

Hydrological modelling of surface run-off using DEMs: an Introduction

Scenario

Hydrological calculations involving run-off are widely used in many environmental management applications. They can be used to understand pollution hazards and their impact on the landscape, such as diffuse pollution from agro-chemicals or point sources of pollution, such as outlet pipes from factories. They are often also used in understanding sedimentation, which can take place during the construction phase of a project and impact the environment. They can also be used to characterise landforms.

In this exercise, we will examine how some of the hydrological tools in ArcGIS work, taking the example of the valley around Broad Chalke in Wiltshire. This valley lies within an Area of Outstanding Natural Beauty (a UK conservation designation) and has numerous chalk streams (<https://www.wwf.org.uk/where-we-work/uk-rivers-and-chalk-streams>) of high conservation value.

Data used in exercise

For this exercise, we have provided you with a 2 metre resolution Lidar-derived DTM, which was downloaded from the Environment Agency web site in geotiff format:

<https://environment.data.gov.uk/DefraDataDownload/?Mode=survey>

Question A: The Environment Agency web site has both Digital Surface Models and Digital Terrain Models. We used a Digital Terrain Model for this exercise. Why? [see last page for answers]

Setup

Unzip the files for the exercise, then start ArcGIS Pro and create a new map, setting up an appropriate working folder and home geodatabase in doing so. Next, add the DTM to your map display. You may be asked if you would like to *Build pyramids* when you open the raster: choose *yes*. Pyramids are a way of optimising the display of a raster. They involve storing pre-calculated, coarser resolution versions of a raster alongside its native grid resolution. The pyramids make it faster to display the raster when you are zoomed out. You may also wish to change the default colour ramp for the DTM to one that is designed for the display of elevation.

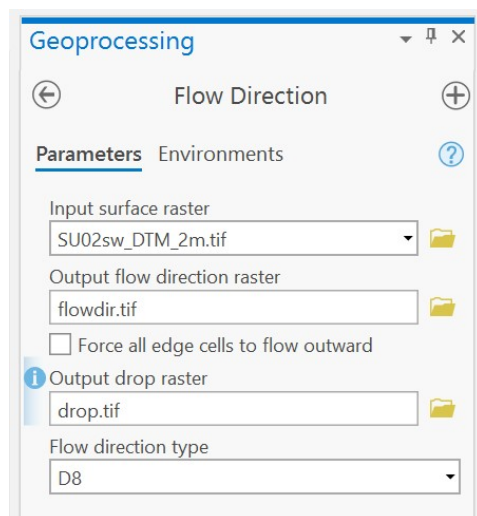
Finally, as we will be working with raster data, you will need to activate the *Spatial Analyst* licence if you have not already done so (Via the *project* menu and *licensing*, then *configure licencing options*).

1.1 Creating a hydrologically conditioned DTM

When using hydrological tools, the first step in working with an elevation model is *hydrological conditioning*. This involves checking the DTM for sinks – areas which are pits, surrounded by higher elevation values. Whilst these can occur naturally, often they result from problems with the input DTM. Bridges or misrepresentation of narrow-sided channels can all result in sinks. Sinks will affect the calculation of hydrologically derived surfaces, so will need to be identified and removed as a first step.

To identify sinks, first head for the *Analysis* menu and choose *Tools*, then search for ‘flow’. Find the tool *Flow Direction* and run it:

- Choose our DTM as the input surface raster

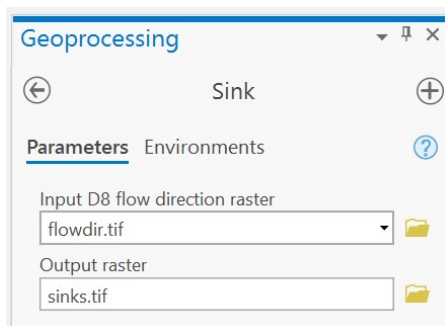


- Choose a name for an output flow direction raster (I normally add a .tif extension to save the output in geotiff format, storing it in a folder rather than geodatabase. This is because ArcGIS can be very ‘fussy’ about folder and file names when using its own native raster format. The format can also be read by other packages such as QGIS). This surface will indicate the direction of steepest slope from each raster pixel and thus surface water flow.
- There is an option to *force all edge cells to flow outward*, though we do not need that here.
- You can also specify an *output drop raster*, which gives an indication of the steepness (as a %) of the flow gradient. Although not required subsequently, I have specified one here, so that we can see what the ‘drop’ looks like across this study site.

