Interpolating Raster Surfaces

Lesson 4: Interpolation

- The basic principles of interpolation
- How to create surfaces using interpolation
- How to control sample points using interpolation
- How to use IDW, Spline, and Kriging
- What the interpolators covered have in common

Introduction to interpolation

- Whether you are concerned with the amount of rainfall, concentrations of pollution, or the differences in elevation, it is impossible to measure these phenomena at every point within a geographic area. You can, however, obtain a sample of measurements from various locations within the study area, then, using those samples, make inferences about the entire geographic area. Interpolation is the process that enables you to make such an inference.
- With spatial interpolation, your goal is to create a surface that models the sampled phenomenon in the best possible way. To do this, you start with a set of known measurements and, using an interpolation method, estimate the unknown values for the area. You then make adjustments to the surface by limiting the size of the sample and controlling the influence the sample points have on the estimated values.

Introduction to interpolation

- The primary assumption of spatial interpolation is that points near each other are more alike than those farther away; therefore, any location's values should be estimated based on the values of points nearby.

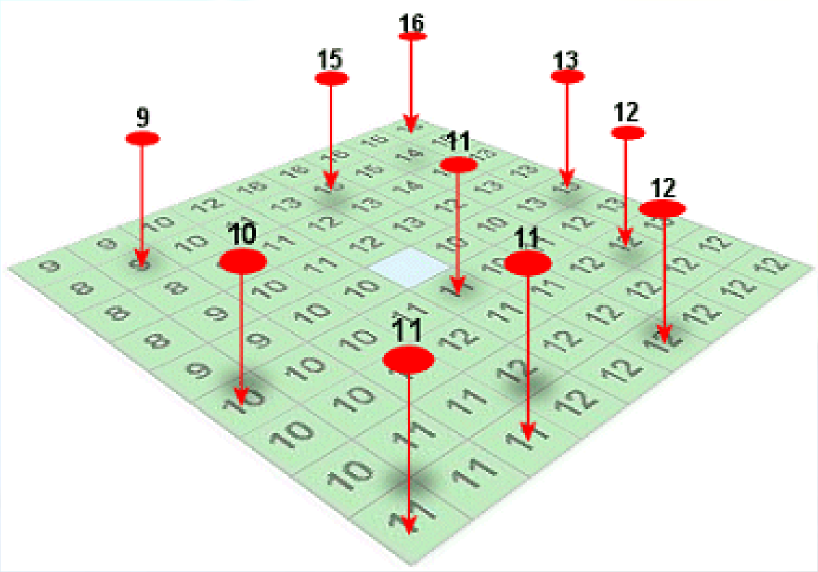

Interpolating the sample points' values creates a surface. As with all of the cells, the unknown value of the light-blue cell in the center will be estimated based on values of the surrounding sample points.

What is interpolation?

- Interpolation is the process of estimating unknown values that fall between known values.

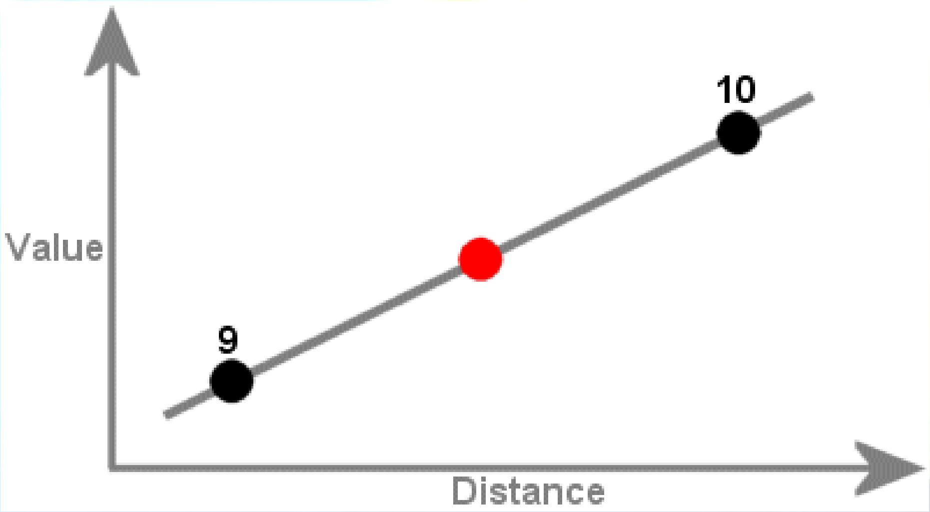

In this example, a straight line passes through two points of known value. You can estimate the point of unknown value because it appears to be midway between the other two points. The interpolated value of the middle point could be 9.5.

What is spatial interpolation?

- Spatial interpolation calculates an unknown value from a set of sample points with known values that are distributed across an area. The distance from the cell with unknown value to the sample cells contributes to its final value estimation.

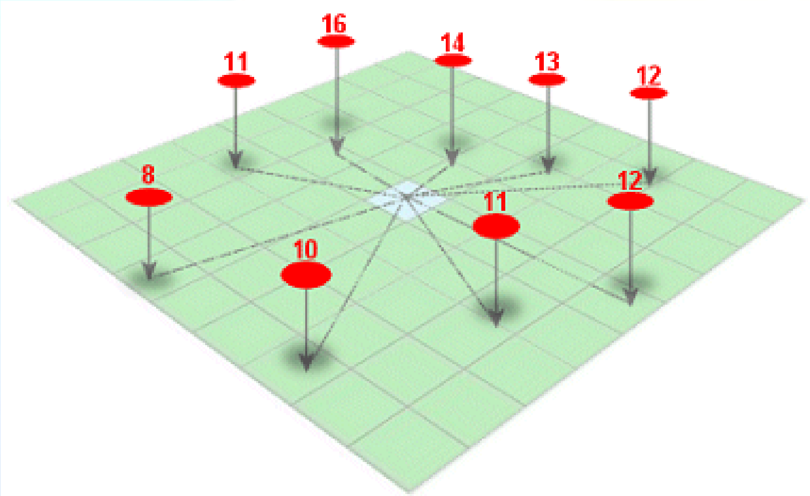

The unknown value of the cell is based on the values of the sample points as well as the cell's relative distance from those sample points.

What is spatial interpolation?

- You can use spatial interpolation to create an entire surface from just a small number of sample points; however, more sample points are better if you want a detailed surface.
- In general, sample points should be well-distributed throughout the study area. Some areas, however, may require a cluster of sample points because the phenomenon is transitioning or concentrating in that location. For example, trying to determine the size and shape of a hill might require a cluster of samples, whereas the relatively flat surface of the surrounding plain might require only a few.

Spatial autocorrelation

- The principle underlying spatial interpolation is the First Law of Geography. Formulated by Waldo Tobler, this law states that everything is related to everything else, but near things are more related than distant things.
- The formal property that measures the degree to which near and distant things are related is spatial autocorrelation. According to this, if it is raining where you are, it is probably raining 10 feet away from you, is less likely to be raining on the other side of town, and might even be clear and sunny 20 miles away.
- Most interpolation methods apply spatial autocorrelation by giving near sample points more importance than those farther away.

