This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

Sky model blends for predicting internal illuminance: a comparison founded on the BRE-IDMP dataset

PLEASE CITE THE PUBLISHED VERSION

<http://dx.doi.org/10.1080/19401490802419836>

PUBLISHER

Taylor & Francis (© 2008 International Building Performance Simulation Association (IBPSA))

VERSION

AM (Accepted Manuscript)

LICENCE

CC BY-NC-ND 4.0

REPOSITORY RECORD

Mardaljevic, John. 2019. "Sky Model Blends for Predicting Internal Illuminance: A Comparison Founded on the BRE-IDMP Dataset". figshare. <https://hdl.handle.net/2134/11577>.

This item was submitted to Loughborough's Institutional Repository (<https://dspace.lboro.ac.uk/>) by the author and is made available under the following Creative Commons Licence conditions.

For the full text of this licence, please go to:
<http://creativecommons.org/licenses/by-nc-nd/2.5/>

Sky Model Blends for Predicting Internal Illuminance: A Comparison Founded on the BRE-IDMP Dataset

John Mardaljevic PhD

Institute of Energy and Sustainable Development
De Montfort University
The Gateway, Leicester, LE1 9BH, UK

e-mail jm@dmu.ac.uk
Phone: +44 (0) 116 257 7972
Fax: +44 (0) 116 257 7981

Abstract: The BRE-IDMP validation dataset contains simultaneous measurements of sky luminance patterns and internal illuminances in two full-size office spaces. This benchmark dataset has been applied previously to test the illuminance predictions from a lighting simulation program under real sky conditions. Sky luminance patterns were mapped into the lighting simulation so that the absolute accuracy of the program could be evaluated without the uncertainties that are introduced when sky models are used. For this follow-on study, the BRE-IDMP dataset is now used to quantify the divergence between the sky model generated luminance patterns and the actually occurring conditions based on the resulting internal daylight illuminances. The internal illuminances were predicted using three ‘narrow-range’ models (CIE overcast, CIE clear and intermediate) and the Perez All-Weather model. Predictions from the narrow-range models were used to investigate formulations for sky model blends. The illuminance effect of arbitrary sky model blends is reproduced in a post-process of the illuminance predictions from the ‘narrow-range’ sky model types. The determination of an optimum sky model blend is described. The findings show that relatively simple blends of just two pure sky models (e.g. CIE overcast and intermediate) may be adequate for the prediction of time-varying illuminances founded on climatic test reference year data.

Keywords: Sky models, daylight, simulation.

1 Introduction

Sky models generate continuous sky luminance patterns. 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 could not account for the random-discontinuous features (i.e. discrete clouds) 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, since it is impossible

to account for the second cause of divergence. Prior to the work reported here, sky model evaluation has largely been carried out in terms of the differences between modelled and measured sky luminance patterns [1][2]. It is noted however that in the vast majority of published work on sky modelling, the stated principle application of the models is to determine “an accurate estimate of the daylight available” [2]. In support of those aims, this study evaluates sky model performance on the basis of a comparison of predictions of internal daylight illuminance with measurements.

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, Figure 1. The four skies were selected from the BRE-IDMP validation dataset [3] to illustrate 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, since in all cases it was based on an independent measurement of direct normal illuminance. 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. The model sky description for these examples was based on the Perez ‘All-Weather’ model [4]. These comparisons are purely illustrative. The sky point luminance is indicated by the height of the surface. 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. The same scaling and rotation were applied to each pair of luminance surfaces.

The four measured skies shown in Figure 1 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 measured skies, though differences are readily apparent in both the underlying pattern (e.g. overcast sky 326_92_11h00) and those caused by the presence of discontinuous features.

2 Illuminance predictions from modelled sky luminance patterns

The BRE-IDMP validation dataset contains simultaneous measurements of sky luminance patterns, solar illuminance and internal illuminance in two full-size mock offices. With this dataset it is possible to specify to an unprecedented degree of precision the daylight conditions at the time of measurement. The dataset comprises 754 simultaneous measurements of internal and external parameters taken from 27 days of monitoring. The 27 days were pseudo-randomly selected by the BRE to cover the range of naturally occurring sky conditions from heavily overcast, through intermediate to clear. On the basis of the Perez clearness index, the validation dataset skies were found to be reasonably representative of the range of skies in the Kew TRY [5]. Thus the 754 skies in the validation dataset can be said contain much of the variation that might be expected for the UK climate.

