Lanai to assess the economic and technical feasibility of increasing the contribution of renewable energy sources on Lanai with a stated goal of reaching 100% renewable energy as part of the Hawaii Clean Energy Initiative. For MECO, enabling reliable installation—and determining the associated technical requirements—of additional renewable resources such as PV systems onto the electric power system is one way to increase renewable energy penetration on Lanai.

The NREL team, with input from interested PV system installers, has been working to create and evaluate viable scenarios to assess potential PV systems that can be cost effectively installed on Lanai. In support of this, NREL installed equipment to

measure variability of the solar resource at several locations on the island. For more than 1 year, NREL has been capturing 3-s time-synchronized solar radiation data at four locations on Lanai that represent some of the possible locations for additional PV systems. Figure A.2 shows the locations of the 3-s solar irradiance data collection points, as well as some of the existing generation and load centers. These data help us understand site resources and variability among sites; however, appropriate system simulation models and assumptions are still needed.

This paper develops a design strategy for the placement of PV generation by maximizing energy produced and minimizing variability—quantified here as power ramp rates. To do this, the planned PV output was simulated from available solar measurements for one complete year, from June 2010 to May 2011. A predetermined amount of PV (e.g., 2MW) was geographically allocated among four locations, with



**Figure A.1.** Map of the island of Lanai, Hawaii, with four sensor locations indicated. The Lanai City and airport load centers are indicated, along with the La Ola PV system, the main island power plant, and a combined heat and power plant located near the coast. *Lanai image courtesy of Forest and Kim Starr.*

differing amounts at each site, and the impact on energy generation and variability was quantified. To formalize this process, a multi-objective optimization scheme was employed to generate Pareto<sup>1</sup>-optimality curves using several variability criteria. Pareto-optimality curves provide information on the trade-off between energy production and variability experienced. These curves are presented as a method for utilities to assess the trade-offs of maximizing energy production while minimizing system variability. The method presented provides a tool for selecting an acceptable level of ramping for an individual system.

### **A.3 Methodology**

#### **A.3.1 Experimental Setup**


The island of Lanai is located at latitude of 20.8°N, where the Walker and Hadley circulations generate northeasterly trade winds. Trade wind cumulus clouds are dominant and cause large and frequent ramps in global horizontal irradiance (GHI) because of the highly variable nature of that cloud type (Figure A.2). GHI was measured at four prospective PV deployment locations (Figure A.2, Table A.1) using Licor LI 200 silicon pyranometers (LI-COR Biosciences, Lincoln, NE, USA) capturing data at a rate of 3 s. The LO station also had regular meteorological measurements of temperature and wind data. Data were collected from April 2010 to September 2011; however, to ensure the summer half year was not weighted more heavily in this study, only data from June 2010 to May 2011 were used.

---

<sup>1</sup> Named for the Italian economist Vilfredo Pareto.

**Table A.1.** Measurement site geographic position and designation

Site	designation	lat. [°N]	lon. [°W]	alt. [m]
Castle & Cook	CC	20.81782	156.92107	464
Waste Water	WW	20.74844	156.89694	110
Challenge LP	CL	20.74102	156.90509	60
La Ola	LO	20.76685	156.92291	381

**Figure A.2.** Global horizontal irradiance on 20 June 2011, showing typical large, cumulus-induced fluctuations on the island of Lanai, Hawaii.

The Castle & Cook (CC) site is located in the center of the island, close to the mountains; the other sites are located in the southern portion of the island surrounded by much flatter terrain. Terrain-induced orographic lifting at the CC site resulted in higher levels of cloud formation and thus a higher frequency of large ramp events. The Waste Water (WW) and Challenge LP (CL) sites are close to each other and thus exhibited stronger correlation than the other two sites. Hereinafter, the four sites are referred to by the designations indicated in Figure A.2 and Table A.1.

### A.3.2 Power output simulation

Power output was simulated using irradiance-to-DC power and DC-to-AC power algorithms that are part of the PVForm performance model (version 3.3; Menicucci 1986,

Menicucci and Fernandez 1988). This model has been widely used since its development, including in the latest version of PVWatts (Marion et al. 2001), and has been shown to have a 5% positive bias and root mean square errors of less than 12% (Perez et al. 1994). The model incorporates a temperature correction for irradiance-to-DC conversion as well as an efficiency reduction for low irradiance levels.

The transposition model used in PVForm (based on that of Perez et al. 1987), to convert horizontal irradiance measurements to plane-of-array irradiance, was not used here. The goal of this study was to develop a methodology to optimize the benefits of geographic smoothing and did not aim to achieve high accuracy in the exact energy produced or ramp rates experienced. Both the energy output and ramp rates would increase if GHI was transposed to plane of array. Importantly, however, the trends would remain the same; that is the basis for the results presented here. An additional consideration is that the latitude is low; therefore, the tilt is close to horizontal, so there is less error in making the horizontal simplification than would occur at higher latitudes.

It should be mentioned that direct performance model conversion of irradiance-to-power output is somewhat misleading. Large PV plants exhibit intraplant smoothing similar to a low-pass filter (Marcos et al. 2011a,b). Both Marcos et al. (2011a) and Lave et al. (2011) developed Fourier-based and wavelet-based methods, respectively, to smooth modeled intraplant power output so that it better matches observed output. Intraplant smoothing was not considered here. The aggregate capacity of the four simulated sites was 2MW, which is relatively small in total area, and it was assumed that direct use of a PV performance model would not impact the utility of the results presented.


### **A.3.3 Multiobjective optimization: maximizing energy generation while minimizing ramp rate**

Using a predetermined installed capacity of PV to be allocated among the four candidate sites, an optimization was performed to determine which weighting (% of total) allocation at each site would provide the most energy with the least variability. To maximize only energy, one would simply place 100% of capacity in the sunniest location; however, this would not provide the smoothest output in general. If variability is a concern, utilities could take advantage of geographic dispersion and place some capacity in less sunny locations. This would result in a potential reduction in total energy produced, but would also lead to reduction in variability. To avoid defining a subjective cost penalty for each parameter being optimized, as is done in standard linear programming, the energy and variability criteria were treated as distinct objectives. This multiobjective optimization problem does not have a single unique solution; instead, it has a range of solutions—known as the Pareto-optimal set—that describe non-inferior alternatives for maximizing energy while minimizing variability. Arbitrarily selecting how to allocate capacity may reduce energy, but not significantly decrease variability; for the Pareto-optimal set, the maximum energy output is achieved for a given level of variability.

Figure A.3 shows an example of the Pareto-optimal set for minimizing the 99<sup>th</sup> percentile of ramp rates while maximizing energy. Moving along this curve—also called the Pareto-optimal frontier—one must trade energy generated for a reduction in variability, characterized here by ramp rates. At any point in the interior, one can increase energy or decrease ramp rate without sacrificing the other. Therefore, any interior point corresponds to an inefficient selection of geographic allocation—denoted as the Pareto-inferior region in the figure. Points along the curve are Pareto-efficient—to



change one quantity, one must sacrifice the other. A system designer can select the maximum level of variability a given electrical system can handle; the corresponding point on the Pareto-optimal frontier will then yield the geographic allocation for maximum energy production.



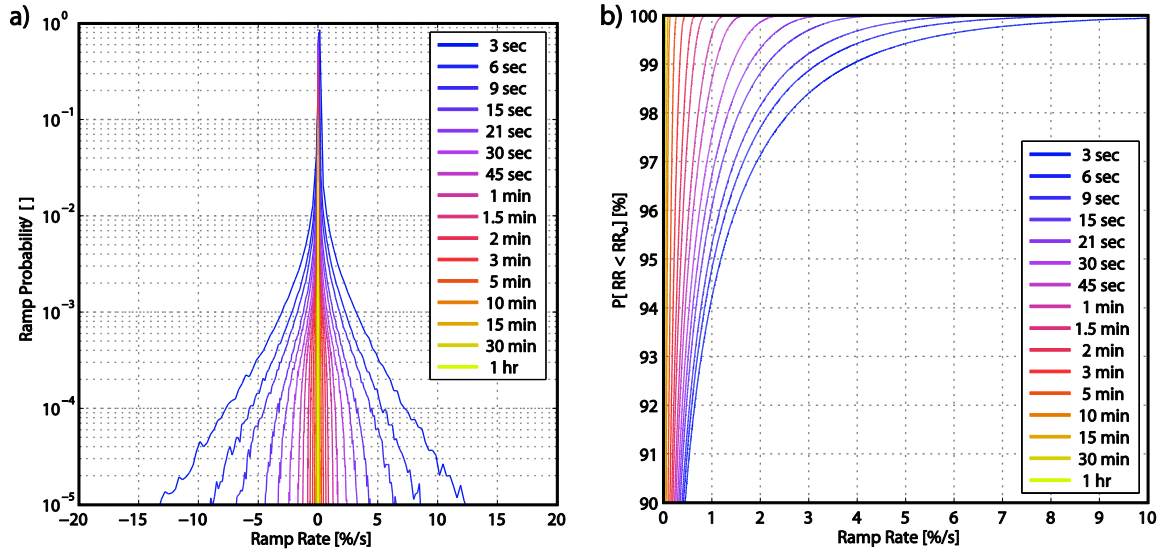
**Figure A.3.** The Pareto-optimal set for maximizing energy while minimizing the 99<sup>th</sup> percentile of ramp rates. The curve is shown to divide two regions: an unattainable set of energy, ramp rate combinations; and an inferior set where individually either energy can be increased or ramp rates decreased without compromising the other.