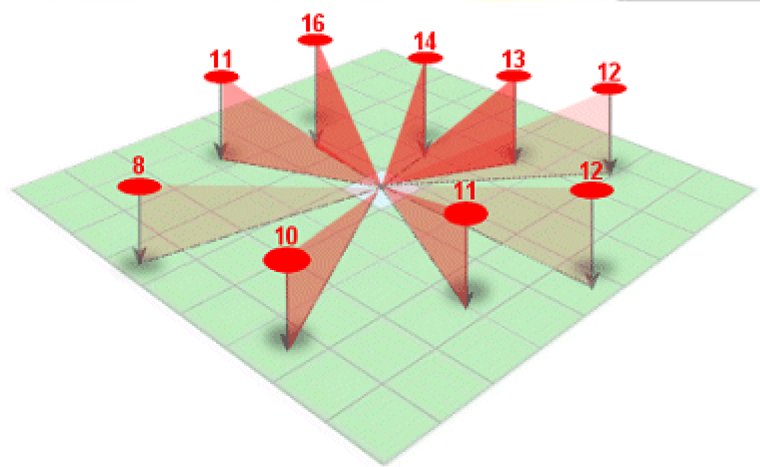

In this graphic, the darkest triangles indicate the most influential sample points.

Sample size

- Most interpolation methods allow you to control the number of sample points used to estimate cell values. For example, if you limit your sample to five points, the interpolator will use the five nearest points to estimate cell values.
- The distance to each sample point will vary depending on the distribution of the points. If you have a lot of sample points, reducing the size of the sample you use will speed up the interpolation process because a smaller set of numbers will be used to estimate each cell value.

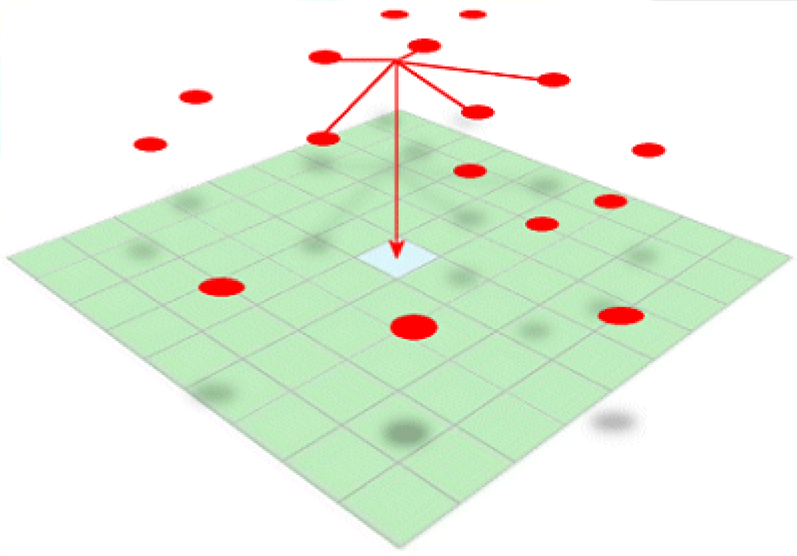

When the sample size is limited to five sample points, as in this case, only the five nearest points are used in the calculation of the estimated cell value. All other points are disregarded.

Sample size radius

- You can also control your sample size by defining a search radius. The number of sample points found within a search radius can vary depending on how the points are distributed. You can choose to use some or all of the samples that fall within this radius to calculate the cell value. A variable search radius will continue to expand until the specified sample size is found. A fixed search radius will use only the samples contained within it, regardless of how many or how few that might be.

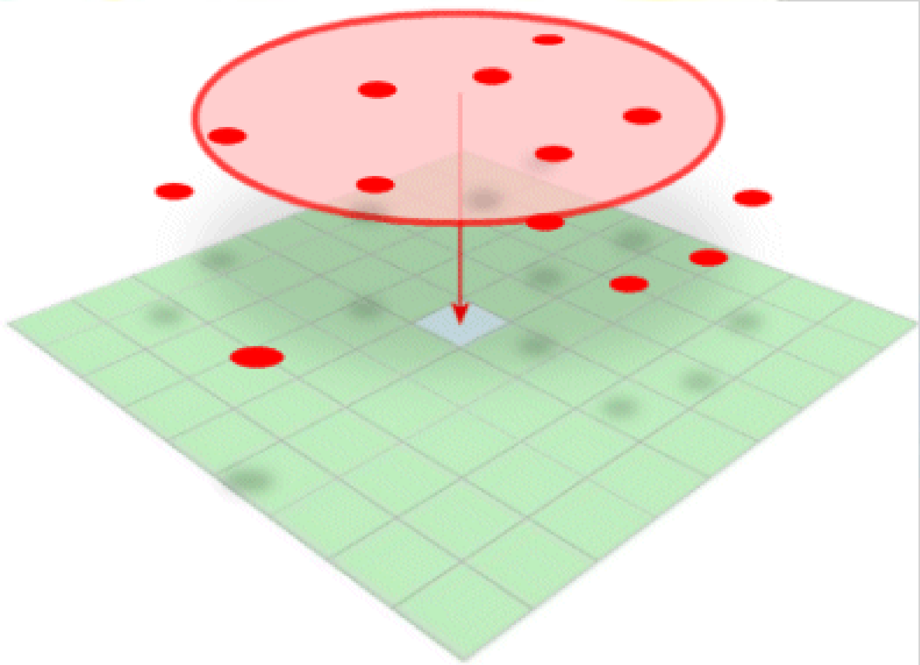

If the search radius in this sample were fixed, only the values of the sample points within the radius would be used to calculate the estimated cell value. If the search radius were variable and the minimum sample size were 8, the search radius would expand until it contained eight sample points.

Interpolation barriers

- The physical, geographic barriers that exist in the landscape, like cliffs or rivers, present a particular challenge when trying to model a surface using interpolation. The values on either side of a barrier that represents a sudden interruption in the landscape are drastically different.

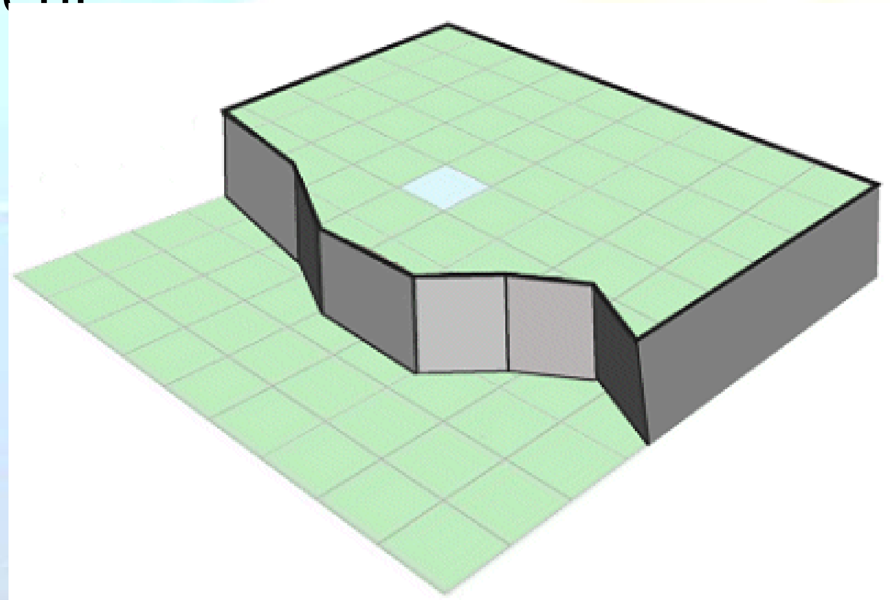

Elevation values change suddenly and radically near the edge of a cliff. When you interpolate a surface with this type of barrier, you can't use known values at the bottom of the cliff to accurately estimate values at the top of the cliff.

