Chapter

5 *Sky Models for Lighting Simulation*

"It's cloudy... and then there's a sun"

ELSPETH

The validation results presented in Chapter 4 have shown that the *Radiance* system can predict internal illuminance to a high degree of accuracy for a wide range of naturally occurring sky conditions. Measured sky brightness data is however, at this point in time, very limited. Long time-series data exist for only a few sites in the world, largely collected as part of the International Daylight Measurement Year.[1] For the majority of lighting scientists and practitioners, non-overcast sky luminance distributions for their locale will have to be derived from measurements of integrated quantities, e.g. irradiance data from weather tapes. This necessitates the use of a theoretical model to generate the sky luminance distribution.

In this Chapter, the performance of a range of sky models is evaluated in terms of their ability to reproduce a sky luminance patterns for the purpose

1. See IDMP website: **http://idmp.entpe.fr/** for a list of the stations.

of illuminance prediction. Performance criteria were based on the error characteristics for predictions of the four total vertical illuminances, and internal illuminances at the six photocell locations in the BRE office (Figure 3-4). Four 'pure' sky models and two sky model blends were evaluated. Sky models and how they are used in *Radiance* was first introduced in Chapter 2 (see Section 2.1.2, 2.3.3, 2.7.1 and 2.7.3).

5.1 Introduction

Sky models generate continuous sky luminance patterns. The discontinuous aspects of skylight - instantaneous cloud patterns - are not addressed. Attempts have been made to devise a theoretical framework which provides for the inclusion of discontinuous brightness features (that is, clouds) on a continuous luminance distribution [Perez 93b]. Here Perez *et al* investigated the possibility of parameterising the magnitude and spatial distribution of discontinuous features based on indices for the sky clearness and sky brightness. The method can be applied to any continuous sky brightness distribution model and may be a way of reconstructing some of the random aspects of daylight from measurements of integrated quantities. It is not, of course, expected to reproduce actual sky brightness configurations observed at a particular instant.

Differences that may arise between measured and modelled sky luminance patterns can result from one or both of the following:

1. The model was unable to reproduce the underlying continuous luminance pattern of the measured sky.
2. The underlying luminance pattern of the measured sky may have been accurately reproduced, but the model did not account for the random-discontinuous features that were present in the measurements.

Evidently, the role of sky model validation is to evaluate the performance of theoretical models based on the first of these causes.[2] Preliminary

comparisons between various sky models and measured sky luminance patterns have recently been reported [Ineichen 94][Littlefair 94].

5.1.1 Real and model skies

Some of the differences, and similarities, between measured and modelled sky luminance patterns are demonstrated in the following examples. The luminance patterns of four measured skies are presented alongside luminance patterns generated by a sky model. The four skies were selected from the BRE-IDMP validation dataset (Section 3.1) to demonstrate something of the diversity in naturally occurring conditions. They cover the range from heavily overcast, through two intermediate skies, to clear sky conditions. The inputs to the sky model generator program were measurements of the direct normal and the diffuse horizontal illuminance recorded at the same time as the scan. The sun description used in both the measured and the theoretical representations was the same for any one sky. The measured and modelled skies are labelled Lumscan and Skymodel respectively. The measured sky luminance patterns were based on the 145 readings taken by the Krochmann sky scanner (Section 3.1.3). The luminance measurements were interpolated to a regular grid compatible with the *Radiance* brightdata format (Section 3.2.4). The model sky description was generated using the **gendaylit** program (Section 2.7.3). This program creates a luminance distribution based on the Perez 'All-Weather' model [Perez 93]. A summary of the specification for the measured and modelled skies is given in Table 5-1.

For these illustrations, the sky brightness distribution is shown as a luminance surface. The height of the surface (z-axis) is proportional to the sky point luminance. *Radiance* was used to generate the 'views' of the measured and modelled skies from which the luminance surfaces were

2. It should be noted that what is considered to be a continuous or discontinuous feature may depend on the angular resolution of the sky scanner. For example, a sky with cirrus of cirrocumulus formations could be recorded as having a very uneven luminance distribution if the scanner managed to resolve the finescale luminance patterns of the cloudlets.

	Lumscan	Skymodel
Sky luminance	Based on measured data	Based on Perez All-weather model
Input parameters	145 measurements of sky luminance	Diffuse horizontal illuminance[a]
***Radiance* pattern type**	brightdata	brightfunc
Luminance at sky point (i.e. pixel	Interpolated from data map	Evaluated from continuous function

Table 5-1. Measured distribution and sky model specification

a. Evaluated from direct normal illuminance and global horizontal illuminance (Section 3.2.5).

derived. Luminance surfaces for the four measured and modelled skies are shown in Figure 5-1 and Figure 5-2. The sky point luminance at the horizon is indicated by the height of the surface's cylindrical 'skirt' (dashed line). The luminance surface is based on an angular fish-eye view of the sky. For this projection, the radial distance (seen here in perspective) from the centre of the surface in the x-y plane is proportional to the zenith angle. For each figure, the same scaling and rotation were applied to the luminance surfaces. Each figure includes the sun position marked on an altitude-azimuth polar plot, and the relative error in the prediction for global horizontal and the four vertical illuminances.

For the modelled sky, the sky point luminance was determined directly from the equation for the sky model. For the measured sky however, the sky point luminance invariably resulted from a bi-linear interpolation of the brightdata datamap. This can be seen in the luminance surfaces for the measured skies: a local high luminance value is shown as a peak rather than as a patch of constant luminance.[3] The four comparisons below are purely illustrative. It is *not* intended that any judgement be drawn on the accuracy of the Perez model from these four cases alone.

3. The scanner measured the average sky luminance across a 'cone' 11˚ wide (Figure 3-6).

Densely Overcast Sky (326_92_11h00)

The measured and modelled luminance patterns are characteristic of those observed for heavily overcast skies - very little azimuthal asymmetry with the zenith luminance greater than that at the horizon, Figure 5-1 (a). The variation in sky luminance with altitude, however, are noticeably different. The luminance scanner recorded a larger luminance ratio from zenith to horizon than that predicted by the model. The lower ratio for the model - indicated by a higher horizon luminance - was the cause of the over-prediction for the vertical illuminances.

Overcast-Intermediate Sky (183_92_10h30)

Cloud cover for this sky was thinner than for the densely overcast sky, Figure 5-1 (b). There was a small component of direct solar radiation, and the maximum sky luminance was at the sun position. Unevenness in the luminance pattern was recorded by the scanner but, of course, does not feature in the distribution generated by the sky model. The vertical illuminance predictions of the Skymodel show the greatest error (over-prediction) for the two surfaces that were illuminated by the sun (i.e. East and South).

Clear-Intermediate Sky 129_92_11h00

The luminance pattern for this sky was dominated by the bright region centred on the sun position, Figure 5-2 (a). However, a fair amount of unevenness resulting from bright patches of cloud was also present in the measurements. These conspicuous cloud patterns were the likely cause of the poor model predictions for the North and West facing vertical illuminances.

Clear Sky 102_92_13h30

The clear sky luminance pattern measured by the scanner does show some slight unevenness, but the form is very similar to that seen in the model sky, Figure 5-2 (b). Both Lumscan and Skymodel exhibit features typical of clear

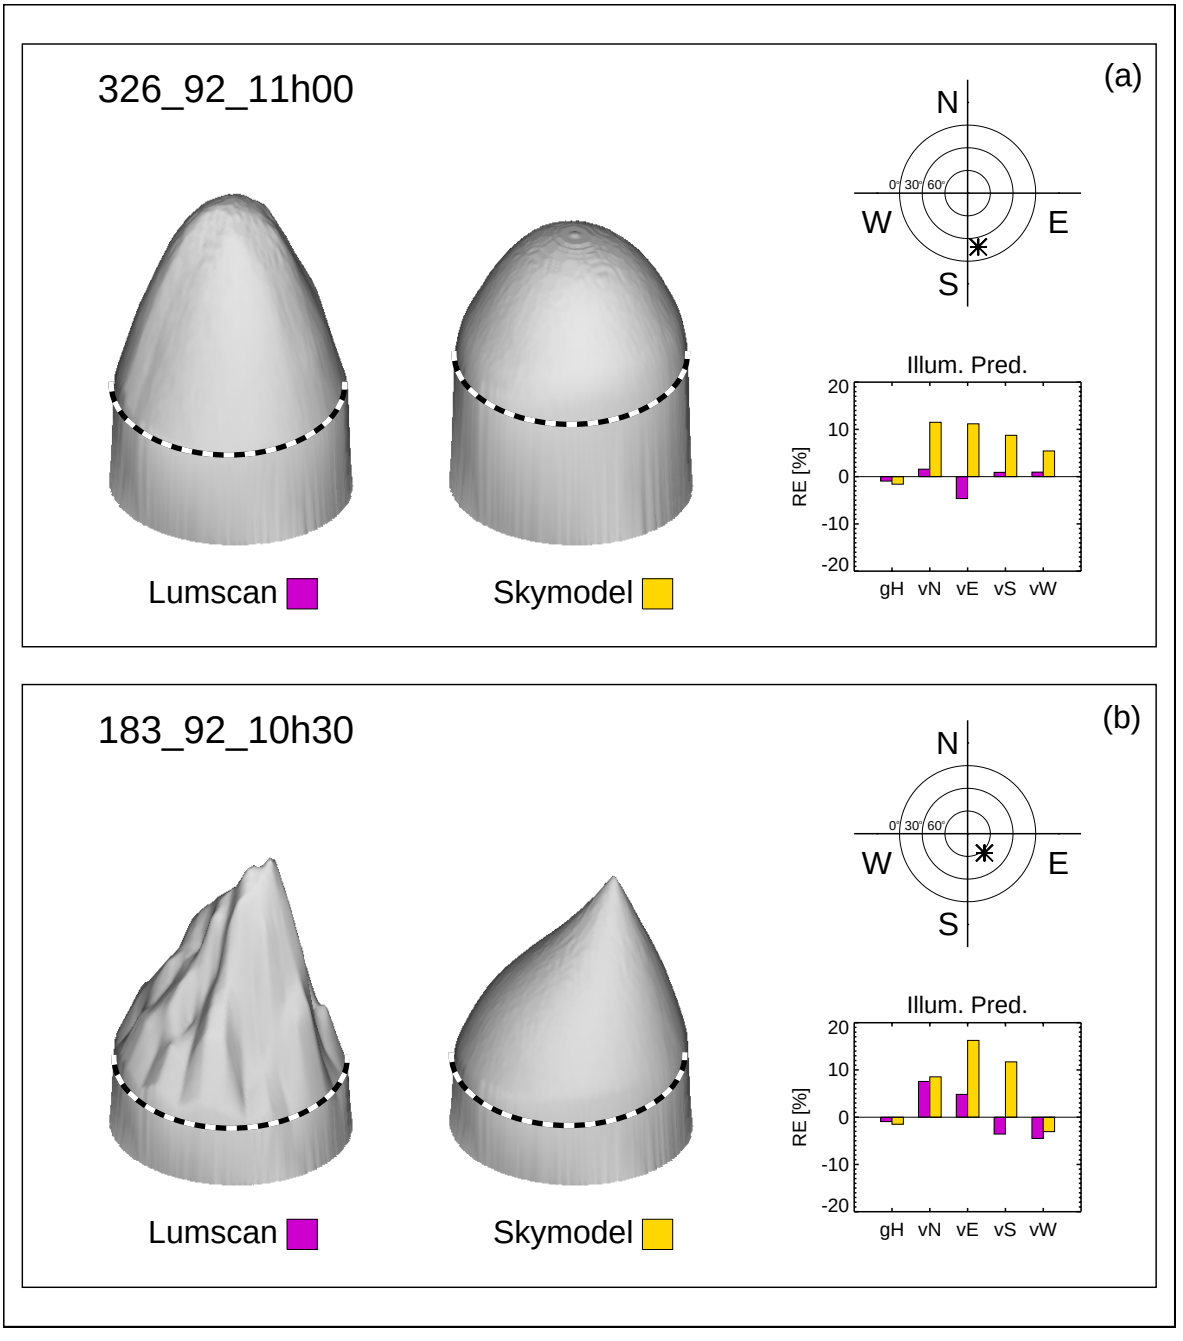

Figure 5-1. Overcast (a) and overcast-intermediate (b) skies

sky conditions - brightening at the horizon and a region of minimum sky brightness about 90˚ away from the sun across the zenith. The model sky however performed poorly for all the vertical orientations. In the absence of conspicuous unevenness in the measured pattern, it is probable that the

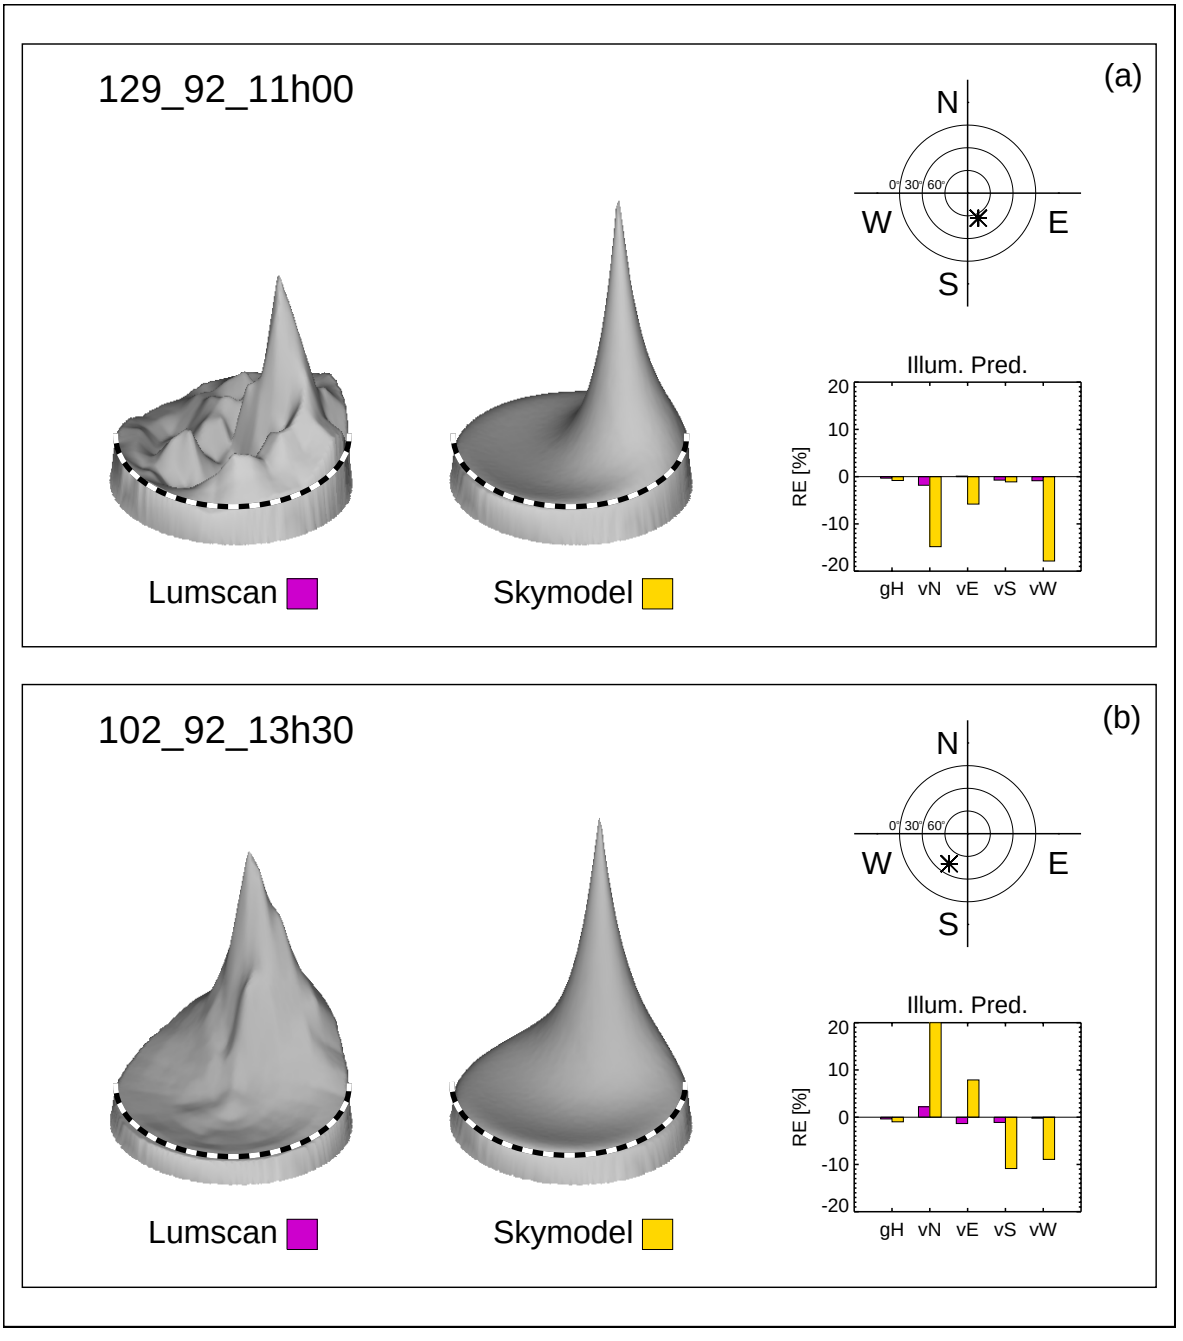

Figure 5-2. Intermediate-clear (a) and clear (b) skies

model was unable to accurately reproduce the underlying luminance pattern for this sky.