- Finally, the *flow direction type* indicates how flow direction will be calculated. The most straightforward algorithm is the *D8* algorithm, which calculates the direction of flow from each pixel to one of its eight neighbouring pixels. The *DInf* algorithm allows for a more nuanced depiction of water flow, whereby flows may be split across more than one pixel, where direction of flow merits this. Since many subsequent tools require D8 flow directions as input, we will use these for now.
- Hit run, then explore the values for flow direction.
- Note that with flow direction, the values either indicate the direction that water flows in (as a compass bearing from 0 to 359 degrees), or with a well-functioning D8 flow direction algorithm, the output values indicate one of eight flow directions (1=east; 2=southeast; 4=south; 8=southwest; 16=west, etc – see: <https://pro.arcgis.com/en/pro-app/tool-reference/spatial-analyst/flow-direction.htm>). Sinks may be assigned different values with the D8 algorithm however, where a DEM has not been hydrologically conditioned.

Now we are in a position to identify our sinks. Head back to the geoprocessing panel, and this time search for and then run 'sink':

- Use your flow direction raster as the input and choose an appropriate output raster name

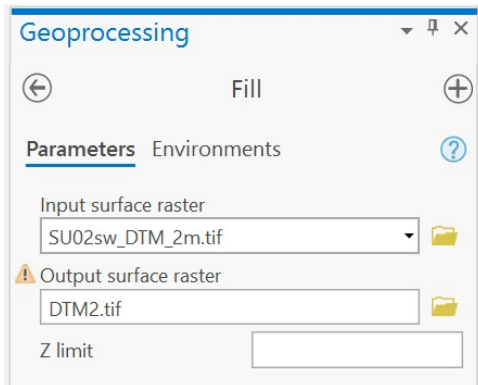


The *Sink* algorithm will identify 'pits' or 'sinks' in the DTM, places where pixels with low elevation values are surrounded by pixels with higher elevation values. Every time the software identifies a patch of such pixels, it gives each patch a unique ID. Run the tool and explore the output layer. You may need to zoom in to see individual 'sinks' and turn off the visibility of other layers. You should see that our DTM does indeed contain some sinks.

ArcGIS Pro also contains a tool for removing sinks, by artificially increasing the elevation values within sink pixels until they are raised to the value of the lowest of their neighbours. This is 'hydrological conditioning'. To run the tool, head back to the geoprocessing panel again and search for, then run 'Fill':

- For this tool, you will need to specify our DTM as the input surface raster and choose an output name.

- Optionally, the *Z limit* sets a maximum difference in elevation (between a sink pixel and its neighbours), above which the sink should not be filled. Normally, this is not needed and can be left blank.



If you run this tool, you should find that you have a very similar-looking DTM output, but with the sinks filled in.

To check that the process has been successful, run the *flow direction* tool again on this layer with the D8 algorithm, creating a new output layer in doing so. Then run the *sink* tool again, using your new flow direction raster as the input. If the process has worked (it did for me), you should find that the 'sink' tool outputs a raster where the minimum and maximum values (shown in the left-hand table of contents panel) are extremely small, rather than being integers (whole numbers). You can remove this layer from your list of layers for display (right-click on it in the left-hand panel and choose remove) if that is the case. Had the process not removed all sinks, we might have to re-run 'fill' again.

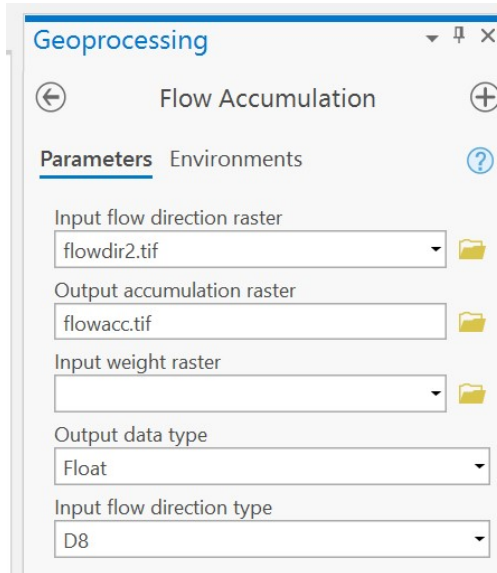
1.2 Delineating stream networks

Now that we have a hydrologically conditioned DTM, we can undertake various analyses. One option is to automatically delineate streamlines and other surface water features. To do this, head back to the geoprocessing panel again, and search for 'flow' once more. This time, choose and run *Flow accumulation*:

- You will need to specify an input flow direction raster. Be sure to input the second one that you produced, from the hydrologically conditioned DTM.
- Choose a name for the output layer that will be produced. For each pixel, this layer will store a count of the number of other pixels that lie upstream of it, the 'flow accumulation', a measure of the area upstream of each pixel. This will be zero for ridges and peaks, but large for streamlines where water accumulates.
- We will not use *input weight raster* here, but it is a useful option. If we had a layer depicting say amounts of nitrogen fertiliser deposited on


fields, instead of counting upstream pixels, the software would sum these fertiliser amounts for each pixel lying upstream of the target pixel. It would output the total amount of fertiliser deposited on land upstream of each pixel therefore. Rainfall values can be used as weights in the same way, as can hazardous land uses (e.g. landfill sites), so this can be a very useful option in pollution control applications.

- Finally, we also need to specify the nature of the algorithm used to calculate the input flow directions, and indicate the type of numbers that we wish to store as the output. *Float* refers to numbers with fractional parts (which we will use here); *integers* are whole numbers.



When you run the tool, the flow accumulation output can sometimes be hard to visualise. One way of making it easier to see patterns is to right-click on the input layer (or highlight the layer in the left-hand table of contents, then click on the *appearance* menu and choose *symbolology*) and choose *symbolology*. If you then choose *classify* from the dropdown in the symbolology panel, you should be able to select *equal interval* as the method of display and increase the number of classes from **5** to say **9**. You may also need to zoom in to see how the flow accumulation values are changing.

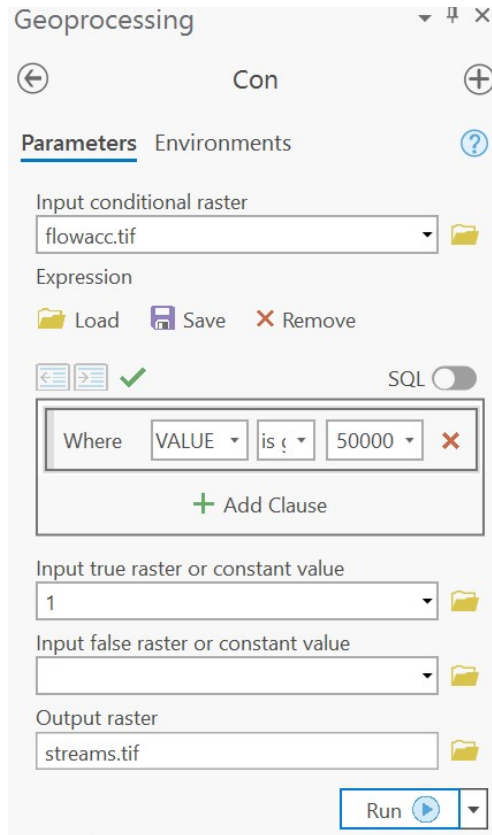
In the example screenshot below, high flow accumulation values in part of our landscape are shown in red, with low values shown in yellows and oranges. Several drainage channels or streamlines can be seen in red. Try clicking around with the *explore* tool (on the *Map* menu ribbon) to see what values have been stored in the 'flow accumulation' output:



To delineate streamlines, a common step is to decide on a threshold value for flow accumulation, then classify everything above this value as a stream. If the threshold value is high, then only major river channels will be identified, as it is lowered, then small seasonal channels will be identified as well as perennial channels. Based on exploring the flow accumulation raster values, I have chosen a value of 50,000 for this (i.e. 50,000 2 x 2m upstream pixels of flow accumulation constitutes a streamline). You may wish to experiment with other values.