The validation exercise described in a previous paper showed that the intrinsic accuracy for illuminances predicted by *Radiance* was typically 10% or better [3]. Thus the accuracy was comparable with that of the instruments used to record the sky, sun and room conditions. That validation exercise was repeated with the sky luminance patterns now provided by sky models. The simulation parameters used were the same as before. Comparable absolute accuracy in the simulation of light transfer can therefore be expected, and any divergence from the predictions obtained using the measured skies can be reliably attributed to the sky-model generated luminance distribution. Illuminance predictions for the 754 unique instances in the validation dataset were obtained using luminance distributions generated by the following sky models:

- the CIE standard overcast sky model;
- the Matsuura Intermediate sky model;
- the CIE clear sky model; and,
- the Perez All-weather sky model.

The first three of the above are referred to as narrow-range models since they were formulated to reproduce luminance patterns for specific 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. Note however that the Perez model was founded on data collected only at one site, i.e. Berkeley, California, USA [4]. The three narrow-range models and the Perez model are called here ‘pure’ sky models because they are distinct formulations. For brevity, the models are occasionally referred to simply as overcast, intermediate, clear and Perez. The absolute value of the sky luminance was, in each case, normalised to the diffuse horizontal illuminance. A separate model sun was based directly on the measurement of direct normal illuminance. Thus, each of the four model skies used the same sun description. Note that, for consistency of process, a model sun was used even when the measurement of direct normal illuminance indicated a sun luminance that was comparable to the background sky luminance (i.e. when conditions were heavily overcast).

Any one of the narrow-range models will be incapable of reproducing the full range of sky conditions in the validation dataset. For these models, the purpose is to generate the basic data from which the illuminance *effect* of a sky model blend may be synthesised in a post-process of the simulation data (this is described in the following section).

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. This effect was noted by Littlefair and an adjustment to the model to prevent this distortion was advised by Perez [2]. A routine examination of the sky model generator program 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 the distortion regardless of the fix.

3 Model sky blends: *ex post facto* synthesis

The usual practice for blending skies is to combine an overcast luminance pattern with one or more non-overcast patterns according to some rule [2]. 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 sky model blends used here are each a composite of an overcast luminance pattern and just one non-overcast luminance pattern. 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:

$$L_{co} = f_{cl}L_{cl} + f_{ov}L_{ov} \quad (1)$$

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. As noted, the luminance for all the model skies was normalised to the measured diffuse horizontal illuminance (E), so that $E = E_{cl} = E_{ov}$. Therefore, setting $f_{ov} = (1 - f_{cl})$ normalises the composite sky to diffuse horizontal illuminance also. Thus:

$$L_{co} = f_{cl}L_{cl} + (1 - f_{cl})L_{ov} \quad (2)$$

and

$$E_{co} = f_{cl}E_{cl} + (1 - f_{cl})E_{ov} \quad (3)$$

Therefore, blending the clear and overcast total illuminances is equivalent to blending the sky components, and then adding the separately computed sun component. The process for blending the intermediate with the overcast was identical to that used for the clear-overcast blend, replacing the subscripts cl with in and co with io where appropriate. For example, the illuminance predicted for the intermediate-overcast blend would be E_{io} . The luminance patterns for the three narrow-range models and the two sky blends are given in Figure 2. The three images in the upper row show 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. $f_{cl} = f_{in} = f_{ov} = 0.5$. 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.

3.1 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. Evidently, the more overcast the actual sky conditions the smaller the contribution of the non-overcast sky 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 of the sky clearness index k over a mixing range bounded by lower and upper values for a fully overcast sky and a fully non-overcast sky respectively. Within the mixing range, the effect of linear and a power-law blending functions were examined. The illuminances synthesised from the narrow-range model predictions using linear and power-law blending functions were compared against measured data. The parameters governing the blend - one for the linear and two for the power-law blend - were set by determining the single value (linear) or combination (power-law) that resulted in a minimum RMSE for the synthesised predictions of vertical illuminance compared against independently measured values. 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 blend 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 one dataset will be both site and sample specific to a greater or lesser degree.