5.1.2 Summary

The four measured skies shown above illustrate something of the range in sky luminance patterns that occur in the UK. From heavily overcast, through intermediate to clear sky conditions, the underlying luminance pattern becomes increasingly anisotropic and dominated by the circumsolar region. Qualitatively, the Perez model representations show similarity with the underlying pattern for the measured skies, even though the vertical illuminances predicted using the model sky were often inaccurate ($|RER| > 10\%$). As previously stated, these were illustrative examples only. In the sections that follow, the sky model formulations available with the standard *Radiance* release are evaluated using all of the 754 skies in the validation dataset.

5.2 *Radiance* generator programs for sky models

The *Radiance* standard release includes the sky model generator program **gensky**. This utility program will produce *Radiance* format sky luminance distributions for four sky model types, with the option to create a sun description for the non-overcast sky models. Another sky model generator program is **gendaylit** (discussed above). This program is not part of the standard release, and so it is not updated with each release of *Radiance*. However, like *Radiance*, it is freely available and it gives the user access to the Perez All-weather model which is not supported by **gensky**.

5.2.1 The models supported by gensky

The **gensky** program can produce sky luminance distributions based on:

- the uniform luminance model;
- the CIE overcast sky model;
- the CIE clear sky model; and,
- the Matsuura intermediate sky model.

The absolute luminance of any of these sky luminance patterns is controlled by supplying the program with either the zenith luminance or the diffuse horizontal illuminance.[4] The clear and intermediate sky models allow the option to automatically create a description of the sun. In which case, the solar luminance is either directly supplied to the program or calculated from horizontal direct illuminance. The sun position can either be defined by altitude and azimuth or calculated by **gensky** from the time and geographical coordinates. The uniform luminance model is unrepresentative of any naturally occurring sky conditions and is therefore excluded from any further consideration.[5] The **gensky** input parameter specification for the remaining three models is described below.

The CIE overcast sky model

The overcast model takes the standard CIE form for this type of sky (see Eq 2-3, Section 2.1.2). The generator command is executed as:

% gensky -ang γ **(180 -** α**) -c -B** I_d

Where the altitude and azimuth[6] are, respectively, γ and α, and I_d is the diffuse horizontal irradiance, which is calculated from global horizontal illuminance, E_{gh}, and direct normal illuminance, E_{dn} using:

$$I_d = \frac{E_{gh} - E_{dn}\sin\gamma}{K_R} \qquad \textbf{(5-1)}$$

where K_R is the *Radiance* luminous efficacy factor (179 lm/W). The solar component, however, is not generated automatically. A *Radiance* description for the sun must be specified manually or generated by another program.

---

4. Actually, as with all *Radiance* programs and descriptions, it is the radiance and/or irradiance that must be specified. See Section 2.3.3.

5. The uniform luminance model is useful however for 'Rights to Light' and other specialist applications.

6. Note, the *Radiance* convention has the azimuth as degrees West of South, rather than degrees East of North.

The description for the sun is:

```
void light solar
0
0
3 Rs Gs Bs
solar source sun
0
0
4 x y z 0.5
```

where **Rs**, **Gs** and **Bs** are, respectively, the red, green and blue spectral radiance values for the sun. The source angle for the sun is 0.5˚, and it is centred on the direction vector (**x,y,z**). In this work, the sun was modelled as an achromatic source and the solar radiance, B_{sun} (**= Rs = Gs = Bs**), was evaluated from the direct normal illuminance, E_{dn} using:

$$B_{sun} = \frac{E_{dn}}{\Delta\omega_s K_R} \quad \textbf{(5-2)}$$

where $\Delta\omega_s$ was the solid angle subtended by the (0.5˚) solar disc. The vector components are computed from the sun altitude and azimuth values. The **gensky** output with the added solar component forms the CIE overcast sky with sun description.

The CIE clear sky model

For the CIE clear sky model, the sky and sun description can be both generated using the **gensky** command. The equation for this sky was given in Eq 2-4, Section 2.1.2. The brightness parameters supplied to **gensky** are diffuse horizontal irradiance and solar radiance. The command used is:

% gensky -ang γ **(180 -** α**) +s -B** I_d **-b** B_{sun}

The "Matsuura intermediate sky" model

This formulation is based on a model that was proposed by Matsuura to describe sky conditions that have a higher turbidity than the CIE clear sky model.[7] At the time that the Matsuura model was implemented into the

7. Private communication - G. Ward, LBL.

gensky program, it was not recognised as a CIE standard.[8] Hereafter, it is referred to as the intermediate sky model. This model takes the form:

$$L(\gamma,\theta) = \frac{L_z \cdot a \cdot b}{2.326} \qquad \textbf{(5-3)}$$

where

$$a = [1.35\sin(5.631 - 3.59\xi) + 3.12]\sin(4.396 - 2.6\xi_s) + 6.37 - \xi$$

$$b = \mathrm{EXP}\{-0.563\theta[(2.629 - \xi)(1.562 - \xi_s) + 0.812]\}$$

L_z = zenith luminance

$\xi = \pi/2 - \gamma$

$\xi_s = \pi/2 - \gamma_s$

θ = angle from sun to sky point

γ = sky point altitude

γ_s = sun altitude

The zenith luminance (actually, radiance) is normalised to the diffuse horizontal irradiance.

In comparison to the CIE clear sky model, the intermediate formulation generally predicts lower luminance for the circumsolar region and slightly higher zenith luminances. Additionally, horizon brightening which can be a prominent feature of the clear sky model, is generally absent. This is discernible in Figure 5-3 which shows plots of the sky luminance versus altitude together with false-colour luminance maps for the clear, intermediate and overcast models. The sky point luminance along an arc from $\gamma = 0^\circ$ (due North), across the zenith to $\gamma = 0^\circ$ (due South) is plotted on the graph (dashed line on the false-colour maps). Each sky model was normalised to the same diffuse horizontal illuminance (30,000 lux). The sun

8. At the time that this thesis was near completion, the CIE announced that an official standard for intermediate skies had been agreed. It is not known if this official formulation is the same as the Matsuura model used here.

altitude and azimuth were 45˚ and 180˚ respectively, though the sun itself was not modelled.

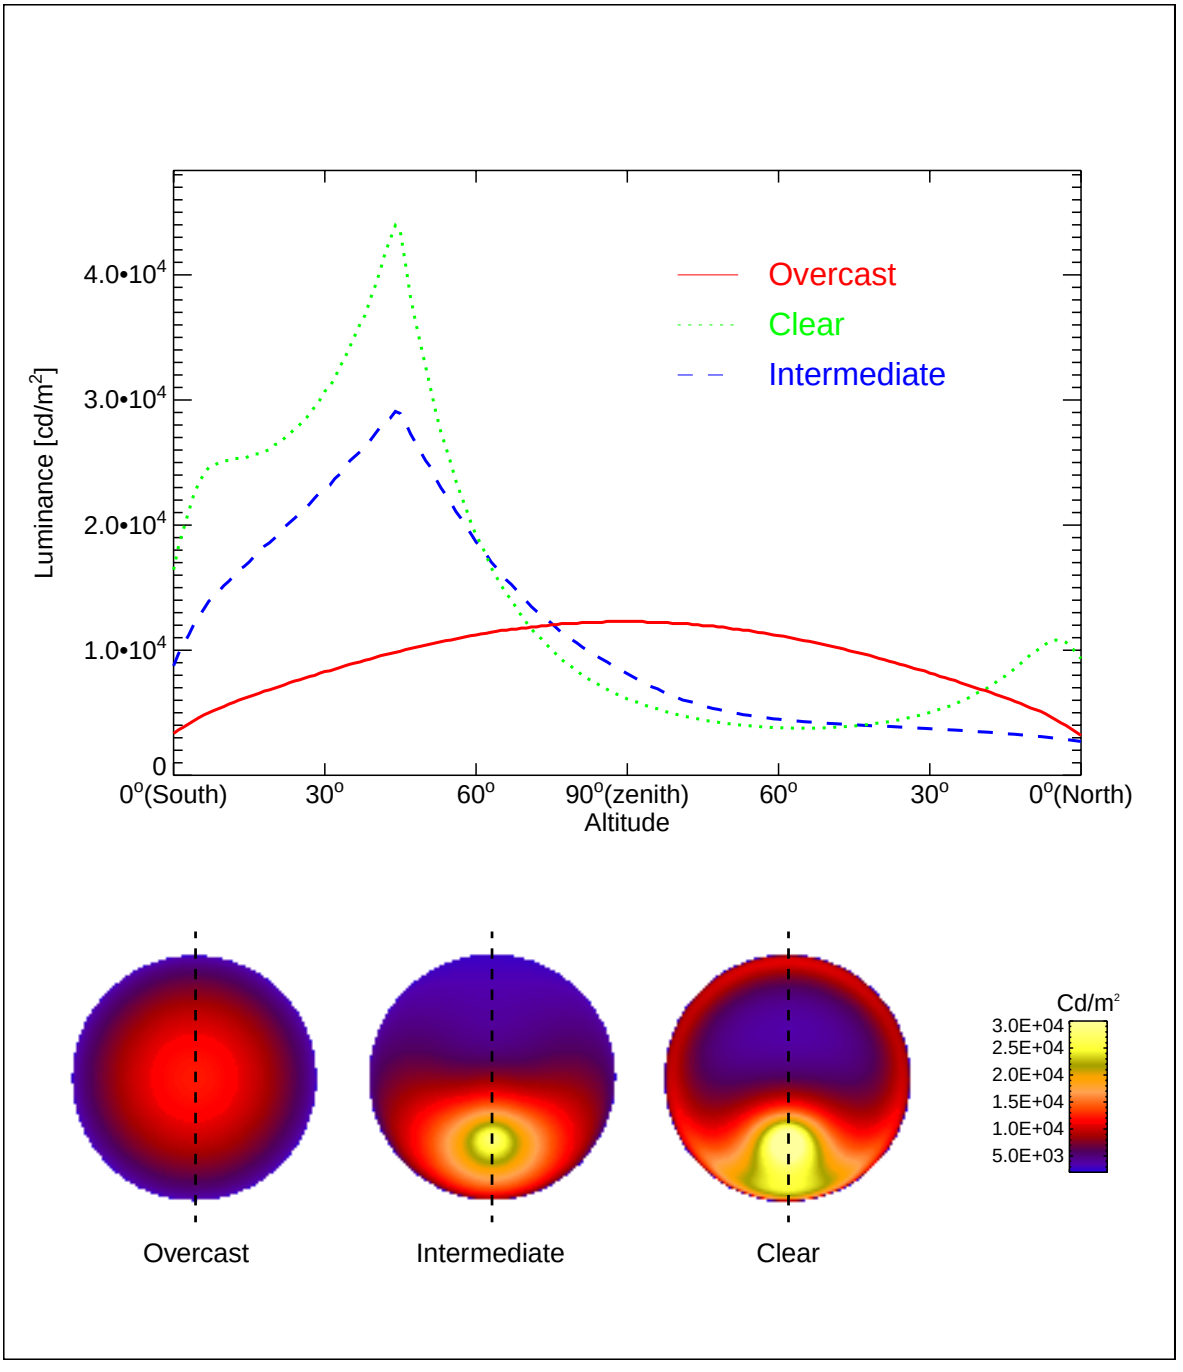

Figure 5-3. Luminance profile and maps for narrow-range sky models

5.2.2 The gendaylit program

The **gendaylit** sky model generator produces a *Radiance* description based on the Perez All-weather model [Perez, 93].[9] This model takes the form:

$$L(\gamma,\theta) = L_z\left[1 + a\mathrm{EXP}\left(\frac{b}{\sin\gamma}\right)\right][1 + c\mathrm{EXP}(\theta d) + e\sin^2\theta] \quad \textbf{(5-4)}$$

Where a, b, c, d and e are adjustable coefficients which depend on solar altitude γ, sky clearness ε and sky brightness Δ. The five coefficients $a, \ldots, e$ are continuous in terms of Δ and γ, and discrete in terms of ε. In other words, the parameters which depend on ε are values held in a look-up table of model coefficients. The model coefficients were derived via least squares fitting of a large data base of ~16,000 sky scans that were recorded at Berkeley (California, USA) between June 1985 and December 1986. Each sky scan recorded 186 measurements of the sky luminance. The **gendaylit** program will evaluate the coefficients $a, \ldots, e$ from diffuse horizontal and direct normal illuminance. In keeping with the conventions of the *Radiance* system, **gendaylit** uses a value of 179 lm/W for luminous efficacy (K_R). The key feature of the Perez model is the potential to generate many sky types, from overcast through to clear, from only the magnitudes of the input parameters. This is in contrast with the different formulations in **gensky** which have to be selected manually by the user.

5.3 Evaluation I: 'Pure' sky models

The illuminance predictions for the validation exercise described in Chapter 3 and Chapter 4 were repeated with the sky luminance patterns now provided by sky models. Illuminance predictions for the skies in the validation dataset were obtained for the following sky models (all with sun):

- the CIE standard overcast sky model;
- the Matsuura Intermediate sky model;.

9. The **gendaylit** program was written by Jean-Jaques Delauney, FhG-ISE, Freiburg, Germany.

- the CIE clear sky model; and,
- the Perez All-weather sky model.

The first three of the above are 'narrow-range' models. These were formulated to reproduce luminance patterns for specific sky conditions. As their names suggest, these sky conditions are: densely overcast with no sun; hazy, thin cloud with sun (intermediate) and clear, sunny sky conditions without clouds. Only the Perez All-weather model was designed to generate luminance patterns for a wide range of sky conditions. These four are, nevertheless, called here 'pure' sky models because they are distinct formulations. For brevity, the models are referred to occasionally simply as overcast, intermediate, clear and Perez.

Any one of the narrow-range models will be incapable of reproducing the full range of sky conditions in the validation dataset. That notwithstanding, the predictions for the narrow-range models serve two purposes. The first is illustrative: just how well do the narrow-range models perform when applied routinely to all the skies of the validation dataset? The second, and more significant purpose, is to generate the basic data from which the illuminance *effect* of a sky model blend may be synthesised (this is examined in Section 5.4). The absolute performance of the narrow-range models should *not* be inferred from the comparison that follows.

5.3.1 Automation of the simulations

For each of the sky models in turn, external and internal illuminance predictions were obtained for each of the 754 skies in the validation dataset. Each sequence of simulations was initiated from an 'executive' IDL program similar to the one described in Figure 3-31. The input parameters for the sky model programs were derived from measurements of the global horizontal illuminance, the direct normal illuminance and the sun position.

The 'executive' program spawned shell scripts that contained commands for the following operations:

1. Generate *Radiance* format skies using **gensky** or **gendaylit** (input parameters were read from temporary files created by the 'executive' program).
2. Create *Radiance* octree for the sun and sky description.
3. Execute **rtrace** to calculate the external illuminances and write to temporary file.
4. Add sun and sky description to the (frozen) octree for the BRE office scene.
5. Execute **rtrace** to calculate the internal illuminances at the six photocell locations and write to temporary file.

Altogether, there were 3,016 (= 754 x 4) executions of **rtrace** for the external illuminances and the same number again for the internal illuminances. The basecase set of ambient parameters was used for all the simulations for the internal illuminance (Section 3.3.2). The input parameters to the generator programs and the measured quantities from which they were derived are listed in Table 5-2.

Model type		CIE overcast	CIE clear	Intermediate	Perez 'All-Weather'
Generator program		**gensky**	**gensky**	**gensky**	**gendaylit**
Sky component	Input	I_{dh}	I_{dh}	I_{dh}	E_{dh}
	Derived from	E_{gh}, E_{dn}, γ, K_R	E_{gh}, E_{dn}, γ, K_R	E_{gh}, E_{dn}, γ, K_R	E_{gh}, E_{dn}, γ
Sun component	Input	B_{sun}[a]	B_{sun}	B_{sun}	E_{dn}
	Derived from	E_{dn}, K_R	E_{dn}, K_R	E_{dn}, K_R	-

Table 5-2. Sky generator program parameters

a. Sun description added by routine in 'executive' program.

