GEM: Spatially Distributed Dynamic Models

Objectives:

This practical has the following objectives:
1. To introduce you to the automation of tasks in ArcGIS Pro
2. To develop understanding of dynamic spatial models
3. to develop understanding of how GIS can be used to solve practical problems (in this case in climatology and soil science)

Scenario:

The Problem: Understanding tropical rainforests and climate

Understanding how tropical rainforests contribute to the global carbon cycle is crucial to climate change modelling. Over the years, climatologists and ecologists have sought to understand the fluxes of carbon moving between the soil, biomass and into the atmosphere as carbon dioxide. Recently, ecologists have suggested that carbon dioxide fluxes from drought-stressed tropical rainforests may be much greater than those from tropical rainforests that are not drought-stressed.

Soil water deficit is commonly used as a measure of drought stress. Unfortunately, virtually no meteorological stations in the tropics measure soil moisture directly. However, there are many meteorological stations that measure precipitation in the tropics. Is it possible that we can map changing patterns of soil water deficit using information about precipitation?

The Model:

Researchers have investigated the relationship between soil water deficit and climate in the field. From their research, it appears that the soil water deficit in the Amazon basin can be approximately estimated as follows (Malhi and Wright, 2004):

$\text{Soil Water Deficit}_{(t)} = 0.8466 * (0.8188 * \text{Soil Water Deficit}_{(t-1)} - \text{Precipitation}_{(t)} + 118)$

Soil water deficit can only be a positive number, so any time that this equation produces a negative number (e.g. when precipitation is very high), the soil water deficit is reset to zero.

Notice that Soil Water Deficit appears twice in this model: once on the left-hand side and once on the right-hand side. The subscripts (e.g. Soil Water $\text{Deficit}_{(t)}$ and Soil Water $\text{Deficit}_{(t-1)}$) tell us which time period we're talking about. Soil Water $\text{Deficit}_{(t)}$ means 'soil water deficit this period', whilst Soil Water $\text{Deficit}_{(t-1)}$ means 'soil water deficit last period'. In other words, this equation suggests that soil water deficit **this month** depends on soil water deficit **last month**.

The Data:

The following map layers are available:

- **prec**: The data provided are part of a bigger set of 36 different precipitation raster maps, representing long-term average monthly precipitation in mm for 1996-98. Each month is indicated by the final digit or two of the name, e.g. **prec_01** = precipitation in January, 1996, **prec_02** = precipitation in February, 1996, **prec_03**, = precipitation in March 1996 etc. until **prec_36** = precipitation in December 1998. We will use just a few of these grids here – **prec_10** (October 1996) through to **prec_18** (Jun 1997).

Notes:

(1) The precipitation data are derived from the CRU TS 2.0 global data set (New et al., 2002; New et al., 2000).

GIS Practical:

Import and inspect your data

Create a new map template and have a look at the various precipitation layers for the different months in ArcGIS Pro. As we will be analysing the raster grids, make sure you activate the *Spatial Analyst* licence (*Project* menu / *Licencing / Configure your licencing options* button).

So far, it is likely that any models that you have designed in ArcGIS Pro have been created through the menus. Here we will consider how we can automate processes using the Model Builder facility in ArcGIS Pro.

At the end, as an extension, we will show you can create a model that processes not just data for a single period of time, but how to make our models **dynamic** and include time within them. We will now attempt this with the soil water deficit model described above.

Create a map for the start of the time series

To start off our time series, we need to set the soil water deficit to an initial level. Our data start at October 1996, so we need to create a grid depicting the soil water deficit as at September 1995 to 'get the ball rolling'. For simplicity, we'll assume that the whole of Amazonia starts off with an initial soil water deficit of 200mm.

To create this initial soil water deficit map layer:

- Head for the *analysis* menu and then *tools*, then search for 'create constant' and then run *create constant raster*
- Create a new raster grid called **initial_water** with the parameters shown in the screenshot below. In the screenshot below, the *output cell size* of 0.5 and the *X and Y minimum and maximum* values have been chosen to match those of our precipitation grids. The X coordinates are all negative, because the represent degrees of longitude west of Greenwich – degrees west and south are normally negative numbers. The minimum Y coordinate is negative and the maximum positive, because our study area straddles the Equator. The *constant value* of **200** represents the initial soil water deficit of 200 mm. When you are done, click on Run.

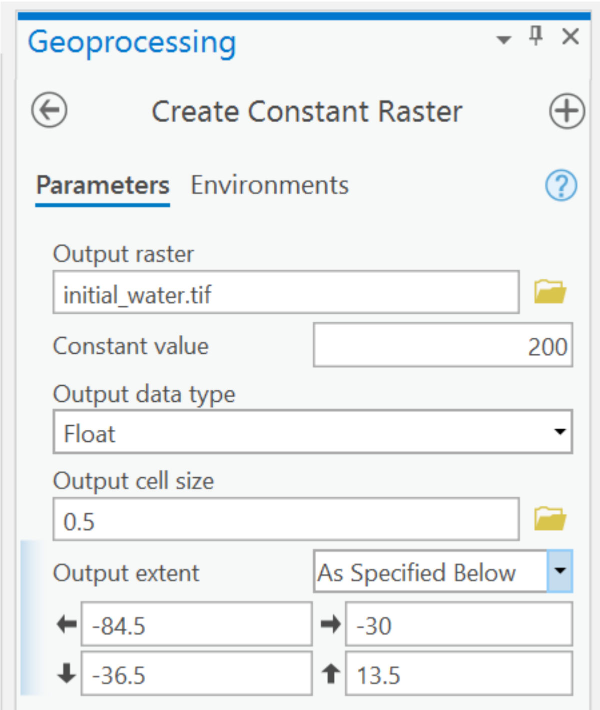

Design a dynamic model using Model Builder

Now that we have created this initial grid, we can start to build up a model of soil water deficit. To do this, we will use the Model Builder facility in ArcView. On the *analysis* menu ribbon, click on the ModelBuilder button:

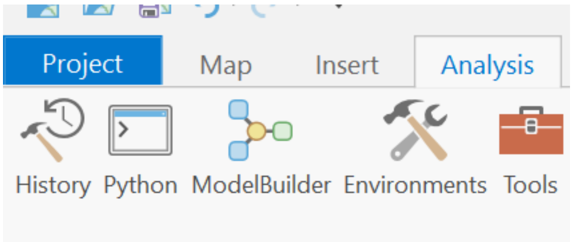

This creates a new window, into which we can add tools for automating processes.