It should be noted that although in this work we use ramp rates exclusively to quantify variability, other metrics—such as ramp size, ramp frequency, or generation shortfall (if load information is available)—can be used and will be the subject of future studies. Ramp rate  $r_j(t, \Delta t)$  at a time  $t$  for an interval  $\Delta t$  is defined as

$$r_j(t, \Delta t) = \frac{p_j(t + \Delta t) - p_j(t)}{\Delta t}, \quad \text{A.1}$$

where  $p_j(t)$  is the power output at site  $j$ . The probability density of ramp rates during the 1 year studied (June 2010–May 2011) for LO is given in Figure A.4 for different  $\Delta t$ . This shows the usual result that ramp rates are largest for the smallest  $\Delta t$ ; thus, for optimization only, ramp rates at the 3-s data capture rate were used. Minimization of 3-s

ramp rates ensures that ramp rates for all other  $\Delta t$  will always be smaller. No ramp rates were excluded in any computation step, even if very small, to ensure illumination of the true probability of occurrence of large ramp events in all statistical results. The cumulative distribution function indicates that 99% of the 3-s ramp rates were below 4% of capacity, and 99.9% of ramp rates were below 10% of capacity for the LO site.



**Figure A.4.** (a) Probability distribution (PDF) and (b) cumulative probability distribution (CDF) of ramp rates for several time steps at the La Ola site. Ramp rates are shown as a percentage of total power output capacity. Small or negligible ramp rates were included; thus, the curves indicate true probability of events. The CDF considers ramp rate magnitudes; the PDF considers the sign of the ramp rates.

The multiobjective optimization employs a genetic algorithm using a tournament selection scheme described in (Deb 2011). The genetic algorithm tries many successive combinations of weights, pushing the Pareto frontier outward until no other variation of weight allocation can extend the curve farther into the optimal territory. During this process, the solution sets in successive iterations are combined (crossover) and “mutated,” and less fit alternatives are removed in an emulation of genetic evolution. The objective functions to be minimized are

$$f(\mathbf{w}) = - \sum_i w_j p_j(t_i) \Delta t, \quad \text{A.2}$$

$$g_n(\mathbf{w}) = P_n(|w_j r_j(t_i, \Delta t)|), \quad \text{A.3}$$

where  $w_j$  is the component of weight vector  $w$  corresponding to the geographic allocation for site  $j$ , and  $P_n(x)$  is the  $n$ th percentile of the argument. Summation on  $j$  is implied for both Eqs. A.2 and A.3. Equation A.2 is the negative of total energy output by the aggregate system for configuration  $w$ , and Eq. A.3 is  $n^{\text{th}}$  percentile of system aggregate ramp rate computed for the entire set of  $i$  ramp rates (i.e., ramp rates from the entire year data set). In Eq. A.3, only the magnitude of the elements in the set is considered, denoted by the vertical bars. The constraints on the optimization are

$$\sum_j w_j = 1, \quad w_j \in [0,1],$$

which requires that all weights sum to unity, and that any given site can be weighted from only 0% to 100%. The notation used earlier contains a subtlety for computational purposes: both power and ramp rates used in Eqs. A.2 and A.3 for individual sites are scaled so that each site is initially at 2MW of capacity (8MW total); when multiplied by the weights  $w$ , the aggregate capacity is the desired 2MW. This implementation and notation must be changed if intraplant smoothing is considered. Power output and thus ramp rates are no longer independent of assigned capacity, and the power time series and thus ramp rates must be adjusted by means of a smoothing filter (Marcos et al. 2011, Lave et al. 2011).

#### A.4 Results

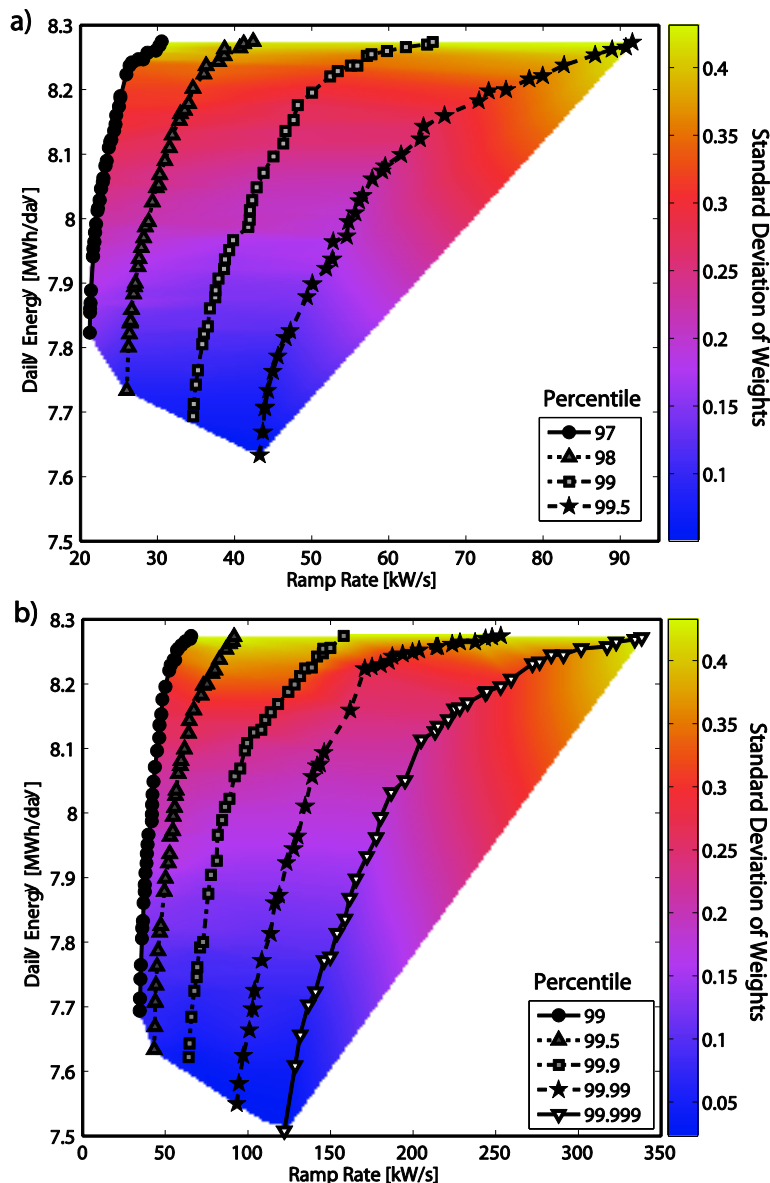
To assess the trade-off between ramp rates and energy production, multiple ramp rate percentiles were used as minimization criteria. Figure A.5a shows the 97<sup>th</sup>, 98<sup>th</sup>, 99<sup>th</sup>, and 99.5<sup>th</sup> percentile Pareto-optimal curves, along with an interpolated color swath to indicate the spread in site weighting. The darker color indicates that the

standard deviation of the four weight values was low. The upper bound of 8.26MWh was the maximum production possible for a 2-MW plant during the 1 year studied; thus, the allocation was 100% at the CL site with corresponding high standard deviation of the four weights. The near verticality of the 97<sup>th</sup> percentile shows that by changing the weights, not much improvement was made in reducing the ramp rates. At higher percentiles, such as 99.5%, significant gains were seen by adjusting the weights, and ramp rates dropped by nearly 40% for a 3% reduction in energy generated.