5.3.2 External illuminance predictions

The first stage of the comparison is a presentation of the relative error in the illuminance predictions for the four vertical illuminances. The relative errors for each orientation are shown as frequency histograms. Each histogram is annotated with the overall MBE and RMSE for the sample. The results were as follows.

CIE Overcast sky model

The distributions showed a low (< 10%) overall bias in predictions for vertical North (VN) and vertical East (vE), Figure 5-4. Illuminances for the vertical South (vS) and vertical West (vW) orientations, however, tended to be under-predicted. The MBE for vS and vW was -15% and -19% respectively. The overall accuracy indicated by the RMSEs was in the range 22% to 28%.

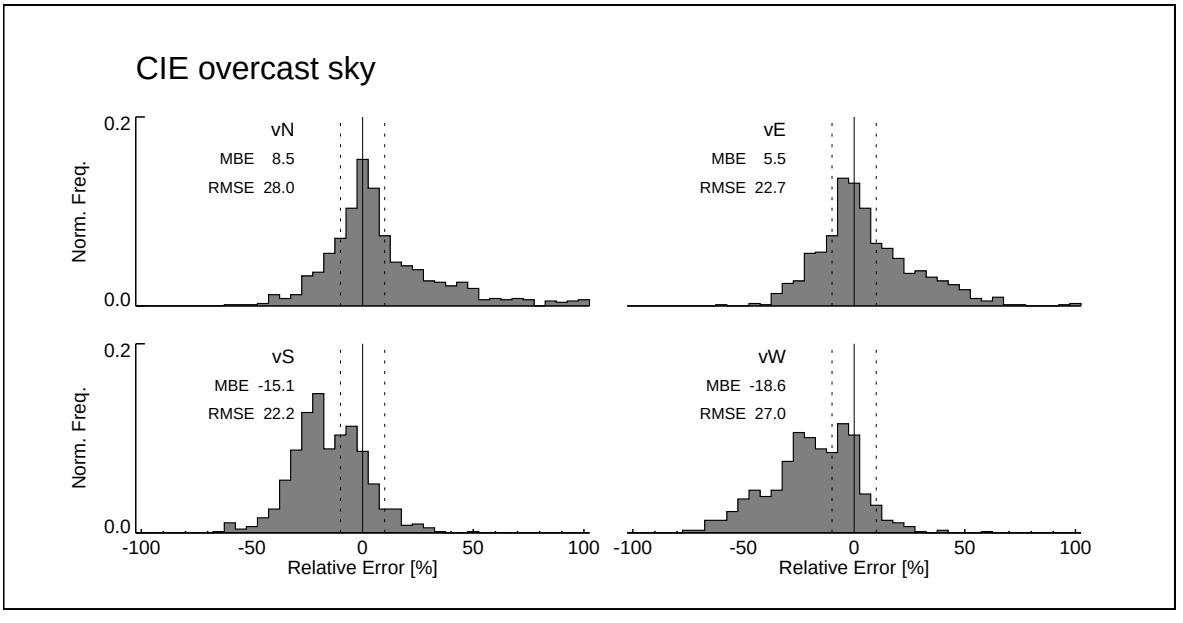

Figure 5-4. CIE overcast sky model

CIE Clear sky model

This model performed fairly poorly for all but the vE orientations, Figure 5-5. The general tendency was to overpredict vertical illuminances,

particularly for the vS and vW orientations where the MBEs were 52% and 58% respectively. The RMSEs were correspondingly very large; in the range 30% (vE) to 94% (vW).

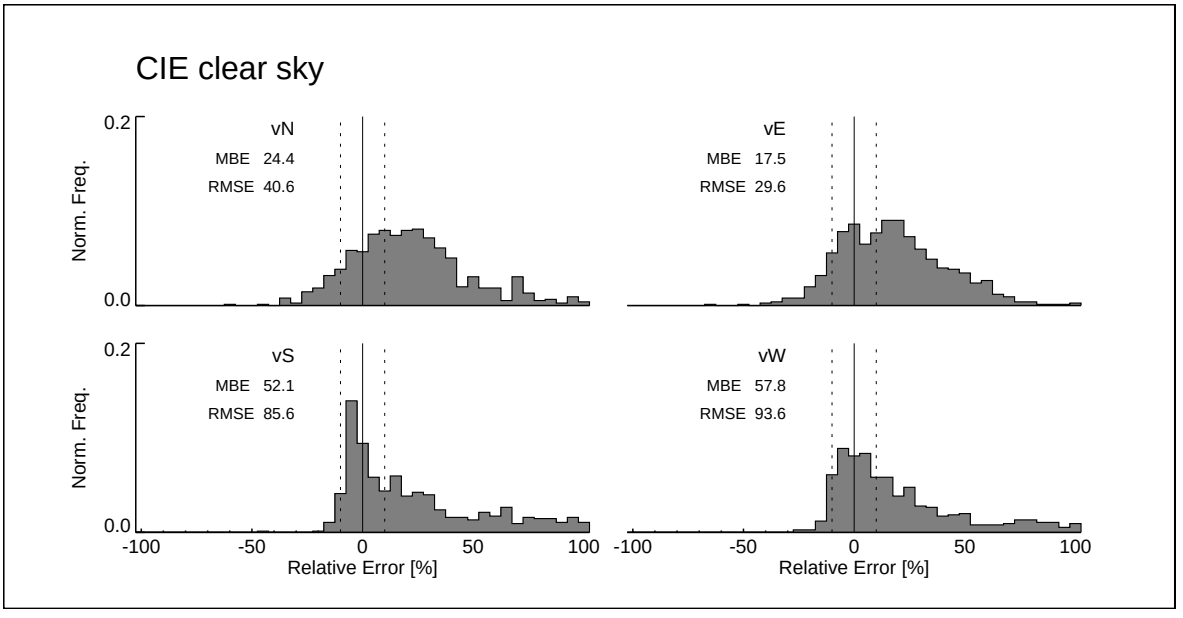

Figure 5-5. CIE clear sky model

Matsuura Intermediate sky model

The performance of the intermediate sky is, in character, similar to that for the clear sky, Figure 5-6. However the bias in the predictions for the intermediate sky was always lower than for the clear sky, and the accuracy greater. Note also that the MBEs for vN and vE were marginally negative (-7% and -13%), whereas with the clear sky model they were markedly positive (24% and 18%).

Perez All-weather model

As is immediately apparent from the distributions, the Perez model performed reasonably well for all vertical orientations, Figure 5-7. For only one orientation (vN) was the MBE greater than 10%, and then only marginally (11%). The MBE was positive for all orientations. The RMSE values were in the range 17% to 42%. The high RMSEs for vS and vW were

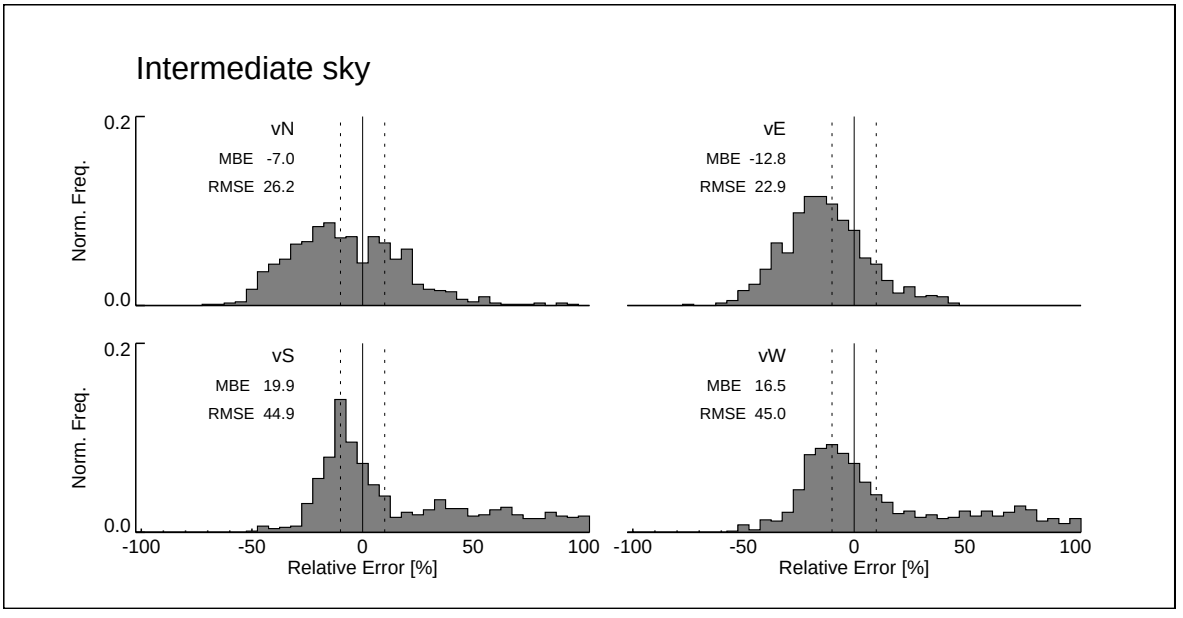

Figure 5-6. Intermediate sky model

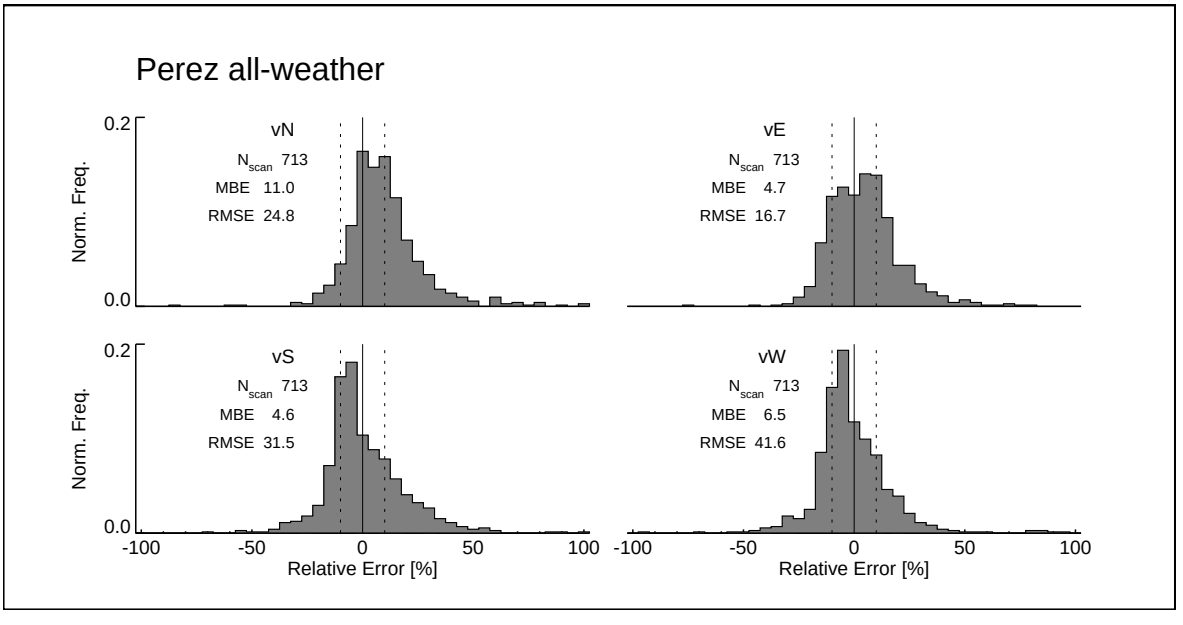

Figure 5-7. Perez all-weather model

caused by a small number of outliers where the RER was greater than 100%. It should be noted that there were 41 skies for which the Perez model description could either not be generated (outside parameter range) or

which produced negative vertical illuminances. These were eliminated from the analysis for this model leaving 713 skies. The negative vertical illuminances resulted from distortions in the sky luminance distribution that can occur unexpectedly for certain combinations of input parameters. These parameter combinations were present in the data collected by the BRE but they were not encountered in the Berkeley data that were used to derive the model.[10] This effect was noted by Littlefair and an adjustment to the model to prevent this distortion was advised by Perez [see Littlefair 94]. A routine examination of the **gendaylit** code showed this fix to be present. This suggests that either the fix (or some other part of the model) was incorrectly coded, or that there are still some parameter combinations that result in distortion, regardless of the fix. The presence of a distortion was taken to be a *negative* value for any of the predicted vertical illuminances. The actual luminance distribution for the sky was not examined. So the possibility remains that some of the other skies may yet have exhibited some distortion. If a distortion did result, but was not sufficient to give a negative vertical illuminance, it would simply be manifest as an under-prediction, and so would *not* be identified as an erroneous sample. In which case, the error metrics for this model may, to some degree, be contaminated.

5.3.3 Analysis of RERs for vertical illuminance predictions

The sky model MBEs and RMSEs for each of the orientations are shown as a histogram chart in Figure 5-8. The results for the measured skies (Lumscan) are included for comparison (taken from Figure 4-1). In terms of bias for the vertical orientations, Perez performed best of the models for vE, vS and vW. Only for vN did the overcast model produce a lower bias than the Perez, and then only marginally. A striking feature of the histogram plots is the similarity in rank order of vN with vE, and of vS with vW. This is apparent for both the MBE and the RMSE. The poor performance of the clear sky model is most apparent for the vS and vW orientations.

10. Private communication - P. Littlefair, BRE.

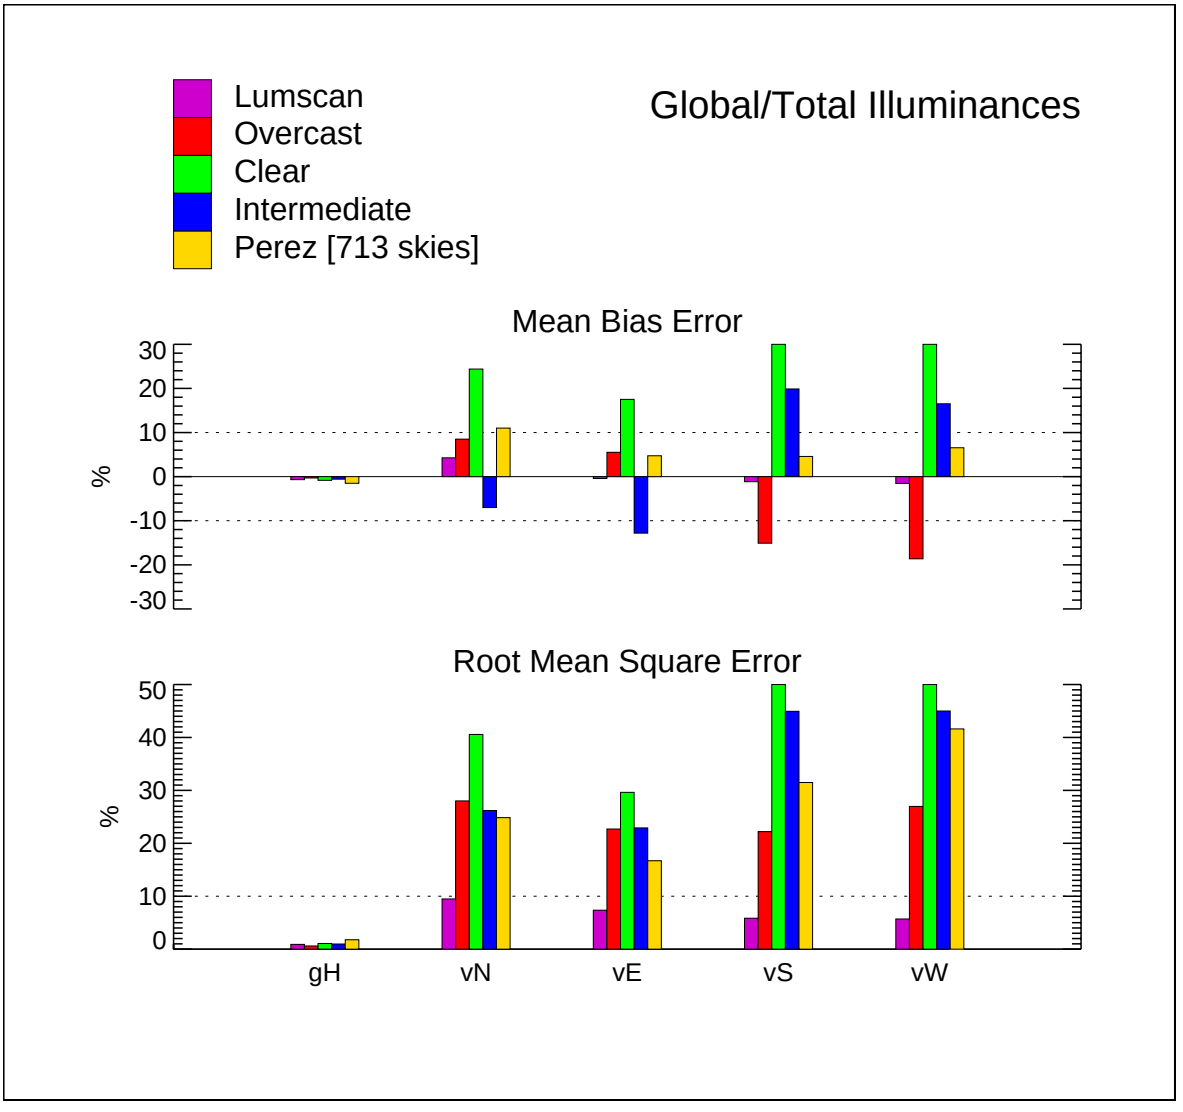

Figure 5-8. Sky models MBE and RMSE

In terms of producing the lowest overall bias for the average (absolute) MBE, the Perez model performed the best, Table 5-3. Next was the overcast, closely followed by the intermediate model. The performance of the clear sky model was markedly worse than the other three. The overall performance of the narrow-range models can be largely attributed to two factors: the characteristic luminance patterns of the sky models; and, the composition of the validation dataset.