Interpolation barriers

- Most interpolators attempt to smooth over these differences by incorporating and averaging values on both sides of the barrier. The Inverse Distance Weighted method allows you to include barriers in the analysis. The barrier prevents the interpolator from using samples points on one side of it.

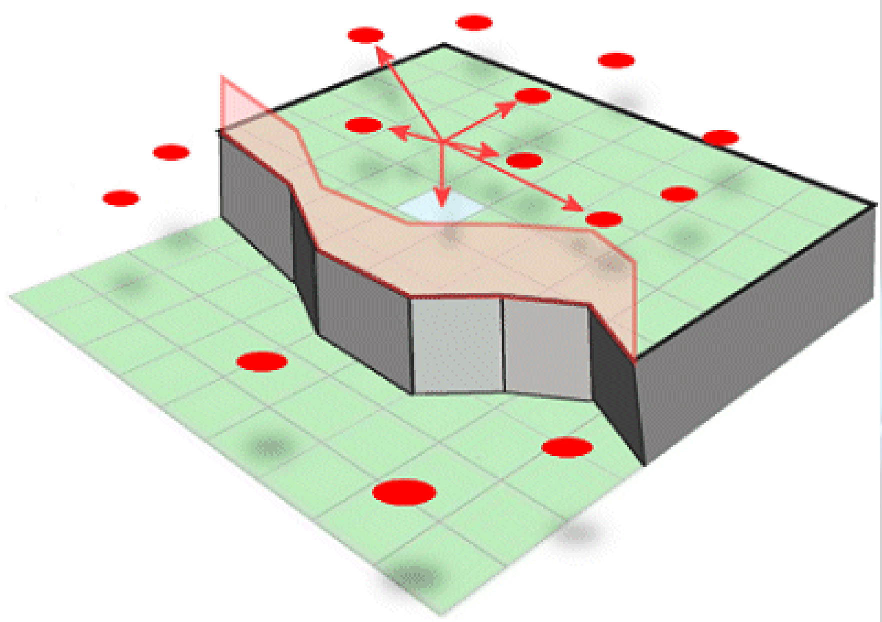

When you use a barrier with interpolation, the estimated cell value is calculated from sample points on one side of the barrier.

Interpolation methods

- **Inverse Distance Weighted** (IDW) takes the concept of spatial autocorrelation literally. It assumes that the nearer a sample point is to the cell whose value is to be estimated, the more closely the cell's value will resemble the sample point's value.
- **Spline** virtually guarantees you a smooth-looking surface. Imagine stretching a rubber sheet so that it passes through all of your sample points.
- **Kriging** is one of the most complex and powerful interpolators. It applies sophisticated statistical methods that consider the unique characteristics of your dataset. In order to use Kriging interpolation properly, you should have a solid understanding of geostatistical concepts and methods.

Inverse Distance Weighted

- IDW works best for dense, evenly-spaced sample point sets. It does not consider any trends in the data, so, for example, if actual surface values change more in the north-south direction than they do in the east-west direction (because of slope, wind, or some other factor), the interpolated surface will average out this bias rather than preserve it.
- IDW interpolation considers the values of the sample points and the distance separating them from the estimated cell. Sample points closer to the cell have a greater influence on the cell's estimated value than sample points that are further away.
- Inverse Distance Weighting cannot make estimates above the maximum or below the minimum sample values. For an elevation surface, this has the effect of flattening peaks and valleys (unless their high and low points are part of the sample). Because the estimated values are averages, the resulting surface will not pass through the sample points.

Inverse Distance Weighted

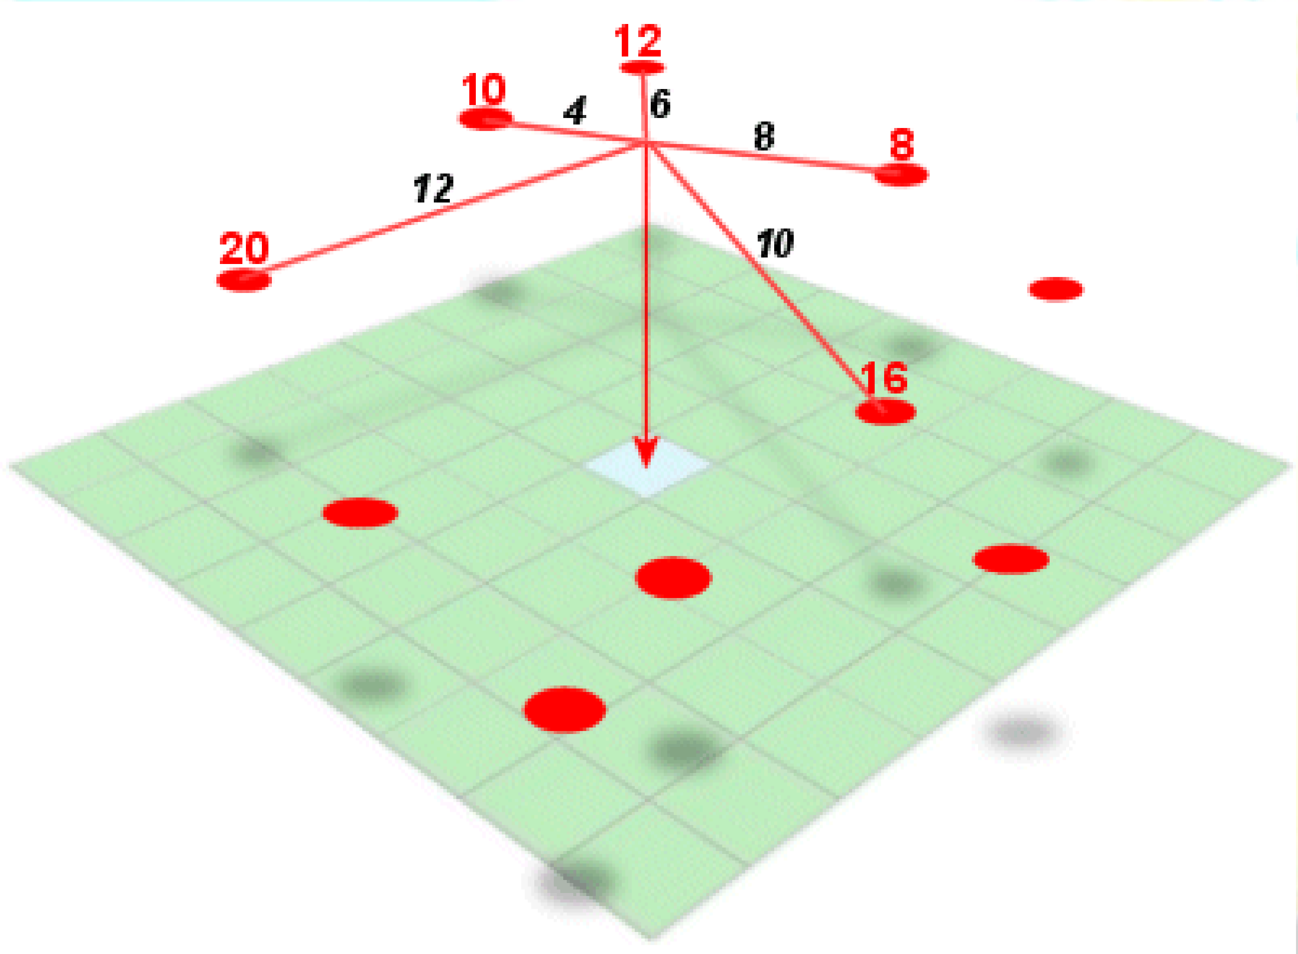

Each of the five sample points in this example have a different value and distance from the estimated cell.