**Table A.2.** Site allocation weight values on the Pareto-optimality curves for minimization of the 99<sup>th</sup>, 99.9<sup>th</sup> and 99.99<sup>th</sup> percentile (pct) of ramp rates, shown as a function of aggregate daily energy produced. Aggregate system ramp rates are also shown.

energy [MWh/day]	99 <sup>th</sup> pct					99.9 <sup>th</sup> pct					99.99 <sup>th</sup> pct				
	weights [%]				ramp rate [kW/s]	weights [%]				ramp rate [kW/s]	weights [%]				ramp rate [kW/s]
	CC	WW	CL	LO		CC	WW	CL	LO		CC	WW	CL	LO	
7.55	–	–	–	–	–	–	–	–	–	–	22	23	29	26	93
7.60	–	–	–	–	–	21	27	30	22	64	20	24	32	24	96
7.70	17	26	36	21	35	20	28	34	18	67	18	25	37	20	103
7.80	13	27	43	17	36	16	27	42	15	73	14	29	41	16	112
7.90	11	27	49	13	38	21	34	41	4	77	12	32	44	12	121
8.00	8	23	60	9	42	16	33	50	1	86	8	34	50	8	135
8.10	5	23	67	5	46	8	29	60	3	98	5	35	55	5	146
8.20	2	19	78	1	52	3	15	81	1	128	2	32	65	1	168
8.26	0	0	100	0	66	0	0	100	0	158	0	0	100	0	253

*Legend: Castle Cook (CC), Waste Water (WW), Challenge LP (CL), La Ola (LO)*



**Figure A.5.** Pareto-optimal sets for the (a) 97th, 98th, 99th, and 99.5th percentiles and (b) 99<sup>th</sup>, 99.5<sup>th</sup>, 99.9<sup>th</sup>, 99.99<sup>th</sup>, and 99.999<sup>th</sup> percentiles. The standard deviation of the geographic allocation (as a fraction of the total) is shown as a color swatch where darker color indicates low deviation and lighter color indicates a single site has a majority allocation.

Taking a more expanded view of the high end of ramp rate space, Figure A.5b shows ramp rate percentiles up to the 99.999<sup>th</sup> percentile. The space between this percentile and the 100<sup>th</sup> percentile is spanned by only 43 data points. More than 1 year of data is needed to provide statistical context for percentiles beyond the “five nines”

level; this was the maximum percentile considered here. At the 99.999<sup>th</sup> percentile, a more than 50% reduction in the magnitude of extreme ramps can be achieved with only a 3% drop in energy produced when site allocations are properly selected. In Figure A.5b, the color shading indicates that for lower ramp rates, geographic allocation was more evenly distributed. As shown, for each energy level, the site weightings appeared somewhat independent of ramp rate percentile. This was a positive result, because it implied that geographic dispersion lowers ramp rates at every level. The multi-objective optimization was run to simultaneously minimize two ramp rate percentiles (99<sup>th</sup> and 99.99<sup>th</sup>) while maximizing energy. The result indicated that the geographic allocation used to reduce ramp rates in one percentile was beneficial to the other. This means that optimal solutions at low ramp rates have similar site allocations, a fact not clear from looking at the standard deviation of weights alone. The three-dimensional plot is not shown because perspective effects make it difficult to interpret on the written page. The site weighting and ramp rates for a selection of points in Figure A.5b are given in Table A.2 as a function of energy and ramp rate percentile.

A system designer is bound by economic constraints and cannot simply minimize ramp rate while disregarding a potentially significant drop in energy output. The exact requirements are system dependent, but an example case was selected where energy was decreased from only the 8.26-MWh maximum to an 8-MWh design point (optimum allocation, OPT). The site weights for this case are given in Table A.3; the resulting ramp rate percentiles are given in Table A.4. For comparison, Table A.4 also includes the ramp rates at different percentiles when all 2MW was allocated to each site separately. As expected, the ramp rates of OPT were smaller at all percentiles, but the important result was that energy was not significantly reduced from the maximum achievable level. In fact, it was significantly higher than could be achieved with 100% allocation at either

CC or LO. Comparing OPT to the equal allocation (EQ) case shows that one can reduce ramps above OPT, but energy drops by nearly 9% below the maximum and 6% below OPT. In the limit of zero correlation of ramp rates between sites, the EQ scenario would be the best geographic allocation for ramp rate minimization. Because there is a small degree of correlation, even at short time scales, a multi-objective optimization minimizing the 99<sup>th</sup> and 99.99<sup>th</sup> percentile and ignoring energy output yielded the lowest overall ramp rates (minimum ramp allocation, MIN; Table A.4). The MIN results marginally improved both ramp rate reduction and energy production above EQ.

**Table A.3.** Site geographic allocation percentage for 7 difference scenarios: CC, WW, CL, and LO allocate all capacity at a single site; OPT uses weighting for the 8MWh design point; EQ equally weights each site; MIN weighting mix shows results from minimization of the 99th and 99.99th percentile ramp rates; and LAN is a lanai case study.

Site	Site geographic allocation [%]			
	CC	WW	CL	LO
CC	100	0	0	0
WW	0	100	0	0
CL	0	0	100	0
LO	0	0	0	100
OPT	9	30	53	8
EQ	25	25	25	25
MIN	22	26	26	26
LAN	13	13	13	61

*Legend: Castle Cook (CC), Waste Water (WW), Challenge LP (CL), La Ola (LO), Optimum Allocation (OPT), Equal Allocation (EQ), Minimum Ramp Allocation (MIN), Lanai case study (LAN)*

**Table A.4.** Energy production and ramp rates at the 99<sup>th</sup>, 99.9<sup>th</sup>, 99.99<sup>th</sup>, and 99.999<sup>th</sup> percentiles (pct) for 7 different weighting scenarios: CC, WW, CL, and LO show results for single site allocation at each respective site; OPT shows results for the 8MWh design point; the EQ case uses 25% weight at each site; MIN uses a site mix to minimize ramp rates irrespective of energy generation; and LAN is a lanai case study scenario.

Site	energy [MWh/day]	ramp rate [kW/s]			
		99 <sup>th</sup> pct	99.9 <sup>th</sup> pct	99.99 <sup>th</sup> pct	99.999 <sup>th</sup> pct
CC	6.93	86	198	314	401
WW	8.19	71	166	266	362
CL	8.27	66	158	253	339
LO	6.72	78	177	276	358
OPT	8.00	41	89	140	189
EQ	7.53	35	65	94	127
MIN	7.55	35	64	93	126
LAN	7.13	49	111	173	225

Legend: Castle Cook (CC), Waste Water (WW), Challenge LP (CL), La Ola (LO), Optimum Allocation (OPT), Equal Allocation (EQ), Minimum Ramp Allocation (MIN), Lanai case study (LAN)

To quantify the level of geographic smoothing of ramp rates, a new metric termed “ramp ratio” was introduced. The ramp ratio  $RR_j^n$  for the  $j^{\text{th}}$  site at the  $n^{\text{th}}$  percentile is defined as

$$RR_j^n = \frac{P_n(|r_j(t_i, \Delta t)|)}{P_n(|w_j r_j(t_i, \Delta t)|)} \quad \text{A.4}$$

where summation over  $j$  is implied by repeated indices. Effectively, ramp ratio compares ramping magnitude when all capacity is allotted to a single location to the aggregate ramping magnitude in a distributed scenario with allocation weights  $w_j$ . For a time period with no variability, such as on a clear day, and assuming all sites are proximal enough that the solar resource is nearly the same, the ramp rates of an individual site are a result of changes in solar geometry. In this case, the ramp ratio has a value close to 1. This corresponds to a high correlation between sites because they increase and decrease in unison with the movement of the sun. When there is variability because of




the presence of clouds, the correlation between sites begins to decrease at short time scales. This results in compensating ramps and ultimately smoothing when the plants are distributed. This characteristic is exhibited as ramp ratios that are greater than 1, indicating that the total system output is smoother on a per capacity basis than the  $j^{\text{th}}$  constituent site. The value of the ramp ratio quantifies the level of increased geographic smoothing—e.g.,  $RR_j^{99} = 2$  indicates that the aggregate system output is two times smoother than site  $j$  at the 99<sup>th</sup> percentile of ramp rates. It may happen that the ramp ratio is less than 1. This implies that a particular site is actually less variable than the aggregate. In this case, the weights allotted to the sites are obviously not optimal, with greater weights allotted to sites with high variability at the expense of low-variability sites. In such a scenario, reallocation of capacity would be beneficial.

