Lesson 51: Surface creation

Table of Contents

Goals

In this lesson, you will learn:

- how surfaces are represented
- how surfaces can be created from sample points
- what interpolation methods are available and how they differ

Topic 1: Surfaces

Some data doesn't lend itself well to being represented as discrete geographic information. Such data can include the spatial distribution of temperature, rainfall, elevation, pollution concentration, and water tables. This type of data is spatially continuous, indicating that a different value can be assigned to each location.

Usually, the distribution of continuous data is not characterized by a sudden change in value, although areas of rapid change are common. Examples of sharp variation include steep slopes, the drop-off in precipitation on the lee side of a mountain, and the change in air temperature in a hot, arid climate near a cool body of water. Usually, most of the data is distributed evenly in space. If displayed graphically, data about continuous phenomena can appear to have a smooth surface, which is why it's often called *surface data*.

Surfaces can be represented by models built from regularly or irregularly spaced sample points on the surface. Interpolation is the mathematical estimation of z values on a surface at unsampled points based on the known z values of surrounding points. Through interpolation, ArcView Spatial Analyst can generate a grid theme from a point theme, thus creating a continuous surface from a limited set of sample data

Density

Spatial Analyst's Density function calculates the number of features within a given area. For example, given a number of customers (points), Density will calculate the number of people per square mile or acre.

Without a population field, Density calculates straight numbers of features per square mile (you can select the units, miles, feet, acres, etc.). In the example above, you can select an attribute field from the customer theme table, like Income or Spending. The output grid would display income per square mile or spending per square mile.

Another example: If you had a point theme of crime locations, you could create a density grid without using a population field. You would end up with a map of crime density. More specifically, you could select a field in the crime table to create a map of density of car thefts or density of break-ins.

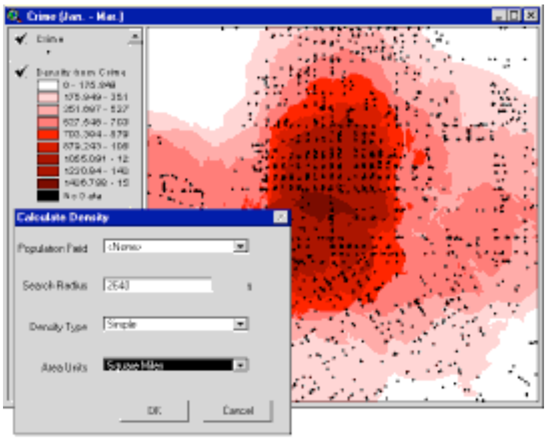

Front: The Density dialog. Back: A theme showing crime density. [Click to enlarge]

Simple and kernel are two standard ways to calculate densities. Kernel uses a smoothing function.

The search radius is a distance from each cell whose points will be used in the density calculation. Area units are the desired unit of measure. Available units are:

- square miles
- square kilometers
- acres
- hectares
- square yards, feet, or inches
- square meters, centimeters, or millimeters

Representing surfaces

ArcView Spatial Analyst can represent surfaces in three common ways: as elevation points, contour lines, and surface grids. Spatial Analyst does not support triangulated irregular network (TIN) datasets.

Point themes have a z value which is used to create contour line themes or surface grid themes. The z value is an attribute like elevation, temperature, or rainfall.

Contour lines are isolines of constant elevation with a specified interval and are a very common way to represent terrain surfaces. Contour accuracy depends on whether the isolines are generated from primary or derived data sources. When contours have been captured directly from aerial photographs as primary data using a stereoplotter, the contours are highly accurate. If the contours have been generated from point data, the location of the contours must be interpolated between known values.

A major drawback of contours is that they only indicate surface value along the isolines. Surface anomalies between contour intervals cannot be represented.

Surface grids can be created from sample points, digital elevation models (DEMs), and other sources. Grids represent information in equally sized square cells arranged in rows and columns. Each grid cell is referenced by its geographic x,y location

Using sample points

Do you need to visit every location in a study area to collect data like elevation or precipitation to create a surface? It would be difficult, or nearly impossible in some cases, to do so. The alternative is to collect the data at sample locations and then use those sample locations to interpolate, or estimate, values for the rest of the surface.

There are various strategies for determining where to locate the sample points. Distance between sample points is an important factor. If the sample distance is large, important variations in the surface may be missed. Smaller sample distances may provide a better representation of the surface, but at the expense of disk space and redundant data. The samples can be regularly or randomly spaced. The more input points and the greater their distribution, the more reliable the results.

The attribute of surface data being measured is called the z value. The amount of rain, level of pollution, and elevation are all examples of z values.

The cell values in the output grid theme are best estimates or interpolated values. Certain assumptions are made when making these estimates. When estimating values, error increases with distance from the samples or known values

Linear interpolation

Given two sample points, linear interpolation estimates values between sample points based on distance.

In the example below, two samples have values of 1 and 2. The data was collected at two rainfall collection stations that are 1 mile apart. The interpolated values are estimated, given the distance between the two sample points. You can estimate that, at a half-mile between each station, the rainfall was 1.5 inches. The other values can also be estimated based on the fraction of distance between the two known samples.

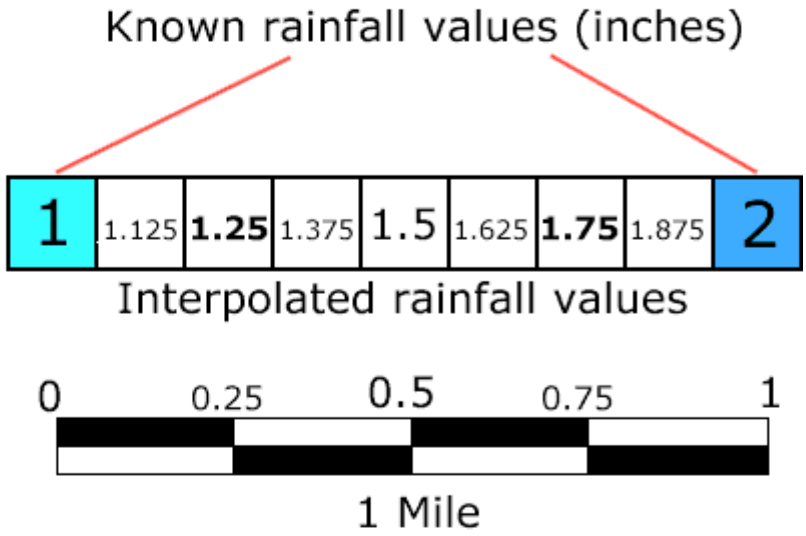

Here, two known rainfall samples are 1 mile apart. Values between the sample points are interpolated based on distance

Surface interpolation

Surface interpolation generates a raster surface from an active point theme in a view. The points may be either regularly or randomly spaced and may contain measurements of elevation, concentration, magnitude, or some other quantity.