Inverse Distance Weighted

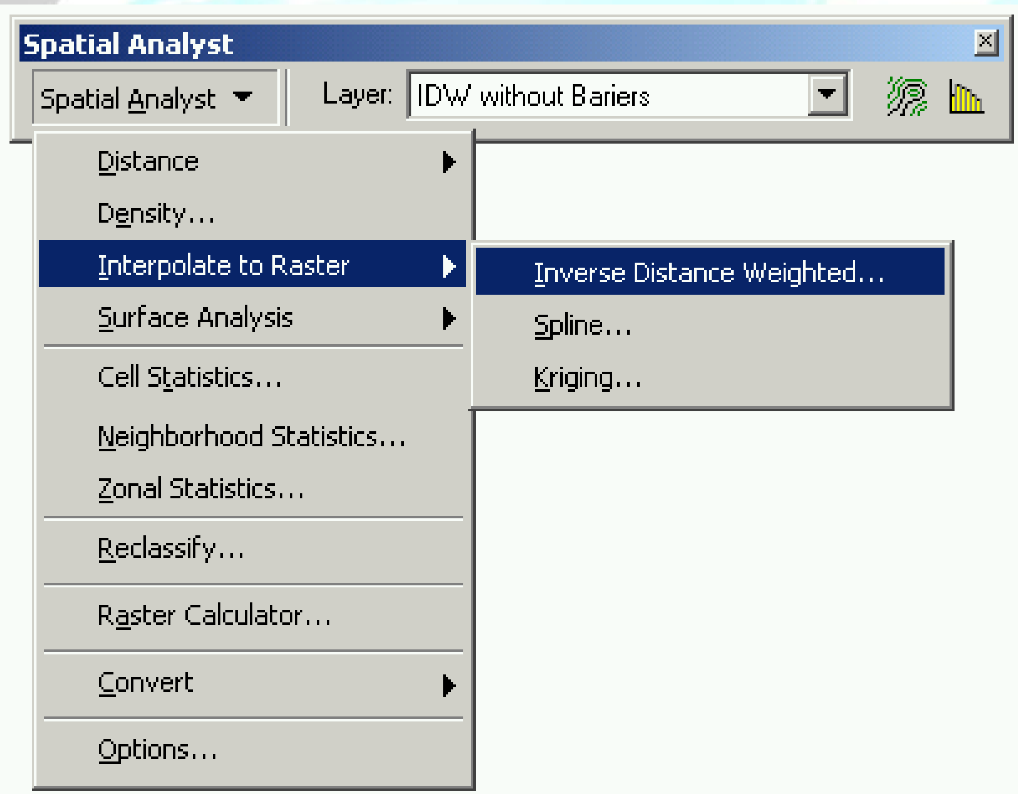

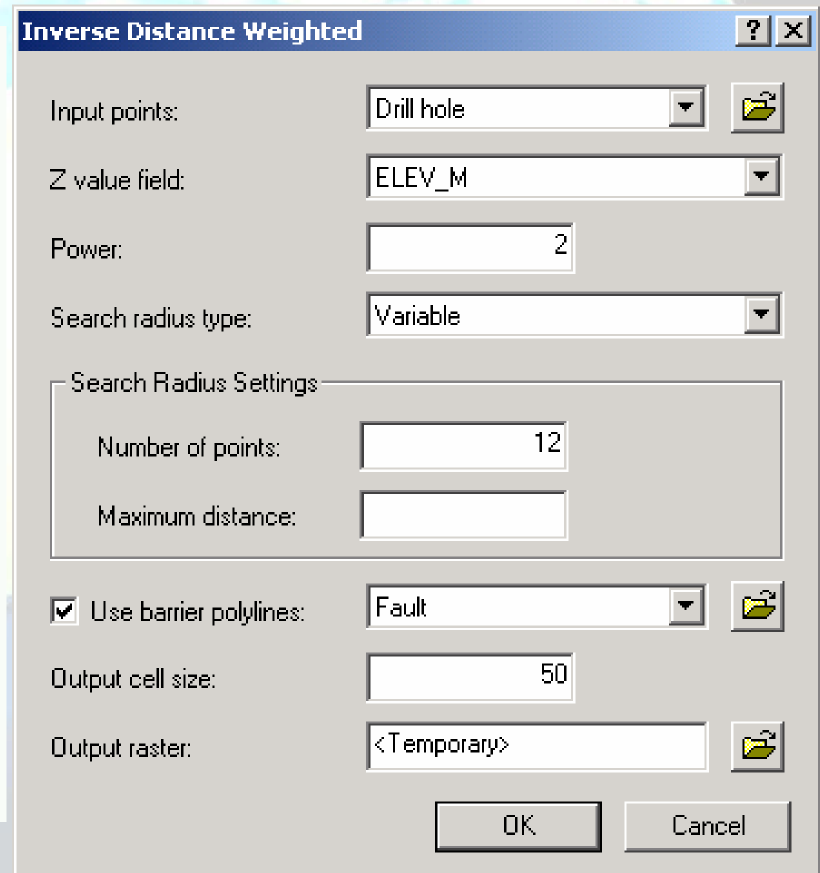

Adjusting Power Settings

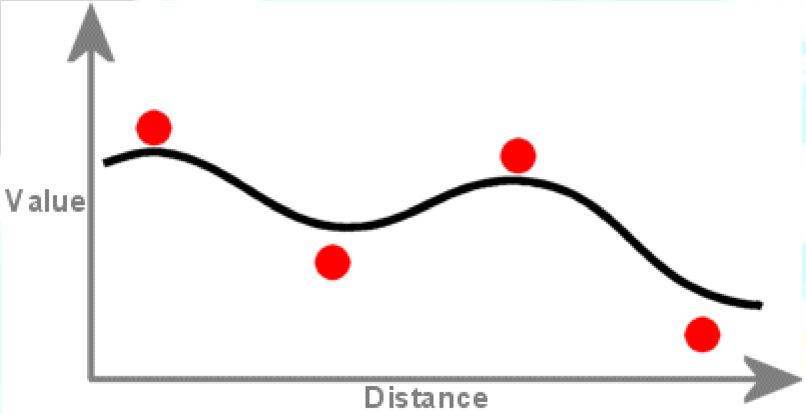

A surface created with IDW will not exceed the value range of the sample points or pass through those points.

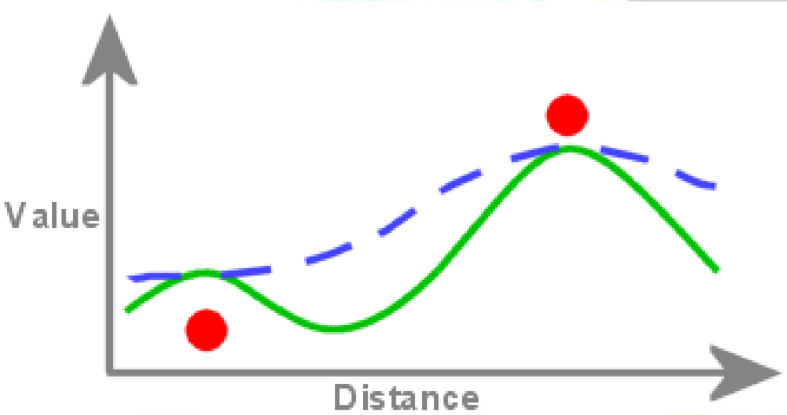

The solid line represents more power and the dashed line represents less power. The higher the power, the more localized an affect a sample point's value has on the resulting surface.