To illustrate this process, go back to the geoprocessing panel again and search for 'con', then run the 'con' tool. This will identify those pixels in an input raster where a given condition (or set of conditions) is true or false:

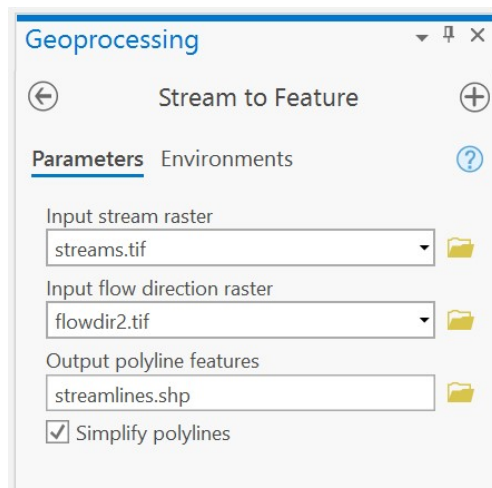
- As the *input conditional raster*, choose your flow accumulation layer.
- Next to 'expression', press the 'add' button, then use the dropdowns to set this to where VALUE (meaning the pixel value field in the flow accumulation raster) IS GREATER THAN, then enter **50,000**.
- For *Input true raster or constant value*, type in **1**. This is the value that the output raster will contain where flow accumulation is greater than 50,000. The *con* tool is very powerful, so whilst we do not need it here, we could specify a raster here instead of the constant value **1**. If we did this, then where the flow accumulation was over 50,000, the output layer would take the values of that raster.
- Leave the *Input false raster or constant value* to blank. This means where the condition is false (flow accumulation is not greater than 50,000), then the output raster will contain null (blank) values. Alternatively, whilst we will not use it here, we could type in say 0 here, in which case pixels with low flow accumulation would be coded as 0.
- Finally, choose a name for the output raster, e.g. 'streams':



When you press 'run', you should hopefully see a new raster layer that depicts what look like stream lines across our study site.

ArcGIS also provides us with a tool for converting our streamlines from raster to vector. To do this, head for the 'geoprocessing' panel once more and search for 'stream', then run *stream to feature*:

- As inputs, choose your output from *con* as the input layer.
- Choose your flow direction raster (the second one, based on the hydrologically conditioned DTM). This will enable the start and end points of polylines to reflect the flow direction.



- Choose an output polyline name. I have saved mine above as a shape file by adding the .shp extension and saving it within a folder rather than geodatabase. You may prefer to use a geodatabase.
- It is a good idea to check the *simplify polylines* option. This will remove unnecessary x, y coordinate pairs from the output line, that are not needed to describe its shape (e.g. because they are halfway along a straight section of stream) through a process called map generalisation.

If you run the tool now, you should be able to generate some surface drainage lines in vector format.

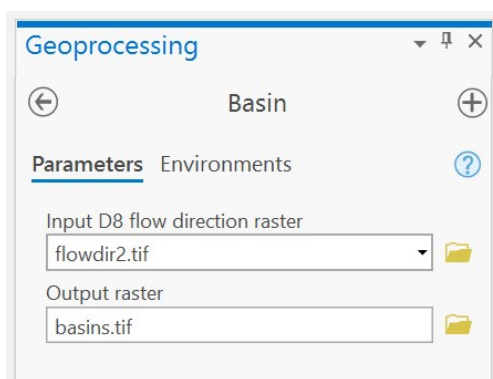
Extensions:

- You could try experimenting with different flow accumulation threshold values for 'con' here to see what impact they have on the resultant stream network.
- It is possible to generate even more understanding of an area's drainage network through these techniques. Take a look at the *stream link* and *stream order* tools. What do they do?

1.3 Delineating drainage basins

It is also possible to delineate drainage basins using flow direction. Head for the *geoprocessing* panel once more and this time search for *Basin*. Run this tool:

- Choose your second flow direction raster as the input layer (the one from the hydrologically conditioned DTM)
- Choose a suitable output name



By following flow direction, the software will now identify patches of pixels which all drain in the same direction – drainage basins (in US terminology) or catchments (UK terminology). Each separate catchment will be given a different unique numeric ID in the output layer. Note that the delineated drainage basins are somewhat dependent on the extent of your raster. At the

edges, catchments that may connect outside our study area extent may be given different identifiers.