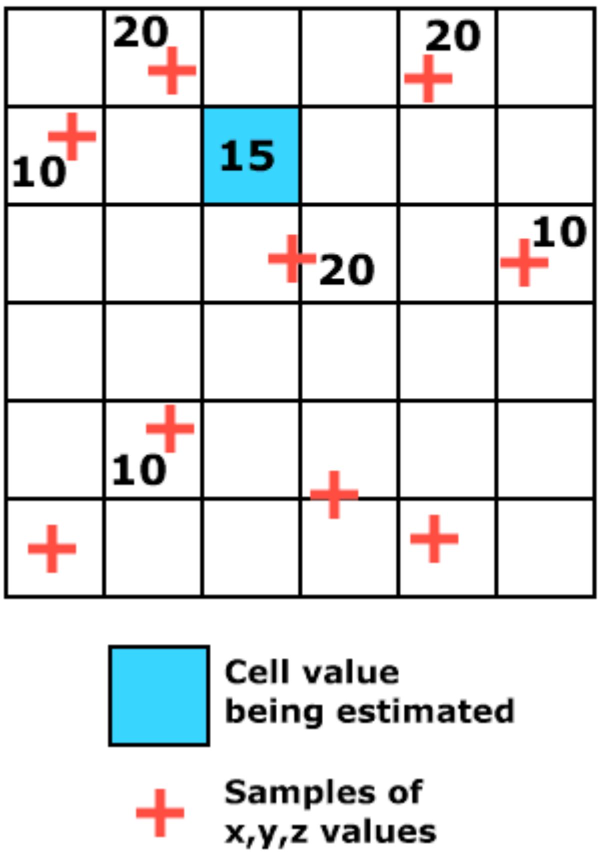

This diagram shows a cell value being estimated from a set of sample points.

Values for each grid cell in the surface are mathematically estimated according to an interpolation method. ArcView Spatial Analyst has four surface interpolation methods that will create a surface from a set of sample points. Spline and IDW (inverse distance weighted) appear as choices in the Interpolate Surface dialog presented to the user when creating a surface. The other two interpolation methods, Trend and Kriging, are accessed using Avenue requests.

Each of the four interpolation methods uses a different approach to determine output cell values with a selected set of sample points. The method you choose depends on the kind of data for which you are creating a surface, the distribution of your sample points, and the phenomenon being studied.

A surface grid can be interpolated from a point theme by choosing Interpolate Grid from the Surface menu. The interpolated grid is a temporary floating point grid. Its default name is "Surface from" followed by the name of the point theme. The grid dataset is written as temporary to the project's working directory, with the name "sface" followed by a unique number.

Grid interpolation is a two-step process. In the first step, you specify the extent, cell size, and mask for the output grid. The extent can be set to that of any theme in the active document. If the active document is a view, the extent can also be set to that of the view or display. Spatial Analyst sets a default cell size and number of rows and columns for the grid. You can change these values manually or set them to match those of any grid theme in the active document.

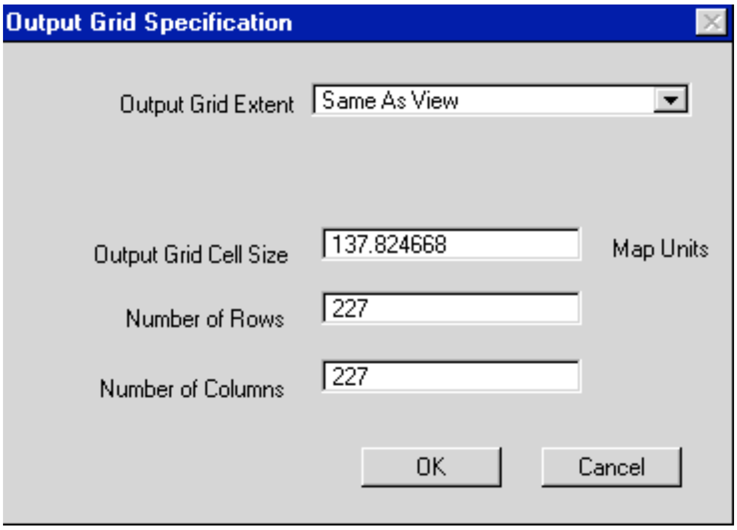

The Output Grid Specification dialog.

In the second step, you choose an interpolation method and the field from the point theme table whose values will be used to create the surface. You can also set various parameters for the interpolation method.

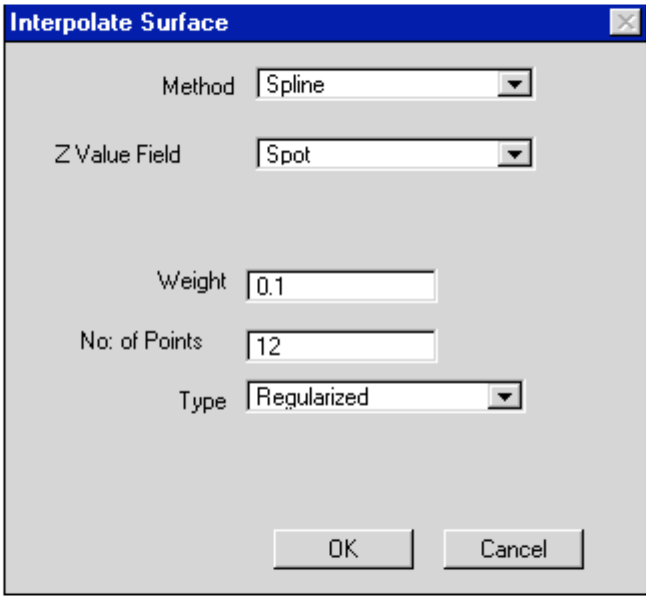

The Interpolate Surface dialog.

The interpolated grid is created and added to the active view. It is always a floating point grid, regardless of whether the input values are integers. By default, it's symbolized with nine classes and a gray monochromatic color ramp, but you can change symbology in the Legend Editor

IDW: Inverse distance weighted

The inverse distance weighted (IDW) interpolation method assumes that each sample point has a local influence that diminishes with distance. In estimating the value for a given cell, it gives greater weight to points closer to the cell than to those farther away. A specified number of points (the default is 12), or, optionally, all points within a radius, are used to determine the value for each cell. The surface being calculated should be a locationally dependent variable.

Use IDW when you have a dense set of points. They should be dense enough to capture the extent of local surface variation needed in your analysis. If you want to capture the high and low surface extremes in your data, make sure that your point dataset includes sample points along these features. If the sampling of input points is sparse or very uneven, the results may not adequately represent the desired surface.

