

Locating the optimal position of wind turbines

Scenario

Wind turbines are seen as a good source of clean energy. After the initial cost of installation a turbine can be expected to last between 20 – 25 years. Modern turbines are typically between 25 – 80 metres high with a rotor diameter of up to 65m. They are large structures which can be seen for many miles.

For the best performance, wind turbines are typically placed in windy places such as the top of valleys, ridges and offshore.

Some people see wind turbines as “blots” on the landscape, considering them as noisy and unattractive structures. Despite being green energy they are not welcomed in areas of outstanding natural beauty such as national parks and due their immense size they can potentially be very hazardous for low flying aircraft and migratory birds.

With so many objections to consider, GIS is an ideal tool for helping a company to locate an optimal site which maximises the energy output whilst considering the wishes of local people and other organisations.

In this exercise you are a manager of a wind farm company looking to site three new wind farms. You will use a mixture of data to locate windy sites and a DEM to calculate a viewshed. A viewshed is a spatial layer indicating which areas are visible to a known location. From this you can identify which towns are able to see a new wind farm and with this knowledge you can tailor any advice given to these populations.

Data used in Exercise

This exercise is supplied with several datasets and you will use these to create a viewshed and answer some basic question about visibility. The supplied data are:

- DEM (a raster binary file)
- Wind speed data (a raster binary file)
- Towns (a shape file)
- Wilderness area boundary (a shape file)

All of these map layers have co-ordinates measured in unprojected degrees of latitude and longitude, based on the WGS1984 datum.

Setup

This exercise comes with zipped data (raster, vector and palette files), it is suggested that you unzip the data to a single location such as **C:\eLearning\Unit5\View**. The output you create can go into this directory.

You will also need to import the raster data for the exercise in order to be able to use it by using the ArcToolBox *conversion tools* under *to raster*

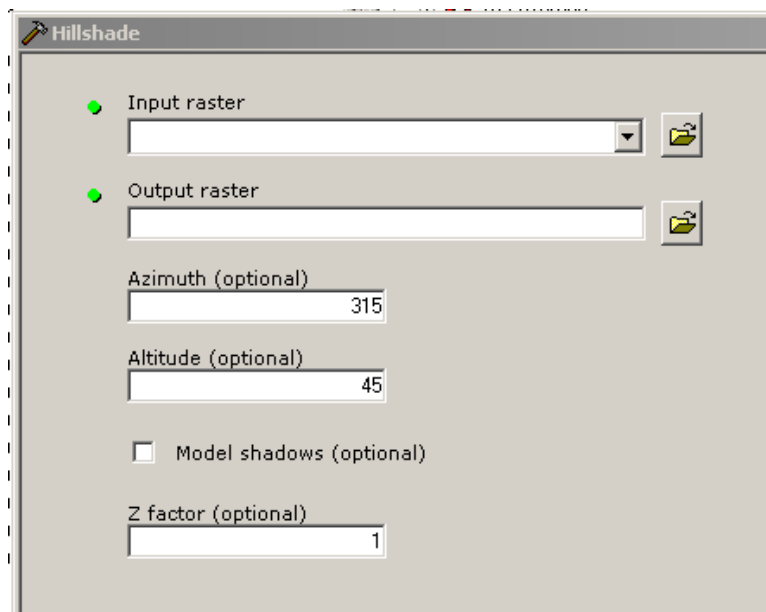
- use *float to raster* to import the **dem** raster map layer

Exercise

1.1 Create a hill shade of your DEM


We will begin by creating a hillshading image of our DEM. To do this:

- go to the ArcToolBox, choose *spatial analyst tools*, then *surface*, and select *hillshade*
- Under *input raster*, select your **dem** raster map layer and provide an appropriate *output raster* name, such as **hillshade**
- To calculate hillshade, you need to specify the location of the sun. This is normally specified as two angles. The *altitude* indicates the sun's position in degrees above the horizon, whilst the *azimuth* indicates the sun's position as a compass bearing. The default settings have the sun at 45° above the horizon and at a bearing 315° (i.e. to the northwest). In most circumstances, these settings provide an appropriate cartographic display.
- The *z factor* enables the effect of terrain to be exaggerated in areas of gentle topography. In effect, the *z factor* scales altitude relative to horizontal distance, so a *z factor* of 3 will triple height differences relative to reality. For now, leave this set to 1.
- By default, *hillshade* only considers the orientation (slope and aspect) of the terrain relative to the sun. It does not take into account the effect that a landform can have on its surroundings, such as the casting of shadows in a valley bottom. The *model shadows* option enables such effects to be modelled. This might be used in remote sensing studies of shadowing, but is typically not used for cartographic purposes, so leave this option unchecked.