Several examples of the ramp ratio  $RR_j^{99}$  are given in Figure A.6 for different allocation scenarios across different months. The EQ case (Figure A.6a) shows that smoothing is not uniform for all months and tends to be higher during the summer. Figure A.6b shows the OPT case. The island of Lanai currently has 1.2MW of PV installed at the LO site, so it is of practical interest to look into the scenario where the remainder of the 2MW is distributed equally among other sites. Figure A.6c shows the Lanai case study scenario (LAN, Table A.3). Overall, compared with Figure A.6b and c, EQ has the highest and most consistent  $RR_j^{99}$ , which is to be expected for nearly uncorrelated signals. During the winter half-year, WW and CL, the sites near each other, have lower  $RR_j^{99}$  in EQ and thus are less variable than CC and LO. This is one reason why the OPT case weighted WW and CL more heavily than the other sites. Figure A.6 b reflects this increase in weighting because the ramp ratios of WW and CL are more uniform throughout the year. These two sites also produce more energy on average and

therefore have a higher weighting in the optimization. The energy generation and ramp rates of the LAN scenario are well within the Pareto-inferior region (Table A.4). Figure A.6c shows the ramp ratio is lower than OPT or EQ, and in fact dips below 1 in November. Ramp ratios below 1 mean that these sites are underutilized in this scenario for their ability to decrease ramping during the winter. The previous discussion for ramp ratio used the 99<sup>th</sup> percentile as the baseline for comparison of monthly and inter-site variability. Use of a different percentile, specifically higher percentiles, provides a different quantification for understanding variability. Extremely large ramp rates—caused by the passage of an individual cumulus cloud where the solar collectors transition to and from the direct solar beam, along with greater than average cloud enhancement just beyond the leading or trailing edge—are infrequent and typically affect only a single site. These extreme ramp events at one site will generally not have any correlation to other locations. Aggregation of sites under these conditions will therefore invariably result in smoothing, which is reflected by higher ramp ratios. The larger the ramp event (and thus higher percentile), the larger the inter-site de-correlation; therefore, the ramp ratio tends to increase as percentile considered increases.

The probability distributions (PDFs) of large ramp events are shown in Figure A.7. Reviewing the design point case OPT, the ramp rates were reduced significantly throughout all four single-site allocation scenarios. The EQ case had the lowest extreme ramp rates by a significant margin. The LAN case reduced variability over the single-site cases, but ramping was larger than OPT by almost an order of magnitude toward the upper tail of the distribution, and it generated more than 10% less energy per annum. The juxtaposition of the OPT case to that of LAN underscores the benefit of designing a system for optimal use of the available solar resource while taking maximum advantage of its spatial de-correlation at the time scale considered.




**Figure A.6.** Monthly ramp ratios for (a) the equal allocation scenario, (b) the 8-MWh design point scenario optimum allocation, and (c) a Lanai, Hawaii, case study reflecting currently installed capacity at La Ola.

## A.5 Conclusions

It is well known that geographic dispersion provides a smoother overall output on a per capacity basis than centralized generation, but the optimal allocation to generate the most energy output while minimizing variability is not widely studied. A multi-objective optimization scheme was employed to investigate if variability can be significantly minimized while maintaining high levels of energy production. It was shown that an optimal set of alternative geographic allocation exists to maximize energy production and minimize variability. With proper geographic allotment of generation,

ramp rates were reduced by 50% whereas energy was reduced by only 3% above the maximum production case. A useful result for system designers is that variability can be tuned to an acceptable level, with the understanding that if the solution is optimal, then energy generation must be compromised, as indicated by the shape of the Pareto-optimal set.



**Figure A.7.** Probability distribution of ramp rates above the 99<sup>th</sup> percentile. The sign of the ramp was not considered. The distribution was scaled such that the probability of all ramps above the 99<sup>th</sup> percentile summed to unity. Allocating all capacity to four sites separately is shown (thin lines), along with three multi-site allocation cases (thick lines).

In addition to showing that there exists a clear energy-variability trade-off, a new metric termed the ramp ratio was introduced, which can be used to assess multisite geographic allocations and their effect on variability. This ratio shows simultaneously how much more smoothing the aggregate system experiences than any individual site, and also if any particular site has significantly less variability per unit capacity than the other sites in the mix. It is important to note that the ramp ratio was developed to assess

optimality of capacity allocation among multiple sites and therefore looking at a single site's ramp ratio in isolation does not provide a complete set of information. The ramp ratio can be extended into a curve in the temporal dimension by segmenting the input data set (as done in this work), or into the magnitude dimension by looking at different percentiles of variability.

The methods applied here will be extended to other variability criteria beyond ramp rates in future work. The metrics of energy and ramp rate are directly derived from resource data, but for power systems application at the distribution level, voltage fluctuations and cost will be the key metrics optimized in the kind of analysis presented here. Additional future work will involve implementing this multi-objective optimization with a feeder simulation where economic benefit is maximized and voltage variations are minimized. In this study, no restriction was placed on the quanta of generation installed at any given site and no site was given a minimum threshold to be considered active. For a realistic feeder simulation, generation siting will be practically limited by space constraints, and by operations and maintenance constraints. These elements need to be included so their impact on the Pareto-optimal set of alternatives can be quantified.

In effect, this paper creates a framework that can be used by cities and municipal utilities that would like to see a high penetration of PV and would be interested in reducing variability impacts on the distribution grid. The development that energy and variability have a clear tradeoff, and that this trade-off can be optimized for unique scenarios, offers an opportunity to communities and utilities to leverage this to their advantage. The methods presented can be incorporated into power system design and simulation tools so that designers of future PV systems can assess the geographical allocations that provide the maximum benefit. This work offers only a glimpse of what

can be achieved in maximizing the utility of the solar resource through optimized geographic allotment.

It is important to note that such studies are valuable before deployment of PV. To conduct such studies, there is a need to collect high-resolution time-synchronized solar radiation data sets for a reasonably extended period of time. Such deployments are relatively inexpensive, but can provide a significant level of benefit when planning high level penetrations of PV on the distribution grid.

### **Acknowledgements**

This work was supported by the US Department of Energy under Contract No. DE-AC36-08-GO28308 with the National Renewable Energy Laboratory. The first author would also like to thank Professor Jan Kleissl for his generous support during the period that this research was conducted.

Publication: Optimizing geographic allotment of photovoltaic capacity in a distributed generation setting. *Progress in Photovoltaics: Research and Applications*, volume 21, issue 6, pp. 1276–1285, 2012. doi: 10.1002/pip.2334

Co-Authors: Manajit Sengupta and James Keller

## References

- Acres B (1895). Some hints on photographing clouds. *Quarterly Journal of the Royal Meteorological Society*, volume 21, issue 95, pp. 160–165. doi: 10.1002/qj.4970219503.
- Anderson M (1964). Studies of the woodland climate: I. The photographic computation of light conditions, *Journal of Ecology*, volume 52, issue 1, pp. 27–41. doi: 10.2307/2257780
- ASTM E1036 - 08, "Standard Test Methods for Electrical Performance of Nonconcentrator Terrestrial Photovoltaic Modules and Arrays Using Reference Cells".
- Bacher P, Madsen H, Nielsen HA (2009). Online short-term solar power forecasting. *Solar Energy*, volume 83, issue 10, pp. 1772–1783. doi: 10.1016/j.solener.2009.05.016
- Barnard ST and Fischler MA (1982). Computational Stereo. *ACM Computing surveys*, volume 14, issue 4, pp. 553–572. doi: 10.1145/356893.356896
- Bayer BE. Color Filter Array, U.S. Patent No. 3,971,065, U.S. Patent and Trademark Office, Washington, DC, 1975.
- Beaubien M, Freedman J, 2001. Comparison of Two Imager-based Methods for Determination of Winds Aloft. *American Meteorological Society Conf.*
- Benjamin S, Brown J, Olson J, Wilczak J, Banta R, DiMego G and Weng F (2010). Improvements in NOAA modeling and data assimilation to improve guidance for the renewable energy industry. *American Meteorological Society Conf.*
- Boland J, Ridley B and Brown B (2008). Models of diffuse solar radiation. *Renewable Energy*, volume 33, pp. 575–584. doi: 10.1016/j.renene.2007.04.012
- Bradbury D and Fujita TT (1968). Computation of height and velocity of clouds from dual, whole-sky, time-lapse picture sequences. *Satellite and Mesometeorology Research Project. Dept. of the Geophysical Sciences. University of Chicago. SMRP research paper, no. 70.*
- Brown DC (1964). An advanced reduction and calibration for photogrammetric cameras. *Air Force Missile Test Center Technical Report No. 59-25, Patrick AFB, Florida.*
- Brown DC (1966). Decentering distortion of lenses. *Photogrammetric Engineering*, vol. 32, no. 3, pp.444–462.
- Brown DC (1971). Close-Range Camera Calibration. *Photogrammetric Engineering*, vol. 37, no. 8, pp. 855–866.
- Brown HE (1962). The canopy camera, Station Paper 72. Fort Collins, CO: U.S. Department of Agriculture, Forest Service, Rocky Mountain Forest and Range Experiment Station.