The form used for the simpler of the two combinations was a straightforward linear mix based on clearness index k , where the fraction of the total due to the clear sky is f_{cl} . The lower bound clearness index was always equal to 1 and k_{ul} was the upper bound, Figure 3. Thus this blend model has only a single parameter k_{ul} .

The search for the upper bound k_{ul} value that gave the minimum RMSE in the prediction of the four vertical illuminances proceeded as follows. Initially, the k_{ul} was set to 1.0 which means that, for all values of clearness index for the 754 skies in the validation dataset, the four vertical illuminance values used were those predicted assuming a pure clear sky in each case. In other words, regardless of the sky conditions, the vertical illuminance predictions were determined using clear sky conditions founded on the diffuse horizontal illuminance. Next, the k_{ul} was set to 1.01. Now, vertical illuminance predictions for those skies with a clearness index between 1 and 1.01 are synthesised from a linear blend of overcast and clear skies using the rule described above. For skies with a clearness index >1.01 , predictions were determined assuming clear sky conditions. Proceeding in this way, all the values of k_{ul} in the range 1 to 3 in increments of 0.01 were tested. Thus, three hundred instances of the simple blend model were tested. This was carried out for the overcast-clear blend and the overcast-intermediate blend - the procedure was the same for both.

A plot of the vertical illuminance RMSEs versus the upper bound for the blended region are shown in Figure 4. Results are presented for both blends. In addition to the curves for the four orientations, the curve showing the average of the four RMSEs (determined at each 0.01 increment) is also shown. Each of the curves show a single stationary (minimum) point - indicated with a vertical dashed line. As expected, the greatest RMSEs are for values of k_{ul} that are close to unity. Here the blend will use the clear sky model formulation when a number of the actually occurring conditions were overcast. Note that, although the minima for the the various orientations are quite different, most of the curves show a marked insensitivity to increases in k_{ul} beyond the stationary point. As noted, the optimum k_{ul} is taken to be that which gives the lowest RMSE in the average of the four predictions of vertical illuminance for the 754 skies in the validation dataset. The k_{ul} optimum values were 1.41 for the clear-overcast blend and 1.10 for the intermediate-overcast blend.

A power-law blend was also tested. This introduced an additional parameter but produced results that were barely distinguishable from the linear

model [5]. Thus, in absence of any compelling reason to use the two parameter power-law model, the simpler linear model was used for both blends to synthesise the illuminance effect of composite sky models. The results using the linear blend models are described in the following section.

4 Results

The internal illuminance predictions that were determined separately for the overcast, intermediate and clear sky models were now combined according to the rules described above to synthesise the illuminance effect of overcast-intermediate and overcast-clear blends. The divergence between the measured internal illuminance and that predicted using the two sky model blends and the Perez model is compared against the illuminances that were predicted using the measured sky luminance distributions. The predictions computed using the measured sky luminance distributions (labelled ‘Lumscan’ in the figures) should be considered ‘benchmark’ values, i.e. those determined using the most complete description of the daylight luminous environment for the test scenario [3]. Prior to the presentation and discussion of the results obtained using the sky models, the following section describes the necessary preparation and filtering of the validation dataset.

4.1 Filtering the benchmark dataset

The BRE-IDMP validation dataset contains one of the most complete set of simultaneous measurements for internal illuminance and sky luminance distribution. Illuminance measurements were taken in a full-size office space under a wide range of naturally occurring conditions. These and the sky luminance measurements were obtained using high-precision, calibrated instruments and observing strict, quality-assured data collection procedures [6]. It is still perhaps the most reliable dataset for the validation of illuminance predictions under daylight conditions, and it provides a far more rigorous evaluation environment than the idealised scenarios devised for standardised tests [7, 8].