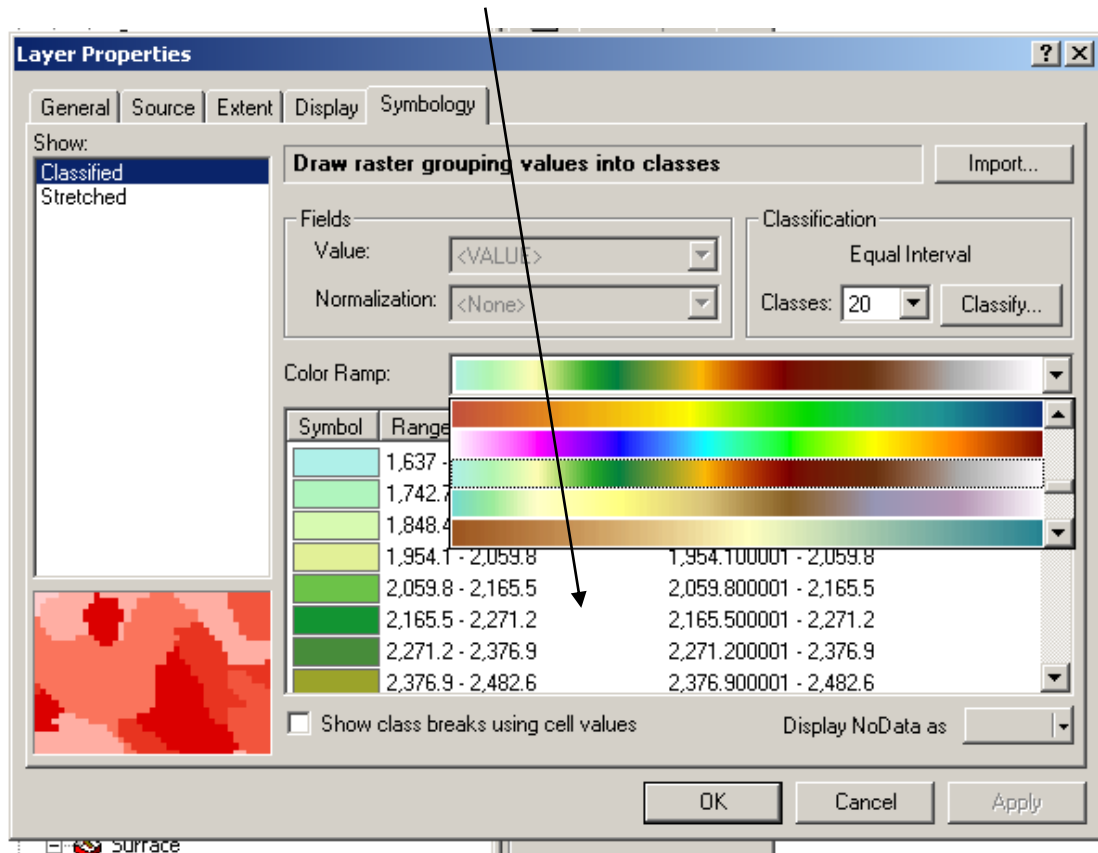
Once you have created your hillshade, you can combine this into a single, integrated depiction of the terrain as follows:

- display your **hillshade** raster grid, so that it appears at the top of the list of layers on the left of the ArcMap screen.
- **Transparency – in ArcGIS version 10 (see below for earlier versions of ArcGIS):** Go to the *windows* menu and click on *Image Analysis* to bring up the image analysis window
- Click on the name of your image in the left-hand panel of this window and then drag the transparency slider part-way across the screen so that it is set at around 60%.



- **[Transparency - in ArcGIS version 9:** right-click on your **hillshade** image and select *properties*, then the *display* tab. Set the *transparent* setting to **60%**, so that you can see through the pattern of shadow to the map layer below. Note that this display method will really only be effective if you use a colour scheme in shades of grey from white to black.]
- Right-click on your **dem** image and click the *symbolology* tab. Under *show* on the left, click on the *classified* option and under *classes* on the right, select a large number of classes. By clicking on the *classify* button, you can also set the method of classification to be *equal interval* (we assume here that you are already familiar with the use of colour in ArcGIS). The *color ramp* shown below is particularly effective at displaying elevation data.