- Brown MZ, Burschaka D, and Hager GD (2003). Advances in computational stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 25, no. 8, pp. 993–1008. doi: 10.1109/TPAMI.2003.1217603
- Cano D, Monget J, Aubuisson M, Guillard H, Regas N and Wald L (1986). A method for the determination of global solar radiation from meteorological satellite data. *Solar Energy* 37, pp. 31–39.
- Cazorla A, Olmo FJ and Alados-Arboledas L (2008). Development of a sky imager for cloud cover assessment, *Journal of the Optical Society of America*, volume 25, number 1, pp. 29–39. doi: 10.1364/JOSAA.25.000029
- Cazorla A, Shields J, Karr M, Olmo F, Burden A and Alados-Arboledas L (2009). Technical Note: Determination of aerosol optical properties by a calibrated sky imager, *Journal of Atmospheric Chemistry and Physics*, volume 9, pp. 6417–6427. doi: 5194/acp-9-6417-2009
- Chow CW, Urquhart B, Lave M, Dominguez A, Kleissl J, Shields JE, Washom B (2011). Intra-hour forecasting with a total sky imager at the UC San Diego solar energy testbed. *Journal of Solar Energy*; volume 85, issue 11, pp. 2881–2893. doi: 10.1016/j.solener.2011.08.025
- Chu Y, Pedro HTC and Coimbra CFM (2013). Hybrid intra-hour DNI forecasts with sky image processing enhanced by stochastic learning. *Journal of Solar Energy*, volume 98, pp. 592–603. doi: 10.1016/j.solener.2013.10.020
- Clayden AW (1898). Photographing meteorological phenomena. *Quarterly Journal of the Royal Meteorological Society*, volume 24, issue 107, pp. 169–180. doi: 10.1002/qj.49702410701
- Collins RD and Crowther KG (2011). Systems-based modeling of generation variability under alternate geographic configurations of photovoltaic (PV) installations in Virginia. *Energy Policy*, volume 39, issue 10, pp. 6262–6270. doi: 10.1016/j.enpol.2011.07.025
- Conrady A (1918). The five aberrations of lens-systems. *Monthly Notices of the Royal Astronomical Society*, volume 79, issue 1, pp. 60–65. doi: 10.1093/mnras/79.1.60
- Conrady A (1919). Decentering Lens Systems. *Monthly Notices of the Royal Astronomical Society*, volume 79, issue 5, pp. 384–390. doi: 10.1093/mnras/79.5.384
- Curtright AE and Apt J. The character of power output from utility scale photovoltaic systems. *Progress in Photovoltaics*; volume 16, pp. 241–247. doi: 10.1002/pip.786
- Deb K. *Multi-objective optimization using evolutionary algorithms*. Wiley : Chichester, England, 2011.
- Debevec P and Malik J (1997). Recovering high dynamic range radiance maps from photographs, *Proceedings of the ACM SIGGRAPH'97*, Los Angeles CA, pp. 369–378.



- Devernay F and Faugueras O (2001). Straight lines have to be straight. *Machine Vision and Applications*, volume 13, issue 1, pp. 14–24. doi: 10.1007/PL00013269
- Dhond UR and Aggarwal JK (1989). Structure from stereo – A review. *IEEE Transactions of Systems, Man and Cybernetics*, volume 19, issue 6, pp. 1489–1510. doi: 10.1109/21.44067
- Dupree W, Morse D, Chan M, Tao X, Iskenderian H, Reiche C, Wolfson M, Pinto J, Williams JK, Albo D, Dettling S, Steiner M, Benjamin S and Weygandt S (2009). The 2008 CoSPA forecast demonstration (collaborative storm prediction for aviation), *Proceedings of the 89th Meeting of the American Meteorological Society, Special Symposium on Weather - Air Traffic*. Phoenix, AZ.
- Dye D (2012). Looking skyward to study ecosystem dynamics. *Eos, Transactions of the American Geophysical Union*, volume 93, issue 14, pp. 141–143. doi: 10.1029/2012EO140002
- Elkholm N (1893). A new instrument for cloud measurements. *Quarterly Journal of the Royal Meteorological Society*, volume 19, issue 85, pp. 38–45. doi: 10.1002/qj.4970198506
- Farr TG, Rosen PA, Caro E, Crippen R, Duren R, Hensley S, Kobrick M, Paller M, Rodriguez E, Roth L, Seal D, Shaffer S, Shimada J, Umland J, Werner M, Oskin M, Burbank D, Alsdorf D (2007). The shuttle radar topography mission. *Review of Geophysics*, volume 45, issue 2, RG2004. doi: 10.1029/2005RG000183
- Faugeras O (1993). *Three-dimensional Computer Vision: A geometric viewpoint*. MIT Press, Cambridge. ISBN 0-262-06158-9
- Feister U and Shields JE (2005): Cloud and radiance measurements with the VIS/NIR Daylight Whole Sky Imager at Lindenberg (Germany). *Meteorologische Zeitschrift*, volume 14, issue 5, pp. 627–639, 2005. doi: 10.1127/0941-2948/2005/0066
- Förstner W and Gülch W (1987). A fast operator for detection and precise location of distinct points, corners and centres of circular features. *Proceedings of the ISPRS Workshop on Fast Processing of Photogrammetric Data*, Interlaken, Switzerland pp. 281–305.
- Fu CL and Cheng HY (2013). Predicting solar irradiance with all-sky image features via regression. *Journal of Solar Energy*, volume 97, pp. 537–550. doi: 10.1016/j.solener.2013.09.016
- Gennery DB (2006). Generalized camera calibration including fisheye lenses. *International Journal of Computer Vision*, volume 68, no. 3, pp. 239–266. doi: 10.1007/s11263-006-5168-1
- Gohari SMI, Urquhart B, Yang H, Kurtz B, Nguyen D, Chow CW, Ghonima M and Kleissl J (2014). Comparison of solar power output forecasting performance of the Total Sky Imager and the University of California, San Diego Sky Imager. *Energy Procedia*, volume 49, pp. 2340–2350. doi: 10.1016/j.egypro.2014.03.248

- Gonzales Y, Lopez C, Cuevas E (2012). Automatic observation of cloudiness: analysis of all-sky images. TECO-2012, WMO Technical Conference on Meteorological and Environmental Instruments and Methods of Observation, Brussels, Belgium, 16-18 October 2012.
- Hamill T and Neerkorn T (1993). A short-term cloud forecast scheme using cross correlations. *Journal of Weather and Forecasting*, volume 8, issue 4, pp. 401–411. doi: 10.1175/1520-0434(1993)008<0401:ASTCFS>2.0.CO;2
- Hammer A, Heinemann D, Lorenz E and Lückehe B (1999). Short-term forecasting of solar radiation: a statistical approach using satellite data. *Solar Energy*, volume 67, issues1-3, pp. 139-150. doi: 10.1016/S0038-092X(00)00038-4
- Hartley RI (1997). In defense of the eight-point algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 19, issue 6, pp. 580–593. doi: 10.1109/34.601246
- Hartley RI and Sturm P (1997). Triangulation. *Computer Vision and Understanding*, volume 68, issue 2, pp. 146–157. doi: 10.1006/cviu.1997.0547
- Hartley R and Zisserman A (2004). *Multiple view geometry in computer vision*. Cambridge University Press. ISBN: 9780511811685, doi: 10.1017/CBO9780511811685
- Heikkilä J and Silvén O (1996). Calibration Procedure for Short Focal Length Off-the-Shelf CCD Cameras. *Proceedings of the International Conference on Pattern Recognition*, pp. 166–170. doi: 10.1109/ICPR.1996.546012
- Heinle A, Macke A and Srivastav A (2010). Automatic cloud classification of whole sky images. *Journal of Atmospheric Measurement Techniques*, volume 3, pp. 557–567. doi: 10.5194/amt-3-557-2010
- Hill R(1924). A lens for whole sky photographs. *Quarterly Journal of the Royal Meteorological Society*, volume 50, issue 211, pp. 227–235. doi: 10.1002/qj.49705021110
- Huo J and Lu D (2009). Cloud determination of all-sky images under low visibility conditions. *Journal of Atmospheric and Oceanic Technology*, volume 26, issue 10, pp. 2172–218. doi: 10.1175/2009JTECHA1324.1
- Huo J and Lu D (2010). Preliminary retrieval of aerosol optical depth from all-sky images. *Advances in Atmospheric Science*, volume 27, issue 2, pp. 421–426. doi: 10.1007/s00376-009-8216-2
- Huo J and Lu D(2012). Comparison of cloud cover from all-sky imager and meteorological observer. *Journal of Atmospheric and Oceanic Technology*, volume 29, issue 8, pp. 1093-1101. doi: 10.1175/JTECH-D-11-00006.1
- Ineichen P and Perez R (2002). A new airmass independent formulation for the Linke turbidity coefficient. *Journal of Solar Energy*, volume 73, issue 3, pp. 151–157. doi: 10.1016/S0038-092X(02)00045-2