IDW is available as an interpolation method on the Surface menu. The Interpolate Grid dialog with the IDW method chosen will allow you to choose either nearest neighbors or fixed radius sampling. With nearest neighbors chosen, you can specify number of neighbors, power and barriers. With fixed radius, you can specify a radius, power, and barriers.

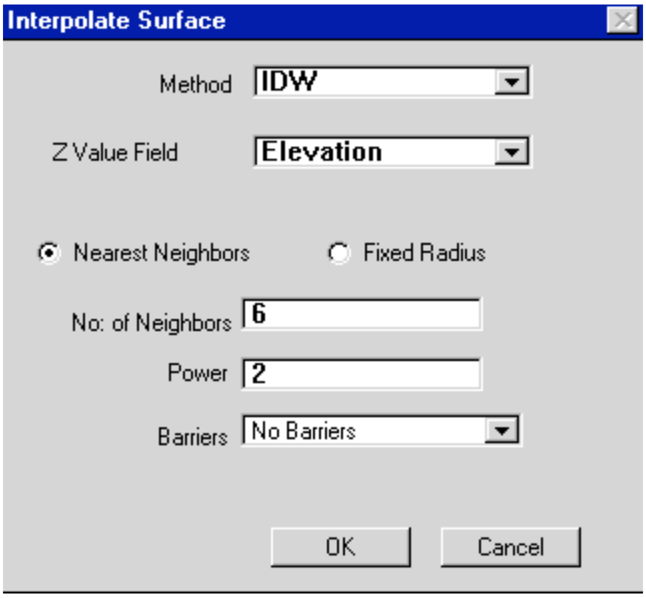

The Interpolate Surface dialog. IDW is the chosen method and Elevation is the z field. Samples will be selected using the six nearest neighbors with a power of 2.

In the example below, the IDW interpolation would estimate a value of 17 for the selected cell. It does not estimate 15 because it weights the closer cells higher. The three samples of 20 have more weight or influence in estimating a value of 17.

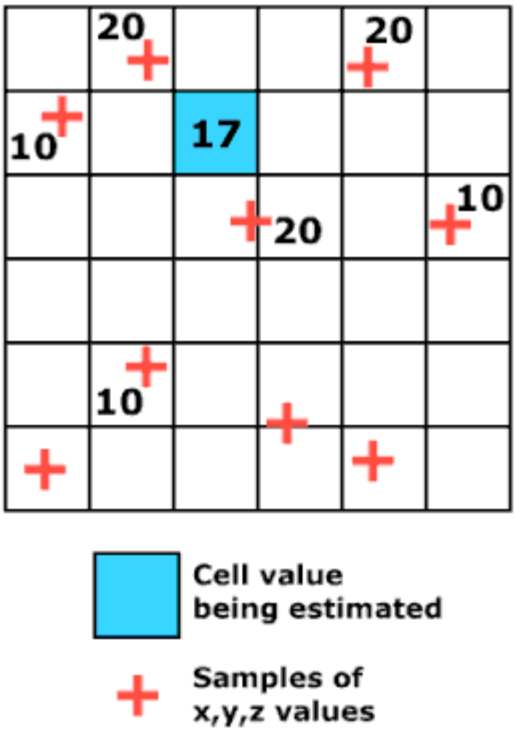

In this example, the IDW interpolation would estimate a value of 17 for the selected cell.

The relative weighting of sample points can be changed by specifying a power (the default is 2). The larger the power, the greater the influence of points close to the processing cell. The power option lets you control the significance of sample points on the estimated values. A larger power means close sample points have more influence on calculating output cell values. Sample points that are farthest away have less influence. If IDW is run with higher powers (greater than 1), it runs with a high degree of local influence, giving the output surface increased detail. If IDW is run with a power of 1 or less, it runs with a global influence, treating each point almost equally to create a smoother output surface.

If IDW is set with higher power values, it is said to be running with a high degree of local influence. If IDW is run with a power of 1 or less, it is said to be run using global influence.

Barriers can also be set to constrain interpolation. A barrier is a line theme that may represent ridges, shoreline, or any other feature that should interrupt interpolation. Barriers limit the number of sample points used to interpolate a given cell's value to those that lie on the same side of the barrier as the cell. The use of barriers significantly slows interpolation time.

For more information on IDW see:

Watson, D. F. and G.M. Philip. 1985. A refinement of inverse distance weighted interpolation. *Geo-Processing*. 2, 315-327

Defining sample subsets for IDW

How do you determine which samples are considered during IDW interpolation? Surfaces potentially have an infinite number of points that can be measured. Obviously, it is impossible to record every point. Consequently, a sampling method must be used to extract representative points to build a model that approximates the surface.

When using the IDW interpolation method, sample subsets of data points may be estimated by using either the nearest number of neighbors or samples within a radius.

Using the nearest neighbor approach, the characteristics of the interpolated surface can be controlled by limiting the input points used in the calculation of the output cell values. You can limit the points by specifying the maximum number of points to be sampled, in which case, the closest ones to the output cell location are selected until the maximum number is reached.

For example, if you specify the six nearest neighbors, the z values of only those six samples will be used in the interpolation. The default number of points used is 12.

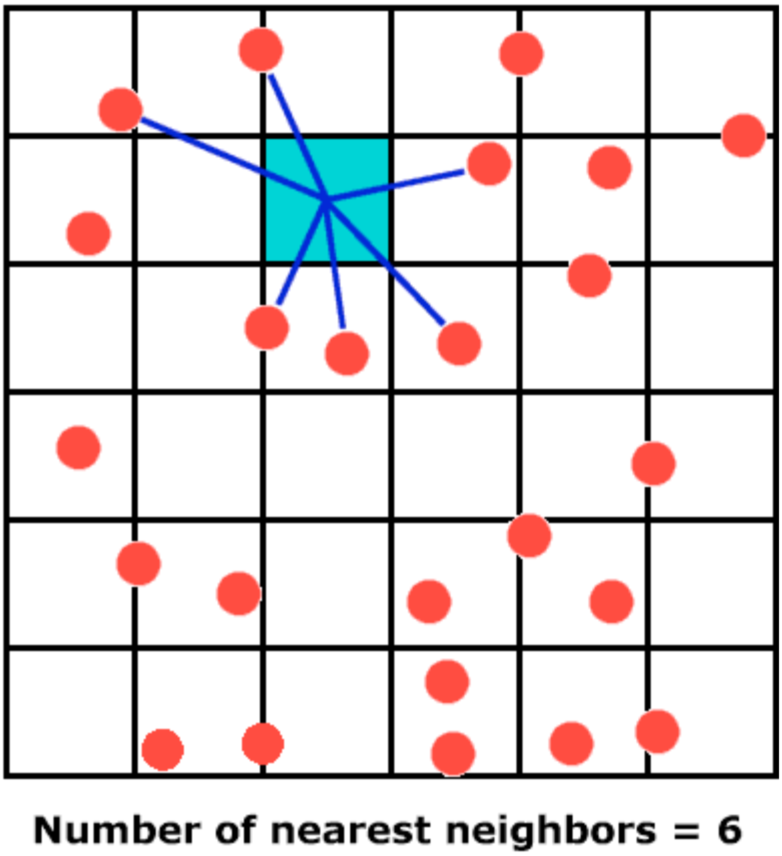

The value of a grid cell is being determined by using the six nearest sample z values.