- You may wish to change the number of decimal places shown in the legend. You can do this by right-clicking somewhere over the legend area of the screen and choosing *format labels*.



- Under *rounding*, if you select *number of decimal places* and then set this to **0**, this should prevent the elevation ranges being displayed with excessive precision.

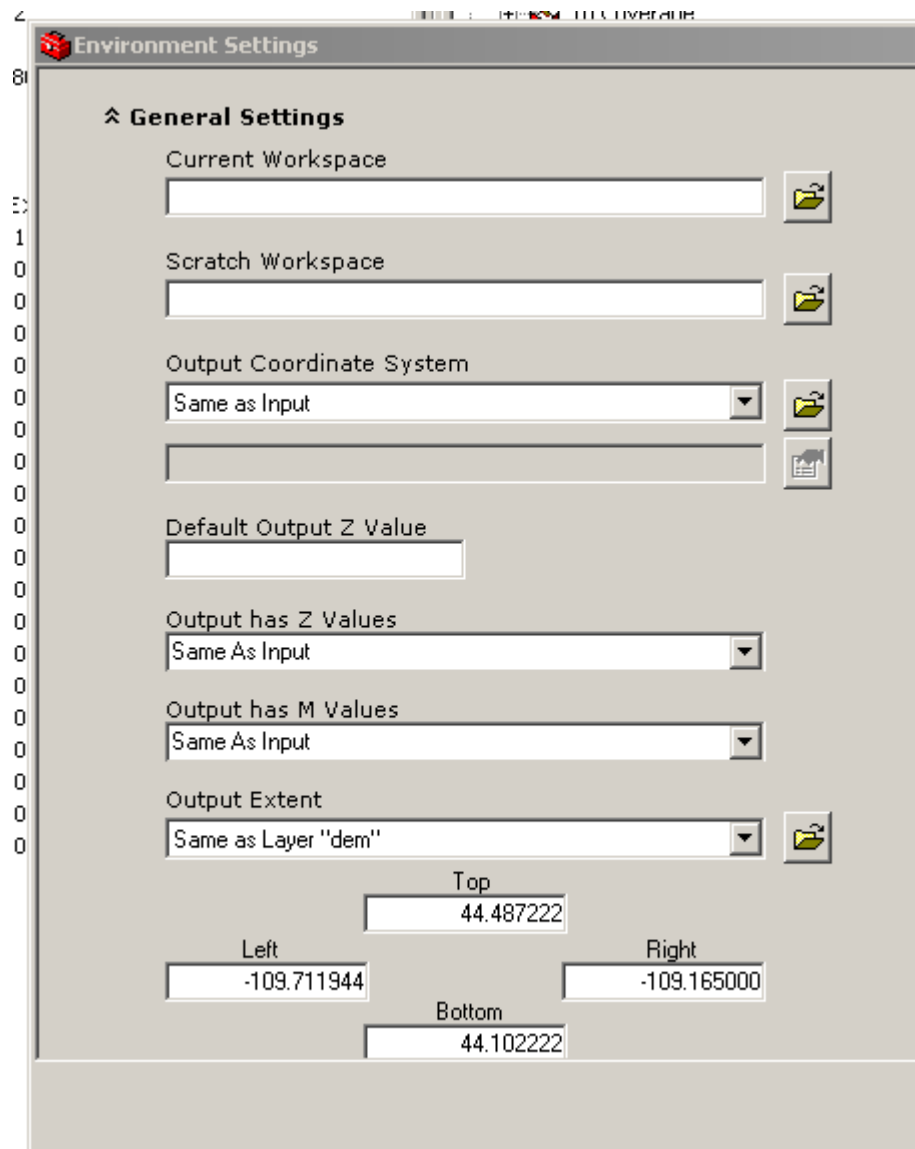
The result should look rather like the display shown below:



1.2 Interpolating a wind speed layer

We now need to identify those areas with high wind speeds. To do this, we will need to interpolate from the point measurements of average wind speed in the file **windspeeds**. To do this:

- Add the **windspeeds** shape file to your map display.
- In the ArcToolBox under *spatial analyst tools*, choose *interpolation* and select *IDW*. IDW uses the interpolation technique called inverse distance weighting (IDW).
- Choose **windspeeds** as the *input point file* and select *data_value* as the *z_value_field* (i.e. the field that contains the wind speed estimates that will form the basis for the interpolation).
- IDW works by calculating a local average wind speed for each output grid cell based on nearby measured wind speeds. Under *search radius settings*, the *number of points* indicates how many of the points in **windspeeds** will be used to calculate this average (**12**). The *search radius* is set to *variable* because ArcGIS will look over a different area for nearby points depending on the density of measurement points (in areas with a dense network of points, the *search radius* will be smaller compared to areas with a sparser network of points).
- Type in an appropriate name for the *output raster*, e.g. **wind_speed**.
- At the bottom of this screen, click on the *environments* button and then on *raster analysis*. This will determine the cell size of your output wind speed raster grid. Set this to be *same as layer 'dem'*.
- Similarly, click on *processing settings* and set the *extent* at the bottom of the screen to be *same as layer 'dem'* too.