- You can adjust the relative influence of sample points. In other words, you can increase how much power the values of sample points have over the interpolation process. Increased power means that the output cell values become more localized and less averaged. Their influence, however, drops off rapidly with distance.
- Lowering the power that sample point values have provides a more averaged output because sample points farther away become more and more influential until all of the sample points have the same influence.

What happens when sample points strongly influence estimated cell values?

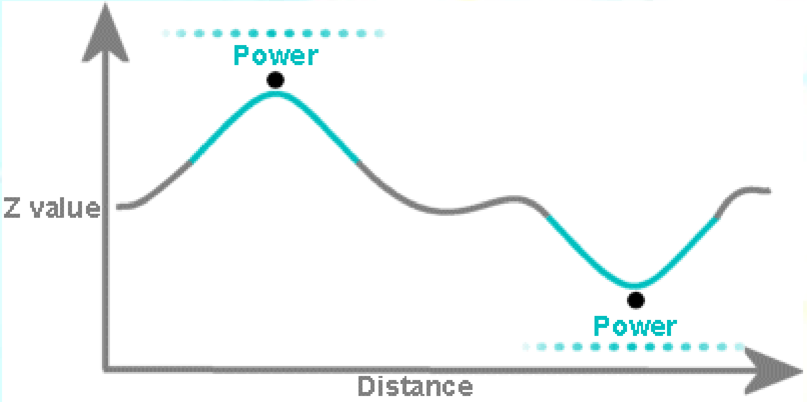

- Sample points are weighted according to their distance from the cell being evaluated. Closer sample points are given more weight than points farther away. As the Power setting increases, the influence of sample points falls off more rapidly with distance. The output cell values become more localized and less averaged. A high power setting will result in bumps and depressions in the surface, which are localized around sample point locations.

Smoothing an IDW Surface

- Decreasing the power will smooth an IDW surface.
- Increase the number of sample points used in the interpolation, or use all sample points within a radius and use a large radius.

Adjusting the Power Settings

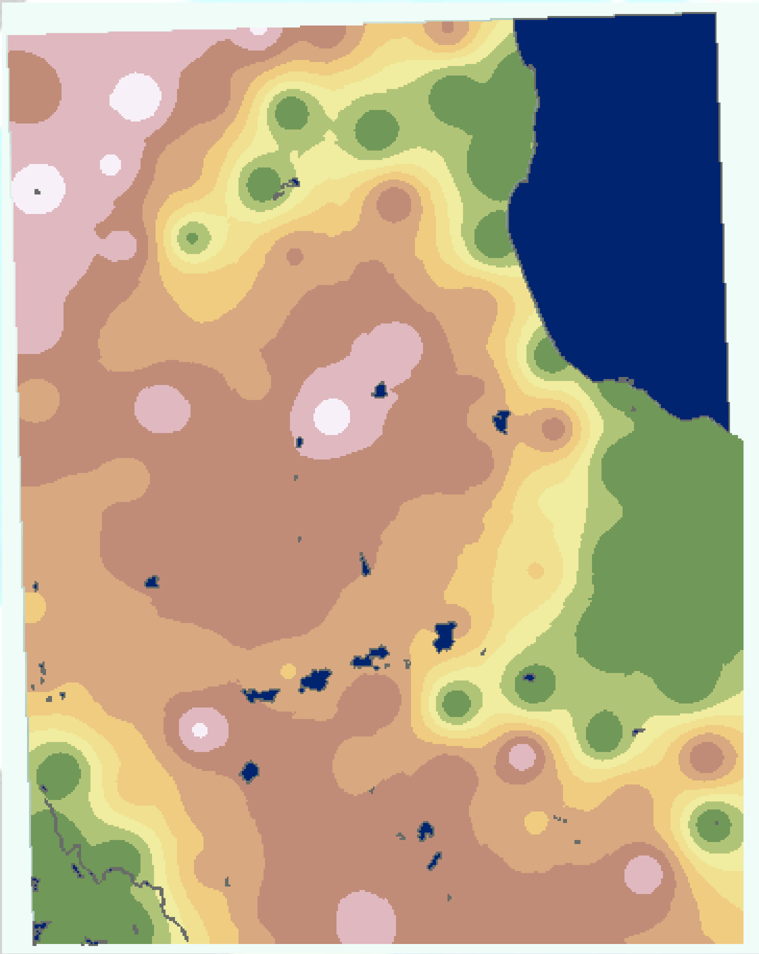

Power of 2

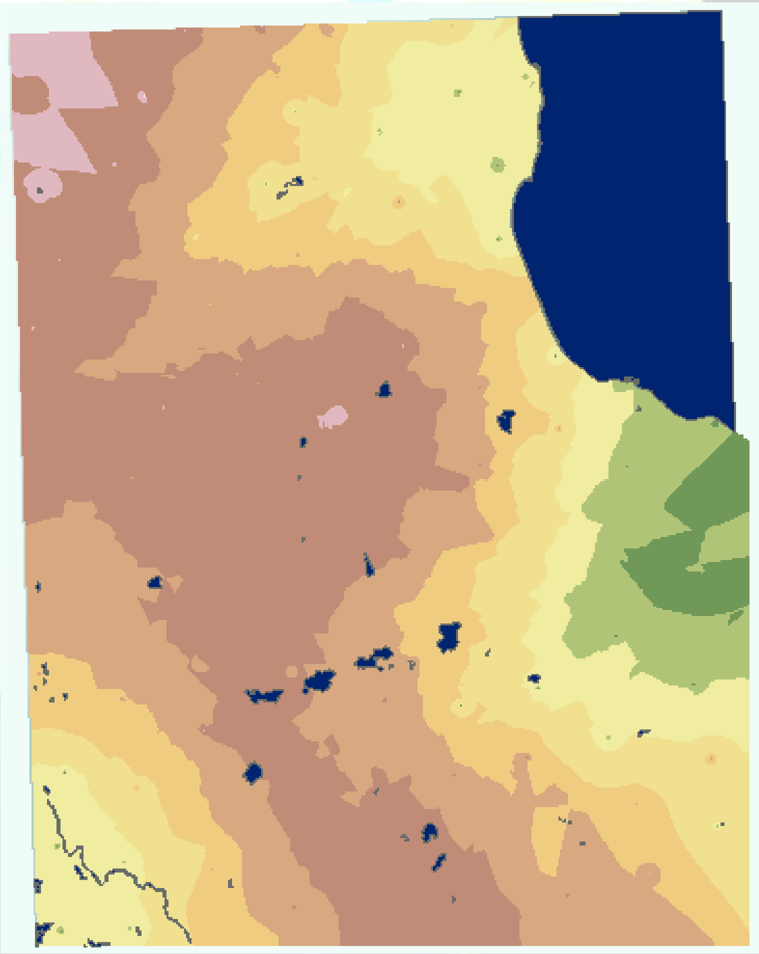

Power of 0.5

Search radius size determination

- It depends on your sample point data. If the points are evenly distributed but far apart, you may need to use a larger search radius with a smaller number of minimum number of points. If the points are evenly distributed but close together, your search radius can be smaller, depending on what you decide for the minimum number of points. If the points are close together in some places and far apart in others, you should consider using a variable search radius.
- You can visually experiment with a search radius by creating a graphic circle (see the Draw toolbar) and setting the width to twice the radius distance. Move the circle to different places on the map, counting the number of sample points contained within it. From this you should be able to determine the appropriate size for a search radius and what the minimum number of sample points should be.