Alternatively, you can specify a radius in map units. In this case, only the input points within the radius distance from the center of the output cell are used unless there are not enough points within the radius. For example, if you specify a radius of 2000 feet and there are 24 points within that radius, the z values of all 24 points are used in the interpolation.

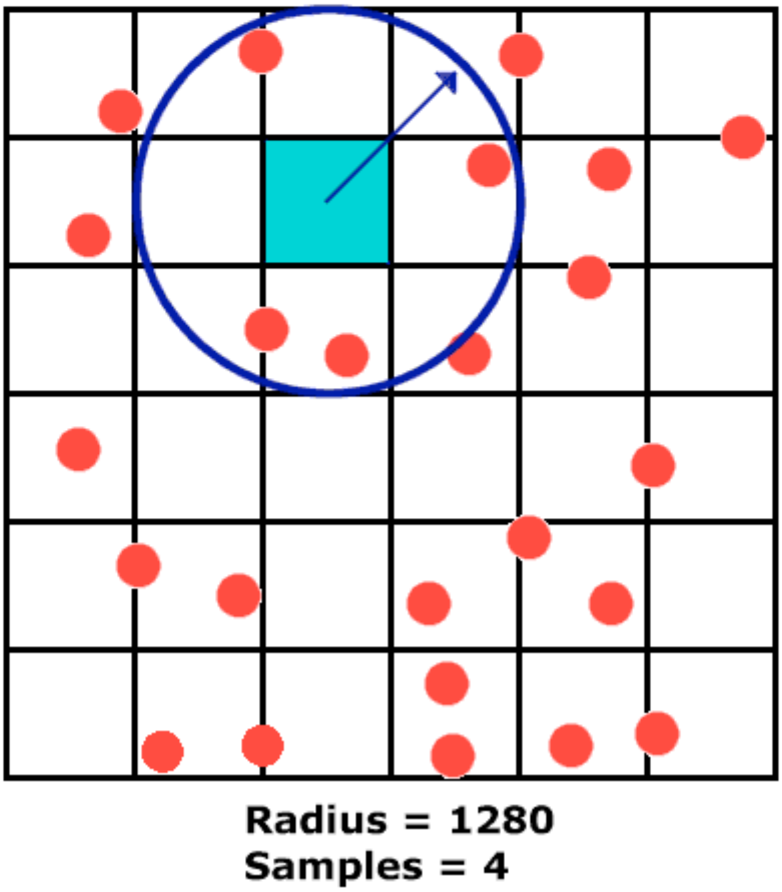

The value of a grid cell is being determined by using the four sample z values within the specified radius.

Because IDW is an averaging technique, the unknown value cannot exceed the highest of high values or the lowest of low values. This means that extreme natural formations like ridges and valleys can't be created unless they have been adequately sampled.

For more information on surfaces see:

Philip, G.M. and D. F. Watson. 1982. A precise method for determining contoured surfaces. *Australian Petroleum Explanation Association Journal.* 22: 205-212

Spline

Spline is a general-purpose interpolation method that fits a minimum-curvature surface to the sample points. The surface passes exactly through the sample points. Like IDW, a surface created with the Spline method will always have the exact value of a sample point at the corresponding surface location. It will also produce a smooth surface because it minimizes curvature. Before the use of computers made it easy to estimate surface values, drafters used flexible rulers to manually fit a surface over the sample points. These rulers were called splines.

Because it generates smooth surfaces, the Spline method is best suited to sample data that varies gently (for instance, elevation or pollution concentrations). It's not appropriate if there are large changes in value within a short horizontal distance.

Spline is available as an interpolation method from the Surface menu. After choosing the Spline method in the Interpolate Surface dialog, you can specify a weight, number of points, and type.

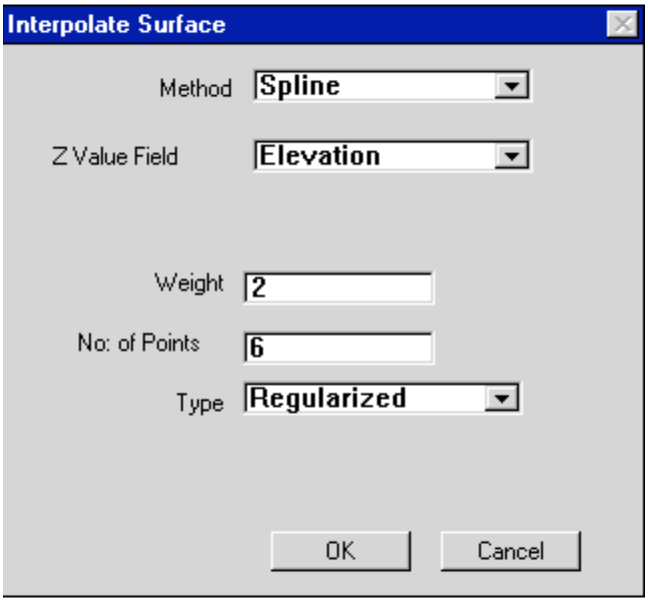

The Interpolate Surface dialog. Spline is the method, and Elevation is the chosen z value.

In the example below, Spline estimates the value of the selected cell as 23. Spline tries to fit a curve using the selected subset of samples--in this case, six samples. The curve would start at one of the cells with a value of 10, go up to a cell with a value of 20, continue up or overshoot, then come down to another 20, and back down to a 10. The cell whose value is being estimated (23) may have been on the upswing of the curve.

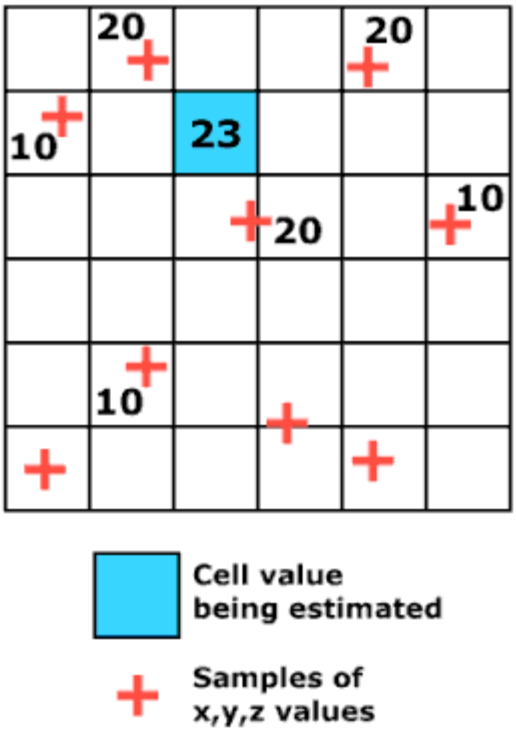

Spline estimates the value of the selected cell as 23 using the selected subset of six samples.