	Overcast	Clear	Intermediate	Perez
Avg(|MBE|)	11.9%	38.0%	14.1%	6.7%
Avg(RMSE)	25.0%	62.4%	34.8%	28.7%

Table 5-3. Vertical illuminance: average |MBE| and RMSE

RERs related to sky model luminance patterns

The luminance patterns for the overcast, intermediate and clear sky models posses characteristic features. Some aspects of the RER distributions for vertical illuminance may be explained in terms of these distinctive features. The plot of model sky point luminance versus altitude (Figure 5-3) is used to illustrate how systematic biases in the prediction of vertical illuminance might arise. Recall that for Figure 5-3, each of the skies were normalised to the same diffuse horizontal illuminance, and the solar altitude and azimuth were 45˚ and 180˚ (due South). Consider the following possibilities for the prediction of the vertical South illuminance (E_S).

1. That the actual sky conditions were consistent with the overcast sky representation. In this event it would be expected that the overcast sky model give a reasonably accurate prediction for E_S. The clear sky model would however, for the same conditions, generate a luminance distribution with a relatively intense peak at the solar position. This is despite the fact that the actual solar luminance (derived from measurement for a heavily overcast sky) was likely to be *less* than the clear sky model luminance at the solar position. It is apparent then that, for the scenario described above, the clear sky representation would always overpredict E_S when actual overcast conditions prevailed. Similarly, the intermediate sky would also overpredict E_S, but to a lesser degree.

2. That the actual sky conditions were consistent with the clear (or intermediate) sky model (sun position as above). Now it is the overcast sky model which gives a poor representation. Since one would expect

some sky brightening about the solar position for actually occurring clear (or intermediate) conditions, the overcast model would then systematically under-predict E_S for this scenario.

There are limits to generalisations elucidated from a relatively small number of specific cases. For example, a bright circumsolar region (when present in the validation dataset) was not always at altitude 45˚ and due South. However, for this dataset, the sun position was 'visible' from the vertical South plane for most of the skies (Figure 3-8). Also, the illuminance is, of course, the cosine (zenith angle) weighted integral of luminance over a hemisphere, not just an arc.[11] With these qualifications kept in mind, the two examples above nevertheless offer some insight. And indeed, the relative sizes of the vertical South MBEs for the overcast and clear skies are consistent with the mechanisms outlined above. For the overcast model, the MBE was -15.1%. Whereas, the clear sky model produced a significantly larger (absolute) MBE of +52.1%. The relative difference of these biases may, in part, be explained by considering also the contribution of direct solar radiation to vertical South illuminance. The clear sky model was a poor representation when the actual conditions were overcast, i.e. when solar radiation was negligible and the sky was the sole contributor to illuminance. In contrast, the overcast model was a poor representation when the actual conditions were clear/intermediate, i.e. when solar radiation was likely to be significant. So although the overcast sky produced less vertical South illuminance than would a clear sky, the solar contribution to E_S was the same for both - the same sun description was used for all the models. Thus, the systematic bias that resulted from modelling actual clear skies with overcast, was less than that from modelling actual overcast conditions with clear skies.

11. Actually, a half-hemisphere for vertical illuminances, where the photocells are shaded from ground luminance.

RERs related to composition of the validation dataset

The 754 skies in the validation dataset contain a large number of overcast skies. Approximately 60% of the skies have a clearness index that falls in to clearness bin number 1, Figure 3-9. It is not surprising therefore that, overall, the overcast sky model performed moderately well. The bias in the distribution of the sun azimuth angle is another factor to consider. Recall that operational factors limited the collection of room illuminance data to after 10:30h (see Section 3.1.4). This means that, for the sample as a whole, only a relatively small number of vE measurements included a component for direct solar radiation. Whereas, the vS and vW planes were often illuminated by direct solar radiation (when present). This bias was the likely reason for the similarity between the distributions for vN and vE, and also between vS and vW. The vN and vE photocells rarely recorded a direct solar contribution, whereas the vS and vW photocells often did. If the distribution in solar azimuth was symmetrical about the N-S line (and without a bias in altitude), the vE and vW RER distributions would then be broadly similar. This would be the case for any of the sky models.

5.3.4 Internal illuminance predictions

The overall sky model MBEs and RMSEs for the prediction of internal illuminance at the six photocell locations are presented as histogram plots in Figure 5-9. For comparison, the results for the measured skies (Lumscan) are shown also (taken from Figure 4-10). These are the results for the entire sample.[12] As one might expect for an office with near-to South facing glazing, the rank order in MBE for the sky models at each photocell is very similar to that for the vertical South illuminance (Figure 5-8). Overall the Perez model gave the lowest bias, closely followed by the overcast model. The clear and intermediate models both performed poorly, though the clear sky model was the worst by a significant margin. The overcast model had a tendency to underpredict internal illuminance. This is in contrast to the

12. That is, all photocell-sky combinations for all 754 skies (713 for the Perez model).

other models which all exhibited a tendency to overpredict internal illuminance, to a greater or lesser degree. The RMSEs for the overcast sky model were, for most photocell locations, markedly lower than for the other three models. The MBE and RMSE for the Lumscan predictions was

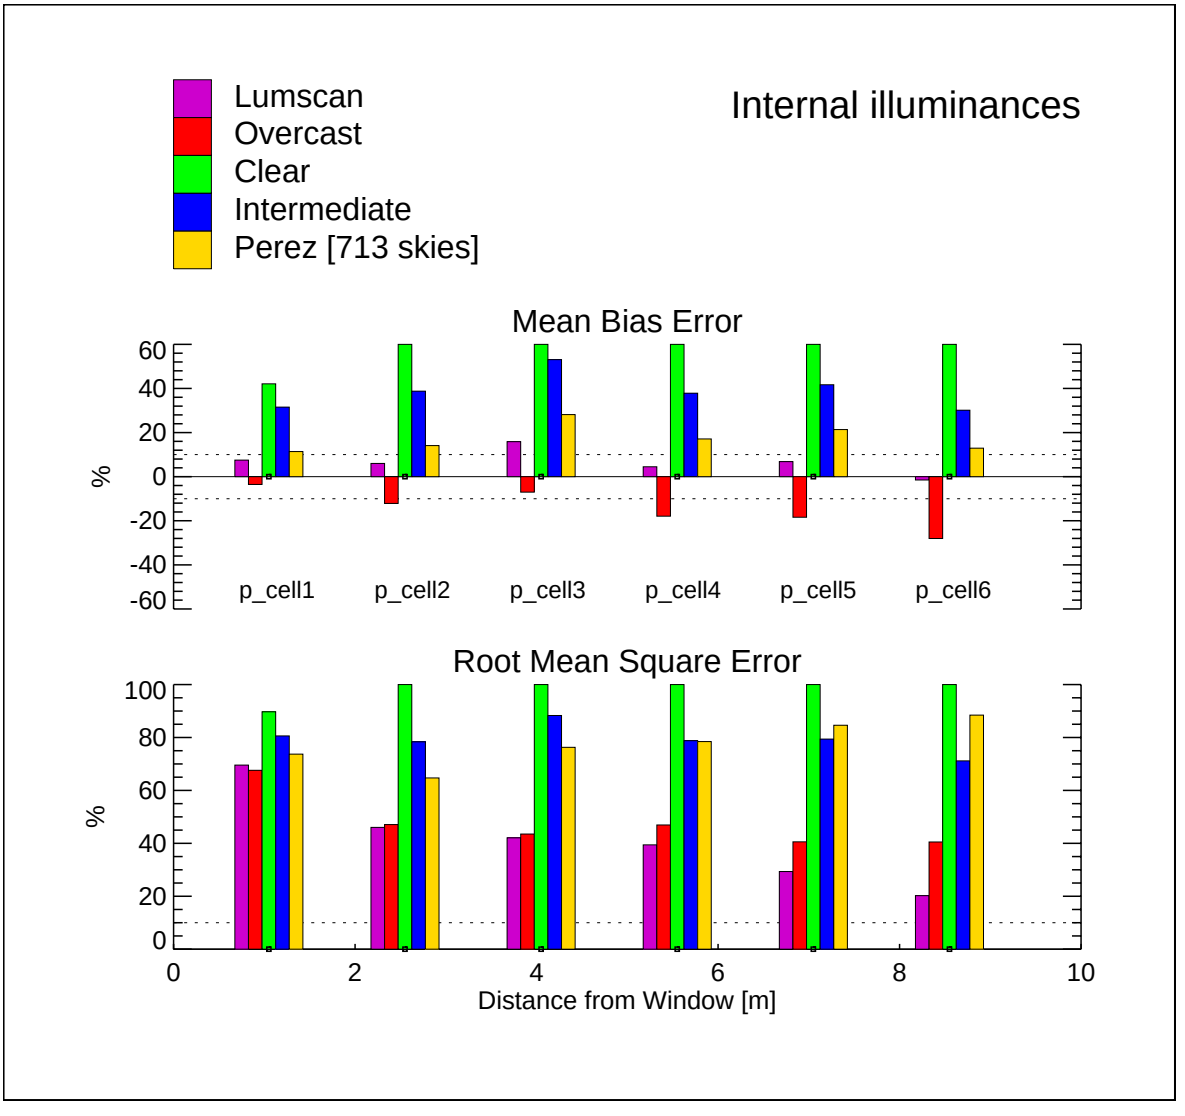

Figure 5-9. Internal illuminance: sky models MBE and RMSE

generally lower than for any of the sky models. Though it is likely that the MBE and RMSE for Lumscan and all the sky models are dominated by

source visibility related errors (Section 3.2.7). The average (absolute) MBE and RMSE across the six photocells are given in Table 5-4.

	Overcast	Clear	Intermediate	Perez
Avg(|MBE|)	18.8	92.2	34.9	12.7
Avg(RMSE)	32.8	155.8	69.1	62.5

Table 5-4. Internal illuminance: average |MBE| and RMSE

Model sky performance based on a percentile analysis

It was demonstrated in Chapter 4 that the BRE-IDMP data contained many occurrences of potentially unreliable photocell-sky combinations (Section 4.5). These were referred to as source visibility related errors (SVRE) and they were identified as resulting from one or more of four error types (Table 3-9). As things stand, there is little scope to correct for any one of these errors with any certainty. Furthermore, because each can have a similar effect on the predictions, it was not possible to dis-aggregate the effect of one error type from the rest. Recall that one of these SVREs (Type B, Table 3-9) resulted from the uncertainty of the sky luminance distribution about the solar position; the sky scanner could not resolve the circumsolar luminance gradients. Sky models for non-overcast skies on the other hand, are designed to reproduce the circumsolar sky luminance.

The partition of the validation dataset, based on visibility of a 6˚ circumsolar exclusion region (CER), eliminated most (possibly all) of the SVRE and elicited a marked improvement in the assessment of the intrinsic accuracy of the illuminance predictions. To do the same for the sky model predictions would remove from the validation dataset those cases where the CER was visible from the photocell. This would have the unfortunate effect of eliminating from the comparison those instances where the illuminance predictions were most sensitive to luminance gradients about the circumsolar region. And indeed those *might* be the instances where the sky model accurately reproduced conditions that the scanner could not

measure. To assess the sky model performance without *a priori* partitioning of the validation data requires a comparison test that is based on something other than MBE and RMSE, because these can be heavily biased by a few outliers. This was achieved using a percentiles-based comparison.

For this, the percentage of the predictions that had a relative error within the range ±*R* were plotted as a function of *R*. These plots are referred to here as percentile-RER plots. The results for the four sky models are shown in Figure 5-10. For comparison, the results for the measured skies are shown

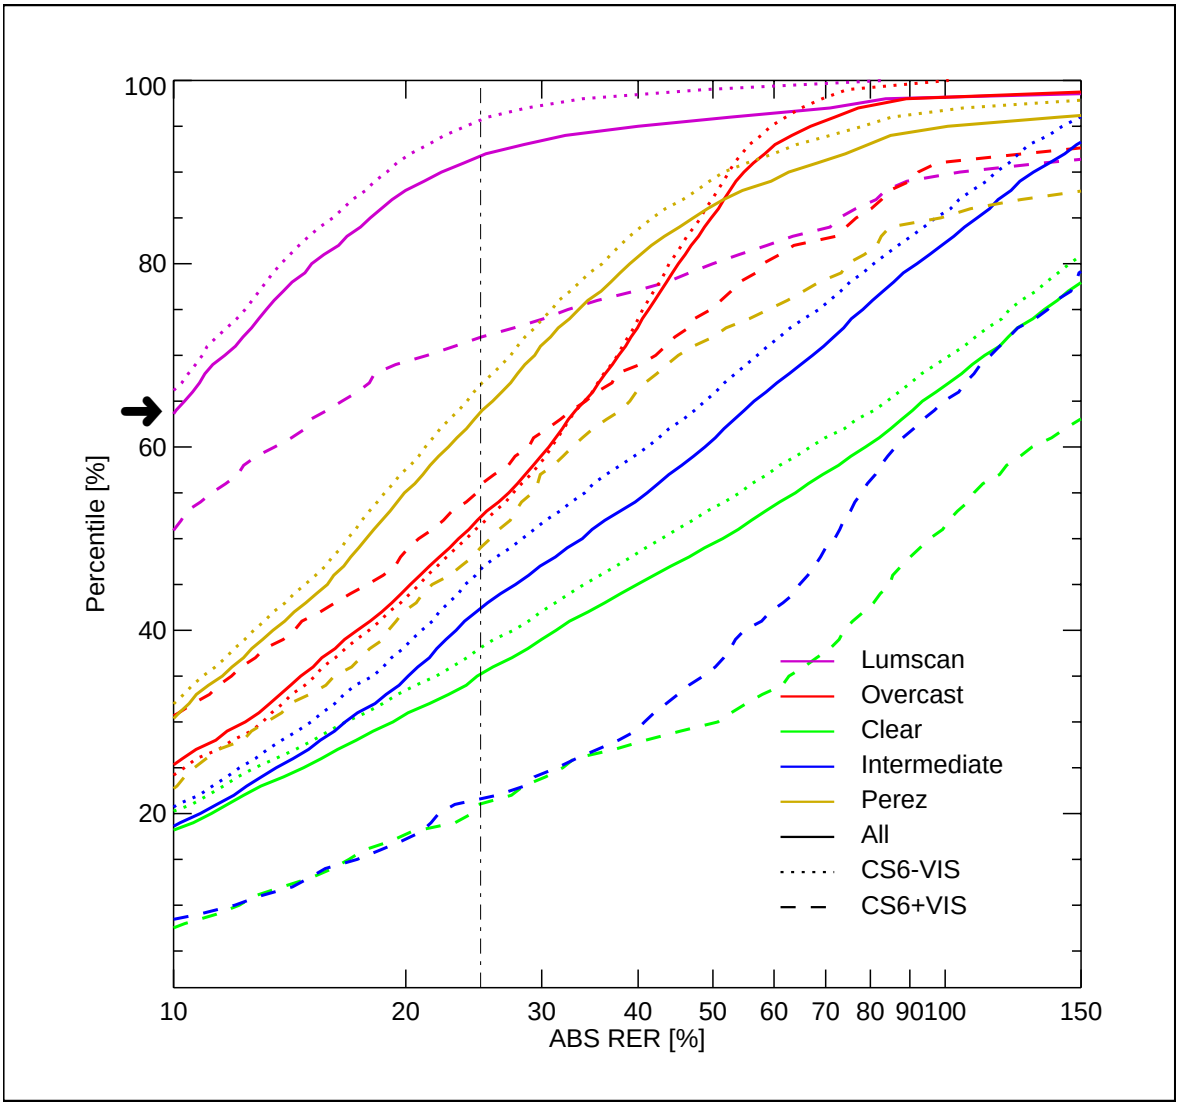

Figure 5-10. Sky model percentiles-RER comparison

also (Lumscan). In addition to showing the percentile lines for all (4,524) photocell-sky combinations for each model, percentiles for each of the partitioned data sets are given also. This figure is to be read as follows: for all photocell-sky combinations (solid lines), ~64% of the illuminance predictions using the measured skies (Lumscan, magenta line) were within ±10% of the measured value (marked by ➜ on graph). Using the same skies, of those instances where the 6˚ CER was *not* visible from the photocell (CS6 -VIS dotted line), ~66% of those predictions were within ±10% of the measured value. Similarly, for *just* those instances where the 6˚ CER *was* visible (CS6 +VIS dashed line), ~51% of the predictions were within ±10% of the measured value. The rank order, best first, for all cases (solid lines) at the ±10% RER line (y-axis) is Lumscan, Perez, Overcast, Intermediate and Clear. For all but the Overcast, the CS6-VIS set had a marginally greater percentage within the RER range ±10%, and the CS6+VIS had a smaller percentage within the RER range ±10%. The rank order of the three sets (All, CS6-VIS and CS6+VIS) for each of the sky models is maintained across the range of (absolute) RER, except for the Overcast.