Click on OK and ArcView should create a new raster grid depicting wind speeds across the study area. On top of this, add in the **towns** and **wilderness** map layers.

1.3 Locating the wind farm

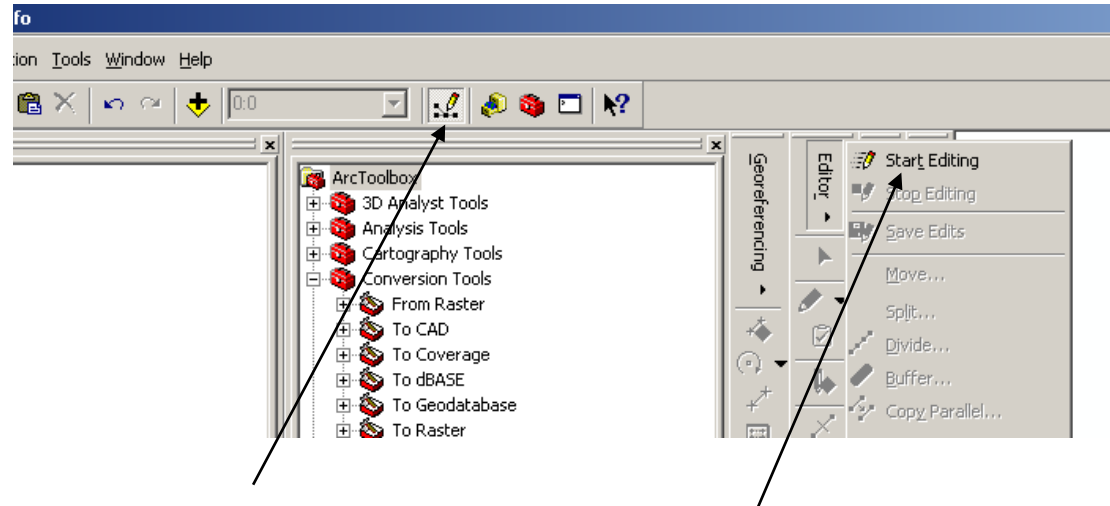
1. You are now ready to locate the three turbine sites. Each point you identify represents a single wind farm. A single farm can have more than one turbine but as a first past you are simply interested in identifying the general farm location.
2. Each wind farm you identify must be:
 - Outside the wilderness area
 - Located in the windiest place
 - Not near another wind farm

We will now create a new personal geodatabase containing a new point feature class of turbine locations:

- In the ArcToolBox, under *data management*, select *workspace* and select *create personal GDB*
- Choose an appropriate folder as an *output location* and then enter a name (e.g. **windfarm**) for your *output personal geodatabase*.
- Next, in the ArcToolBox, under *data management*, select *feature class* and select *create feature class*.
- Select as the *output location* the personal geodatabase that you have just created and enter an appropriate name for the *output feature class* (e.g. **turbines**)
- Set the *geometry type* to *point* and choose *OK*.

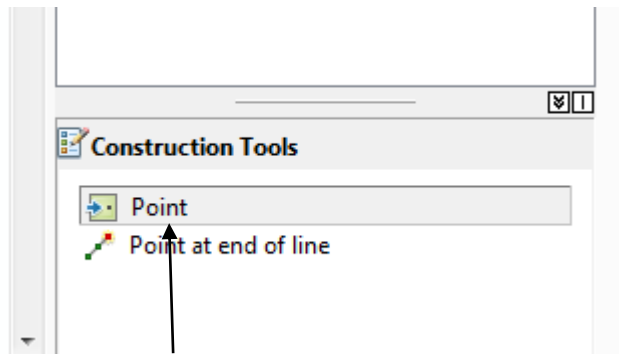
You should now see your new point feature class appear on the left-hand side of the ArcMap screen. We now need to add some points to this new feature class. To do this, we need to start editing the **turbines** map layer:

- Make sure that the editor toolbar is visible by clicking on the appropriate button (see below)
- Click on *editor* and choose *start editing* (note that this toolbar may appear in a different place on your screen, depending on how your computer is configured).