We're going to start by doing the bit of the equation shown in bold and underlined:

Soil Water Deficit$_{(t)}$ = 0.8466 * (**0.8188 * Soil Water Deficit$_{(t-1)}$** – Precipitation$_{(t)}$ + 118)

Via the tools button, search for 'times' and then double-click on the *times* tool to add this to your Model builder canvas:

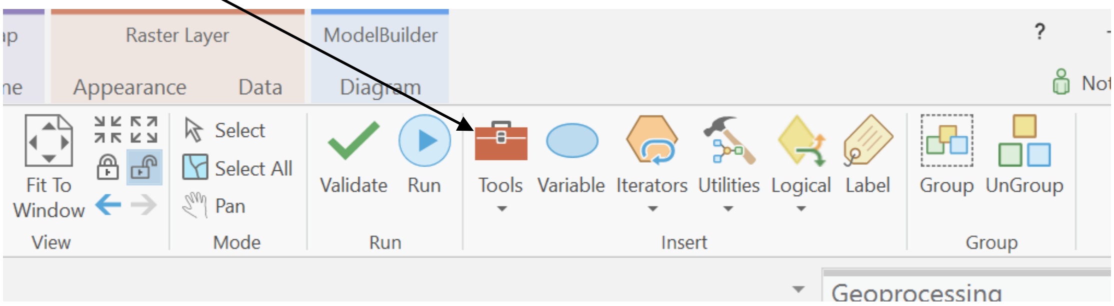

This tool will multiply 2 rasters together, or multiply a raster by a constant number. If you right-click on the tool and choose *open*, you can select our **initial_water** raster (the opening water balance for the first period) as the first layer. We can then then enter the constant from the equation above as a second input (0.8188), and specify a name for the output (I normally add a .tif extension, because I find ArcGIS is less 'fussy' over file and folder names when working with geotiff rasters, but this is optional):

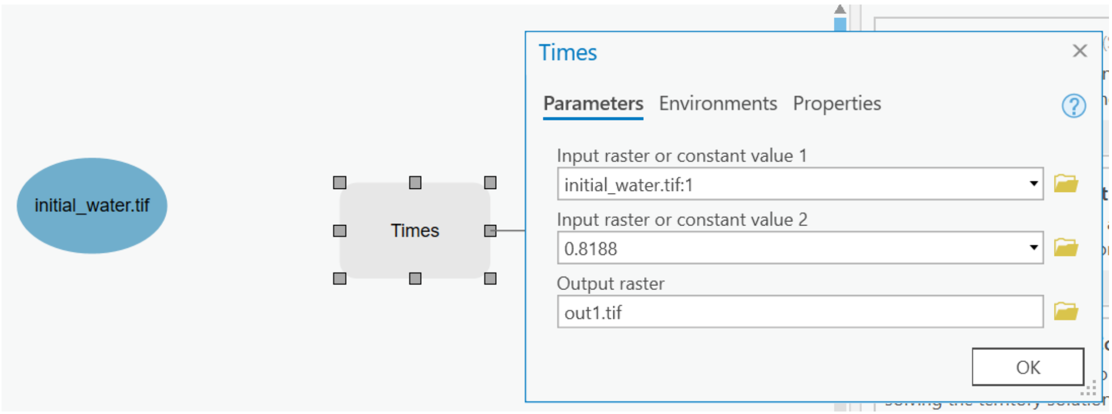

- Your diagram should now look something like this (you may need to move the symbols around slightly to tidy it up):

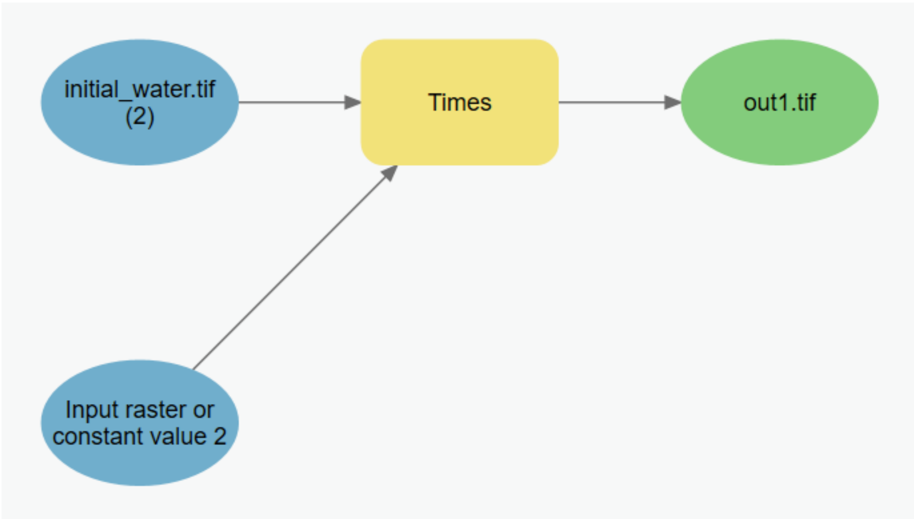

Now for the next bit - the + 118 part underlined:

Soil Water Deficit$_{(t)}$ = 0.8466 * (**0.8188 * Soil Water Deficit$_{(t-1)}$** - Precipitation$_{(t)}$ + 118)

- Find the ‘plus’ tool via the ‘tools’ button on the ribbon, then add this to your model by double-clicking on it. Again, right-click and choose *open*, then set the input values (i.e. the constant of 118 and the raster output from your *Times* tool):

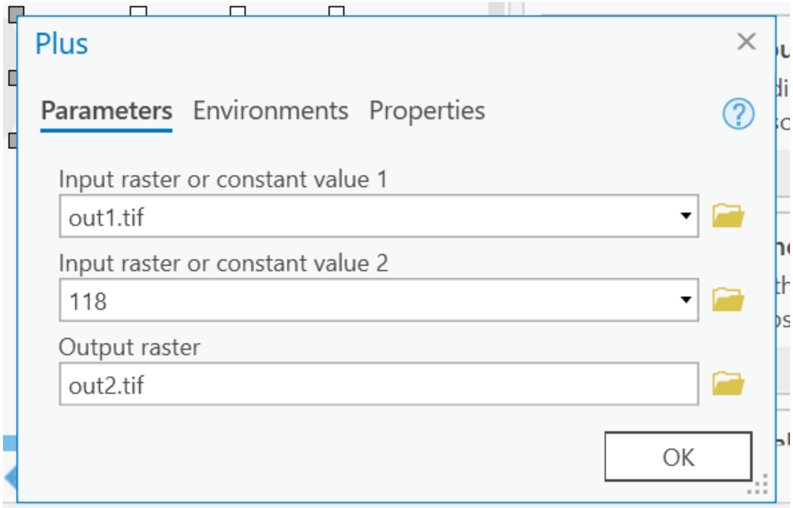

Hopefully, now you should have a Model Builder diagram which looks something like this:

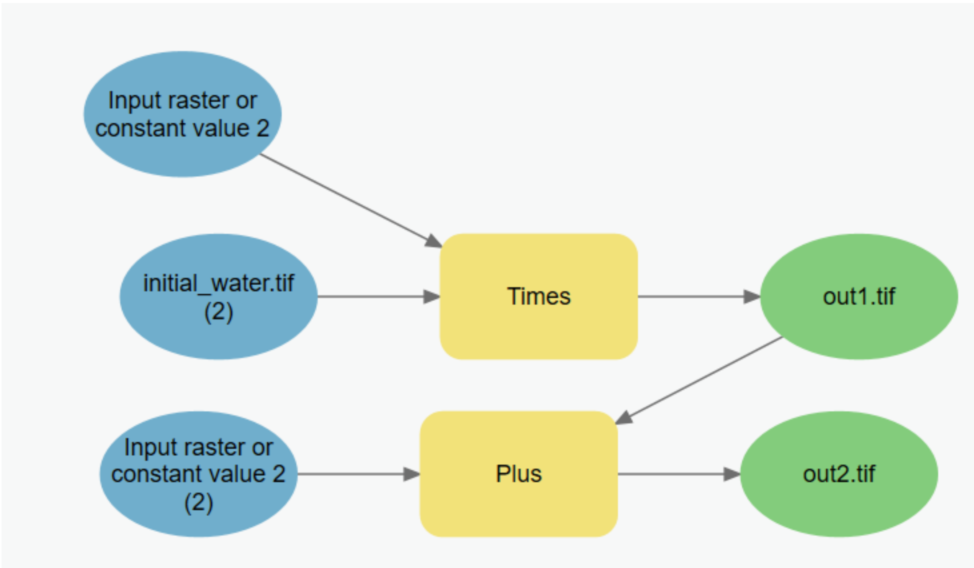

Now we need to calculate the precipitation part (underlined):

Soil Water Deficit$_{(t)}$ = 0.8466 * (**0.8188 * Soil Water Deficit**$_{(t-1)}$ **- Precipitation**$_{(t)}$ **+ 118**)

- Find the *minus* tool from within *tools* and double-click to add it into the Model Builder window.
- Let us explore another way of adding inputs to tools. Drag your mouse from inside **out2** (or whatever you called the output from 'plus') to inside the *minus* box:

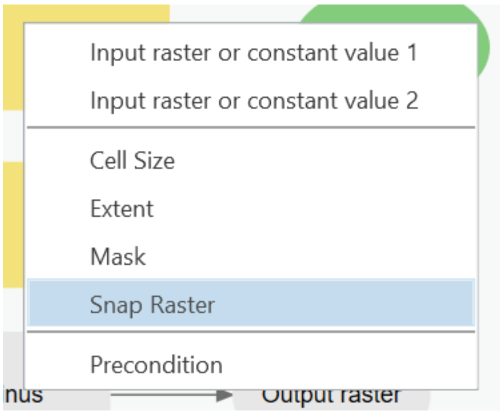

- ArcGIS will ask you how this layer should be used via the pop-up menu above. You can select *input raster or constant 1* to make the raster one of the two required inputs to the tool. **Note:** We will not use these, but the options *cell size, extent, mask, and snap raster* are all optional *environments settings,* accessible via 'environments' when you open up the settings dialog box for the tool. *Cell size* for example specifies the raster grid size (e.g. 100 x 100m) of the output from the tool. *'Mask'* is a way of restricting the tool's calculation to only part of a study area, specified in an optional layer.

- If you right-click on the *minus* tool, you should be able to add **prec_**10 as the *input raster or constant 2*. Change the name of the *output raster* to **step3**. Your model should now look like the diagram below:

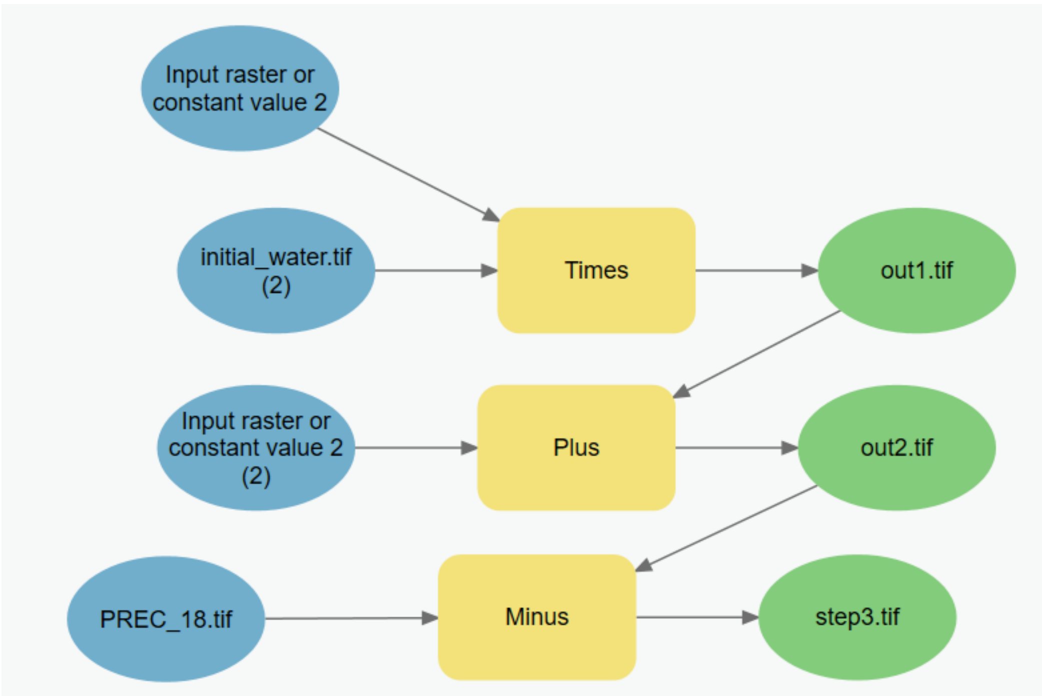

We've now done all of the parts of the equation shown in bold. We simply need to do the final part, multiplying the result of our calculation so far by 0.8466:

Soil Water Deficit$_{(t)}$ = <u>0.8466 *</u> (**0.8188 * Soil Water Deficit**$_{(t-1)}$ **- Precipitation**$_{(t)}$ **+ 118**)