Confounding expectation, the rank order, best first, for the Overcast model at the ±10% line is CS6+VIS, then All, then CS6-VIS. The cause for this is revealed in Figure 5-11. These plots are similar to those given in Figure 4-

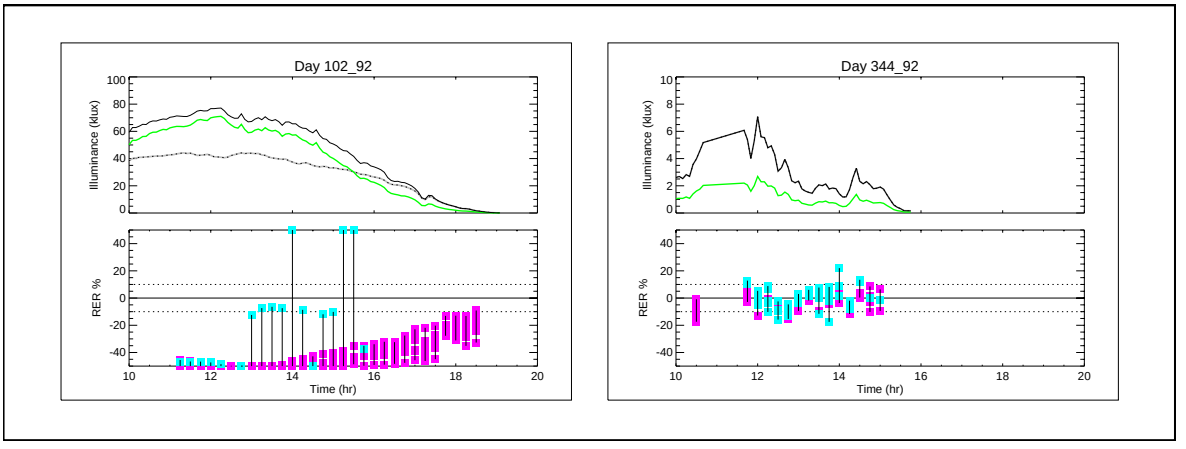

Figure 5-11. RER time-series for overcast (with sun) model - CS6-VIS(■) and CS6+VIS(■)

34 to Figure 4-37, only here the relative error for the predictions using the Overcast model are shown. The instances where a photocell did not 'see' the circumsolar region (CS6-VIS) are shaded magenta (■), and where it did (CS6+VIS) are shaded cyan (■). For days such as 344_92, the Overcast model performed well regardless of the visibility or otherwise of the CER, because of course there was negligible sun. For clear sky days however (102_92), the CS6-VIS and CS6+VIS sets can have very different RER characteristics. For those cases where the CER was not visible (■), there was consistent under-prediction. This is expected because the maximum luminance for the Overcast sky is at the zenith, rather than at the circumsolar position which was mostly in the South for this day (see Figure 5-3). This was so even when the circumsolar region was not directly visible. Where the CER was visible (■), there was significant under-prediction, over-prediction *and* a run of very accurate values. The accurate values result from those occasions when a photocell was, correctly predicted, to be directly illuminated by the sun. When this happened, the illumination from the sky had a relatively marginal effect - though note how these (mostly) accurate predictions all share a small negative bias. The under-prediction may result when a photocell predicted to be in shade, was actually in sun. And vice-versa for over-prediction. It is generally the case that the *extremes* in the RER are greater for the CS6+VIS set than those for the CS6-VIS set. This explains the peculiar behaviour of the three percentile lines for the Overcast model.

It may well have been the case that, for a number of skies, the circumsolar region was better reproduced by one or more of the non-overcast sky models than by the measured-interpolated distributions described in Chapter 3. If so, the analysis described above was insufficient to reveal this. It may have been worthwhile to repeat the percentiles analysis using only clear skies had a much larger sample of non-overcast skies been available.

5.4 Evaluation II: Sky model blends

Three of the four sky models evaluated above were devised to be applicable to a limited range of possible sky conditions, that is, overcast, intermediate and clear. It is only the Perez model that was specifically formulated to represent all (or at least the majority) of naturally occurring skies. It is not surprising therefore that the one model designed with 'wide-range' applicability performed better than the other three 'narrow-range' models. However, wider application of the 'narrow-range' models might be achieved by blending the luminance patterns of two or more of them in response to meteorological conditions. This would produce a sky luminance pattern which is a blend of the component model patterns. The following sections describes the formulation and testing of two sky model blends.

5.4.1 Model sky blends: *ex post facto* synthesis

The common practice for blending skies is to combine an overcast luminance pattern with one or more non-overcast patterns according to some rule [Littlefair 94]. A number of ways of achieving this are currently in use. For the investigation described here, just two simple sky model blends are evaluated in terms of their ability to reproduce sky conditions for the purpose of internal illuminance prediction. The 754 skies in the BRE-IDMP dataset were sufficient for the validation work described in Chapter 3 and Chapter 4. This number of skies however is too small to warrant an exhaustive examination of complex sky models blends, since certain sky conditions were represented by a relatively small number of measurements (see Figure 3-9). Accordingly, the sky model blends used here are each a composite of an overcast luminance pattern and just one non-overcast luminance pattern.

An optimum sky model blending function for the validation dataset was determined for each of the sky model blends. The *effect* of a composite sky was synthesised by combining the existing illuminance predictions for the narrow-range sky models. Proceeding in this way, any arbitrary blending of

the sky models can be investigated without calling for additional illuminance predictions.

The illuminance predictions for the overcast, intermediate and clear models were re-used to synthesise the illuminance effect of an overcast-intermediate blend and an overcast-clear blend. The theoretical basis for this is described below using the clear overcast blend as an example. In terms of sky luminance, the resultant sky point luminance for a clear-overcast blend L_{co} would be:

$$L_{co} = f_{cl}L_{cl} + f_{ov}L_{ov} \quad \textbf{(5-5)}$$

Where L_{cl} and L_{ov} are, respectively, the sky point luminances for the clear and overcast models. The weighting given to the components are f_{cl} for the clear sky, and f_{ov} for the overcast. The applied weighting was constant across the sky vault. Therefore, the resultant diffuse horizontal illuminance from a sky will vary in proportion to the weighting factor, e.g. for clear sky:

$$\int_{2\pi} f_{cl}L_{cl}(\gamma,\alpha)\sin\gamma d\Omega = f_{cl}\int_{2\pi} L_{cl}(\gamma,\alpha)\sin\gamma d\Omega = f_{cl}E_{cl}^{dh} \quad \textbf{(5-6)}$$

where $f_{cl}E_{cl}^{dh}$ is the diffuse horizontal illuminance for the clear sky distribution $L_{cl}(\gamma,\alpha)$ weighted by f_{cl}. Recall that the luminance for all the models skies was normalised to the diffuse horizontal illuminance, so that $E^{dh} = E_{cl}^{dh} = E_{ov}^{dh}$. Therefore, setting $f_{ov} = 1 - f_{cl}$, normalises the composite sky to diffuse horizontal illuminance also. Thus,

$$L_{co} = f_{cl}L_{cl} + (1 - f_{cl})L_{ov} \quad \textbf{(5-7)}$$

and

$$E^{dh} = f_{cl}E^{dh} + (1 - f_{cl})E^{dh} \quad \textbf{(5-8)}$$

The illuminance predicted using the pure sky models was the total illuminance; the component illuminances from the sun and sky were not calculated separately.[13] For all the sky models however, the illuminance from the sun was the same; they all shared the same description for the

sun. Therefore, blending the clear and overcast total illuminances is equivalent to blending the sky components (Eq 5-8), and then adding the sun component E^{sun}:

$$f_{cl}(E^{sun} + E^{dh}) + (1 - f_{cl})(E^{sun} + E^{dh}) = E^{sun} + f_{cl}E^{dh} + (1 - f_{cl})E^{dh} \quad \textbf{(5-9)}$$

The luminance patterns for the three narrow-range models and the two sky blends are given in Figure 5-12. This figure shows luminance surfaces for the intermediate, the overcast and the clear sky models, all normalised to the same horizontal diffuse illuminance. Below, are two “half-and-half” blends for an intermediate-overcast blend and a clear-overcast blend, i.e. $f_{in} = f_{ov} = f_{cl} = 0.5$. Both of the sky model blends would produce the same diffuse horizontal illuminance as the ‘pure’ sky luminance patterns. The same scaling and rotation were used to display each luminance surface.

Although the blended luminance surface shows what a particular composite pattern would look like, in the analysis that follows, they were never actually generated. To recap, their illuminance effect was synthesised from the existing illuminance predictions for the narrow range models.

5.4.2 The blending functions

The weighting factor for the non-overcast sky f_{cl} (or f_{in}) should depend in some way on the clearness of the sky.[14] Evidently, the more overcast the actual sky conditions the smaller f_{cl} should be. For fully overcast skies, f_{cl} should equal zero. Conversely, for progressively clearer skies, f_{cl} should tend to unity. The factor f_{cl} therefore should be some function $\mathbf{f}(\varepsilon)$ of the sky clearness index ε (Eq 3-1) over a mixing range bounded by lower and upper values for ε. Within the mixing range, the effect of a linear and a power-law blending function was examined. The illuminances synthesised

13. These, or any other illuminance components, could of course be calculated separately if desired (see Section 4.4).

14. The clear-overcast blend is used to illustrate the linear and power-law blends. The blending function equations for the intermediate-overcast blend are essentially the same, and for brevity, they are not reproduced.

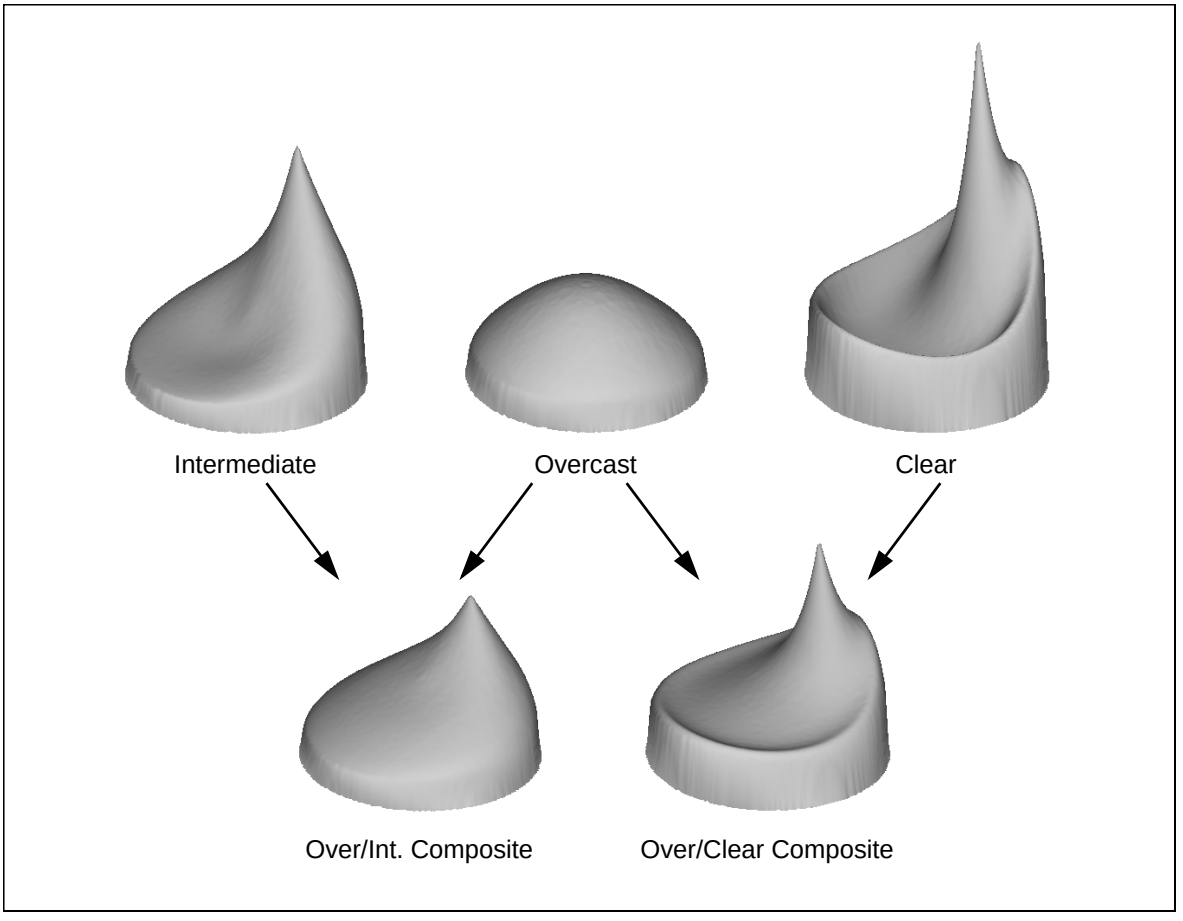

Figure 5-12. Example composite skies

from the narrow-range model predictions using linear and power-law $\mathbf{f}(\varepsilon)$ were compared against measured data. The optimization for $\mathbf{f}(\varepsilon)$ was as follows. The parameter (linear) or parameter combination (power-law) that resulted in a minimum RMSE for the synthesised predictions of vertical illuminance was selected as the optimum $\mathbf{f}(\varepsilon)$. Since the goal was the comparison of predictions for internal illuminance with measurement, it could be argued that either:

1. the RMSE for predictions of vertical South should be minimised since the room has approximately South facing glazing; or, taking this reasoning one step further,
2. the RMSEs for internal illuminance should be minimised.

Both these approaches were rejected because they limit the generality of the $\mathbf{f}(\varepsilon)$ to either a specific orientation (1) or a specific orientation *and* an actual room configuration (2). Nevertheless, it remains the case that any mixing function elucidated from this one dataset will be both site and sample specific to a greater or lesser degree.

Linear mixing function

The form used for this, the simpler of the two combinations, was a straightforward linear mix based on clearness index ε, where the fraction of the total due to the clear sky is

$$f_{cl} = \begin{cases} \left(\dfrac{\varepsilon - \varepsilon_l}{\varepsilon_{ul} - \varepsilon_l}\right) & \varepsilon \le \varepsilon_{ul} \\ 1 & \varepsilon > \varepsilon_{ul} \end{cases} \qquad \textbf{(5-10)}$$

The lower bound clearness index ε_l was always equal to 1 and ε_{ul} was the upper bound, Figure 5-13 (a).

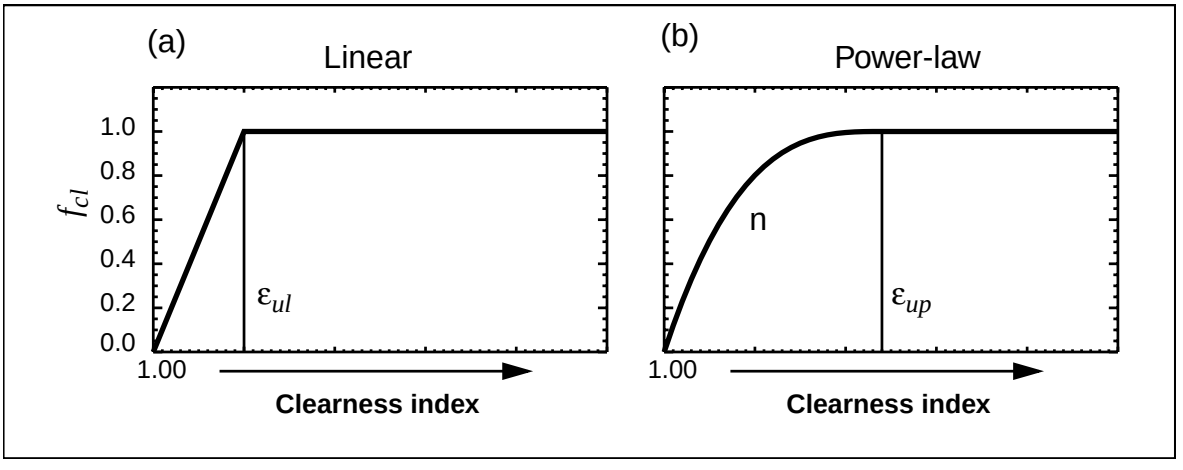

Figure 5-13. Linear and power-law blending functions

The vertical illuminance RMSEs versus ε_{ul} are shown in Figure 5-14 for clear-overcast and intermediate-overcast linear blends. The average RMSE versus ε_{ul} is also shown. Each of the curves show a single stationary (minimum) point. The minimum of the average RMSE is taken to be the

optimum value. The minima for all the curves are also marked in Figure 5-14. The optimum (that is, average RMSE minima) values for ε_{ul} were 1.41 for the clear-overcast blend and 1.10 for the intermediate-overcast blend. For both of the blends, the curves show marked insensitivity to increasing ε_{ul} beyond the stationary point.