Extension:

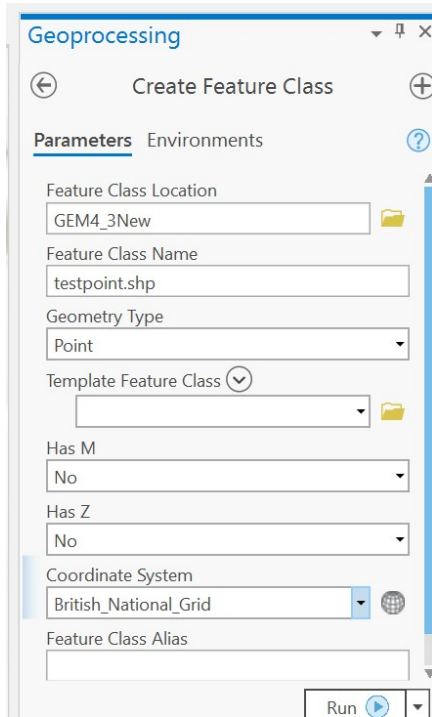
- By searching the geoprocessing panel for tools once more, you may also wish to try calculating *flow length*. This is the longest distance water will flow, either upstream or downstream, from a given pixel. It can also be constructed from flow direction. This measurement is widely used in understanding erosion and sedimentation, for example in its upstream form as the parameter L (slope length) in the Universal Soil Loss Equation (see <https://milford.nserl.purdue.edu/weppdocs/overview/usle.html>).

1.4 Hydrological processing of cloud-based DEMs

ArcGIS Pro also supports the processing of DEMs held in the cloud (i.e. on a remote server, rather than on a local drive), rather than locally via its *Ready to use* tools. We will briefly explore this functionality here too.

Note: you need to have an ArcGIS Online account to run 'ready to use' tools – you will only be able to do this part if you have such an account.

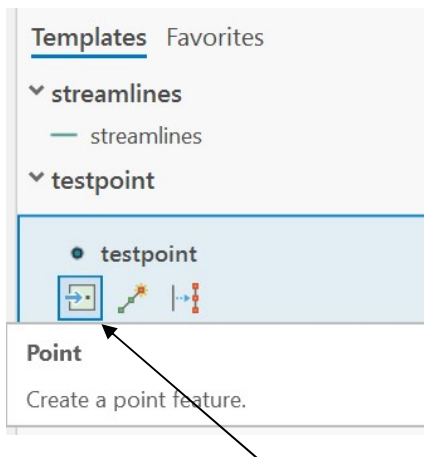
To do this, first we will need to create a new, empty point vector layer, then digitise a sample point into this layer. To do this, head for the right-hand catalog panel, then right-click in an appropriate location and choose *new* and either shape file, geopackage or feature class (depending on your preference – I have gone with a shape file below):



- Set the *coordinate system* to be the same as the *current map*, which should be British National Grid coordinates.
- Set the geometry type to be *point* rather than multipoint (note: multipoint refers to a group of points that share a single set of attributes, e.g. a business that has three different premises, but only one owner or name).
- Set *Has M* and *Has Z* to be No (M means 'measure coordinates' – specialist coordinates that describe locations along lines; Z means include elevation as well as X and Y coordinates in describing point locations).

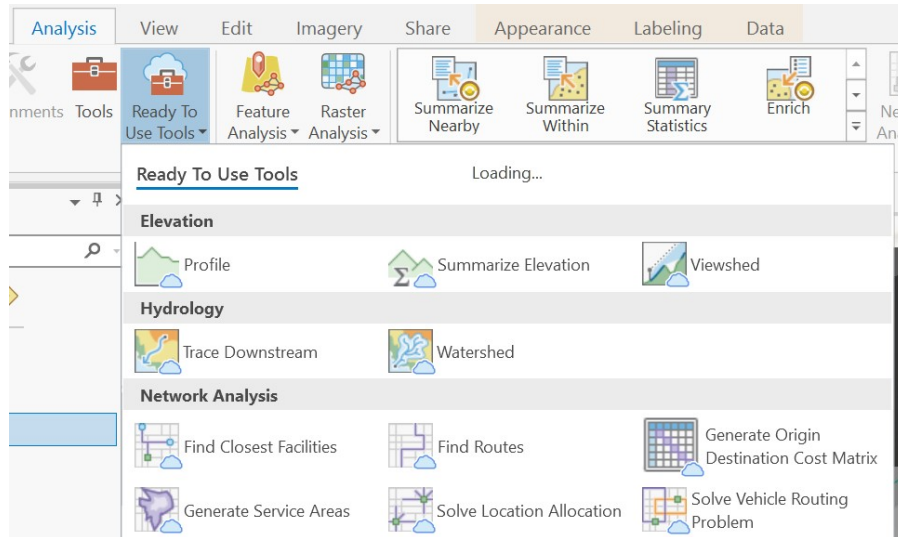
Once we have created our new layer, head for the *Edit* menu to digitise a new point within it, then hit the *create* button on the ribbon there.