Increasing the number of points creates smoother surfaces, but also increases computation time.

There are two types of splines that can be used to interpolate a surface: regularized and tension. Different values are entered into the spline algorithm to vary the tautness/looseness of the splines. These values are various different rates of change (derivatives) calculated from the local trends shown by the input data points. A local trend can be described as a constant and regular change in the data in a particular direction. The first derivative reveals the change of z value, which corresponds to the magnitude of the gradient or slope in a localized area. The second derivative corresponds to the rate of change of slope in an area or the curvature of that part of the surface. The third derivative reflects the changes in or extremes of the curvatures.

A regularized spline incorporates the first derivative (slope), the second derivative (rate of change in slope), and the third derivative (rate of change in the second derivative) into its minimization calculations. Regularized spline offers a looser fit, but may have overshoots and undershoots. The weight parameter when using a regularized spline makes a smoother surface. Higher weight values smooth more than lower values.

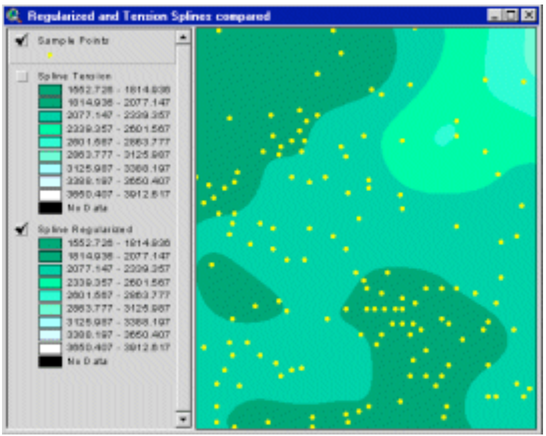

A regularized spline. A looser fit is produced but possibly at the expense of overshoots and undershoots. [Click to enlarge]

The tension spline uses only first and second derivatives. It forces the curve. The weight parameter when using a tension spline makes a coarser surface. Higher weights produce coarser surfaces than lower weights.

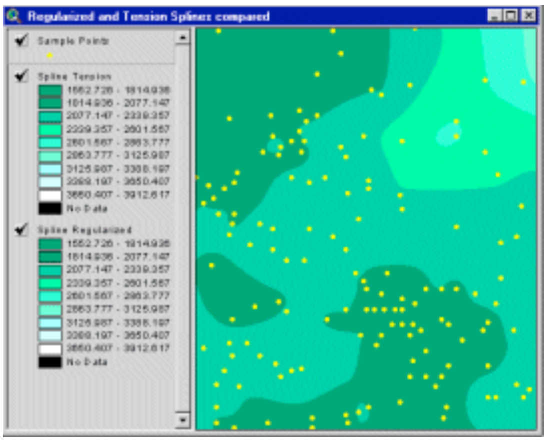

A tension spline. A tension spline generally makes a coarser surface. It forces the curve. [Click to enlarge]

To help visualize Spline, think of it working like a sheet of plastic. Regularized spline is like thin plastic and tension spline is like thicker plastic

Exercise

Create surfaces with the IDW and Spline interpolation methods

This exercise familiarizes you with ArcView Spatial Analyst surface interpolation functions available from the Surface menu. You will create different surfaces using the surface generation methods of IDW (inverse distance weighted) and Spline. You will then compare the results from the different functions.

You will be interpolating elevation for an area roughly six miles square, centered on the Crafton Hills. The Crafton Hills are due east of the city of Redlands, California. You will be using a cell size of 330 feet. The small area and large cell sizes were chosen for the exercise so that the interpolators would run faster.

If you have not downloaded the exercise data for this module, you should [download the data now](#).

Step 1 Start ArcView

Start ArcView and load the Spatial Analyst extension.

Note: If you are running ArcView GIS 3.1, you see a Welcome to ArcView GIS dialog. Click Cancel to close this dialog.

If ArcView is already running, close any open projects.

Step 2 Load the CellTools sample extension

This exercise requires the Celltools sample extension. You'll load it now.

To load the CellTools sample extension, make the Project window active, then choose Extensions from the File menu. Check the box next to Cell Tools (sample).

If Cell Tools (sample) is not listed in the Extensions dialog, you'll need to copy the **celltool.avx** file from ArcView's samples\ext directory to ArcView's ext32 directory. These directories are located in ArcView's installation directory. A typical location (the default) for the installation directory is **c:\esri\av_gis30\arcview**.

If you cannot find this path, try searching for the file **arcview.exe** to help you determine the correct path.

Once you've done this, the sample extensions will be listed in the Extensions dialog the next time you open it. In that dialog, check the sample extension(s) you want to load and click OK.

If you don't want this sample extension to appear in the Extensions dialog permanently, you may load it temporarily by running a simple script. It will be available only for your current ArcView session.

To run this sample script in ArcView, perform the following steps:

1. In ArcView's Project window, click the Scripts icon, then click New to create a new script.

2. Type the following line into the script window:

```
system.SetEnvVar( "USEREXT",
"$AVHOME\samples\ext".AsFilename.GetFullname)
```

3. Compile and run the script.

Once you've done this, the sample extensions will be listed in the Extensions dialog the next time you open it. In that dialog, check the sample extension(s) you want to load and click OK.

Step 3 Open the project

From the File menu, choose Open Project. Navigate to the **surfalsa\lesson1** directory and open the project **l1_ex1.apr**.

Note: If you are running ArcView GIS 3.1, you see an Update l1_ex1.apr message box. Click No to dismiss this box.

When the project opens, you see an Elevation view containing an elevation grid theme and a point theme of sample elevations.

VIEW RESULT

Step 4 Examine the sample points

Start by examining the sample point theme attributes. Make the Vipcov.shp theme active, then click the Open Theme Table button  to examine the attributes of the sample point theme.

The Spot field contains the elevation of the sample points.

VIEW RESULT

Close the theme table when you have finished examining it.

REVIEW CONCEPT

Step 5 Interpolate a surface using IDW

The inverse distance weighted (IDW) interpolation method interpolates by giving greater weight to the sample points closest to the cell whose value is being interpolated. As the distance between the input point and the cell increases, the significance of the sample decreases. IDW will not interpolate values beyond the vertical extremes of the provided samples. That is, it cannot find ridges or valleys if they are not represented in the sample set. The best results are obtained with very dense sampling.