- to do this, drag and drop the *times* tool into the Model Builder window once again.
- Draw a line from **step3** (or whatever you called the output from *minus)* into this tool, setting it to be *input raster or constant 1*.
- Double-click on the new *times* tool and set the output raster to be say **step4**.
- Set the *input raster or constant 2* to 0.8466. Your model should now look something like this:

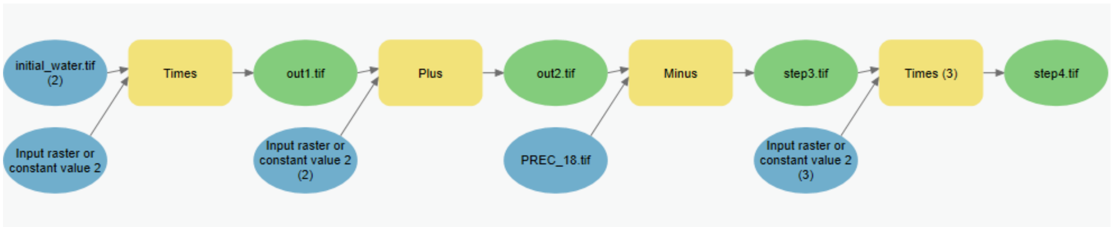

You may well be finding that your Model Builder window is getting quite messy at this point. If you would like to automatically tidy up your boxes, you can press this button:

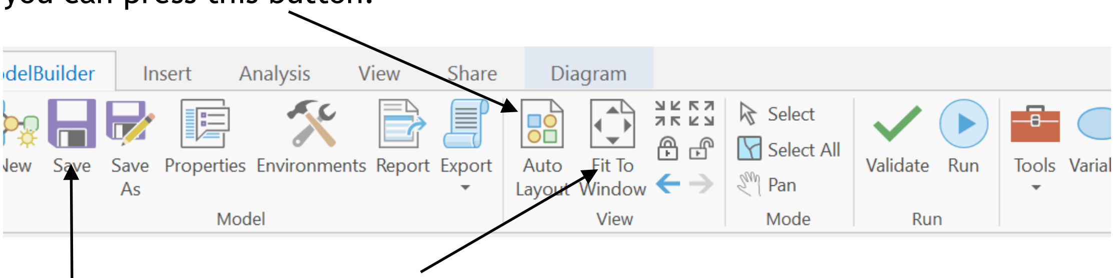

Also very useful is *fit to window*, which enables you to see your whole model. As things are getting quite complex, now may also be a good time to save your model too. Notice also that if you press *Export*, you can save your model either as Python (or more specifically ArcPy, the variant of Python used within ArcGIS) or as a picture (e.g. a jpg), which is very useful, should you wish to document a workflow in a coursework report.

We still have one more job to do. If you re-read the model description on page 1, you will notice that our model can only take positive numbers (a positive soil water deficit number means that our soil water column is saturated and has water sitting on top of it). We therefore need to reset any negative numbers to be zero. To do this:

- Within the *tools* on the ribbon, search for 'greater' and then double-click on the *greater than* tool to drop it into your Model Builder window.
- Draw a line from the **step4** oval to the *greater than* box, setting this to be *input raster or constant 1*.

- Open up the *greater than* tool and call the *output raster* **step5**.
- set the *input raster or constant 2* to be **0**.
- Finally, double-click on the *times* tool from within *tools* and add it into your Model Builder window.
- Draw lines linking this *times* box to **step4** and **step5**, making them inputs to the tool. Call the *output raster* **water_deficit**.
- Hopefully, you should now have a formidable-looking model like the diagram below, which you should now save (via *save* on the *model* menu).

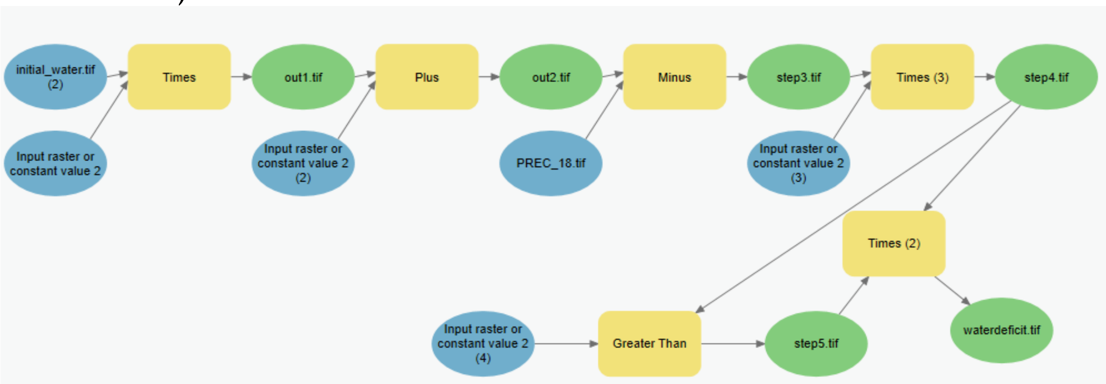

You may wish to use *Save as* on the menu ribbon to save your work with a more meaningful name (the label below is a bit too long – perhaps just call it 'water deficit model'):

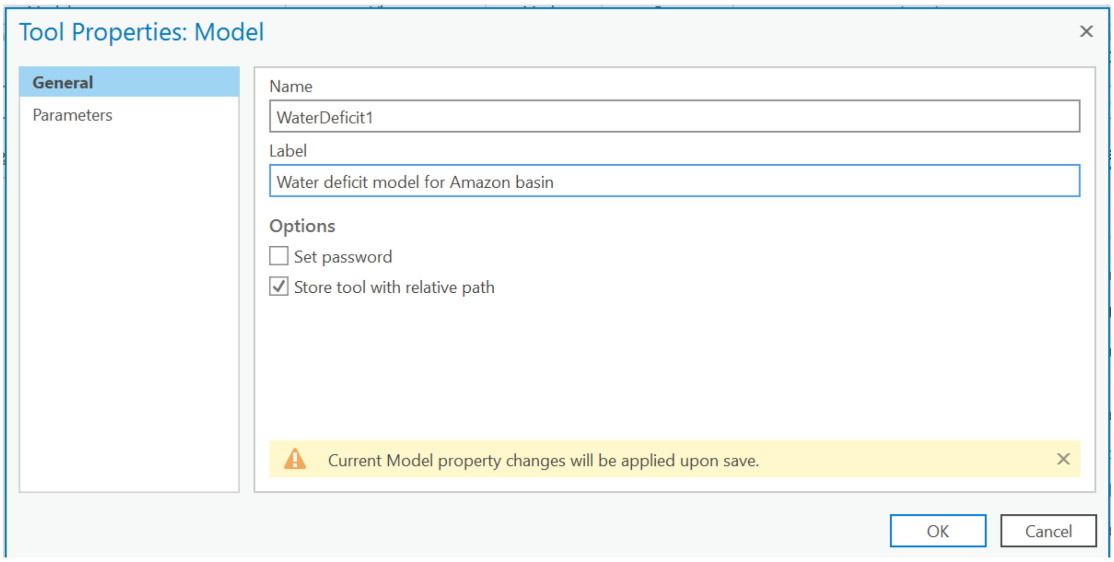

Notice also the setting in the side panel parameters: we will learn how to set these up in just a moment.