The BRE-IDMP dataset however is fairly complex, and its use for exacting validation purposes proved to be quite demanding. It was discovered that the dataset contained occurrences of potentially unreliable photocell-sky combinations [3]. It should be noted that the principle causes for these errors were entirely due the finite resolution of the various measuring instruments

used. Furthermore, the errors could not reasonably have been reduced by any significant amount using other means. The causes were referred to as source visibility related errors (SVREs), and they were identified as resulting from one or more of four error types. The two most significant of the SVREs relate to imprecision in the model representation of the real scene. For example, significant errors in the illuminance prediction can result from small geometric and/or orientation differences between the simulation model and reality: misalignment of just one millimetre can produce large relative errors (RERs) when there are shadows cast on - or near to - the photocell by the window frame bars. The office glazing had several window bars, and although they were measured individually to an accuracy of $\sim 2\text{mm}$, positional errors of 1 to 2 cm relative to the overall scale of the room were possible. It was reasonable to assume therefore that at least some of the high RERs might be due to a mis-match between the modelled geometry and that of the actual office. Note that it is virtually impossible to conclusively attribute any one specific high RER to positional misalignment alone. Another possible SVRE is the uncertainty of the brightness distribution about the solar position. When clear sky conditions prevail, there will be large luminance gradients within the acceptance angle of the scanner at and near to the solar position. To protect against sensor damage, the scanner did not record the luminance at the patch closest to the solar position, and this value had to be estimated. In any case, the potentially huge luminance gradients within the 11° acceptance angle could not have been reliably estimated even if the mean luminance across the circumsolar region had been recorded. The overall sky luminance was normalised to diffuse horizontal illuminance, but it remains that the luminance gradients about the the circumsolar are unknown. This could be the cause of errors in the prediction of illuminance when the photocell has a direct “view” of the circumsolar region, since the direct sky component of illumination can be a major part of the total illuminance.

As things stand, there is little scope to correct for any one of these SVRE 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. However it was possible to identify “at-risk” entries in the validation dataset which could give large RERs in illuminance predictions due to the SVREs. The potential for inaccuracy in the illuminance predictions resulting from all four SVRE types - acting independently or in combination - is greatest for sunny conditions when the sun (and circumsolar region) come into view from the photocell position. Thus the photocell-sky combi-

nations at-risk are those where the photocell can “see” or nearly “see” the sun position. A test was carried out to determine which of the 4,524 unique photocell-sky combinations (i.e. 754 skies and the six photocell locations) had a direct view of a 6° circumsolar region centred on the sun position [3]. This patch of sky was called the circumsolar exclusion region (CER). The illuminance predictions at each of the six photocells were partitioned into sets designated as either ‘reliable’ or ‘potentially unreliable’ depending on the visibility of the CER.

This was proven to be a highly effective discriminator: the vast majority of the high RER predictions (i.e. $> 50\%$) for the measured sky luminance patterns were those where the photocell had a direct view of the CER. By separating out these ‘unreliable’ data from the ‘reliable’ entries, it was possible to arrive at a true assessment of the absolute accuracy of the simulation program (using measured sky luminance patterns). Using only the ‘reliable’ data with the measured sky luminance patterns, 67% of predictions were within $\pm 10\%$ of the measured values, and 96% were within $\pm 25\%$. Thus, the simulation program can be considered to be capable of very high accuracy, considerably better than anything yet demonstrated for physical (i.e. scale) modelling [9].

A few outliers however did remain in the ‘reliable’ set. These could be the result of rapidly changing illumination conditions during the period of measurement, chance reflections of sunlight from rainwater puddles, and so forth. Given the number of potentially confounding factors at work, and the practicalities of sky metrology, there seems little scope for eliminating every single occurrence of inaccurate prediction caused by imprecise model representation. Furthermore, it is evident that just a few outliers can greatly influence the MBE and RMSE for a sample of otherwise accurate predictions.

4.2 A comparison of illuminance predictions based on percentiles

It is proposed that a percentiles-based evaluation gives a more insightful analysis for datasets that may contain outliers with occasional very high RERs. For the evaluation of the accuracy of the illuminance predictions determined using sky models, the percentage of the predictions that had a relative error within the range $\pm R$ were plotted as a function of R . These plots are referred to here as percentile-RER plots. Percentile-RER curves are given for illuminance predictions computed using:

- The ‘benchmark’ measured sky luminance patterns (Lumscan).
- The intermediate-overcast blend (Int/Ovc).
- The clear-overcast blend (Clr/Ovc).
- The Perez All-Weather model (Perez).