Press this button to view editor toolbar ...here is the toolbar

- When you start editing, ArcView will ask you which map layer you wish to edit. Choose the *windfarm* geodatabase and the *turbines* feature class.
- In **ArcGIS 10** [see instructions below for earlier versions of ArcGIS], the *create features* window will now appear to the right of your screen. Click on the name of your map layer in this right-hand panel and you should see the construction tools appear in the bottom-right corner



- Click on the *point* tool and then use this to start digitising within your map window.
- **[If you are using ArcGIS 9.3 or earlier:** To create a new point feature, make sure that on the editor toolbar, you have the *task* option set to *create new feature* (see below).]




- Digitise the locations of three separate sites for your wind turbines using this tool and making sure that you follow the guidelines described earlier.
- When you are finished, choose *stop editing* on the editor toolbar and save your edits.

1.4 Construct a viewshed

We are now in a position to calculate the area from which the new turbines will be visible. To do this:

- from the ArcToolBox, choose *spatial analyst tools*, *surface*, and then *viewshed*.
- Select **dem** as your input raster and choose your **turbines** as your *input point or polyline observer features*.
- Call the *output raster* an appropriate name, such as **viewshed**.
- The fairly daunting looking *refractivity coefficient* is used in specialist applications involving non-visible bits of the spectrum, such as siting mobile phone masts. We can ignore it here.
- Hit OK and you should now see the areas from which your proposed turbines are visible (note that this may take some time, as the calculations are fairly computer-intensive).

The resultant raster grid indicates how many of the input features (wind turbines) are visible from each grid cell. A value of zero means that the grid cell cannot be seen by any of the turbines, a value of one means that the grid cell is visible from one turbine, two means that the grid cell is visible from two turbines, and so on. You can see this by clicking on the viewshed layer using the 'identity' tool.



Note: It is possible to take into account the height of both an observer and an object superimposed on a landscape. You can add new fields to the attributes of the **turbines** that will be used in the calculation (by opening up the map layer's attribute table, and then selecting 'add field' from the button in the top left of this screen). A numeric field called OFFSETA can be added to include an observer height in the calculations (e.g. 1.6 metres, the typical height of a person), whilst a numeric field called OFFSETB can be added to include an object height (e.g. for the turbines, which can vary in height depending on capacity and circumstances). Other fields can be added to the attribute table to control other aspects of the calculation, such as the maximum distance or radius over which calculations are made for each feature.

Somewhat harder to handle are obstacles, such as buildings, vegetation and so forth. Although Digital Surface Models (as derived using LiDAR) will include object heights, some 'bare earth' elevation products – particularly those coarser spatial resolution products used over wider areas - will not include these objects. One solution is to 'extrude' a bare earth elevation product by using tools like the 'math / trigonometric / plus' tool within Spatial Analyst to add on approximate object heights for buildings, tree stands, and so forth.

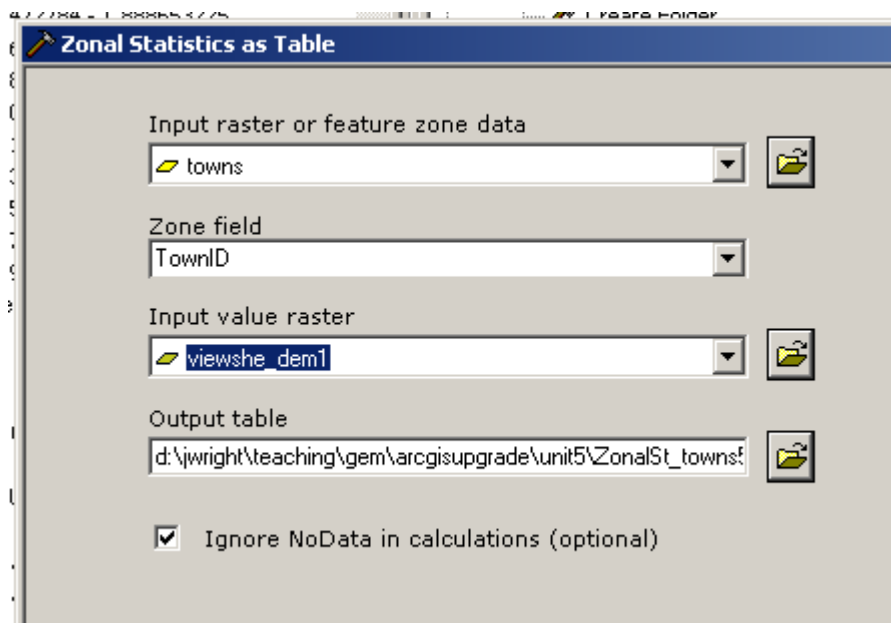
1.5 Which towns can see the wind farms and how much?