You'll use the IDW method to interpolate a surface of elevation using the Vipcov.shp sample point theme.

Make the Vipcov.shp theme active. From the Surface menu, choose Interpolate Grid. In the Output Grid Specification dialog, set the Output Grid Extent to Same as Hillcgrd, and the Output Grid Cell Size to Same as Hillcgrd.

Click OK.

In the Interpolate Surface dialog, set the Method to IDW and choose Spot as the

Z Value Field.

Click the Nearest Neighbors radio button, set No. of Neighbors to **10** and specify a Power of **2**.

Click OK.

When the Surface from Vipcov.shp theme appears, rename it **Idw1grd** and turn it on.

VIEW RESULT

A visual comparison of the two grids is not very useful--a more statistical approach is called for. In the next step, you'll use the Map Calculator to subtract the Idw1grd grid from the Hillcgrd grid.

REVIEW CONCEPT

Step 6 Compare grids

From the Analysis menu, choose Map Calculator. Enter the following expression:

[Hillcgrd] - [IDW1grd]

Click Evaluate.

When the Map Calculation 1 theme appears, rename it **IDWt1grd** and turn it on. Close the Map Calculator.

VIEW RESULT

The lighter areas are where the IDW-interpolated values are below the original surface; the darker areas are where they are above it.

Step 7 Examine the test grid

Click the CellValue tool and examine a random selection of light and dark cells in the IDWt1grd test grid.

Now you'll create a histogram of the test grid.

Make IDWt1grd active and click the Histogram button .

Next, you'll examine the statistics for the test grid.

In the Table of Contents, double-click on IDWt1grd to open the Legend Editor. In the Legend Editor, click the Statistics button.

The Statistics dialog appears.

VIEW RESULT

Notice that the mean difference is –12 and that the standard deviation is about 50. That means that about 62 percent (the definition of standard deviation) of the cells are within the range of 62 feet below the original surface to 38 feet above it.

Click OK to close the Statistics dialog. Close the Legend Editor.

In the next step, you'll try using a smaller sample size, and see what effect it has.

Step 8 Interpolate a surface using IDW with a smaller sample size

You will now use the IDW method again, but this time using a smaller sample size.

Make the Vipcov.shp theme active. From the Surface menu, choose Interpolate Grid. In the Output Grid Specification dialog, set the Output Grid Extent to Same as Hillcgrd and the Output Grid Cell Size to Same as Hillcgrd.

Click OK.

In the Interpolate Surface dialog, set the Method to IDW and choose Spot as the Z Value Field.

Click the Nearest Neighbors radio button, set No. of Neighbors to **5** and specify a Power of **2**.

Click OK.

When the Surface from Vipcov.shp theme appears, rename it **Idw2grd** and turn it on.

VIEW RESULT

In the next step, you'll use the Map Calculator as before to subtract the Idw2grd grid from the Hillcgrd grid.

Step 9 Compare grids

From the Analysis menu, choose Map Calculator. Enter the following expression:

[Hillcgrd] - [Idw2grd]

Click Evaluate.

When the Map Calculation 1 theme appears, rename it **IDWt2grd** and turn it on. Close the Map Calculator dialog.

VIEW RESULT

The lighter areas are where the IDW interpolated values are below the original surface, and the darker areas are where they are above it.

Step 10 Examine the test grid

Use the CellValue tool to examine a random selection of light and dark cells in the IDWt2grd test grid.

Make IDWt2grd active and click the Histogram button.

Next, you'll examine the statistics for the test grid.

Double-click on IDWt2grd to bring up the Legend Editor. In the Legend Editor, click the Statistics button.

VIEW RESULT

The range of differences is greater, but the mean is closer to 0 and the standard deviation is smaller, meaning more cells have values closer to those of the original elevation grid.

Further experimentation with the sample size or the power value may yield a closer fit. The lesson here is that surface generation is usually an iterative process, where you experiment with various combinations of parameters and sample sizes to find the interpolation that works best with a particular set of samples.

Click OK to close the Statistics dialog. Close the Legend Editor.

Next, you will use the Spline interpolation method.

Step 11 Interpolate a surface using a regularized spline

The Spline method uses a two-dimensional splining interpolation and results in a smooth surface that passes through all the sample points. Generally, the greater the number of sample points included in the interpolation, the smoother the surface.

You'll use the Spline method to create an interpolated surface using the default Regularized method, then compare the new surface to the original.

Make the Vipcov.shp theme active. From the Surface menu, choose Interpolate Grid. In the Output Grid Specification dialog, set the Output Grid Extent to Same as Hillcgrd and the Output Grid Cell Size to Same as Hillcgrd.

Click OK.

In the Interpolate Surface dialog, set the Method to Spline, and choose Spot as the Z Value Field. Set the type to Regularized.

Click OK.

When the Surface from Vipcov.shp theme appears, rename it **Spl1grd** and turn it on.

VIEW RESULT

Again, a visual comparison of the two grids is not very useful. In the next step, you'll use the Map Calculator to subtract the Spl1grd grid from the Hillcgrd grid.

REVIEW CONCEPT

Step 12 Compare grids

From the Analysis menu, choose Map Calculator. Enter the following expression:

[Hillcgrd] - [Spl1grd]

Click Evaluate.

When the Map Calculation 1 theme appears, rename it **SPLt1grd** and turn it on. Close the Map Calculator.

VIEW RESULT

The lighter areas are where the Spline-interpolated values are below the original surface, and the darker areas are where they are above it.

Step 13 Examine the test grid

Use the CellValue tool to examine a random selection of light and dark cells in the SPLt1grd test grid.

Make Splt1grd active and click the Histogram button.

Next, you'll examine statistics for the test grid.

Double-click on SPLt1grd to open the Legend Editor. In the Legend Editor, click the Statistics button.

VIEW RESULT

Based on the comparison, these results are a bit better than those of IDW, in that the mean is closer to zero.

Click OK to close the Statistics dialog. Close the Legend Editor.

Next, you will try the Tension option.

Step 14 Interpolate a surface using a tension spline

You'll now create an interpolated surface using the tension spline method, then compare the new surface to the original.

Make the Vipcov.shp theme active. From the Surface menu, choose Interpolate Grid. In the Output Grid Specification dialog, set the Output Grid Extent to Same as Hillcgrd and the Output Grid Cell Size to Same as Hillcgrd.

Click OK.

In the Interpolate Surface dialog, set the Method to Spline and choose Spot as the Z Value Field. Set the Type to Tension.

Click OK.

When the Surface from Vipcov.shp theme appears, rename it **Spl2grd** and turn it on.

VIEW RESULT

Again, a visual comparison of the two grids is not very useful. In the next step you'll use the Map Calculator to subtract the Spl2grd grid from the Hillcgrd grid.