Colour is used to distinguish between the four types, and for each one separate curves are given for: the complete dataset (solid line), the ‘reliable’ set only (dotted line) and the ‘unreliable’ set only (dashed line), Figure 5. Thus the plot shows twelve curves.

As expected, the predictions obtained using the ‘reliable’ set from the measured sky luminance patterns (magenta line) are the most accurate. Considering just the ‘reliable’ set for each of the sky model blends and Perez, the clear-overcast blend performs marginally better than the intermediate-overcast followed by the Perez. The difference in performance between the sky models is markedly less than the difference in performance between the best sky model (i.e. clear-overcast blend) and the measured sky luminance patterns. For example, 41% of illuminance predictions for the clear-overcast blend were within $\pm 10\%$ of measurements compared to 32% of predictions from the Perez model. The percentile lines (All, ‘Reliable’ and ‘Unreliable’) for the measured skies (Lumscan) were markedly better than the corresponding lines for the two blend models up to $\sim 100\%$ where they converge.

The instruments used to collect the measurements that comprise the BRE-IDMP dataset (i.e. the illuminance meters for the internal and external conditions and the sky scanners) have accuracies in the range $\pm 10\%$ to $\pm 25\%$ [6]. The error analysis described in [5] indicates that the *Radiance* simulation program is capable of delivering illuminance predictions with an accuracy of $\pm 10\%$. Achieving this accuracy however requires input data of comparable quality and rigour to that of BRE-IDMP dataset. Lower quality datasets and/or less rigorous evaluation methodologies may produce misleading assessments [10]. Given that there is not yet an established performance ‘target’ for lighting simulation, it is suggested that the assessment of performance should be based on the percentage of predictions that fall within $\pm 10\%$, $\pm 25\%$ and $\pm 50\%$ of the measured values. The percentile values for the predictions falling within those three ranges determined using the ‘reliable’ subset of the measured data are given in Table 1.

4.3 Time-series plots of the relative error in illuminance prediction

Although summary metrics, such as the three percentile values described above, are needed to characterise overall performance, an examination of the time-series of the relative error in predictions reveals how trends in performance are related to prevailing sky conditions. A full exposition for all twenty seven days of monitoring is not possible within the confines of a journal paper (see [5] for this), however key trends are revealed in the examples given below for just two of the days.

The first figure shows measurements and predictions a clear sky day, Figure 6. The topmost plot shows a time-series of measurements for global horizontal illuminance (solid line) and diffuse horizontal illuminance (dotted line) recorded at 5 minute intervals. The middle plot shows the relative errors (RERs) in the internal illuminance predictions for the six photocells where the sky luminance pattern was based on the scanner measurements. At each time-step, the relative error in the six illuminance predictions are plotted in a vertical line. Predictions were made every 15 minutes when a scan of the sky was carried out. The photocell sky combinations deemed ‘reliable’ are plotted using a downwards pointing triangle shaded magenta (▼). Those deemed ‘unreliable’ (i.e. where the photocell could “see” the circumsolar region) are indicated by an upwards pointing triangle shaded cyan (▲). Predictions with RERs outside the range are plotted at the range limits. The majority of predictions for the reliable set are within the $\pm 10\%$ range, and all are within $\pm 25\%$ range. The mechanism to detect potentially ‘unreliable’ entries is clearly successful: the five identified all have RERs outside of the range. The plot at the bottom of this group shows the RERs in the illuminance predictions where the sky conditions were generated using a clear-overcast blend (◀) and the Perez All-Weather model (▶). The symbols are plotted either side of the vertical line that covers the range in RERs. The three plots are time-aligned, and it is easy to identify on the Skymodel plot the predictions that correspond to the large RER points in the Lumscan plot above. For this clear sky day the difference in performance between the blend model and the Perez is readily apparent. The clear-overcast blend performs consistently better than the Perez until around 17h00, after which the blend model tends to under-predict illuminances. In contrast, over the period 11h00 to 17h00 the Perez tends to under-predict illuminances, after which there is some over-prediction. The clearness index range over which the overcast and clear sky