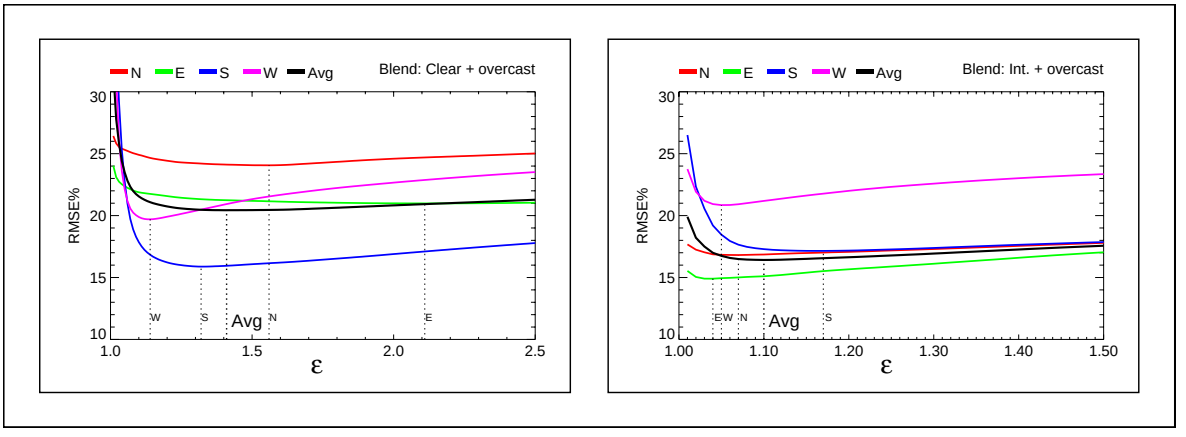

Figure 5-14. Plots of RMSEs for linear blends

Power-law mixing function

This form uses a parabola-like function for the mixing range:

$$f_{cl} = \begin{cases} 1 - \left\{1 - \left(\dfrac{\varepsilon - \varepsilon_l}{\varepsilon_{up} - \varepsilon_l}\right)\right\}^n & \varepsilon \le \varepsilon_{up} \\ 1 & \varepsilon > \varepsilon_{up} \end{cases} \qquad \textbf{(5-11)}$$

Where, for $n > 1$, the transition from pure overcast to pure clear sky with increasing ε is more gradual, and, arguably, more physically realistic than for the linear combination, Figure 5-13 (b). For $n = 1$, the power-law form reduces to the linear form. The average RMSE (for the four vertical illuminances) was evaluated for all the parameter combinations covering the range $1 \le \varepsilon_{up} \le 2$ and $1 \le n \le 3$. The average RMSE for the clear-overcast and intermediate-overcast blends are shown as false-colour maps, Figure 5-15. Also shown on each map is the trajectory of the RMSE minima for each value of n (gray line). The minimum of these minima is marked by

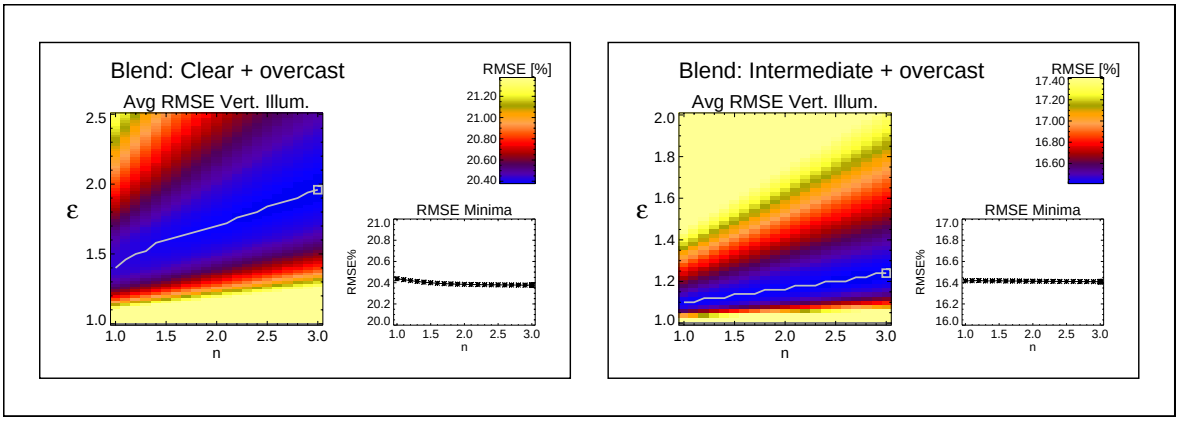

Figure 5-15. Maps of RMSEs for power-law blend

a gray box. The inset plot is of the RMSE minima versus n. Although the optimum is indicated at $n = 3$ for both blends, the difference in minima between $n = 3$ and $n = 1$ (linear form) is very slight indeed. Accordingly, there is little to choose between the linear blend and the (arguably) more physically realistic $n = 3$ power-law blend for either of the blend models. Invoking Ockham's Razor once again, the simpler linear blend model is used for both models to synthesise the illuminance effect of composite sky models. The results using the linear blend models are described in the following section.

5.4.3 Illuminance predictions for sky blends

The relative errors in the predictions for vertical illuminance are presented as frequency histograms. The results for the clear-overcast and the intermediate-overcast blends are shown in Figure 5-16 and Figure 5-17, respectively. As expected, both blend models offer a significant improvement in performance over the any of the narrow-range models alone (see Figure 5-4, Figure 5-5 and Figure 5-6). The MBEs and RMSEs for the sky blends are compared against those for the Perez models and the measured skies (Lumscan) in Figure 5-18. It is worth noting that, of the two blends, the clear-overcast blend performed better for the South and West orientations, whilst the intermediate-overcast blend performed better for

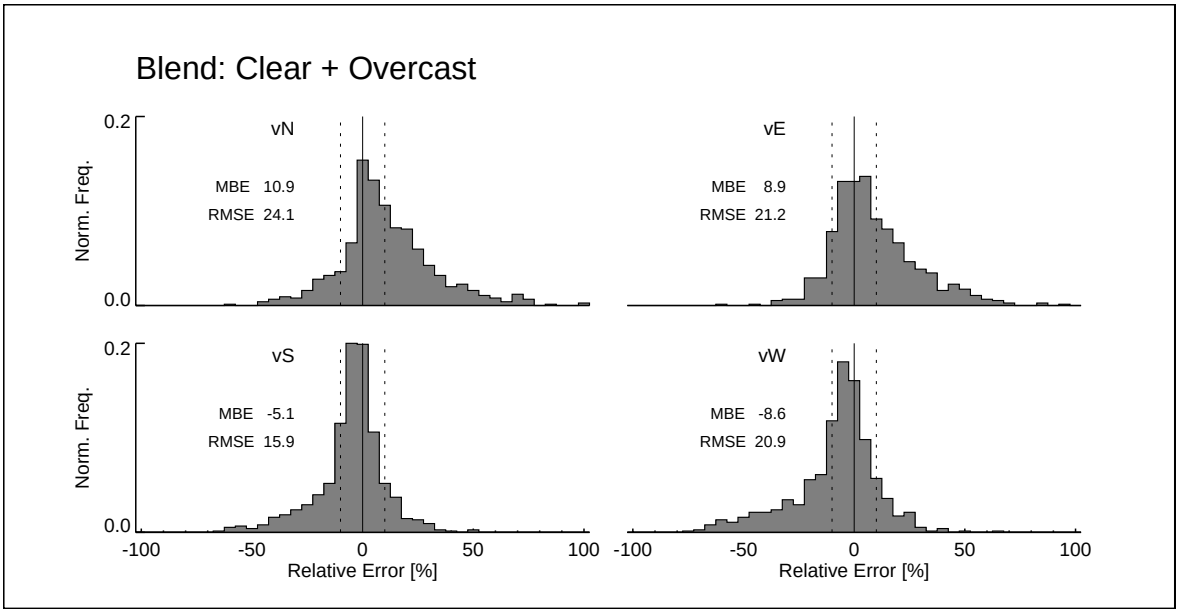

Figure 5-16. Clear + overcast sky composite

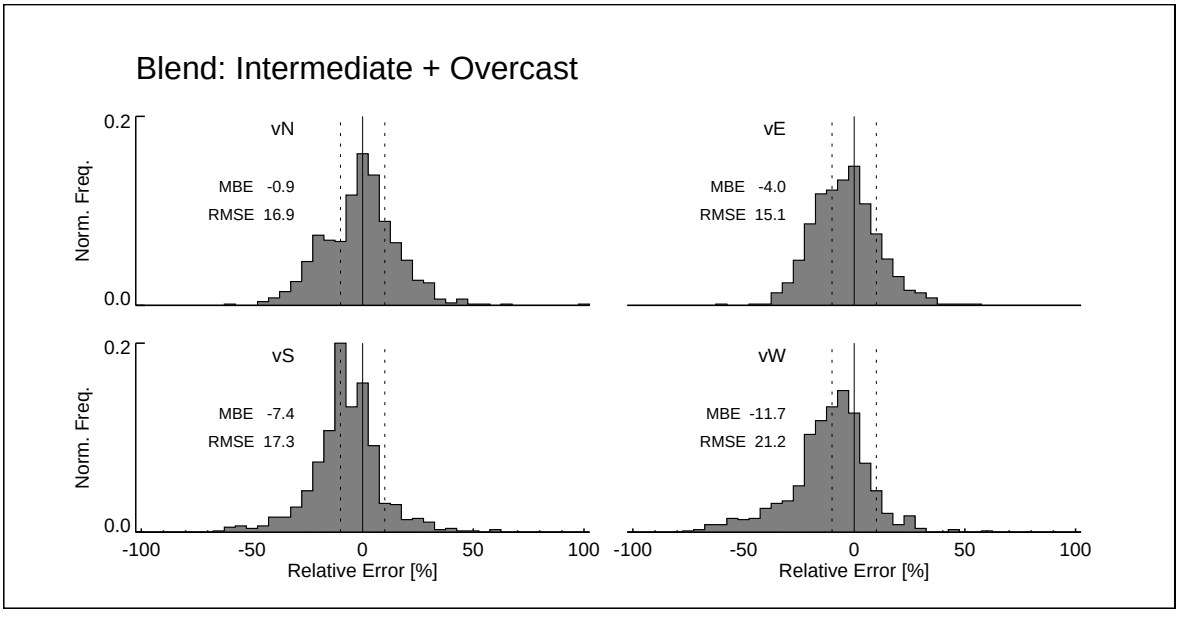

Figure 5-17. Intermediate + overcast sky composite

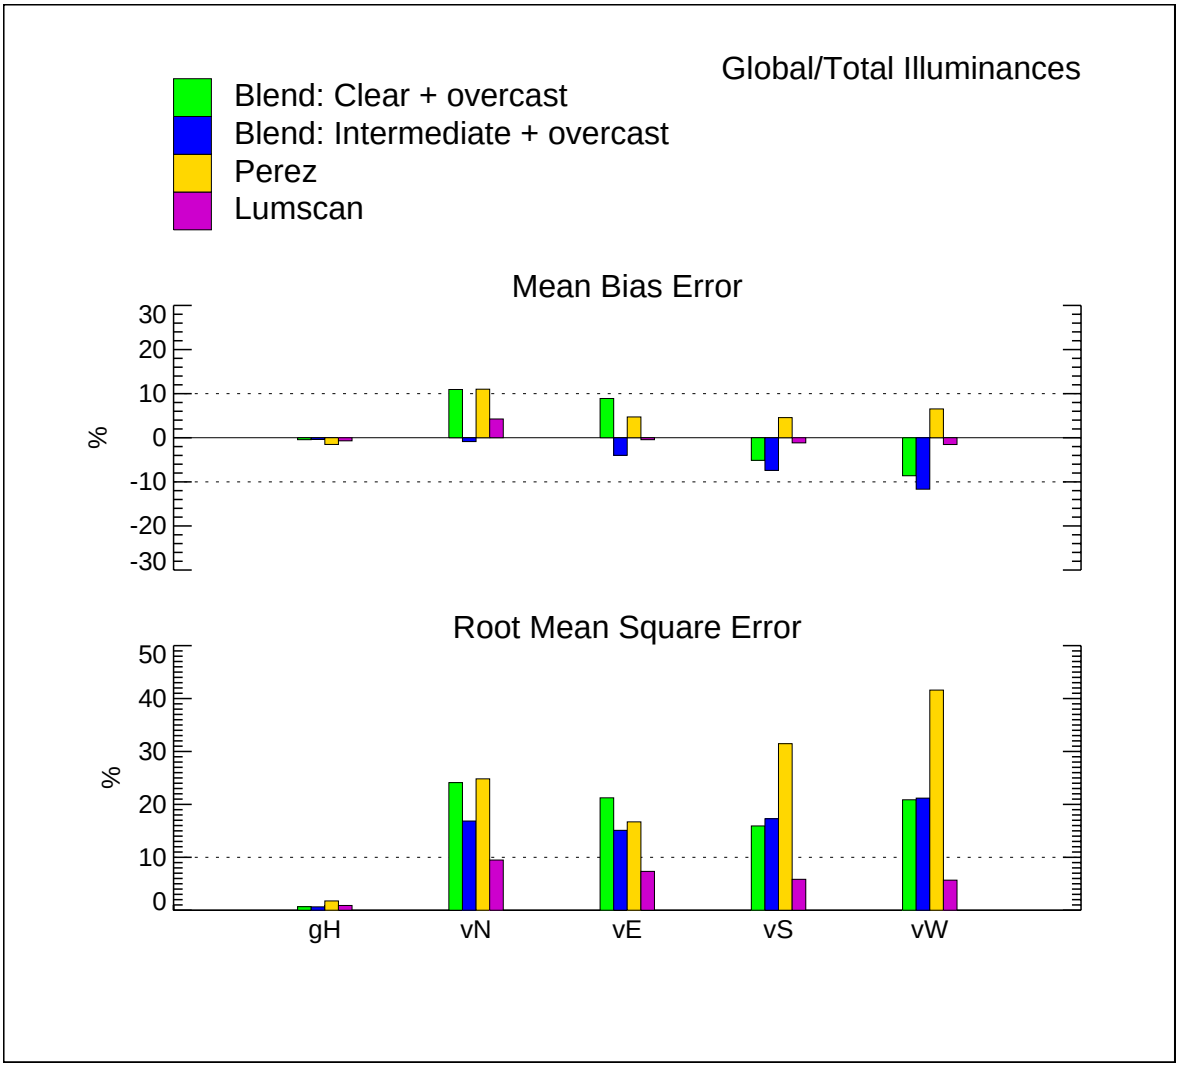

Figure 5-18. Sky models and composites MBE and RMSE

the North and East orientations. Recall that the azimuth angle for the sun position was contained within the range 160˚ to 310˚ (Figure 3-8). Only occasionally, therefore, did direct sun make a contribution to the illuminance for the North and East orientations. It would appear to be the case then, that the clear-overcast blend performed best for orientations South and West where direct sun was often a major contributor to total illuminance. Whilst the intermediate-overcast blend performed best for orientations North and East where direct sun was rarely a significant contributor to total illuminance. This observation may be evidence,

admittedly slight, that a three component model - overcast, intermediate and clear - would perform better than either of the two-component models considered here. As stated previously, it was felt that the number of skies in the BRE-IDMP validation dataset did not warrant testing the configuration of sky blends based on more than two component skies.

Percentile plots for the sky blends

The percentiles analysis (Section 5.3.4) was repeated, but now the narrow-range models (overcast, intermediate and clear) are replaced by the two sky model blends, Figure 5-19. The Lumscan and Perez results are shown once again for comparison. At the ±10% RER line (y-axis), the percentage of the complete sample of skies that achieved this accuracy was: Lumscan, ~63%; clear-overcast blend, ~39%; intermediate-overcast blend, ~34% and Perez, ~30%. The percentile lines (All, CS6-VIS and CS6+VIS) for the measured skies (Lumscan) were markedly better than the corresponding lines for the two blend models up to |RER| ~100% where they converge. Note that the rank order in performance for the percentile lines up to |RER| = 50% is the same as the rank order in RMSE for the prediction of vertical South illuminance (Figure 5-18).