Step 15 Compare grids

From the Analysis menu, choose Map Calculator. Enter the following expression:

[Hillcgrd] - [Spl2grd]

Click Evaluate.

When the Map Calculation 1 theme appears, rename it **SPLt2grd** and turn it on. Close the Map Calculator.

VIEW RESULT

The lighter areas are where the Spline-interpolated values are below the original surface, and the darker areas are where they are above it.

Step 16 **Examine the test grid**

Use the CellValue tool to examine a random selection of light and dark cells in the SPLt2grd test grid.

Make Splt2grd active and click the Histogram button.

Double-click on SPLplt2grd. In the Legend Editor, click the Statistics button.

VIEW RESULT

Another improvement. Again, the mean has gotten closer to zero, and the standard deviation is smaller.

Click OK to close the Statistics dialog. Close the Legend Editor.

Step 17 **Close the project**

Close the project without saving any changes.

You have completed this exercise.

This comparison of IDW and Spline has really proved only one thing: For this set of samples of this specific elevation data, the Spline method is best. Under other circumstances or with different data another interpolator might give better results.

Remember that other interpolation methods, like Kriging and Trend, are available as Avenue requests. Kriging is the most powerful and sophisticated of ArcView Spatial Analyst's interpolators. It requires an advanced level of knowledge to use properly. Like all of ArcView Spatial Analyst's interpolators, it is employed iteratively, varying its parameters until the statistical "best fit" has been found for the data. To give you an idea of the power of Kriging, if you had used the MakeKriging request with this set of sample points, you would have achieved a mean of nearly 0 and a smaller standard deviation than the other two methods explored here.

When you perform surface interpolation of non-topographic data, you should use all the interpolation functions and experiment with each function's parameters and sample size to discover the interpolation technique that works best with your data. You must consider the phenomenon you are modeling, the error in the samples, and distance effects between samples

Topic 2: **Interpolation methods using Avenue**

In the last topic, you learned about the two interpolation methods available from the Spatial Analyst user interface: IDW and Spline. There are two additional interpolation methods available using Avenue: Kriging and Trend. These interpolation methods, especially Kriging, are much more difficult to master, so only the general principles of Kriging and Trend are covered here.

Kriging is the most powerful and sophisticated of the grid interpolators. It requires an advanced level of knowledge to use properly. Like all the grid interpolators, Kriging should be employed iteratively, varying parameters until the statistical "best fit" has been found for the data. It is introduced here in case you want to explore its use on your own.

It's important to note that the IDW and Spline interpolation methods are also available through Avenue. You can find detailed information on each method in the ArcView online help under the topics for the MakeIDW, MakeSpline, MakeKriging, and MakeTrend requests

Kriging

Kriging is a powerful statistical interpolation method used in diverse applications, such as health sciences, geochemistry, and pollution modeling. The technique was co-developed by Georges Matheron (a French geomathematician) and D. G. Krige (a South African mining engineer) to analyze geochemical data to find gold deposits (it worked).

Kriging is usually classified as one of the exact interpolation methods because of its ability to retain the original sample point values on the resultant surface, but can also remain in a class of its own. This is because Kriging is an optimal interpolator, meaning that it finds the best possible estimate in an unbiased manner.

Kriging is based on the theory of regionalized variables, which states that naturally occurring attributes tend to reflect random quality, but at the same time, continuous and spatially dependent traits. The theory also states that objects closer together will tend to reflect dissimilar values to some range of the distance.

It can be stated that Kriging has the ability to model reality well because it encompasses this theory with its emphasis on stochastic spatial variation (randomness of z values).

Kriging generates the resultant surface by incorporating local weighted averaging in its estimates. Each estimation is based on a small sample of points surrounding the estimation point, called the subset, which are weighted according to their distances (and sometimes directions) from each other and from the estimation point.

ArcView Spatial Analyst supports two types of Kriging: Ordinary and Universal. In Ordinary Kriging, a subset's members surrounding the estimation point are weighted according to the distance they are away from the estimation point and from each other. In Universal Kriging, the subset surrounding the estimation point is weighted according to the distance they are away from the estimation point, from each other, and their directions or location to each other.

Kriging is performed in ArcView Spatial Analyst using the MakeKriging request.

Avenue syntax:

```
Interp.MakeKriging (aType, aRadius, aBarrierFTab, outVarGridFN)
```

- aType is the type of mathematical function used to model the semi-variance.
- aRadius controls how searching for input points is done. It can be a Radius or Nil. If it is Nil, a variable radius of 12 points with no maximum distance is used.
- aBarrierFTab is used as a limit within the Radius to control searching for input points across such features as ridge lines. It can be a line FTab or Nil. If it is Nil, no limit is used within the radius when searching for input points.
- outVarGridFN is the filename of the optional output grid dataset of predicted semi-variance values for each cell in the grid created when this Interp object is passed to the Grid.MakeByInterpolation or anFTab.Interpolate requests. outVarGridFN can be a filename or Nil. If it is Nil, a grid dataset isn't created.

This discussion is merely an introduction to Kriging. Kriging is a complex procedure that requires greater knowledge about spatial statistics than can be conveyed in this module. Before using the MakeKriging request, you should have a thorough understanding of the fundamentals of Kriging and have assessed the appropriateness of your data for modeling with this technique. If you would like to learn more about Kriging, some excellent references about Kriging and statistical interpolation are listed below.

- Burrough, P.A. 1986. *Principles of Geographic Information Systems for Land Assessment*. New York: Oxford University Press.
- Heine, G.W. 1986. A controlled study of some two-dimensional interpolation method. *COGS Computer Contributions*. **3**, 2: 60-72.
- McBratney, A.B., and R. Webster. 1986. Choosing functions for semi-variograms of soil properties and fitting them to sampling estimates. *Journal of Soil Science*. 37: 617-639.
- Oliver, M.A. 1990. Kriging: a method of interpolation for geographical information systems. *International Journal of Geographic Information Systems*. **4**, 4: 313-332.
- Press, W.H. et al. 1988. *Numerical Recipes in C, The Art of Scientific Computing*. New York: Cambridge University Press.
- Royle, A.G., F.L. Clausen, and P. Frederiksen. 1981. Practical universal kriging and automatic contouring. *Geoprocessing*. 1: 377-394

Trend

The Trend surface interpolator creates surfaces that show the regional tendency of the variable in question (the z value). Instead of producing a surface that fits reality well, the Trend surface interpolator was designed to display gradual long-range variations in the data. Because this interpolator approximates the surface, it regards the data points as guidelines to which to fit a surface. As a result, the surface will lie above and below the input data points.

The Trend surface interpolator uses a polynomial regression to fit a least-squares surface to the input points. This finds the single best-fit equation to generate the entire surface. Trend surface interpolation creates smooth surfaces. The surface generated isn't likely to pass through the original data points because it performs a best fit for the entire surface and therefore is often classified as an approximate interpolation method. Use Trend when you don't want to model a surface precisely but are more interested in finding trends in the input sample data.