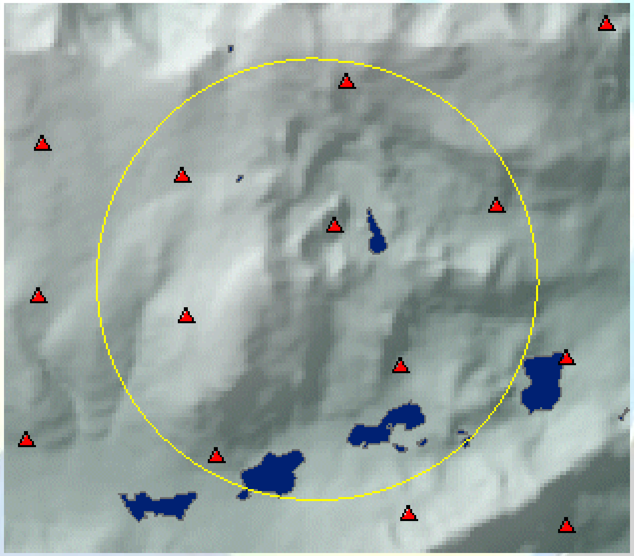

IDW-Key Points

- IDW is a good interpolator for a phenomenon whose distribution is strongly correlated with distance. A classic example is noise, which falls off very predictably with distance.
- IDW does less well with phenomena whose distribution depends on more complex sets of variables because it can account only for the effects of distance.
- One potential advantage of IDW is that it gives you explicit control over the influence of distance; an advantage you don't have with Spline or Kriging.
- You can create a smoother surface by decreasing the power, increasing the number of sample points used, or increasing the search radius. To create a more locally influenced surface, do the opposite.
- You may be able to improve the accuracy of an IDW surface by using line layers as barriers. On elevation surfaces, barriers can represent abrupt changes in elevation, such as cliffs.

Spline

- Instead of averaging values, like IDW does, the Spline interpolation method fits a flexible surface, as if it were stretching a rubber sheet across all the known point values.

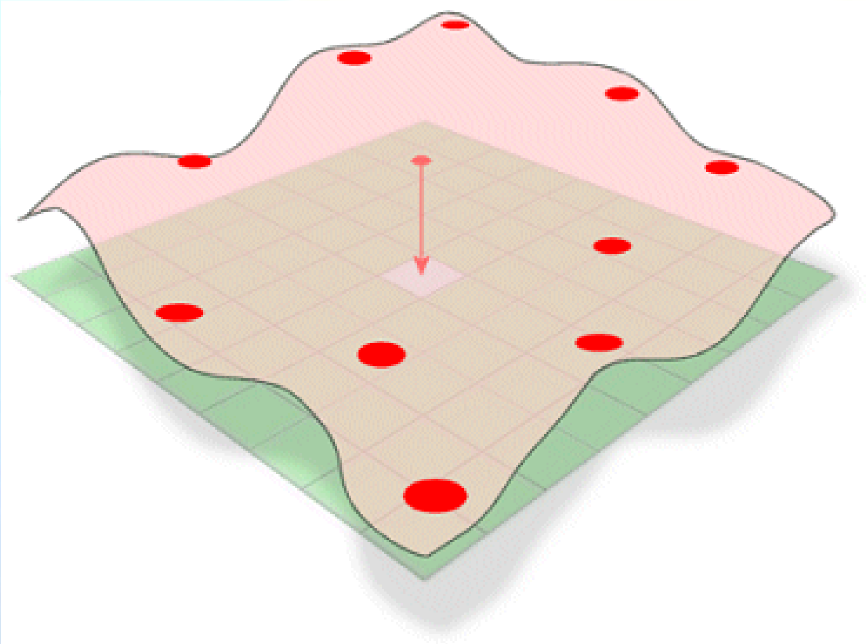

The Spline method of interpolation estimates unknown values by bending a surface through known values.

Spline

- This stretching effect is useful if you want estimated values that are below the minimum or above the maximum values found in the sample data. This makes the Spline interpolation method good for estimating lows and highs where they are not included in the sample data.

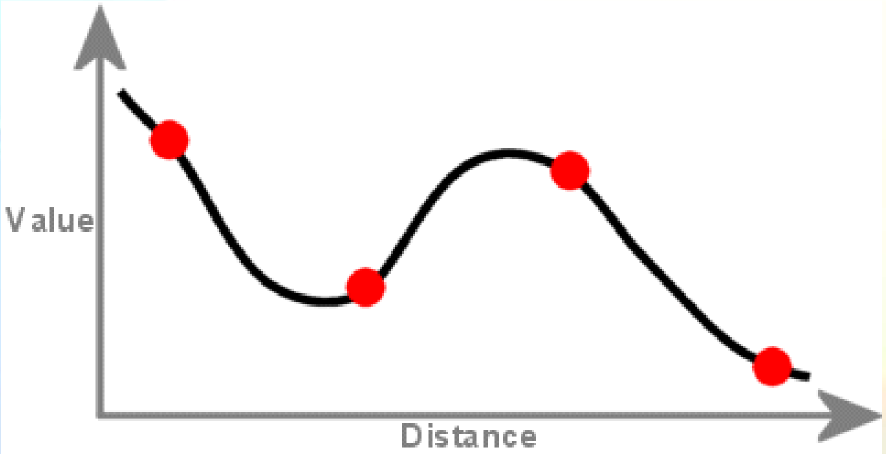

A surface created with Spline interpolation passes through each sample point and may exceed the value range of the sample point set.

Spline Limitations

- When the sample points are close together and have extreme differences in value, Spline interpolation doesn't work as well. This is because Spline uses slope calculations (change over distance) to figure out the shape of the flexible rubber sheet.
- Phenomena that cause surface values to change suddenly, such as a cliff face or a fault line, are not represented well by a smooth-curving surface. In such cases, you might prefer to use IDW interpolation, where barriers can be used to deal with these types of abrupt changes in local values.

Types of Spline

- Regularized and Tension
- A Tension Spline is flatter than a Regularized Spline of the same sample points, forcing the estimates to stay closer to the sample data. You might say that the Tension Spline method produces a surface more rigid in character, while the Regularized Spline method creates one that's more elastic.

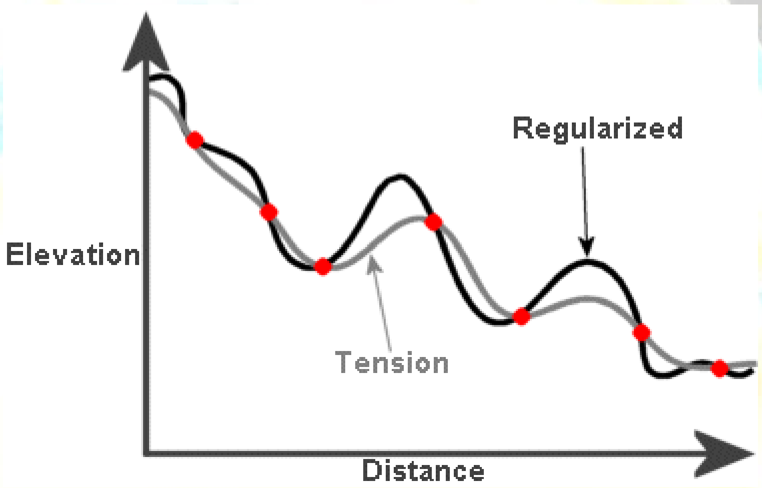

Notice that the tension curve is flatter than the regularized curve. The estimates are forced to stay closer to the sample data. You might say that the Tension Spline method produces a surface more rigid in character, while the Regularized Spline method creates one more elastic in character.