Run your model:

Try running your model and see what happens. You can do this by frst clicking on the *validate model* button to check for any obvious errors and then the *run* button:

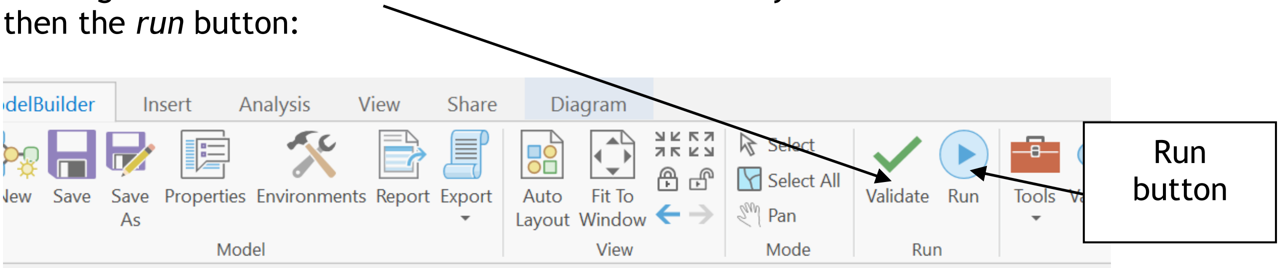

If you head across to your map display (via the tabs above the model canvas), you should be able to display the resultant model output by adding from the catalog panel (you may need to right click on your working folder and press *refresh* if the model output is not visible initially).

Take a look at the output raster grid, **water_deficit**, and see how it looks. You should find that it looks something like the screenshot below:

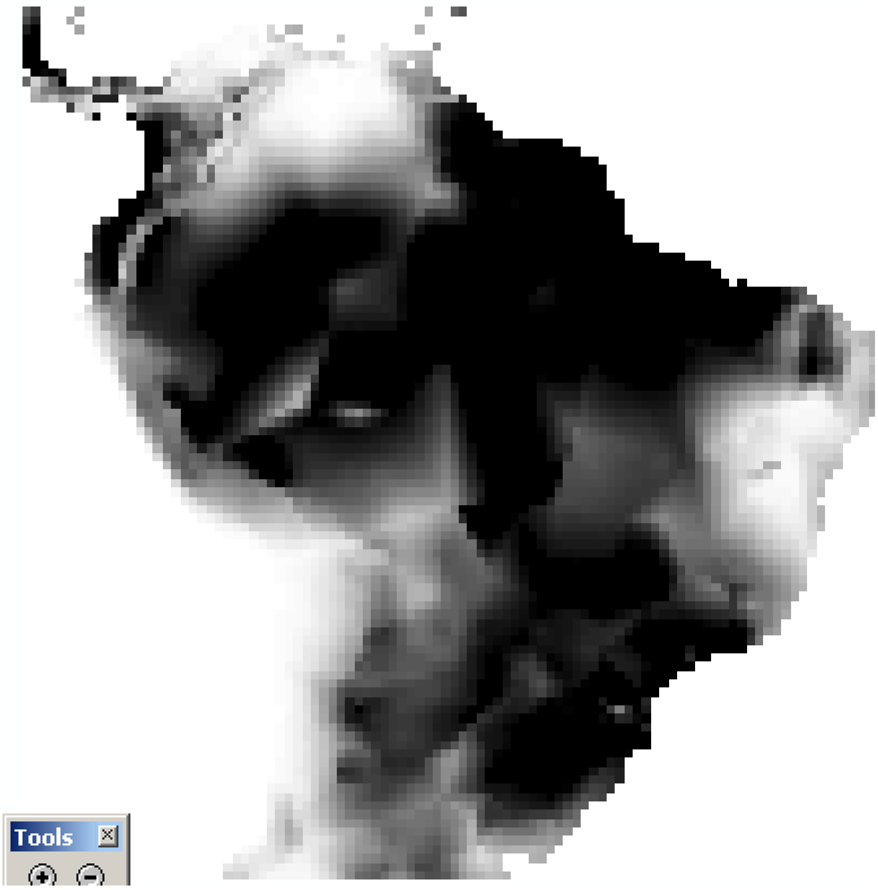

Parts of the Brazilian coast, Venezuela, and the western coast along Peru and Chile are running up large soil water deficits and appear in lighter colours.

Running the model for different time periods:

So far, we have run our model for the first month (October 1996) in our set of 9 months of precipitation data. We can, however, run it for the next two months (November and December 1996) as well. To do this:

- Go back from your map window to the model window via the tabs at the top of your map display.
- Right-click on the **initial_water** oval and in the pop-up box that appears, select *parameter*. A parameter is an input or output value for a model that is set at runtime by the user.
- Do the same for both **prec1** and **water_deficit** to make these model parameters too. You should now see a 'P' for 'parameter' appear next to each of these 3 boxes:

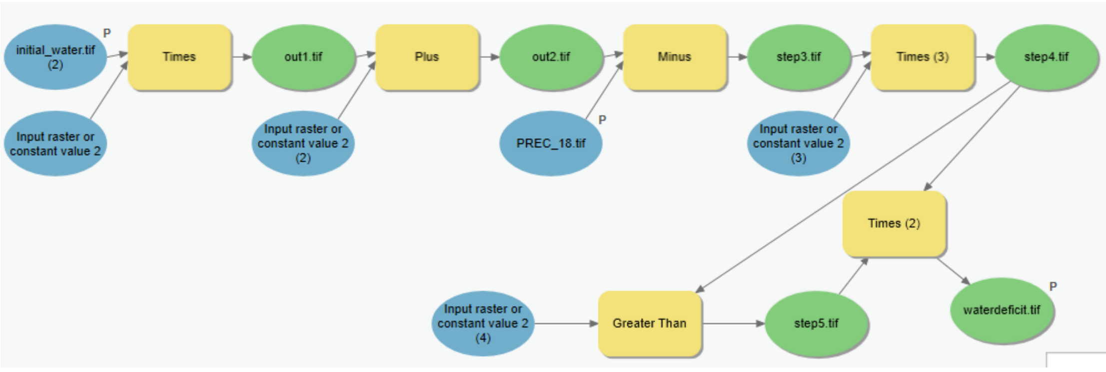

- Save your model and then close the Model Builder window.