Trend may be the most appropriate tool for interpolating grids of local head elevation from water well data. Groundwater hydraulics are poorly understood, and the sample data is often quite sparse, so it is probably inappropriate to use a more sophisticated interpolator.

The Trend surface interpolator is available in ArcView Spatial Analyst using the MakeTrend request.

Avenue syntax:

```
Interp.MakeTrend (anOrder, useLogistic)
```

The only parameter in Trend that is not common to the other surface generation functions is anOrder. The Trend request allows you to control the order of the polynomial used to fit the surface. Trend is easy to understand by considering a first-order polynomial. A first-order trend surface interpolation simply performs a least-squares fit of a plane to the set of input points. The default is 1.

How do you determine the appropriate order to use? As the order of the polynomial is increased, the surface being fitted becomes progressively more complex. As illustrated in the diagram below, a second-order polynomial will have one curve in it, a third-order will have two curves, etc.

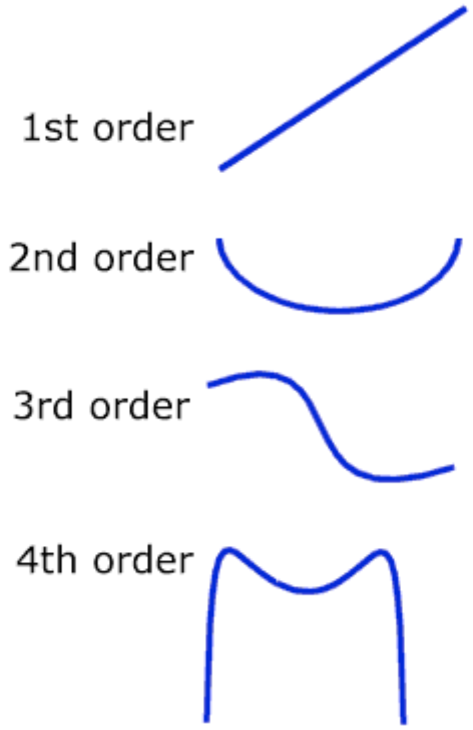

A first-order polynomial is a flat tiled surface, a second-order has one curve, a third-order has two curves, a fourth-order has three curves, and so on.

A higher-order polynomial will not always generate the most accurate surface. The best order is dependent upon the data. The lower the RMS error reported, the more closely the interpolated surface represents the input points. The most common order of polynomials is 1 through 3.

The UseLogistic option determines what type of trend surface analysis is performed. If useLogistic is true, a logistic trend surface analysis is performed, generating a continuous probability surface. This option is appropriate only in situations where the data is binary or logical, hence the name logistic.

An example would be a set of irregularly distributed sample points where a particular species was either found or not found. The output grid would be the probability of detecting the species in each cell. If useLogistic is false, a polynomial regression is performed to fit a least-squares surface to the set of input points. This is more suited for continuous data.

Trend's advantages are:

- Trend surfaces can be used as one part of a two-step process to reveal small scale/local deviations not witnessed when using another interpolator. The trend surface values are subtracted from the input point dataset attribute values, as per x,y location, and the resultant values are used by another interpolator to display the local components.
- It is highly useful to derive a generalized surface from an otherwise "noisy" input point dataset. Maverick or poorly measured data can be smoothed out of the data, but the degree of eccentricity can have ill effects on the resultant surface.

Trend also has many limitations affecting its validity as an accurate surface interpolator:

- Trend does not represent reality well. As you attempt to create a more reliable surface with higher orders, meanwhile being constantly reassured by lower RMS error values that the surface is fitting the input point dataset more closely, you also must realize that the surface is deviating farther from reality.
- All local detail is lost to the smoothing algorithm.
- Clustering and edge effects can hamper the validity of the interpolation. Clustering will create locational bias in the resultant surface, and anomalous edge values can be created where boundaries are poorly represented due to the extrapolation involved with Trend

Summary

In this lesson, you learned that surfaces are geographic phenomena represented as a set of continuous data. ArcView Spatial Analyst can represent surfaces with elevation points, contour lines, and surface grids.

Surfaces can be represented by models interpolated from sample points. Inverse distance weighted (IDW) and Spline are the interpolation methods available from the Spatial Analyst user interface. Avenue provides two additional interpolation methods through the MakeKriging and MakeTrend requests.

IDW interpolation uses an inverse distance weighting method. Sample subsets for use with IDW are defined using either the nearest neighbor or fixed radius sampling method.

The Spline interpolation method fits a minimum-curvature surface through the sample points. A regularized spline offers a looser fit, but may have overshoots and undershoots. A tension spline generally produces a coarser surface as it forces the curve.

Kriging is the most powerful of the interpolation methods and is based on the theory of regionalized variables. It assumes the data being modeled is statistically homogeneous throughout the surface. Trend is not used to model a surface precisely, but rather to find trends in the input sample data. Trend fits one polynomial equation to the entire surface

This is the **Introduction to Surface Analysis with ArcView Spatial Analyst - Lesson 1 Self test**.

Please watch your time—you have 2 hours to complete this test.

Use the knowledge you have gained in *Introduction to Surface Analysis with ArcView Spatial Analyst* to answer the following questions. You will need to correctly answer 7 of the following questions to pass.

Netscape Users: Do not resize this browser window. This can cause the page to reload and generate new questions.

GOOD LUCK!

---

1. Sample points must be regularly spaced in order to create a surface.

   True

   False

2. A surface created with the Spline interpolation method will pass exactly through the sample points.

   True

- False

3. Contours can indicate surface value only along the isolines.
   - True
   - False

4. The attribute of surface data that is being measured is referred to as the:
   - Q value
   - Z value
   - Y value
   - X value

5. What is the default power used by IDW to weight sample points?
   - 1
   - 2
   - 3
   - 4

6. Which of the following interpolation methods makes use of a polynomial order as one of its parameters?
   - IDW
   - Kriging
   - Spline
   - Trend

7. By default, an interpolated grid is symbolized with:
   - 9 classes and a gray monochromatic color ramp
   - 9 classes and a red monochromatic color ramp
   - 13 classes and a gray monochromatic color ramp
   - 13 classes and a red monochromatic color ramp

8. Which of the following most likely produces the smoothest surface?
   - A tension spline with a lower weight value
   - A tension spline with a higher weight value
   - A regularized spline with a higher weight value
   - A regularized spline with a lower weight value

9. An interpolated grid is:
   - Integer if the input values are integer
   - Floating point only if the input values are floating point

- Always integer
- Always floating point

10. A surface created with the Spline method will always have the exact value of a sample point at the corresponding surface location.

- True
- False