of a full year are shown as (tiny) shaded rectangles arranged in a 365 (days of the year) by 24 (hours of the day) matrix. The diffuse horizontal irradiance is the energy from the sky that is incident on an unobstructed horizontal surface. The direct normal irradiance is the energy from the sun and circumsolar region incident on a surface that is normal to the direction of the sun. The shading in Figure 19.3 represents the magnitude of the irradiation with zero values shaded light-gray. Presented in this way it is easy to appreciate both the prevailing patterns in either quantity and their short-term variability. Most obvious is the daily/seasonal pattern for both irradiances: short periods of daylight in the winter months and longer in summer. The hour-by-hour variation in the direct normal irradiance is clearly visible, though it is also present to a lesser degree in the diffuse horizontal irradiance (i.e. from the sky). Of course, both diffuse and direct irradiances will, in reality, vary over periods shorter than an hour. However, the hourly datasets are the most generally available and they do exhibit much of the variation in conditions that might be expected (e.g. in the absolute magnitudes of the two quantities, the occurring sun positions, etc.). Furthermore, these standard datasets provide definitive yardstick quantities for modelling purposes.[2]

It is relatively a straightforward matter to generate sky and sun conditions from, respectively, the diffuse horizontal and the direct normal irradiance quantities. It is understood that it is impossible to recreate an actually occurring sky brightness pattern from a measurement of diffuse horizontal irradiance because a real sky will exhibit unique brightness configurations resulting from cloud patterns and so on. However, it is possible to achieve reasonable approximations to actually occurring conditions using theoretical sky models that generate idealized sky brightness patterns from the basic irradiance quantities found in climate datasets. The character of the sky brightness pattern (e.g. overcast, intermediate or clear) can be inferred from the relative values in the diffuse horizontal and the direct normal irradiances, and the sun position is calculated from the 'time-stamp'.

To recap, the theoretical models and the basic meteorological data exist to generate hour-by-hour descriptions for the sky and sun that are representative of actual conditions for the majority of locales in the developed world. Then how might we make use of these to assess solar access? One possibility is to make a quantitative measure of illumination for an urban setting at one or more times of the year using sky and sun conditions generated from the meteorological data. Unfortunately, this approach will give only a 'snapshot' of solar access, as indicated by the resulting illumination, for the times examined.

Whilst it can be informative to determine, say, a monthly average for a scalar quantity such as temperature, illumination is strongly dependent on the directional character of the incident light. Associated with every data value in Figure 19.3 are the solar altitude and azimuth which, of course, vary from hour to hour.[3] In terms of providing a basis for predicting solar access, the notion of 'average' days is less than useful because an 'average' sun position would give entirely misleading patterns of illumination. In fact, simple averages could result in sun positions that never occur in reality, for example sun position due south and at a lower altitude than that for the winter solstice. Illumination parameters, therefore, are generally not suited to manipulations such as averaging. Sub-sampling the meteorological dataset (e.g. taking only the first day of each month) will reduce the number of hours to consider. However, the action will inevitably introduce biases because equally valid but most likely quite different sky and sun conditions/positions would be excluded. Indeed, it seems that the only way to avoid the pitfalls of averaging or sub-sampling is to consider the meteorological dataset in its *entirety*. That is, all the hourly sky and sun conditions for a period of a full year. Only this can capture the full range of both the short- and long-term variations in the sky and sun conditions. A schema to evaluate solar access based on a full year's meteorological data is described in the next section.

The new schema

There are essentially two ways to base an evaluation of solar access on the illumination provided by all of the hourly meteorological conditions that can be generated from a Test Reference Year (TRY) dataset. One is to examine the time series of illumination, hour by hour, that results from all of the unique sky and sun conditions. Necessarily, this would involve the manipulation and analysis of large amounts of data, approximately 4000 values (i.e. daylight hours) for every location evaluated. The other, much more straightforward method, is to base the measure of solar access itself on the *cumulative* effect of the illumination that results from *all* of the unique sky and sun conditions. In other words, on a measure of all the light energy from the sun and sky that is incident on a surface over a period