patterns are mixed is fairly narrow. Tests using similar blends with inputs taken from annual UK climate data (i.e. containing $\sim 4,400$ hourly values) resulted in the generation of a mixed sky luminance pattern for only $\sim 1/8$ of the time [5]. For the rest of the time, the blending function selected 100% contribution from either the overcast or the non-overcast sky. It is likely that, between 16h00 and 17h00, the sky model blend switched from pure clear sky to pure overcast. For practical purposes, little importance should be given to the RERs for Skymodel predictions after 18h00 since the absolute values of internal illuminance were very low. Nevertheless, it is worth noting that the Lumscan predictions maintain high accuracy irrespective of the absolute levels of illumination.

The second example day was overcast throughout, as is evident from the topmost plot where the time-series for global and diffuse horizontal illuminance are identical, Figure 7. As before, the potentially ‘unreliable’ cases are identified, but here these are not associated with any conspicuously high RER predictions since the SVREs generally do not manifest under totally overcast sky conditions. The clear-overcast blend performs comparably well for this overcast sky as it did for the clear sky conditions in the previous example, though it should be noted that here the blend would have selected a pure (CIE) overcast sky throughout the day since the direct normal contribution was always zero. The predictions using the Perez All-Weather model are now generally greater than the measurements. Since the clear-overcast blend - giving a pure CIE overcast sky for this day - delivered mostly accurate predictions, the reason for the relatively poor performance of the Perez model is perhaps revealed in Figure 1. The measured fully overcast sky (label 326_92_11h00) shows a zenith luminance approximately three times that of the horizon, i.e. very similar to the CIE overcast formulation. In contrast, the overcast pattern generated by the Perez model exhibits a smaller ratio - the zenith luminance is approximately two times that of the horizon. Since the model skies were normalised to the same diffuse horizontal illuminance, the effect of the smaller ratio for the Perez overcast condition is to elevate the low altitude sky luminance relative to the CIE overcast sky. This in turn will lead to a relative over-prediction of internal illuminances using the Perez compared to the CIE overcast model, since the ‘view’ of the luminance pattern through the window (i.e. that at low altitude) is the most significant contributor to internal illuminance. This seems to be a general characteristic of the Perez model when used to generate luminance patterns for heavily overcast sky conditions. This feature of the Perez model may have resulted

from the calibration dataset if ‘overcast’ skies in Berkeley are systematically different from those recorded in more mainland locales. This seems plausible as Berkeley has a Mediterranean climate with a high occurrence of foggy conditions.

5 Conclusion

The results presented here have demonstrated how sky models can be evaluated based on predictions of internal illuminance. 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 noticeably the worst. This is perhaps to be expected since the sky model blends were ‘tuned’ to the validation dataset whereas the Perez model was not. However it is possible that the Perez model is also site specific to some degree since the model coefficients were based on fits to sky luminance data recorded at one location [4].

Note that the actual blended part of the overcast-intermediate blend is more physically plausible than that for the overcast-clear blend. At least for the half-and-half blend shown in Figure 2 where the overcast-clear blend has the unlikely combination of a slight zenith ‘hump’ (contribution from the overcast sky) and horizon brightening (contribution from the clear sky). Nonetheless, the overcast-clear blend performed marginally the best suggesting that either the validation dataset is comprised more of two distinct populations of largely overcast and largely clear skies with a relatively small number of skies occupying the blended region, or that the effect has little significance for the purpose of predicting illuminance.