You might wonder where your model has been saved. It has been saved into the working folder you set for your project and automatically added to the set of tools available to you. If you now head for the *Analysis* menu and choose *tools*, try searching for 'water' (or use words from the label you gave your model if you did not use the term 'water' when you saved it). You should be able to see your model appearing, like any other tool.

- If you right-click on your tool, you can choose *edit* to open it up and return to the Model Builder window and change the workflow for the tool. Close your model down for now and go back to the geoprocessing panel.
- If you right-click on your tool again, *locate* displays where the tool is (in your project folder).

You can also double-click on your tool in the geoprocessing panel and run it. You will then be prompted to enter the three parameters for it that you specified earlier:

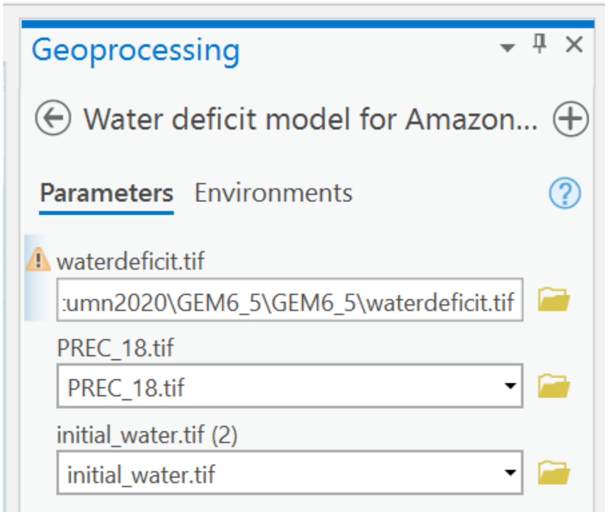

- What we can now do is to feed in our results from the first month into the next month's calculations. To do this:
    - Set the starting raster grid (initial_water) of soil water deficit to be **water_deficit**. In effect, we are feeding in the soil water deficit that we calculated from last time around into the next month's calculations here.
    - Set the precipitation raster grid to be **prec_11**. Here, we are now feeding in the next month's precipitation figures.
    - Set the output *water deficit* grid to be **water_def2** – we will use a different name for the output, because it represents a different month's soil water deficit.

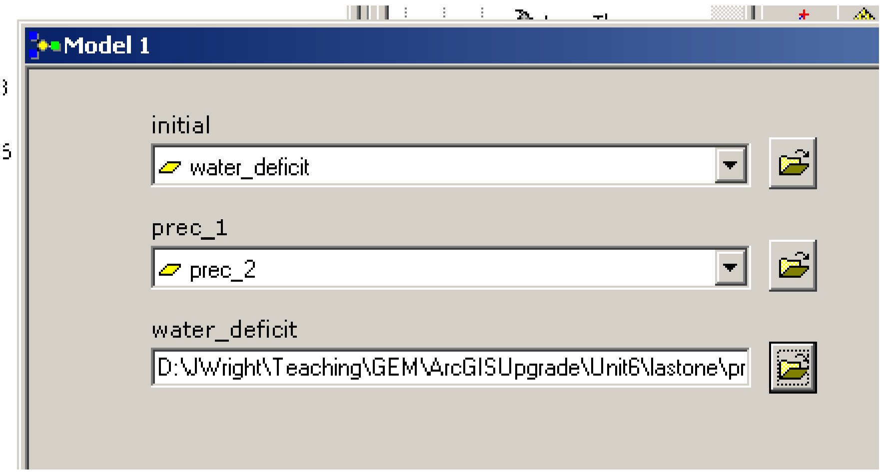

- We can now run this model by hitting OK, but this time it will be for the next month in the time series – November 1996.
- Take a look at your **water_def2** map layer (you may need to display this using a *classified* symbology) – it should look something like this:

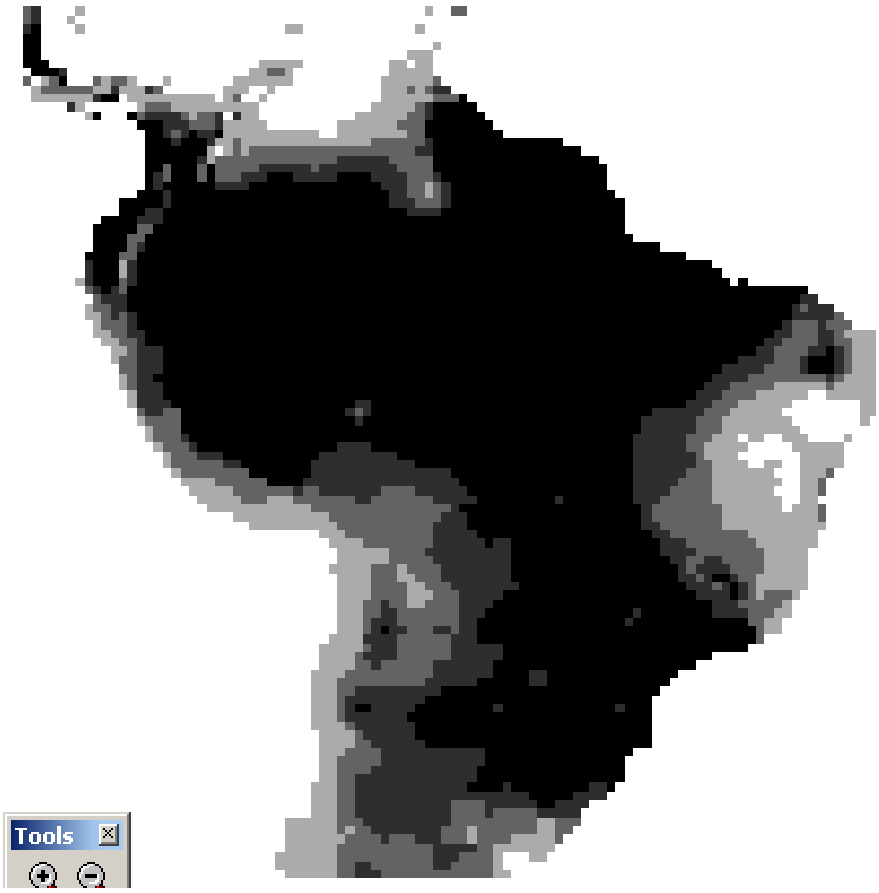

Again, the lighter areas are in eastern Brazil, where soil water deficit is high.