In the final part of this exercise you will introduce a layer representing towns in the region. You will identify which towns can see the wind farms and how much of the town is affected. *Be aware that your results will differ in detail from this example as your viewshed depends entirely upon where you located your wind farms.*

Have a think about how you might work out which towns can see the wind turbines before turning the page.

One way of finding out the towns from which the wind turbines are visible is to use ArcView's 'zonal statistics' feature:

- in the ArcToolBox under *spatial analyst tools*, select *zonal* and choose *zonal statistics as table*.
- Choose **towns** as the *input raster or feature zone data*. These are the zones for which you will summarise values from a raster image.
- **TownID** – the only integer (whole number) field in this layer's attributes – should automatically be selected as the *zone field*.
- Select your raster viewshed layer that you have just created as the *input value raster* and choose an appropriate name for the *output table*, such as **turbine_vis**



Click on OK and ArcView will overlay the town boundaries onto the viewshed raster grid and calculate summary statistics, based on the grid cell values in the viewshed raster. The results will be stored in a new table.

To view this table, you will need to open it up and then right-click on it and choose *open*.

What does the information in the table mean?

Each row in this table represents a different town (from our 'feature zone data' in the input screen). The first column, **OID** (see illustration below), is simply a unique serial number for each town (town 0, town 1, town 2, and so on). The second column, **value**, contains the values of **TownID**, copied over from our original map layer of towns that we fed into Viewshed when we specified a 'zone field'. Based on this logic, the row with Value = 2 represents the town in the southeast corner of our towns map layer with **TownID** = 2, whilst the row with Value = 5 represents the town up in the northwest part of the study site.

OID	VALUE	COUNT	AREA	MIN	MAX	RANGE	MEAN	STD
0	2	17112	0.001320	0	0	0	0	0
1	3	10913	0.000842	0	2	2	0.533309	0.840432
2	4	12161	0.000938	0	2	2	0.528986	0.623132
3	5	9623	0.000743	0	1	1	0.060376	0.238182

Moving along the columns, **count** tells us how many grid cells in our viewshed raster grid lay within the boundary of that particular town. The **area** column is an estimate of the area of the town, though in our case the numbers here are areas in squared degrees of latitude and longitude, so should not be considered as meaningful (this is because our original DEM's reference system was measured in degrees of latitude and longitude).

The relevant information for us comes in the following set of columns. These are a summary of the numbers stored in the raster viewshed for each town. '**Min**' is the smallest number stored in any of the grid cells that lay inside the boundary of a particular town, whilst '**Max**' is the largest number in any of these grid cells. Looking at the first row in the example shown above, you should be able to see that the town with **value** = 2 (i.e. the one in the southeast corner) had a value for both '**Min**' and '**Max**' of 0. Since 0 means 'not visible', this means that *none* of the turbines could be seen from anywhere in this town. In contrast, looking at the town with **value** = 5 (i.e. the one up in the northwest of the study area), the column '**Max**' contains a 1. This means that in part of the town, 1 of the turbines was visible. In the remaining two towns, the **max** column tells us that 2 of the turbines were visible in at least parts of both towns. N.B. If your table looks different to the example above, do not be alarmed. You are likely to have placed your wind turbines in slightly different positions when drawing them on-screen.

The other columns of information – **range**, **mean**, and so on – contain other summary measures of the values of our viewshed raster for each town (e.g. the **mean** column contains the average number of turbines visible within each town, averaging across all the pixels within the town's boundary). However, for our purposes here, they are less useful than the **max** and **min** columns.

This final calculation may seem excessively complicated. A simpler alternative, for example, might be to overlay the towns map layer on top of our viewshed raster and simply look at how visible each town is by eye. The

advantage of using 'zonal statistics', however, would come if we were looking at many different locations (e.g. several hundred houses, not just 4 towns) and wished to automatically calculate inter-visibility for all of them using ArcView.

Additional background

RenewableUK

Web link: <http://www.renewableuk.com/>