- Ineichen P (2006). Comparison of eight clear sky broadband models against 16 independent data banks. *Journal of Solar Energy*, volume 80, issue 4, pp. 468–478. doi: 10.1016/j.solener.2005.04.018
- Johnson RW, Koehler TL and Shields JE (1988). A Multi-Station Set of Whole Sky Imagers and A Preliminary Assessment of the Emerging Data Base. *Proceedings of the Cloud Impacts on DOD Operations and Systems - 1988 Workshop*, pp. 159-162, 1988.
- Johnson RW, Hering WS and Shields JE (1989). Automated Visibility and Cloud Cover Measurements with a Solid-State Imaging System. University of California, San Diego, Scripps Institution of Oceanography, Marine Physical Laboratory, SIO Ref. 89- 7, GL- TR-89-0061, NTIS No. ADA216906, 1989.
- Johnson R, Shields JE and Koehler T (1991). Analysis and Interpretation of Simultaneous Multi-Station Whole Sky Imagery. University of California, San Diego, Scripps Institution of Oceanography, Marine Physical Laboratory, SIO 91-3, PL-TR-91-2214.
- Kalisch J and Macke A (2008). Estimation of the total cloud cover with high temporal resolution and parametrization of short-term fluctuations of sea surface insolation. *Meteorologische Zeitschrift*, volume 17, issue 5, pp. 603–611. doi:10.1127/0941-2948/2008/0321
- Kannala J and Brandt SS (2006). A generic camera model and calibration method for conventional, wide-angle, and fish-eye lenses. *Pattern Analysis and Machine Intelligence, IEEE Transactions*, volume 28, issue 8, pp.1335–1340. doi: 10.1109/TPAMI.2006.153
- Kassander AR and Sims LL (1957). Cloud photogrammetry with ground-located K-17 aerial cameras. *Journal of Meteorology*, volume 14, issue 1, pp. 43–49. doi: 10.1175/0095-9634-14.1.43
- Kassianov E, Long C and Christy J (2005). Cloud-base-height estimation from paired ground-based hemispherical observations. *Journal of Applied Meteorology*, volume 44, issue 8, pp. 1221–1233. doi: 10.1175/JAM2277.1
- Kasten F and Young A (1989). Revised Optical Airmass Tables and Approximation Formula. *Applied Optics*, volume 28, issue 22, pp. 4735–4738. doi: 10.1364/AO.28.004735
- Klebe DI, Blatherwick RD and Morris VR (2014). Ground-based all-sky mid-infrared and visible imagery for purposes of characterizing cloud properties. *Atmospheric Measurement Techniques*, volume 7, pp. 637-645. doi: 10.5194/amt-7-637-2014
- Kleissl J (2013). *Solar Energy Forecasting and Resource Assessment*. Academic Press, Boston, 2013. ISBN 978-0-12-397177-7, doi:10.1016/B978-0-12-397177-7.00010-3

- Lalonde JF, Narasimhan SG and Efros AA (2010). What Do the Sun and the Sky Tell Us About the Camera? *International Journal of Computer Vision*, volume 88, issue 1, pp 24–51. doi: 10.1007/s11263-009-0291-4
- Lara-Fanego V, Ruiz-Arias JA, Pozo-Vazquez D, Santos-Alamillos F and Tovar-Pescador J (2011). Evaluation of the WRF model solar irradiance forecasts in Andalusia (southern Spain). *Journal of Solar Energy*, volume 86, issue 8, pp. 2200–2217. doi:10.1016/j.solener.2011.02.014
- Lave M, Stein J, Ellis A, Hansen C, Nakashima E and Miyamoto Y. Ota City: Characterizing output variability from 553 homes with residential PV systems on a distribution feeder. Sandia National Laboratories: Albuquerque, NM, 2011; report SAND2011-9011.
- Lave M, Kleissl J and Arias-Castro E (2012). High-frequency irradiance fluctuations and geographic smoothing. *Journal of Solar Energy*, volume 86, pp. 2190–2199. doi: 10.1016/j.solener.2011.06.031
- Lave M and Kleissl J (2013). Cloud speed impact on solar variability scaling - application to the wavelet variability model. *Journal of Solar Energy*, volume 91, pp. 11–21. doi: 10.1016/j.solener.2013.01.023.
- Lewis JP (1995). Fast Template Matching. *Proceedings of Vision Interface 95*, Canadian Image Processing and Pattern Recognition Society, Quebec City, Canada, May 15-19, 1995, pp. 120-123. (Also see the expanded and corrected version "Fast Normalized Cross-Correlation").
- Liebowitz D and Zisserman A (1998). Metric rectification for perspective images of planes. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 482-488. doi: 10.1109/CVPR.1998.698649
- Linke, F., 1922. Transmissions-Koeffizient und Trübungsfaktor. *Beitr. Phys. fr. Atmos.* 10, 91–103.
- Long CN and DeLuisi JJ (1998). Development of an automated hemispheric sky imager for cloud fraction retrievals. *Proceedings of the 10th Symposium on Meteorological Observations and Instrumentation*, Phoenix, Arizona, American Meteorological Society, pp. 171–174.
- Long CN, Sabburg JM, Calbó J and Pagès D (2006). Retrieving cloud characteristics from ground-based daytime color all-sky images. *Journal of Atmospheric and Oceanic Technology*, volume 23, issue 5, pp. 633–652. doi: 10.1175/JTECH1875.1
- Longuet-Higgins HC (1981). A computer algorithm for reconstructing a scene from two projections. *Nature*, volume 293, pp. 133–135. doi: 10.1038/293133a0
- Lorenz E, Hurka J, Heinemann D and Beyer HG (2009a). Irradiance forecasting for power prediction of grid-connected photovoltaic systems. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, volume 2, issue 1, pp. 2–10. doi: 10.1109/JSTARS.2009.2020300

- Lorenz E, Remund J, Müller S, Traunmüller W, Steinmaurer G, Pozo D, Ruiz-Arias J, Fanego V, Ramirez L, Romeo M, Kurz C, Pomares L and Guerrero C (2009b). Benchmarking of different approaches to forecast solar irradiance. 24th European Photovoltaic Solar Energy Conference, Hamburg, Germany, 21 –25.
- Luoma J, Kleissl J and Murray K (2011). Optimal inverter sizing considering cloud enhancement. *Journal of Solar Energy*, volume 86, issue 1, pp. 421–429, 2011
- Malkus JS and Ronne C (1954). On the structure of some cumulonimbus clouds which penetrated the high tropical troposphere. *Tellus*, volume 6, issue 4, pp. 351–366. doi: 10.1111/j.2153-3490.1954.tb01130.x
- Mann S and Picard RW(1994). On being 'undigital' with digital cameras: extending dynamic range by combining differently exposed pictures. Technical Report 323, MIT Media Lab Perceptual Computing Section, 1994. Also in *Proceedings of Imaging Science and Technology 48th annual conference*, May 1995, 422–428, 1995.
- Marcos J, Marroyo L, Lorenzo E, Alvira D and Izco E (2011). Power output fluctuations in large scale PV plants: one year observations with one second resolution and a derived analytic model. *Progress in Photovoltaics: Research and Applications*, volume 19, issue 2, pp. 218-227. doi: 10.1002/pip.1016
- Marcos J, Marroyo L, Lorenzo E, Alvira D and Izco E (2011). From irradiance to output power fluctuations: the pv plant as a low pass filter. *Progress in Photovoltaics: Research and Applications*, volume 19, issue 5, pp. 505-510. doi: 10.1002/pip.1063
- Marcos J, Marroyo L, Lorenzo E and Garcia M. Smoothing of PV power fluctuations by geographical dispersion. *Progress in Photovoltaics: Research and Applications*, volume 20, issue 2, pp. 226-237. doi: 10.1002/pip.1127
- Marion B, Anderberg M, George R, Gray-Hann P and Heimiller D. PVWATTS Version 2 – Enhanced Spatial Resolution for Calculating Grid-Connected PV Performance. National Renewable Energy Laboratory: Golden, CO 2001, report NREL/CP-560-30941.
- Marquez R and Coimbra CFM (2011). Forecasting of global and direct solar irradiance using stochastic learning methods, ground experiments and the NWS database. *Solar Energy*, volume 85, issue 5, pp. 746–756. doi: 10.1016/j.solener.2011.01.007
- Marquez R, Gueorguiev VG and Coimbra CFM (2012). Forecasting of global horizontal irradiance using sky cover indices. *Journal of Solar Energy Engineering*, volume 135, issue 1, pp. 001017-1–001017-5. doi: 10.1115/1.4007497
- Marquez R and Coimbra CFM (2013). Intra-hour DNI forecasting based on cloud tracking image analysis. *Journal of Solar Energy*, volume 91, pp. 327–336. doi: 10.1016/j.solener.2012.09.018.
- Martins F, Pereira E and Abreu S (2007). Satellite-derived solar resource maps for Brazil under SWERA project. *Journal of Solar Energy* volume 81, pp. 517–528. ISSN 0038-092X, doi: 10.1016/j.solener.2006.07.009