Try running the model one more time using the next month's precipitation data:

- Double-click on your new model tool in the toolbox once again:
  - Set the starting raster grid of soil water deficit to be **water_def2**.
  - Set the precipitation raster grid to be **prec-12**. Here, we are now feeding in the next month's precipitation figures.
  - Set the output *water deficit* grid to be **water_def3**.
  - Run the model again.

We have run our model once again, but this time for December 1996. We could carry on in this way, running the model again and again for the rest of the months in our time series.

Automating time series processing - ideas for taking this practical further (optional):

There are other ways of building models over time using the Model Builder feature within ArcGIS. In particular, you can have your model run on a time series 'stack' of raster grids, rather than entering in the relevant month's data at each time step. Such an approach uses what is called an iterator – a means of repeatedly running the same calculation within a loop.

We will illustrate how this works by using iteration to add up the total amount of precipitation in our monthly time series. To do this:

- First, outside of Model Builder, use the *create constant raster* tool to create a raster with the value zero, having the same dimensions (extent and cell size) as the other grids in the time series.
- Next, go back to the *Analysis* menu again and choose the Model builder icon to create a second, new model.
- This time, in your model, head for the *iterators* menu on the ribbon and choose *iterate rasters*
- Identify the folder that contains the precipitation rasters for this exercise and enter that as the *workspace:*

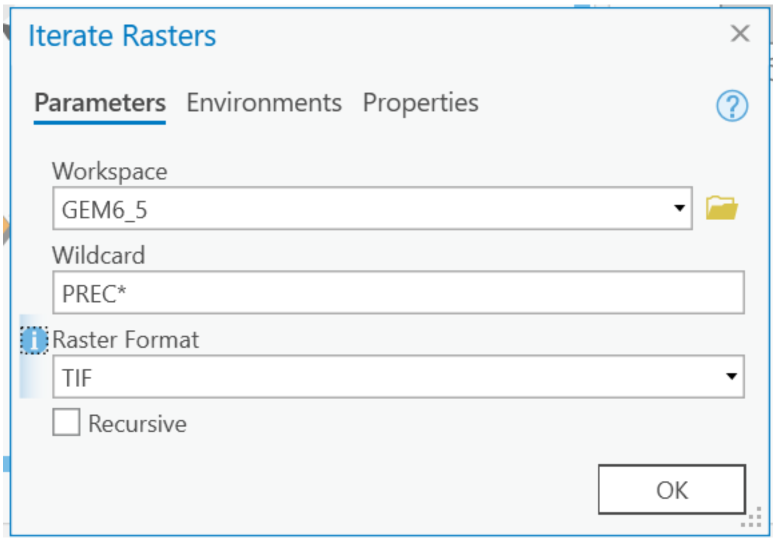

- You may well have multiple raster files by now in your working folder, so it may be a good idea to specify a wildcard (**prec***). The wildcard means that only files with names beginning **prec…** will be processed. Similarly, setting raster format to **TIF** means only geotiff format rasters will be processed.
- Notice that *iterate rasters* produces two outputs. One is the raster to be processed and by default, once you have set up the options for the tool, the name of the first raster layer meeting any criteria you set (wild cards, file formats) will be displayed.
- The second output is *name*. This is the name of each raster being processed, in a format suitable for use in other tools. Within iterators, you can set up names for outputs (e.g. output map layers or tables) such that a different name is used on each pass through a loop. This is achieved through what ArcGIS calls inline substitution. For example, as the output name of a *plus* operation within an *iterate rasters* loop, we could write 'plus%name%'. For each pass

through the loop, the part inside the % signs would take the name of the raster currently being processed via *iterate rasters*. Thus, the name of the output raster would change from 'plusprec_10', 'plusprec_11', etc in our case by means of this inline substitution. That way, we could process many rasters in the same way, saving the outputs for each one. %Name% is just one of several options for setting up inline substitution[1].

- Next, add in the *plus* tool via the *tools* menu. Set it up so that it takes as inputs:
    - the raster filled with zeros that you created just before starting this model
    - the output raster from *iterate rasters*.
- As the output raster name, we will use an example of inline substitution. We will use an inline substitution format %n% - this is a way of indicating the number of passes through the loop. %n% is 1 for the first raster, 2 for the second raster in 'iterate rasters', and so on. If you specify the output name as 'output%n%.tif', the software will call your output file 'output1.tif' when processing the first raster, 'output2.tif' when processing the second raster, and so on:

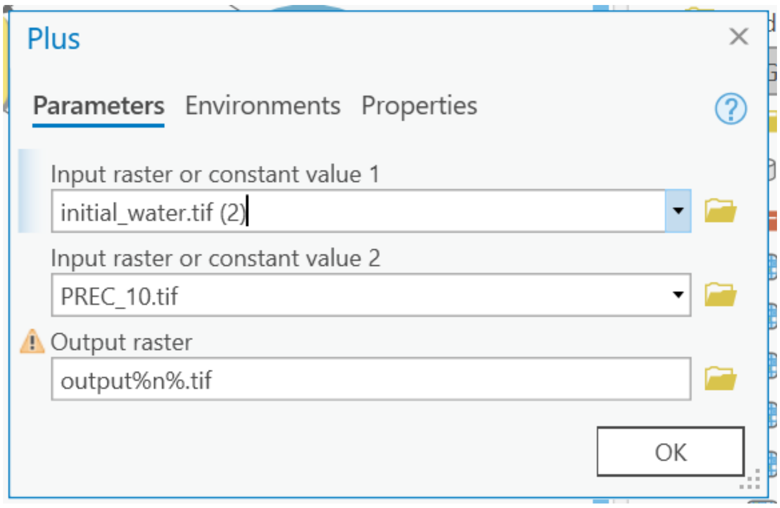

Next, head for the *tools* menu again, but this time, search for 'copy' and add the *copy raster* tool to your model. Use this tool to make a copy of your **output%n%** raster called something like **finalsum**:

[1] A complication with %name% is file extensions, such as .tif. Where file names are specified, often we have to strip out the extension part of the name such as .tif. There is a tool on the 'utilities' menu called 'parse path'. This splits out the folder pathway from the file name from the extension name, so can be used like any other tool in a model to simplify file names.