It should also be noted that the clear-overcast blend performed better than the intermediate-overcast blend, if only marginally, even though the luminance distribution of the composite form is somewhat unrealistic.[15] It is likely therefore that the slightly better performance of the clear-overcast blend over the intermediate-overcast blend resulted from a number of sunny sky conditions where the contribution of the overcast sky was zero for both blends. Evidence that this was indeed the case is given in the time-series RER plots for two clear sky days, Figure 5-20. For both of these days, the clear-overcast blend resulted in more accurate illuminance predictions than the intermediate-overcast blend for most of the skies.

15. The half-and-half clear-overcast sky has pronounced horizon brightening and a relatively low circumsolar luminance (see example in Figure 5-12). Yet, horizon brightening is generally associated with clear sky conditions when the circumsolar luminance is large.

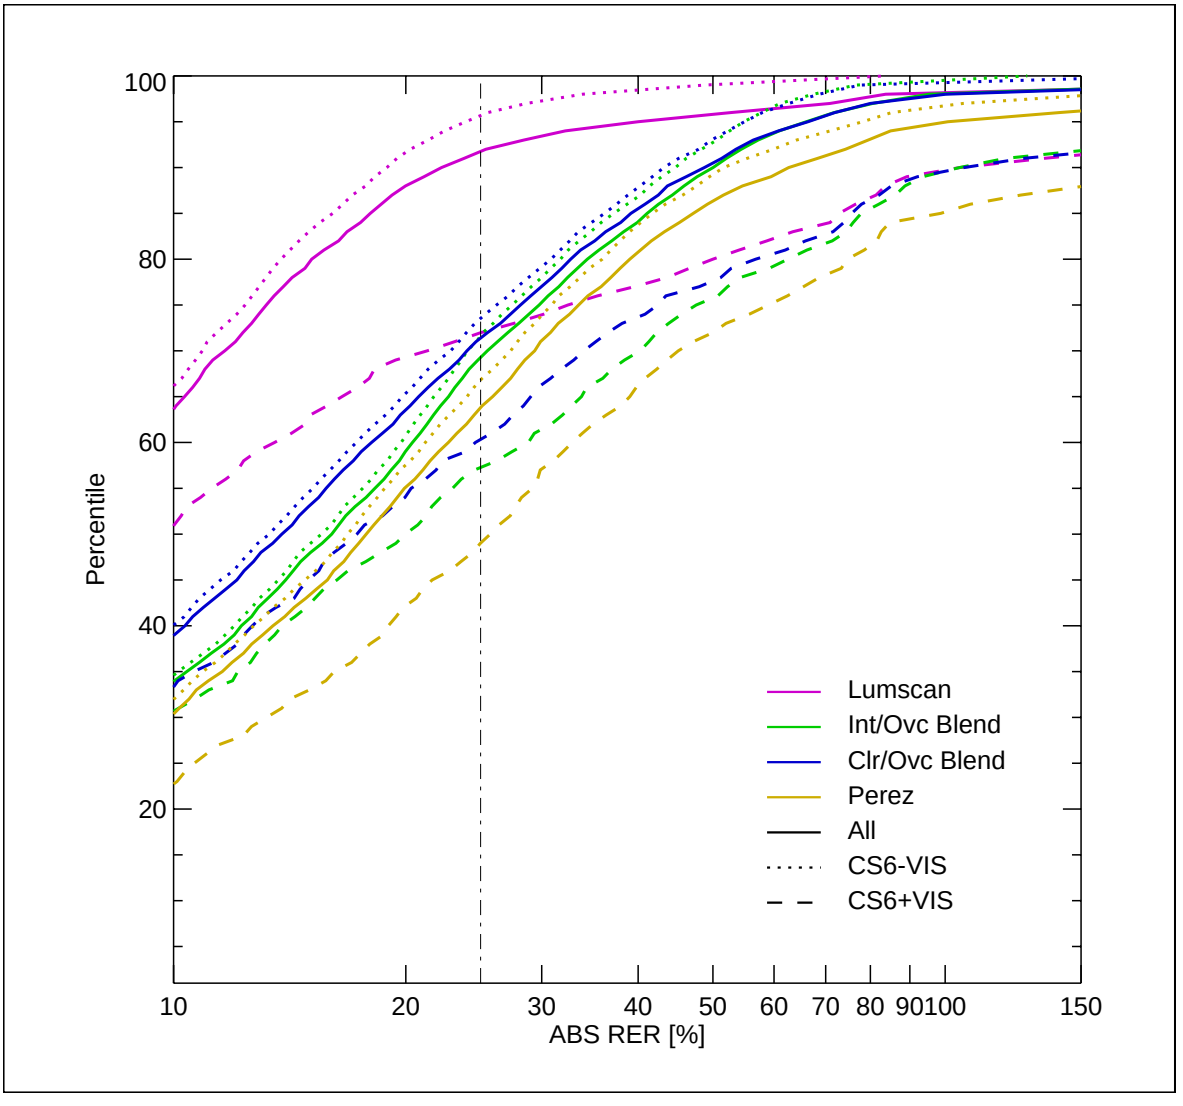

Figure 5-19. Percentile sky blends

Time-series plots for the clear-overcast blend and Perez

With time-series plots it is possible to determine patterns in the RER for predictions of internal illuminance that are difficult to discern from summary metrics. It is practical to show results for only two cases per plot, any more and the plots are too cluttered. Although it would be instructive to show time-series plots for the two blend models, the Perez model, and, for comparison, Lumscan, it would require a large number of plots and lead to

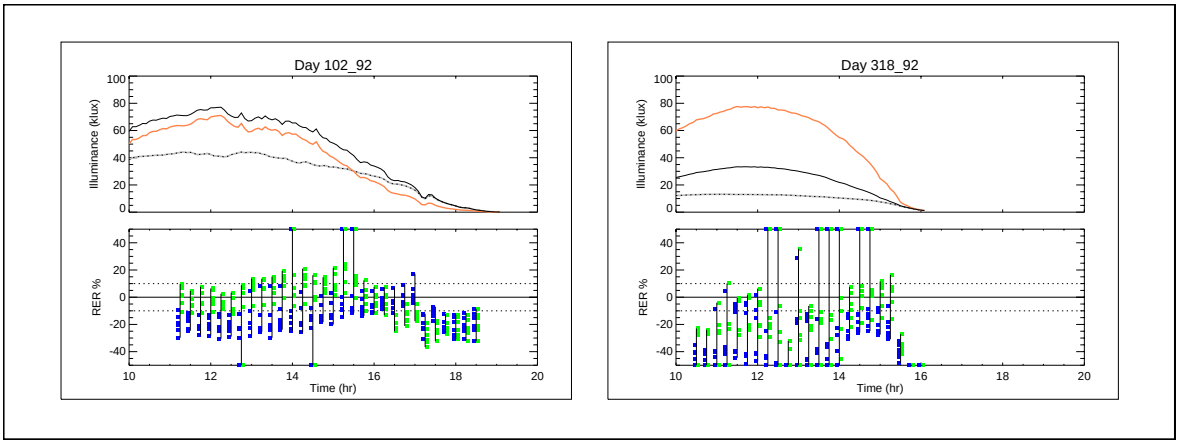

Figure 5-20. Time-series RER for clear-overcast (■) and intermediate-overcast (■) blends

some repetition. For brevity, just one set of time-series plots for all skies in the validation dataset are shown in this chapter.

The two blend models performed identically for overcast conditions because, of course, they both used the same overcast sky model. Also, the Lumscan results were presented in the previous chapter. Accordingly, the two cases shown here are the best performing model - the clear-overcast blend - and the Perez model. The relative error in the illuminance predictions for these two models is shown alongside the corresponding time-series for global horizontal, diffuse horizontal and vertical South illuminance in Figure 5-21 to Figure 5-24. The format for these plots is the same as that used in Chapter 4, only here all the predictions for both models are shown.

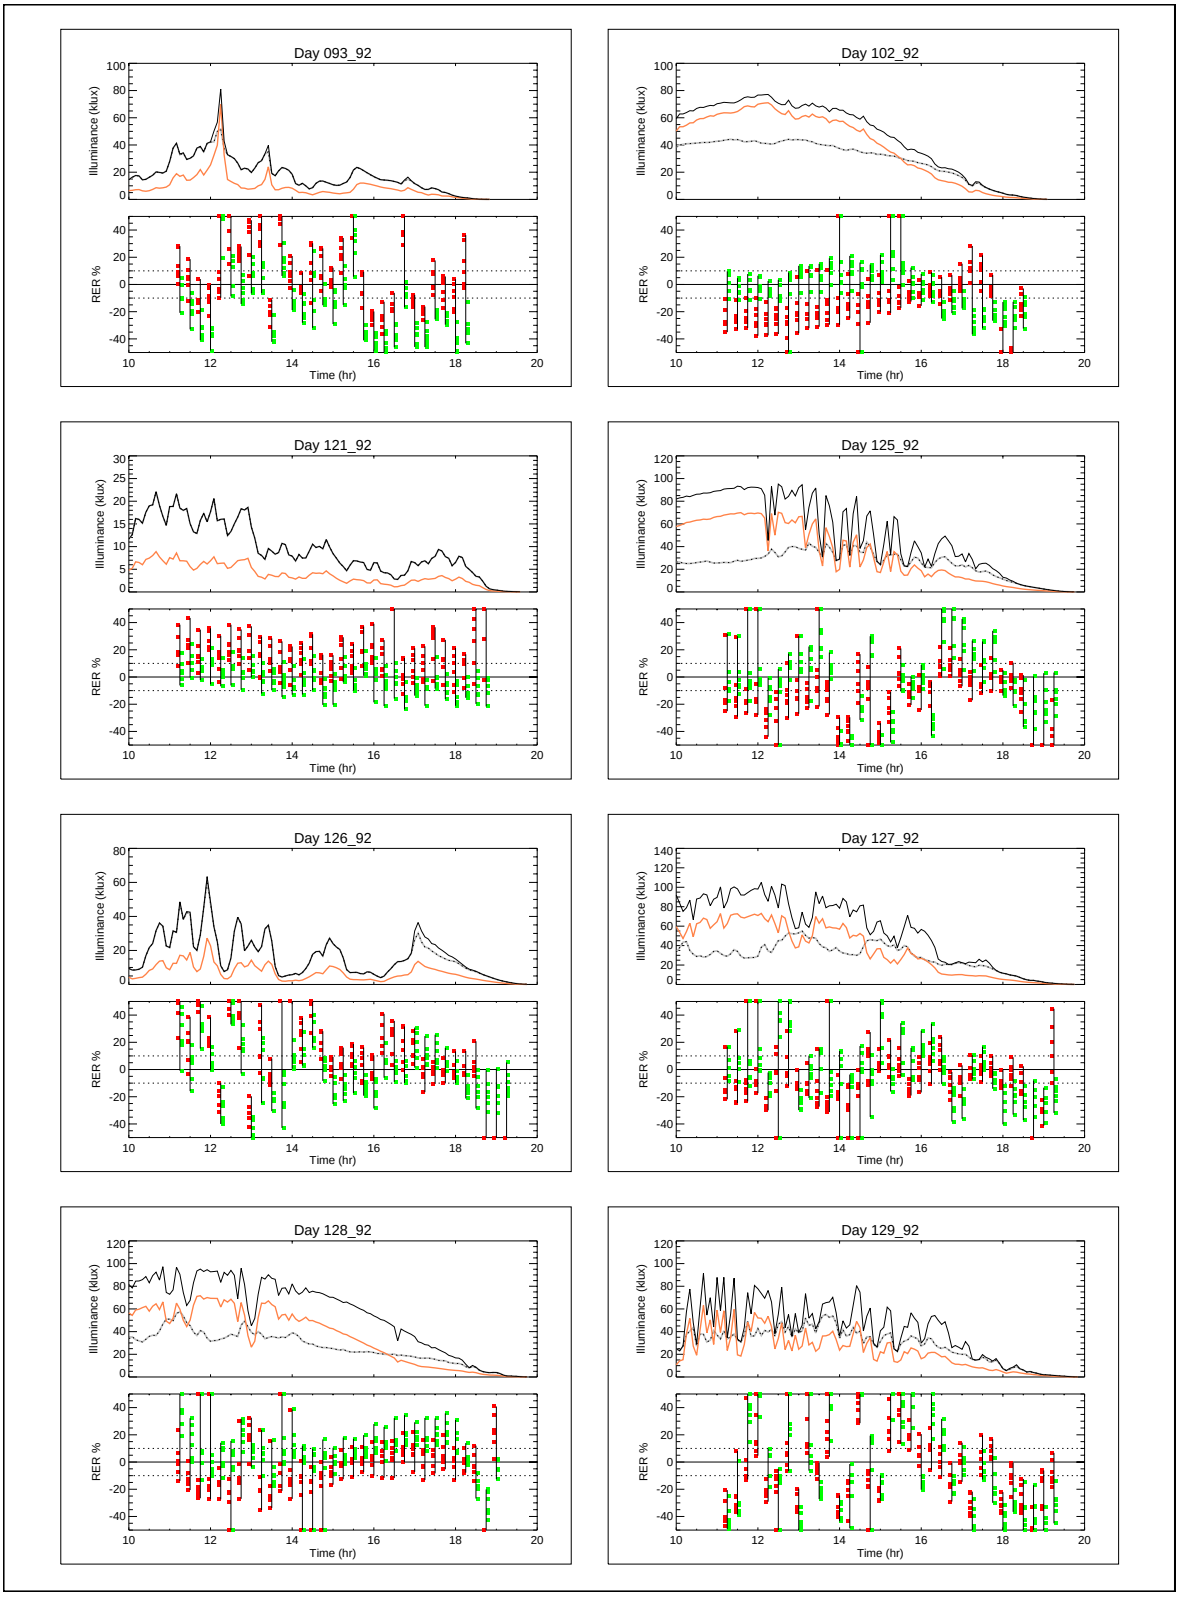

Figure 5-21. Time-series RER for clear-overcast blend (■) and Perez(■)