Are the sky model blends described here ‘fit for purpose’? As noted, the often stated purpose of sky models is to generate luminance distributions for daylight prediction under realistic (i.e. climate-based) conditions. In other words, daylight prediction under realistic non-overcast (with sun) conditions as well as overcast skies where appropriate. It is proposed that the only meaningful daylight analysis is one that is carried out using climatological datasets, i.e. 365 days of hourly values of irradiance (or illuminance) determined from measurements taken at or near the site of interest (e.g. city). A period of a full year at short time-step is needed to capture in the simulation the full range of both short-term and long-term (i.e. seasonal) variations in the sky and sun conditions. The pattern of hourly values in a climate file is unique and, because of the random nature of weather, it will never be

repeated in precisely that way. However, climate files are representative of the prevailing conditions at the site, and they do exhibit the full range in variation that typically occurs. Consideration of any period shorter than a full year would select only some of the possible range in sky and sun conditions, and introduce unpredictable biases into the analysis. Furthermore, it is not possible to make meaningful daily or monthly averages of daylight quantities because they typically exhibit large changes in direction and magnitude throughout the course of the day. This being so, there seems little, if any, value in examining illumination under realistic sky/sun conditions for just a handful of snapshot instances since only a small part of the actually occurring range in conditions would be sampled. Thus for quantitative daylight modelling, the only option appears to be an annual simulation using realistic, climate-based sky and sun conditions.

Climate-based modelling of daylight has been demonstrated by a number of researchers [11][12], and shown to be effective in practical application [13]. For the prediction of annual daylighting profiles described in those two papers, Mardaljevic used the overcast-intermediate blend described here, whereas Reinhart used the Perez model. More recently, the CIE Standard General Sky has been proposed as a scheme that can reproduce the underlying luminance distribution for a wide range of naturally occurring sky types [14]. The answer to the question “which sky model(s) to choose?” is not straightforward. For climate-based prediction of illuminance, the two simple sky-blends performed only marginally less well than the benchmark Lumscan model which was based directly on measured sky luminance patterns. For those that favour simplicity in addition to faithfulness, the blend model may well be ‘fit for purpose’. The tuning of the sky blends needs to be repeated using other datasets and preferably for a larger number of skies. However it seems a reasonable proposition that the simple sky blends described here will be applicable to many other climates, possibly with the need of some ‘tuning’ that is locale-specific. The type and number of sky models that are used for a daylight coefficient based annual simulation should not affect the overall computational cost. The computation of the daylight coefficients is usually several orders of magnitude more demanding than the generation of, say, 4,400 (i.e. number of daylight hours in a year) sky luminance distributions (and the corresponding internal illuminances) which is usually an interactive process. Thus, computational overhead is no reason to favour simple sky model blends over more complex formulations. Rather, the issue is one of avoiding undue complexity where it offers no tangible (or testable) improve-

ment over simplicity. This should be a consideration when sky models are used in lighting simulation.

There is as yet no consensus on the selection of sky model types or the use of sky model blends for lighting simulation. This was the case when the number of commonly used sky model types was just a handful, i.e. far fewer than the fifteen types offered by the CIE General Sky. The large number of types offered by the CIE General Sky certainly does not simplify the matter of selection. However, a recent study by Tregenza suggests that a subset of the fifteen types will most likely suffice for most climates [15]. That study examined the probability of occurrence of the various sky types using data collected at the BRE. Tregenza found that:

five sky types account for nearly 80% of the scanned data sets; some types are rarely applicable or not used at all. This suggests that the daylight climate could be characterized by a small subset of standard types without significant loss of accuracy.

What is not yet clear however is how to select the most suitable sky type and/or blend on a time-step basis from the data contained in climate files. A recent paper by Dumortier and Kobay gives the first results for a possible solution in which the Perez All-Weather model is used as means to select suitable types from the CIE General Sky [16]. It is however early days and further work along these lines is needed.

Acknowledgements Sky luminance and room data were kindly provided by Paul Littlefair (currently), Maurice Aizlewood (formerly) and Howard Porter (formerly) of the UK Building Research Establishment. This research is presented in support of the efforts of CIE TC 3.37. The comments of a sharp-eyed referee are gratefully acknowledged.