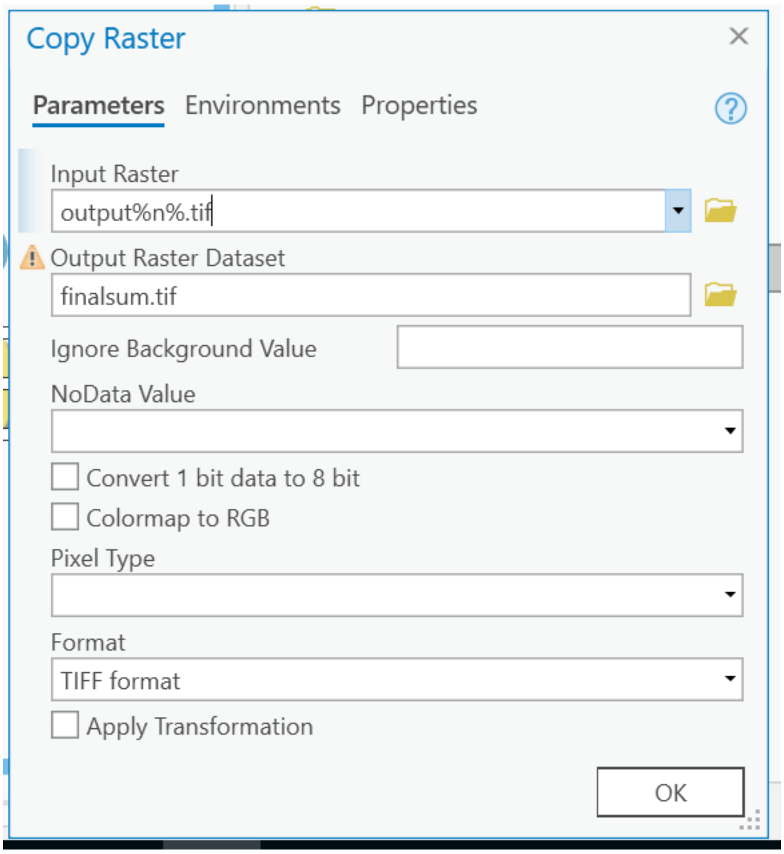

We will now introduce another idea for iterators – feedback loops. Sometimes, with an iterator calculation, we need the output from one time step to feed into the calculation for the next time step. Here, we are trying to add up total precipitation across all of our monthly raster layers. If we have added up precipitation for January, February and March via three passes through the loop, we need the fourth pass through the loop for April to use the running total of precipitation from the preceding three months. We can do this by setting up a feedback loop.

To do this, click and hold down your mouse button somewhere inside the output raster from 'plus', then drag your mouse over to your raster filled with zeros. You should hopefully see a blue dotted line.

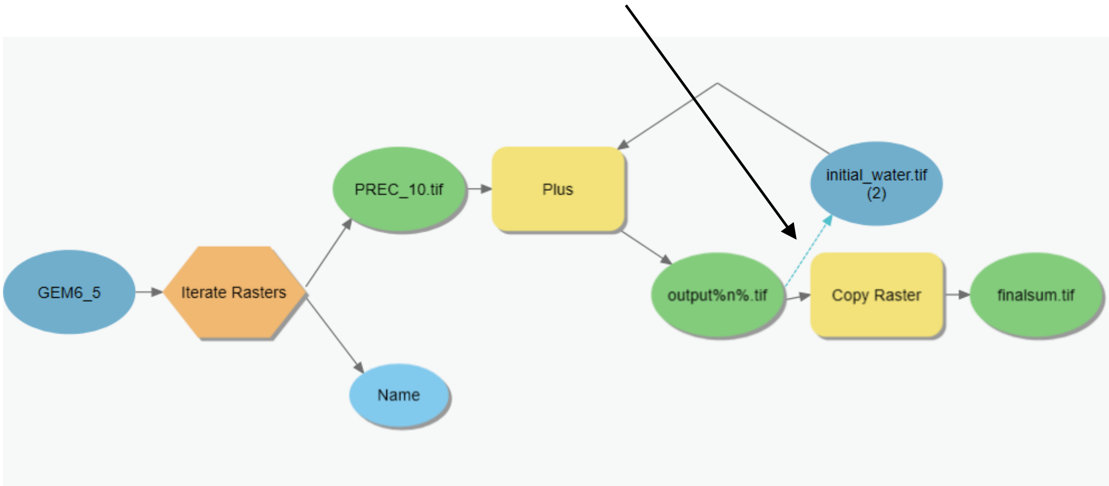

This blue dotted line is a feedback loop. After the very first pass through the loop, the **output%n%** layer will take the place of our initial raster on each pass through the loop, enabling us to carry over a running precipitation total in our calculation.

One more point: when you run a model, only the final products are saved permanently, not intermediate data (i.e. calculations that are stepping stones to a final product). As far as Model Builder is concerned, our **output%n%** layer looks like intermediate data and not like a final product. Typically, a final product does not form an input to any other tool or feedback loop. This is why we therefore had to add in ‘copy raster’, so that there was a final product, which would be permanently saved at the end of the calculation.

Finally, right click on **finalsum** and then choose *add to display*. This will mean that on each pass through the loop, this layer is added to your map display.

If you now save your model and validate it, you should be able to run it. Notice that as it runs, you should see the tool that is running highlighted in red, with a print-out onscreen of the model’s progress.

All being well, when it finishes, you should find that **finalsum** contains the total precipitation over the nine months for the Amazon basin, summed across all nine input rasters.

Extension question (if you wish to stretch yourself): Could you take our soil water deficit model and use ‘iterate rasters’ to implement it across all nine months’ precipitation data via a new model?

Summary

The ‘iterators’ functionality provides a much more sophisticated way of handling time within the Model Builder environment. Note that there are other programming constructs that can be used within the Model Builder environment too, such as ‘branching’ (where a part of a model is only processed if certain initial conditions are met). These skills, though challenging, are in high workplace demand.

References:

You do not have to read any of these references, but you may find them useful background.

The equation that predicts soil water deficit in the Amazon basin is taken from this paper:

Malhi Y and Wright JA (2004): 'Spatial patterns and recent trends in the climate of tropical forest regions'. *Philosophical Transactions of the Royal Society, Series B – Biological Science* **359**: 311-329

The climate data that are used in this exercise were created by the University of East Anglia Climate Research Unit. Details of how these maps were created are available in:

New, M., Hulme, M. and Jones, P.D., (2000): 'Representing twentieth century space-time climate variability. Part 2: development of 1901-96 monthly grids of terrestrial surface climate'. *Journal of Climate* **13**: 2217-2238