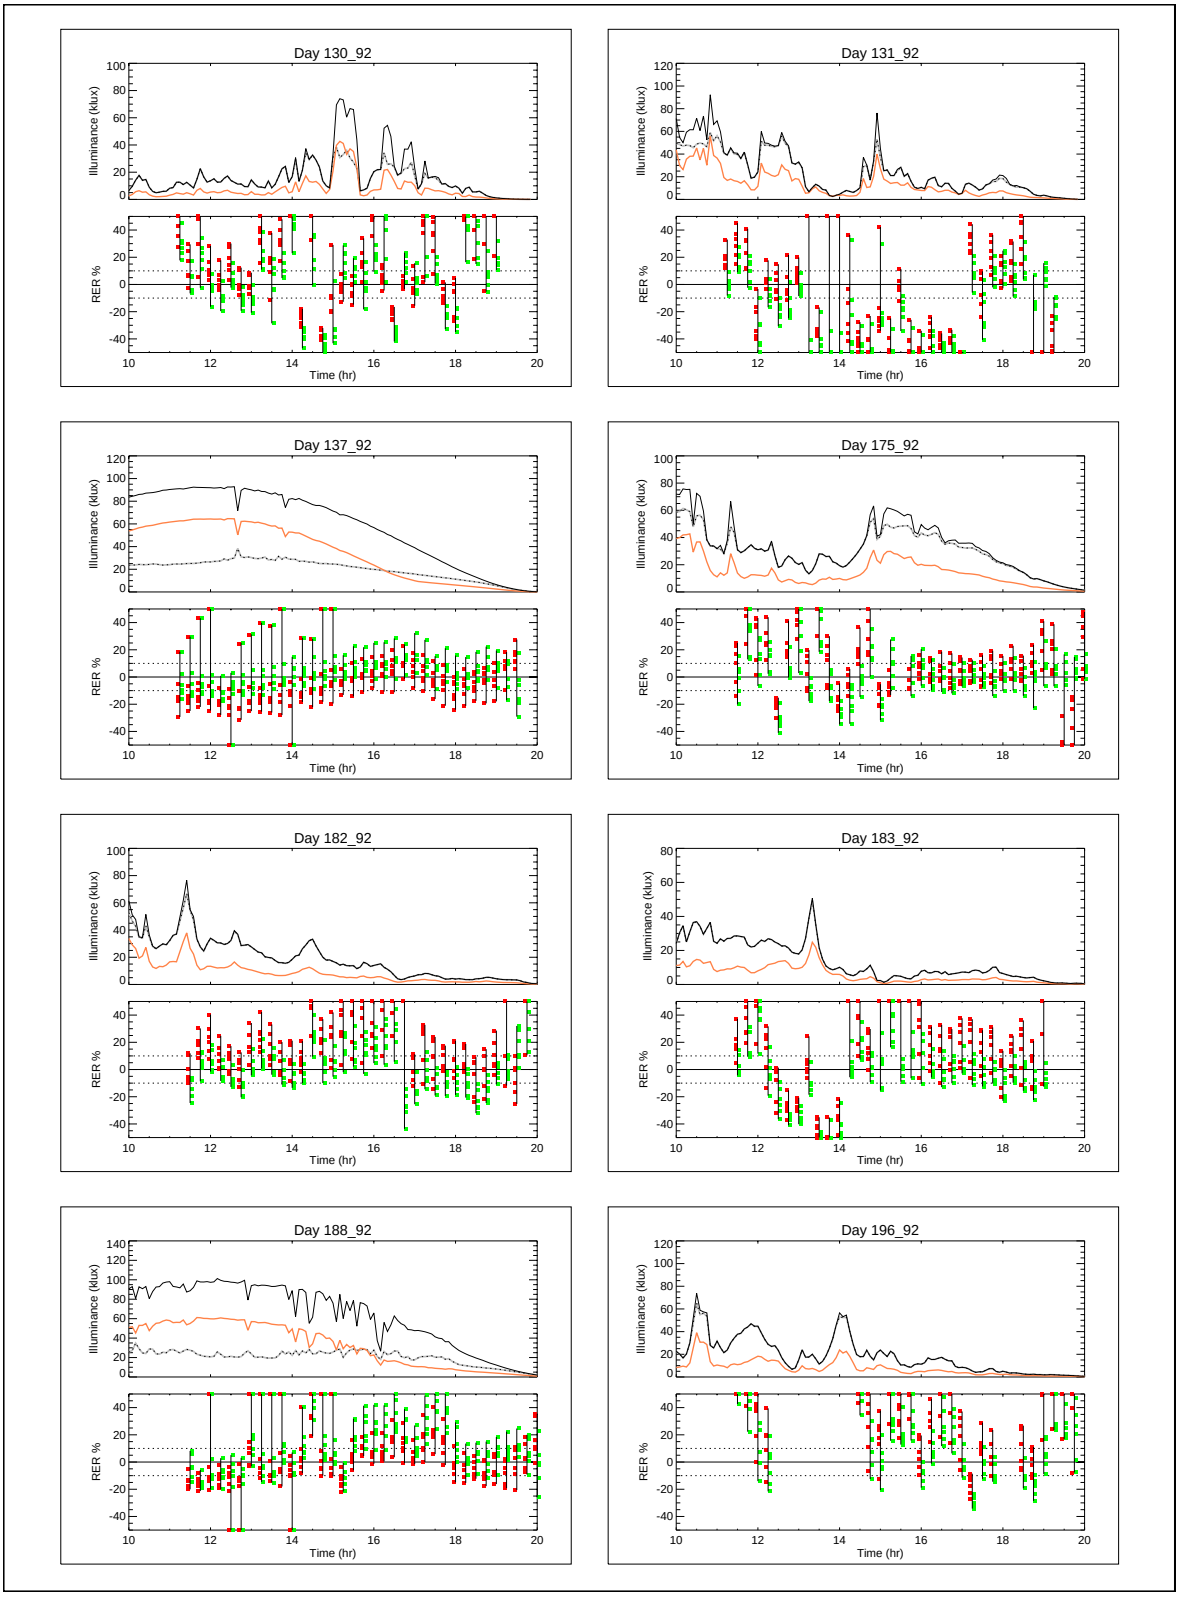

Figure 5-22. Time-series RER for clear-overcast blend (■) and Perez(■)

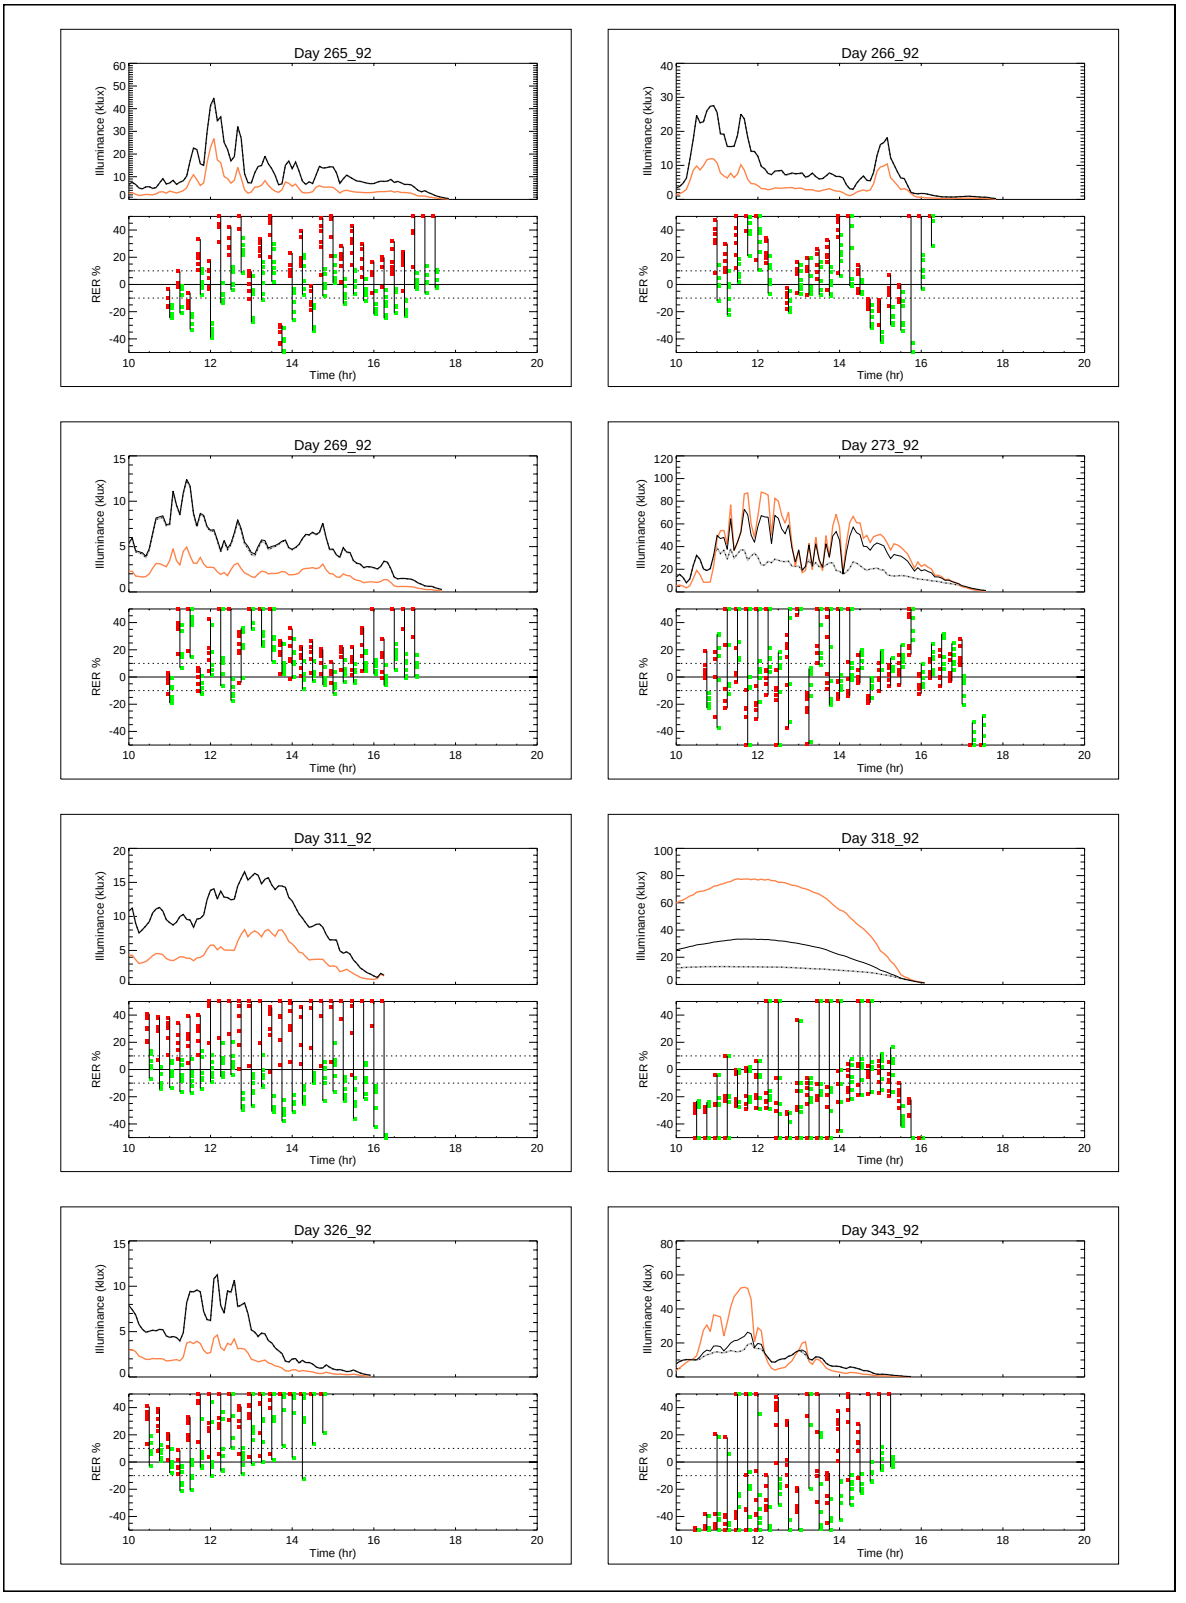

Figure 5-23. Time-series RER for clear-overcast blend (■) and Perez(■)

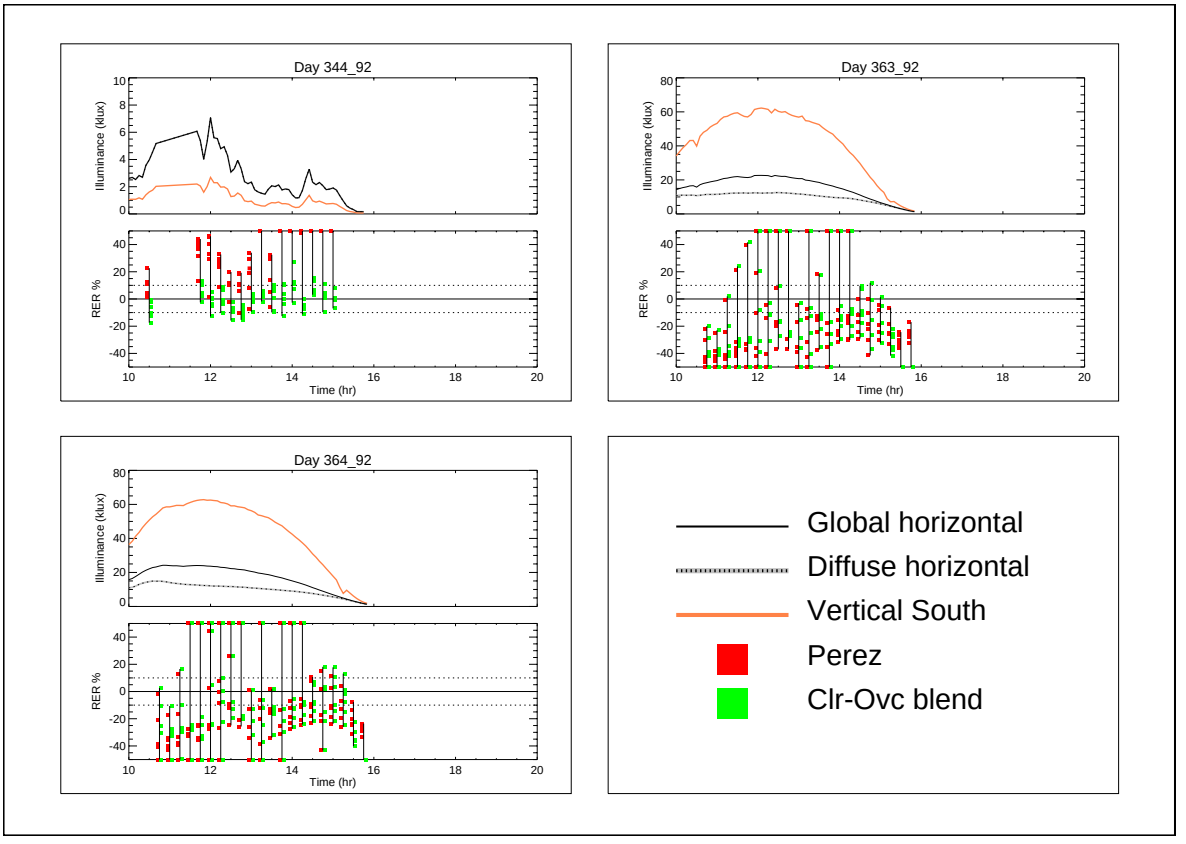

Figure 5-24. Time-series RER for clear-overcast blend (■) and Perez(■)

The following are noted:

- For clear sky days 102_92, 128_92 and 137_92, the clear-overcast blend generally performed better than the Perez model which tended to under-predict illuminances.
- There was a general tendency for the Perez model to overpredict illuminances for overcast skies (e.g. 121_92, 265_92 and 311_92).
- Clear sky conditions in winter (318_92, 363_92 and 364_92) resulted in RERs for both models very similar to that achieved using measured sky luminance distributions (see Figure 4-36 and Figure 4-37 for the Lumscan RERs).

- For several of the overcast days, both models performed relatively poorly compared to the measured skies (e.g. 093_92, 131_92, 183_92 and 196_92). Here it was very likely that, although overcast, these skies contained 'lumpy' sky luminance patterns. Measured by the scanner, these 'lumpy' patterns could not, of course, be reproduced by either of the sky models.

The time-series plots have revealed consistent differences in performance for the two models. It would appear that the low zenith-to-horizon luminance ratio noted in Figure 5-1(a) is a feature of the Perez model for overcast skies, and was the cause of the regular over-prediction for these conditions. There was also some tendency for the Perez model to regularly under-predict for clear sky conditions, though this was not always the case throughout the entire day.

5.5 Conclusion

The results presented here have demonstrated how sky models can be evaluated based on predictions for internal illuminance. Four 'pure' sky models and two sky model blends were examined. Illuminance predictions for the sky model blends were synthesised from the illuminance predictions for the narrow-range 'pure' sky models.

Routine application of the narrow-range models for all 754 skies resulted in poor performance overall. The two blend models and the Perez model performed reasonably well, with the clear-overcast blend marginally the best of the three, and the Perez model marginally the worst. It is not possible to generalise these findings without further work because, of course, the sky model blends were 'tuned' to the validation dataset whereas the Perez model was not. It has already been noted that the Perez model is also site specific to some degree (Section 5.2.2). Nevertheless, the blend models and the Perez model could be applied with reasonable confidence to, say the Kew TRY since this has a similar composition to the validation dataset (Section 3.1.4).

The internal illuminances (predicted and measured) were generally more sensitive to the luminance of the visible part of the sky rather than the much larger part of the sky that was not 'seen'. Thus, the results presented here have demonstrated more the ability of the sky model/blend to reproduce luminance patterns for those parts of the sky that were visible from the photocell locations, than for the sky in total. Which is, of course, as it should be if the purpose of the sky model is to provide daylight illumination for internal spaces.

This investigation is, as far as the author is aware, the first comparison of sky model performance against measured sky luminance patterns that was based on predictions of internal illuminance. The breadth of the analysis was commensurate with the number of skies in the validation dataset. A larger sample, particularly for clear sky conditions, would have allowed a more thorough investigation. Accordingly, conclusions drawn for the performance of these sky models/blends needs to be made with caution.

A more comprehensive evaluation of sky model performance, based on internal illuminances, should examine the effects of glazing orientation. Also, more than 754 skies needs to be used. But how many skies should be considered a representative sample? The answer will depend, to a degree, on the intended use for the sky model. For daylight illumination, one goal is the prediction of time-varying internal illuminances using realistic models for the sky and the sun conditions. This could be carried out in conjunction with dynamic thermal analysis, using the same period and timestep as the thermal simulation [Clarke 98]. This invariably means a simulation period of a full year at a timestep of one hour with meteorological conditions derived from test reference year data.

Any analysis that is based on a TRY time-series would be computationally very demanding; approximately 4,000 unique skies (i.e. daylight hours in the year) would need to be modelled for each case. Using the 'standard calculation', that is modelling the illuminance effect of each individual sky

for thousands of unique skies, could take days or even weeks of computer processor time. A potentially more efficient approach to predicting the internal illuminance for a large number of unique sky and sun configurations - measured or modelled - is described in the following chapter.