You should then see the create tools in the right-hand panel:



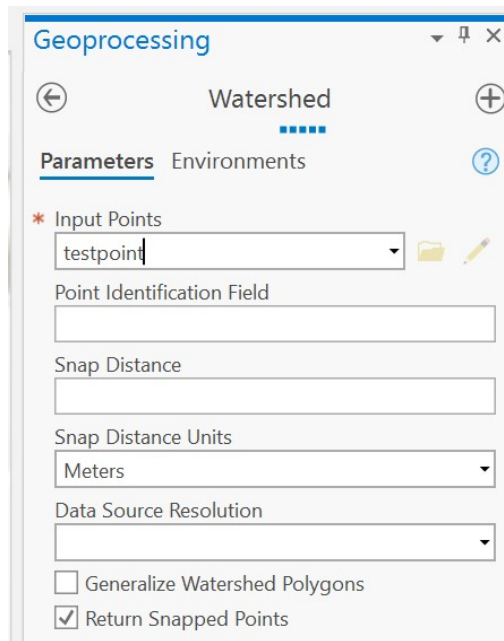
Choose the *create point* tool, then use this to digitise a point near or on one of your streamlines. When you have digitised your point, press *save* on the ribbon.

Now we are ready to run the 'ready to use' tools, so head for the *Analysis* menu, then choose *Ready to use tools*:




Try running the *Watershed* tool, which will work out the area upstream of our digitised point, then return this area as a polygon layer. When you run the tool, if you are not already signed in to ArcGIS Online, you will need to do so now (the tool requires some analysis credits to run).

Try running the tool:



- As input points, choose the new point layer that we have just created.
- We do not need it here with just one input point, but optionally, if we had multiple points, we could specify a *point identification field*. This field would be used to assign IDs to output polygons, so we would know which area was upstream of which point.
- Optionally, we can also specify a *snap distance*, though we do not need to do so here. This would be the maximum distance our input points would be moved ('snapped') to align with a river network held in the cloud by ESRI. Checking *return snapped points* means that the tool will return a layer showing the adjusted locations of the input points.
- Leave the *data source resolution* blank for the default value for now. This means that the software will use 3 arc second (90 metre) resolution SRTM DEM data, held in the cloud by ESRI. We could alternatively set this to **finest**, which would mean the highest available resolution of DEM data held by ESRI would be used for our study area.

The tool may take a while to run. When it finishes, you should see something like this:



The tool shows us the polygon upstream of our point, though with quite coarse boundaries, reflecting the 90 metre cloud-held DEM used to delineate it. Had we checked the *generalise watershed polygons* option, the stepped boundary of the catchment (reflecting the 90 metre pixels) would appear smoother. We can also see how the input has been moved towards a river channel in this coarse resolution cloud-held DEM.

Note that there is a very similar tool called *Watershed*, which you could run with a DEM held locally, rather than one in the cloud.

Extensions:

- You may wish to explore the other ready to use tools for elevation in the group, such as *trace downstream* and *viewshed*.

Questions:

- What environmental applications could you see for tools such as *trace downstream* and *watershed*?
- What kind of operations do ESRI typically implement as ready to use tools in ArcGIS Pro?

Exercise

Answers:

- A. A Digital Surface Model would include objects such as vegetation. For hydrological modelling, we need a 'bare earth' model – a Digital Terrain Model – from which such objects have been excluded.
- B. There are obvious pollution control applications of these tools. For example, what areas might be impacted by an accidental pollutant release event? Where we have detected a pollutant at a monitoring point, which areas should we investigate to identify the potential source of the pollutant? The tools could also be used for applications such as invasive species modelling and control (e.g. where might seeds or rhizomes from an invasive plant such as Japanese knotweed be dispersed through surface flow in the landscape?)
- C. 'Classic' ready to use tools are ones where you hold very simple data locally (e.g. points representing pollutant spill events), then these are combined with data that are much more complex to prepare for analysis. For example, here we had to hydrologically condition our own DTM before analysis. With the 'ready to use' tools, this has been done for us. Other ready to use tools use complex road network data with traffic parameters in a similar manner.