- Mathiesen P and Kleissl J (2011). Evaluation of numerical weather prediction for intra-day solar forecasting in the continental United States. *Journal of Solar Energy*, volume 85, issue 5, pp. 967–977. doi: 10.1016/j.solener.2011.02.013
- Mathiesen P, Collier C and Kleissl J (2013). A high-resolution cloud-assimilating numerical weather prediction model for solar irradiance forecasting. *Journal of Solar Energy*, volume 92, 47–61. doi: 10.1016/j.solener.2013.02.018
- McCartney E (1976). *Optics of the Atmosphere: Scattering by Molecules and Particles*. John Wiley & Sons, pp. 216.
- Meeus J (1998). *Astronomical Algorithms* (2nd edition). Willmann-Bell Inc. Richmond Virginia, USA, 1998. ISBN: 978-0943396354
- Menicucci DF and Fernandez JP (1988). User's Manual for PVFORM: A photovoltaic simulation program for stand-alone and grid interactive applications. Sandia National Laboratories: Albuquerque, NM, 1988; report SAND85-0376 · UC-276.
- Menicucci DF (1986). Photovoltaic array performance simulation models. *Solar Cells*, volume 18, pp. 383–392.
- Miyamoto K (1964). Fish Eye Lens. *Journal of the Optical Society of America*, volume 54, issue 8, pp. 1060–1061. doi: 10.1364/JOSA.54.001060
- Monteiro C, Bessa R, Miranda V, Botterud A, Wang J and Conzelmann G (2009). *Wind Power Forecasting : State-of-the-Art 2009*. Argonne National Laboratory, Decision and Information Sciences Division, report ANL/DIS-10-1. doi:10.2172/968212
- Murata A and Otani K (1997). An analysis of time-dependent spatial distribution of output power from very many PV power systems installed on a nationwide scale in Japan. *Solar Energy Materials and Solar Cell*, volume 47, pp. 197–202. doi: 10.1016/S0927-0248(97)00040-8
- Nakajima T, Tonna G, Rao R, Boi P, Kaufman Y and Holben B (1996). Use of sky brightness measurements from ground for remote sensing of particulate polydispersions. *Journal of Applied Optics*, volume 35, issue 15, pp. 2672–2686. doi: 10.1364/AO.35.002672
- Ohta Y and Kande T (1983). Stereo by intra- and inter-scanline search using dynamic programming. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 7, issue 2, pp. 139–154. doi: 10.1109/TPAMI.1985.4767639
- Öktem R, Prabhat, Lee J, Thomas A, Zuidema P and Romps DM (2014). Stereophotogrammetry of oceanic clouds. *Journal of Atmospheric and Oceanic Technology*, volume 31, issue 7, pp. 1482–1501. doi: 10.1175/JTECH-D-13-00224.1
- Orville HD and Kassander AR (1961). Terrestrial photogrammetry of clouds. *Journal of Meteorology*, volume 18, issue 5, pp. 682–687. doi: 10.1175/1520-0469(1961)018<0682:TPOC>2.0.CO;2

- Otani K, Minowa J, Kurokawa K. Study on a real solar irradiance for analyzing areally-totalized PV systems. *Solar Energy Materials and Solar Cells*, volume 47, pp. 281–288. doi: 10.1016/S0927-0248(97)00050-0
- Page J (2003). The role of solar radiation climatology in the design of photovoltaic systems, chapter X in Markvart T and Castaner L, editors, *Practical handbook of photovoltaics: fundamentals and applications*. Elsevier, Oxford, 2003, pp. 5–66.
- Pedro HTC and Coimbra CFM (2012). Assessment of forecasting techniques for solar power production with no exogenous inputs. *Solar Energy*, volume 86, issue 7, pp. 2017–2028. doi: 10.1016/j.solener.2012.04.004
- Pereira E, Martins F, Abreu S, Couto P, Stuhlmann R and Colle S (2000). Effects of burning of biomass on satellite estimations of solar irradiation in Brazil. *Journal of Solar Energy*, volume 68, pp. 91–107. ISSN 0038-092X, doi: 10.1016/S0038-092X(99)00044-4
- Perez R and Seals R (1987). A new simplified version of the Perez diffuse irradiance model for tilted surfaces. *Journal of Solar Energy*, volume 39, issue 3, pp. 221–231. doi: 10.1016/S0038-092X(87)80031-2
- Perez R, Doty J, Baily B and Stewart R (1994). Experimental evaluation of a photovoltaic simulation program. *Journal of Solar Energy*, volume 52, issue 4, pp. 359-365. doi: 10.1016/0038-092X(94)90142-2
- Perez R, Ineichen P, Moore K, Kmiecik M, Chain C, George R and Vignola F (2002). A new operational model for satellite-derived irradiances: description and validation. *Journal of Solar Energy*, volume 73, issue 5, pp. 307–317. doi: 10.1016/S0038-092X(02)00122-6
- Perez R, Kivalov S, Schlemmer J, Hemker Jr. K, Renné D and Hoff T (2010). Validation of short and medium term operational solar radiation forecasts in the US. *Solar Energy*, volume 84, issue 12, pp. 2161–2172. doi: 10.1016/j.solener.2010.08.014
- Perez R, Cebecauer T and Šúri M (2013a). Semi-empirical satellite models, chapter 2 in: Kleissl J, editor: *Solar Energy Forecasting and Resource Assessment*. Elsevier; 2013. ISBN: 978-0-12-397177-7. doi: 10.1016/B978-0-12-397177-7.00002-4
- Perez R, Hoff TE (2013b). *SolarAnywhere Forecasting*, chapter 10 in: Kleissl, J., editor: *Solar Energy Forecasting and Resource Assessment*. Elsevier; 2013. ISBN: 978-0-12-397177-7, doi:10.1016/B978-0-12-397177-7.00010-3
- Perez R, Lorenz E, Pelland S, Beauharnois M, Van Knowe G, Hemker Jr. K, Heinemann D, Remund J, Müller SC, Traunmüller W, Steinmayer G, Pozo D, Ruiz-Arias JA, Lara-Fanego V, Ramirez-Santigosa L, Gaston-Romero M and Pomares LM. (2013c). Comparison of numerical weather prediction solar irradiance forecasts in the US, Canada and Europe. *Solar Energy*, volume 94, pp. 305–326, doi: 10.1016/j.solener.2013.05.005

- Pfister G, McKenzie R, Liley J and Thomas A (2003). Cloud coverage based on all-sky imaging and its impact on surface solar irradiance, *Journal of Applied Meteorology*, 42, 1421–1434.
- Ramalingam S, Sturm P and Lodha SK (2005). Towards complete generic camera calibration. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pp. 1093–1098. doi: 10.1109/CVPR.2005.347
- Reda I and Andreas A (2004). Solar position algorithm for solar radiation applications. *Journal of Solar Energy*, volume 76, issue 5, pp. 577–589. doi: 10.1016/j.solener.2003.12.003
- Remund J, Perez R and Lorenz E (2008). Comparison of Solar Radiation Forecasts for the USA. *Proceedings of the 23rd European PV Conference*, September 1-4, Valencia, Spain.
- Rigollier C, Lefèvre M and Wald L (2004). The method Heliosat-2 for deriving shortwave solar radiation from satellite images. *Journal of Solar Energy*, volume 77, issue 2, 159–169, ISSN 0038-092X, doi: 10.1016/j.solener.2004.04.017
- Rogers M, Miller S, Combs C, Benjamin S, Alexander C, Sengupta M, Kleissl J and Mathiesen P (2012). Validation and analysis of HRRR insolation forecasts using SURFRAD. *Proceedings of the American Solar Energy Society*, Raleigh, NC, 2012.
- Román R, Antón M, Cazorla A, de Miguel A, Olmo FJ, Bilbao J and Alados-Arboledas L (2012). Calibration of an all-sky camera for obtaining sky radiance at three wavelengths. *Journal of Atmospheric Measurement Techniques*, volume 5, 2013–2024, 2012. doi: 10.5194/amt-5-2013-2012
- Roy S. Optimal planning for utility generation by photovoltaic sources spread across multiple sites. *IEEE Transactions on Energy Conversion*, volume 21, issue 1, pp. 181–186. doi: 10.1109/TEC.2005.845452
- Sebag J, Krabbendam VL, Claver CF, Andrew J, Barr JD and Klebe D(2008). LSST IR camera for cloud monitoring and observation planning. *Proceedings of the International Society of Photonics and Optics, Ground-based and Airborne Telescopes II conference (SPIE 7012)*, volume 7012, 2008. doi: 10.1117/12.789570
- Schmetz J (1989). Towards a surface radiation climatology: retrieval of downward irradiances from satellites. *Atmospheric Research*, volume 23, issues 3-4, pp. 287–321. doi: 10.1016/0169-8095(89)90023-9
- Schroedter-Homscheidt M, Hoyer-Klick C, Rikos E, Tselepis S and Pulvermüller B (2009). Nowcasting and forecasting of solar irradiance for energy electricity generation, *SolarPACES Conference*.
- Seiz G, Shields J, Feister U, Baltsavias E, Gruen A (2007). Cloud mapping with ground-based photogrammetric cameras. *International Journal of Remote Sensing*, volume 28, issue 9, pp. 2001–2032. doi: 10.1080/01431160600641822