References

- [1] P. Ineichen, B. Molineaux, and R. Perez. Sky luminance data validation: Comparison of seven models with four data banks. *Solar Energy*, 52(4):337–346, 1994.
- [2] P. Littlefair. A comparison of sky luminance models with measured data from garston, united kingdom. *Solar Energy*, 53(4):315–322, 1994.
- [3] J Mardaljevic. The BRE-IDMP dataset: a new benchmark for the validation of illuminance prediction techniques. *Lighting Research and Technology*, 33(2):117–134, 2001.
- [4] R. Perez, R. Seals, and J. Michalsky. All-weather model for sky luminance distribution—preliminary configuration and validation. *Solar Energy*, 50(3):235–245, 1993.
- [5] J. Mardaljevic. *Daylight Simulation: Validation, Sky Models and Daylight Coefficients*. PhD thesis, De Montfort University, Leicester, UK, 2000.
- [6] M. E. Aizlewood. Innovative daylighting systems: An experimental evaluation. *Lighting Research and Technology*, 24(4):141–152, 1993.
- [7] F Maamari and M Fontoynt. Analytical tests for investigating the accuracy of lighting programs. *Lighting Research and Technology*, 35(3):225–239, 2003.
- [8] F. Maamari, M. Fontoynt, A. Tsangrassoulis, C. Marty, E. Kopylov, and G. Sytnik. Reliable datasets for lighting programs validation—benchmark results. *Solar Energy*, 79(2):213–215, 2005.
- [9] S. W. A. Cannon-Brookes. Simple scale models for daylighting design: Analysis of sources of error in illuminance prediction. *Lighting Research and Technology*, 29(3):135–142, 1997.
- [10] J Mardaljevic. Verification of program accuracy for illuminance modelling: Assumptions, methodology and an examination of conflicting findings. *Lighting Research and Technology*, 36(3):217–239, 2004.
- [11] J. Mardaljevic. The simulation of annual daylighting profiles for internal illuminance. *Lighting Research and Technology*, 32(3):111–118, 2000.

- [12] C. F. Reinhart and S. Herkel. The simulation of annual illuminance distributions - a state-of-the-art comparison of six radiance-based methods. *Energy and Buildings*, 32(2):167–187, 2000.
- [13] J. Mardaljevic. Examples of climate-based daylight modelling. *CIBSE National Conference 2006: Engineering the Future, 21-22 March, Oval Cricket Ground, London, UK*, 2006.
- [14] S. Darula and R. Kittler. CIE General Sky Standard Defining Luminance Distributions. *Proceedings eSim 2002, Montreal, Canada, September 11-13*, 2002.
- [15] P. Tregenza. Analysing sky luminance scans to obtain frequency distributions of cie standard general skies. *Lighting Research and Technology*, 36(4):271–279, 2004.
- [16] D Dumortier and M. Kobav. Use of the Perez all weather sky luminance model to obtain the frequency of CIE standard sky types. *Lux Europa, Berlin*, pages 137–140, 2005.

Sky model type	Percentage of predictions within:		
	$\pm 10\%$	$\pm 25\%$	$\pm 50\%$
Lumscan (measured)	67	96	99
Clear/Ovc blend	41	74	94
Int/Ovc blend	35	72	93
Perez	32	67	90

Table 1: Percentage of predictions that fall within $\pm 10\%$, $\pm 25\%$ and $\pm 50\%$ of the measured values for the reliable subset of the validation data

Figure 1: Illustration of measured (Lumscan) and modelled (Skymodel) sky luminance patterns

Figure 2: Sky model blends

Figure 3: Linear blending function

Figure 4: RMSE minima for clear/overcast and intermediate/overcast blends

Figure 5: Percentile plots for the absolute relative error in illuminance prediction

Figure 6: Time series comparison of internal illuminance predictions - clear sky day. The upper plot shows a 5 minute interval time-series of global horizontal illuminance (solid line) and diffuse horizontal illuminance (dashed line). The middle plot shows the relative error in the predicted illuminance at each of the six photocells where the sky luminance distribution was based directly on measurements from the sky scanner. Measurements were taken at 15 minute intervals, and different shades are used to identify the reliable ($\blacktriangledown$) and potentially unreliable ($\blacktriangle$) entries. The lower plot shows the relative error in illuminance predictions where the sky luminance distribution is now provided by sky models: clear overcast blend ($\blacktriangleleft$) and the Perez All-Weather ($\blacktriangleright$)

Figure 7: Time series comparison of internal illuminance predictions - over-cast day (otherwise the caption is the same as Figure 6)

Figure 8: Legend for Figures 6 and 7