- Shah S and Aggarwal J (1996). Intrinsic Parameter Calibration Procedure for a (High-Distortion) Fish-Eye Lens Camera with Distortion Model and Accuracy Estimation. *Journal of Pattern Recognition*, volume 29, issue 11, pp. 1775–1788. doi: 10.1016/0031-3203(96)00038-6
- Shields JE, Johnson RW and Koehler TL (1993a). Automated Whole Sky Imaging Systems for Cloud Field Assessment. Proceedings of the Fourth Symposium on Global Change Studies, January 17–22, American Meteorological Society, Boston.
- Shields JE, Johnson RW and Karr ME (1993b). Automated Whole Sky Imagers for Continuous Day and Night Cloud Field Assessment. Proceedings of the Cloud Impacts on DOD Operations and Systems Conference.
- Shields JE, Karr ME, Tooman TP, Sowle DH and Moore ST (1998a). The Whole Sky Imager –A Year of Progress, Proceedings of Eighth Atmospheric Radiation Measurement (ARM) Science Team Meeting, Tucson, Arizona, March 23–27.
- Shields JE, Johnson RW, Karr ME, Wertz JL (1998b). Automatic day/night whole sky imager for field assessment of cloud cover distributions and radiance distributions. Proceedings of the 10th Symposium on Meteorological Observations and Instrumentations, Phoenix, AZ, American Meteorological Society, Boston, 165–170.
- Shields J, Karr M, Burden A, Johnson R and Hodgkiss W (2007). Continuing Support of Cloud Free Line of Sight Determination including Whole Sky Imaging of Clouds. Final Report for ONR Contract N00014-01-D-0043 DO #13, Marine Physical Laboratory, Scripps Institution of Oceanography, University of California San Diego, Technical Note 273.
- Shields J, Karr M, Burden A, Johnson R, Mikuls V, Streeter J and Hodgkiss W (2009). Research toward Multi-Site Characterization of Sky Obscuration by Clouds. Final Report for Grant N00244-07-1-009, Marine Physical Laboratory, Scripps Institution of Oceanography, University of California San Diego, Technical Note 274.
- Shields JE, Karr ME, Johnson RW and Burden AR (2013). Day/night whole sky imagers for 24-h cloud and sky assessment: history and overview. *Journal of Applied Optics*, volume 52, issue 8, pp. 1605–1616. doi: 10.1364/AO.52.001605
- Slama CC, Theurer C and Henriksen SW (1980). *Manual of Photogrammetry*, fourth edition. American Society of Photogrammetry. ISBN: 0-937294-01-2
- Strachey R and Whipple GM (1890). Cloud Photography conducted under the Meteorological Council at the Kew Observatory. *Proceedings of the Royal Society of London*, volume 49, no. 296–301, pp. 467–480.
- Stumpf J, Tchou C, Jones A, Hawkins T, Wenger A and Debevec P (2004). Direct HDR capture of the sun and sky. AFRIGRAPH 2004, 3rd International Conference on Computer Graphics, Virtual Reality, Visualization and Interaction in Africa, Cape Town, South Africa — November 3-5. doi: 10.1145/1029949.1029977

- Sturm PF and Maybank SJ (1999). On plane-based camera calibration: a general algorithm, singularities, applications. Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition, volume 1, pp. 432–437. doi: 10.1109/CVPR.1999.786974
- Tardif JP, Sturm P and Roy S (2006). Self-calibration of a general radially symmetric distortion model. Appearing in, Leonardis A, Bischof H, Pinz A, editors: Lecture Notes in Computer Science, Computer Vision - ECCV 2006, Part IV of the Proceedings of the 9th European Conference on Computer Vision, Graz, Austria, May 2006, pp. 186–199. ISBN: 978-3-540-33838-3, doi: 10.1007/11744085\_15
- Tsai R (1987). A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. IEEE Journal of Robotics and Automation, volume 4, issue 4, pp.323–344. doi: 10.1109/JRA.1987.1087109
- Urquhart B, Chow CW, Nguyen D, Kleissl J, Sengupta M, Blatchford J and Jeon D (2012). Towards intra-hour solar forecasting using two sky imagers at a large solar power plant. Proceedings of the American Solar Energy Society, Denver, CO, USA, 2012.
- Urquhart B, Ghonima M, Nguyen D, Kurtz B, Chow CW and Kleissl J (2013). Sky imaging systems for short-term solar forecasting, chapter 9 in: Kleissl J, editor: Solar Energy Forecasting and Resource Assessment. Elsevier; 2013. ISBN: 978-0-12-397177-7, doi: 10.1016/B978-0-12-397177-7.00009-7
- Urquhart B, Kurtz B, Dahlin E, Ghonima M, Shields JE and Kleissl J (2014a). Development of a sky imaging system for short-term solar power forecasting. Atmospheric Measurement Techniques Discussion, volume 7, pp. 4859–4907. doi:10.5194/amtd-7-4859-2014
- Urquhart B, Kurtz B and Kleissl J (2014b). Sky camera geometric calibration using solar observations. IEEE Transactions on Pattern Analysis and Machine Intelligence, *in submission*.
- Urquhart B and Kleissl J (2014). Dense cloud stereo photogrammetry using dynamic programming. *In preparation*.
- Warner C, Renick JH, Balshaw MW and Douglas RH (1973). Quarterly Journal of the Royal Meteorological Society, volume 99, issue 419, pp. 105–115. doi: 10.1002/qj.49709941910
- Wei G and Ma S (1993). A complete two-plane camera calibration method and experimental comparisons. Proceedings of the Fourth International Conference on Computer Vision, pp. 439–436. doi: 10.1109/ICCV.1993.378183
- Weng J, Cohen O and Hernriou M (1992). Camera calibration with Distortion Models and Accuracy Evaluation. IEEE Transactions on Pattern Analysis and Machine Intelligence, volume 14, issue 10, pp. 965–980. doi: 10.1109/34.159901

- Whipple GM (1890). A brief notice respecting photography in relation to meteorological work. *Quarterly Journal of the Royal Meteorological Society*, volume 16, issue 75, pp. 141–146. doi: 10.1002/qj.4970167502
- Wiemken E, Beyer HG, Heydenreich W and Kiefer K (2001). Power characteristics of PV ensembles: experiences from the combined power production of 100 grid connected PV systems distributed over the area of Germany. *Journal of Solar Energy*, volume 70, issue 6, pp. 513–518. doi: 10.1016/S0038-092X(00)00146-8.
- Wood-Bradley P, Zapata J and Pye J (2012). Cloud tracking with optical flow for short-term solar forecasting, *Proceedings of the 50th Conference of the Australian Solar Energy Society*, Melbourne, December 2012.
- Yang H, Kurtz B, Nguyen D, Urquhart B, Chow CW, Ghonima M and Kleissl J (2014). Solar irradiance forecasting using a ground-based sky imager developed at UC San Diego. *Journal of Solar Energy*, volume 103, pp. 502–524. doi: 10.1016/j.solener.2014.02.044
- Zehnder JA, Hu J and Razdan A (2007). A stereo photogrammetric technique applied to orographic convection. *Monthly Weather Review*, volume 135, issue 6, pp. 2265–2277. doi: 10.1175/MWR3401.1
- Zelenka A, Perez R, Seals R, Renné D (1999). Effective accuracy of satellite-derived irradiance. *Theoretical and Applied Climatology*, volume 62, issues 3-4, pp. 199–207.
- Zhang Z (2000). A Flexible New Technique for Camera Calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 22, issue 11, pp. 1330–1334.