Spline Settings

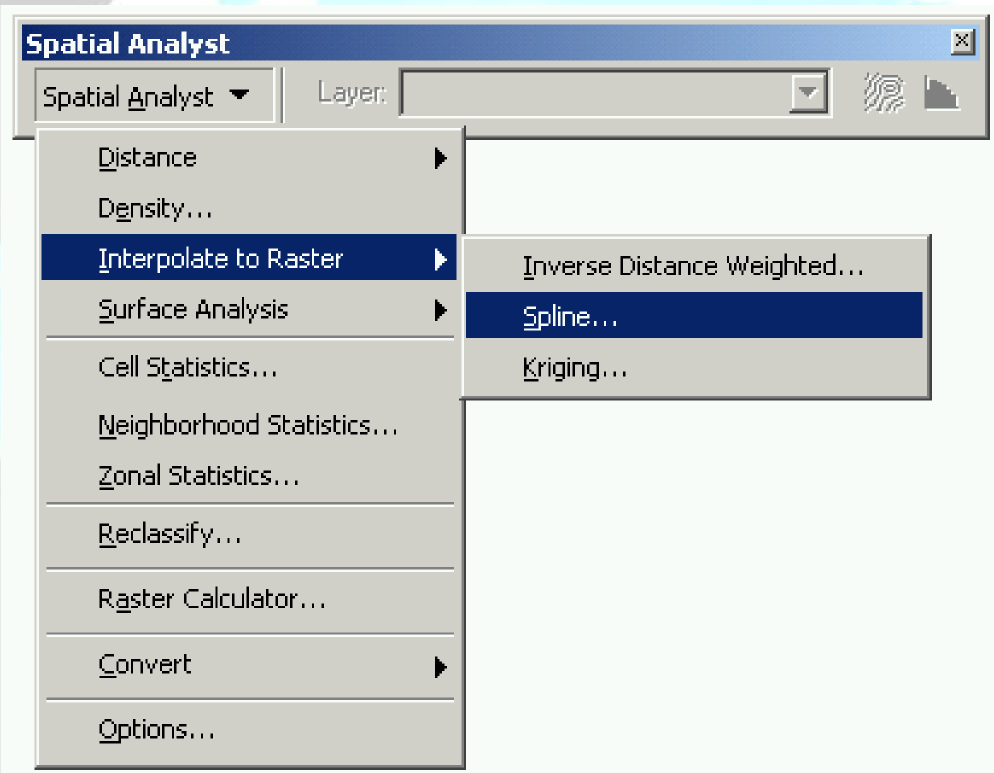

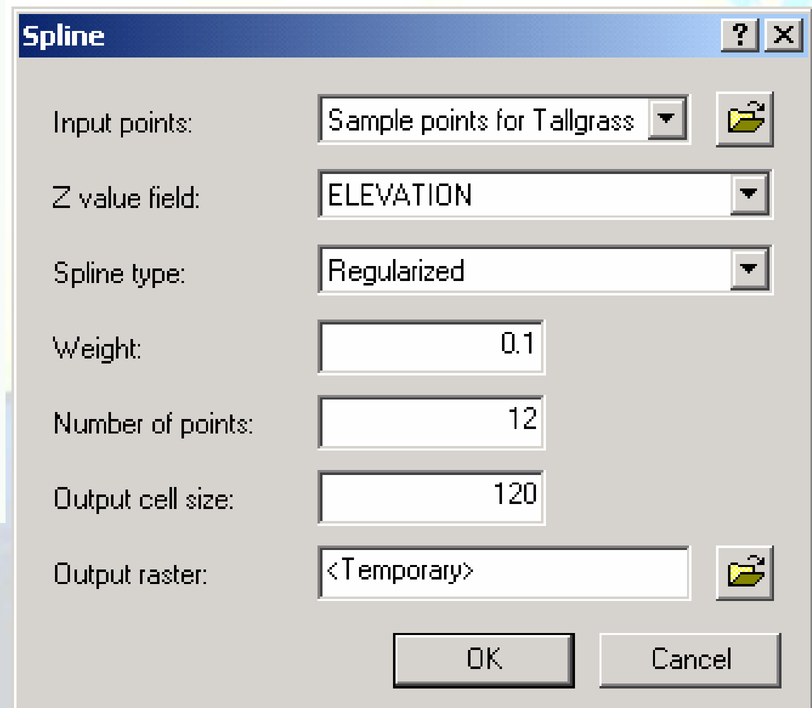

Spline Weights

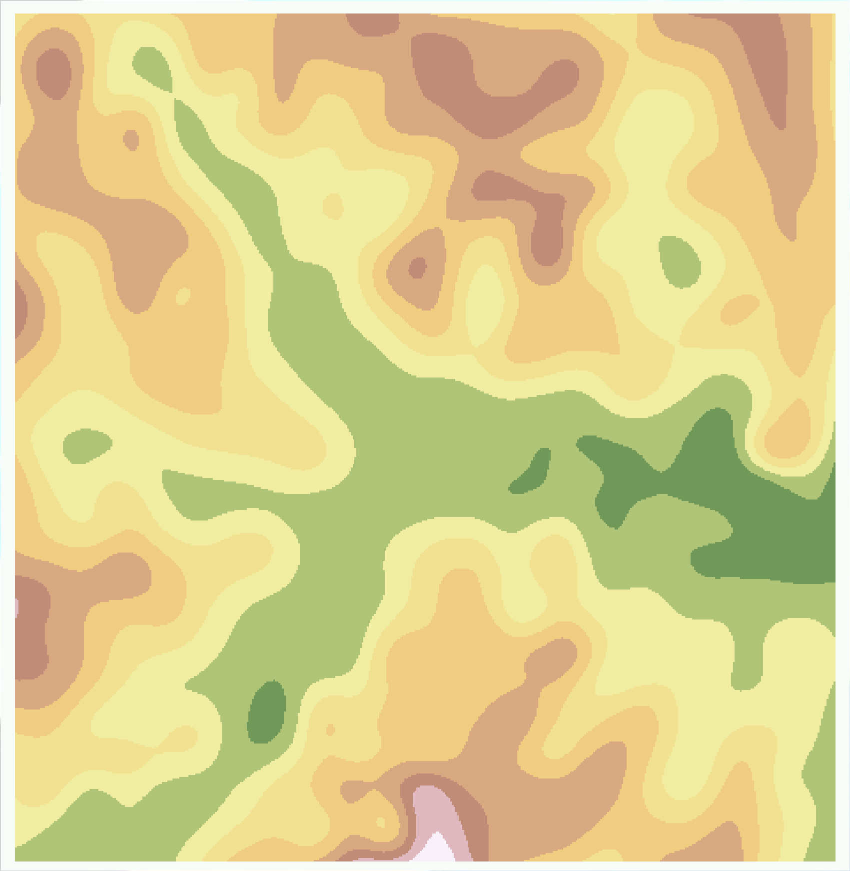

Spline Regular 0.1

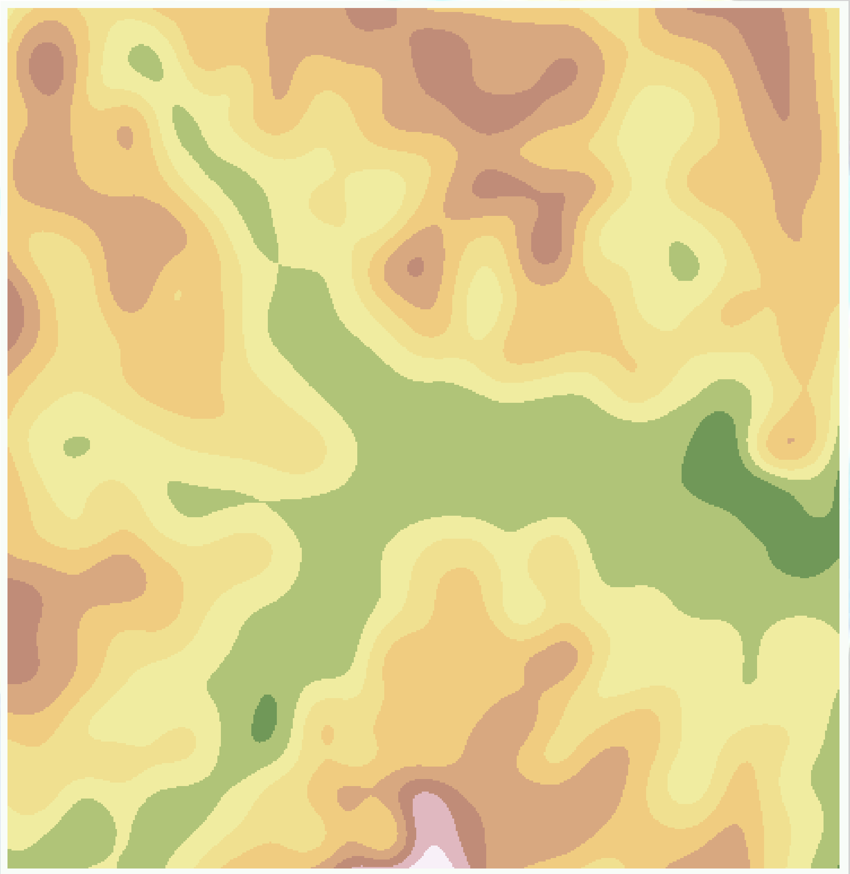

Spline Regular 1

Spline Types

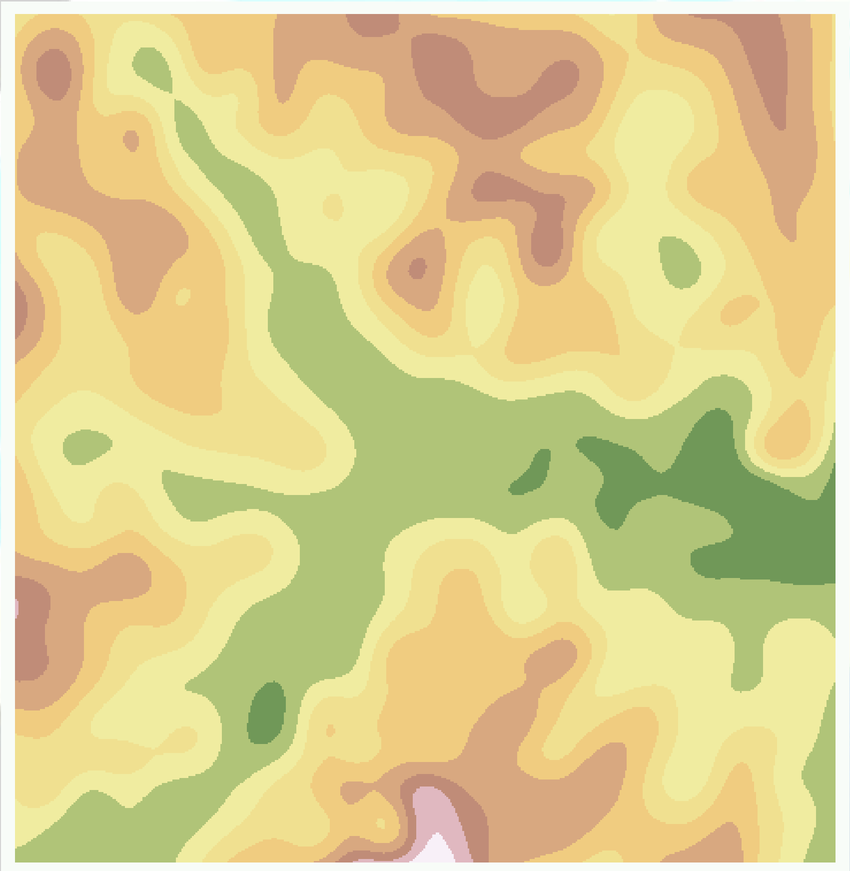

Spline Regular 0.1

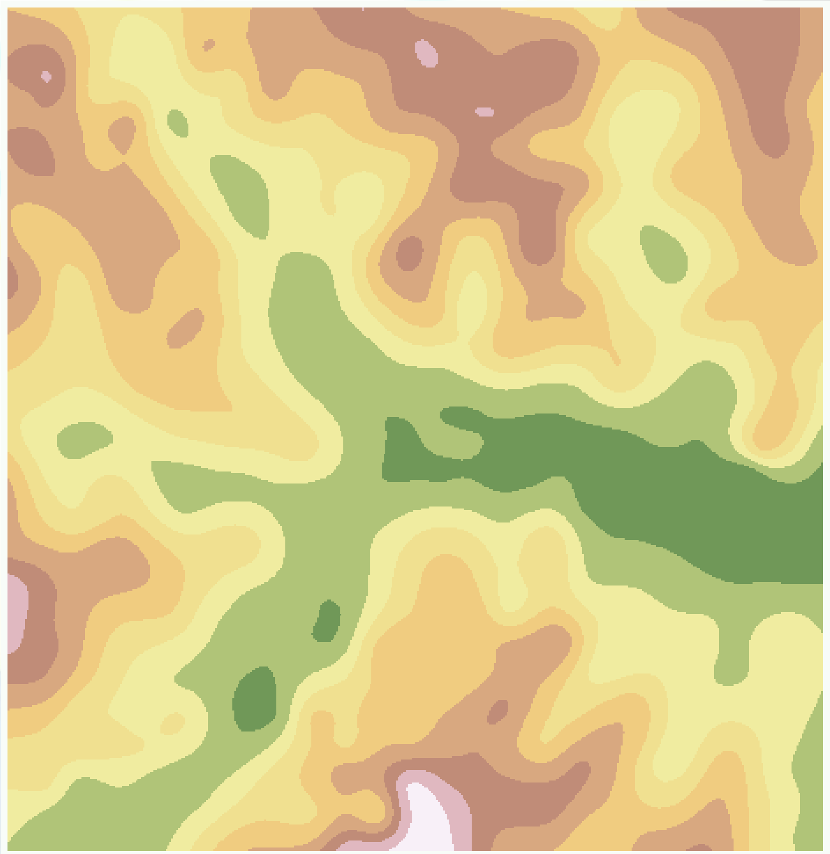

Spline Tension 0.1

Spline – Key Points

- An advantage of the Spline interpolator is that it can make estimates outside the range of input sample points.
- There are two types of Spline interpolators. The regularized spline creates a more elastic surface. The tension spline creates a less flexible surface.

Kriging

- Kriging (like IDW) is a weighted average technique, except that the weighting formula in Kriging uses much more sophisticated math. Kriging measures distances between all possible pairs of sample points (that’s right, all of them) and uses this information to model the spatial autocorrelation for the particular surface you're interpolating.
- In other words, Kriging tailors its calculations to your data by analyzing all the data points to find out how much autocorrelation

Kriging

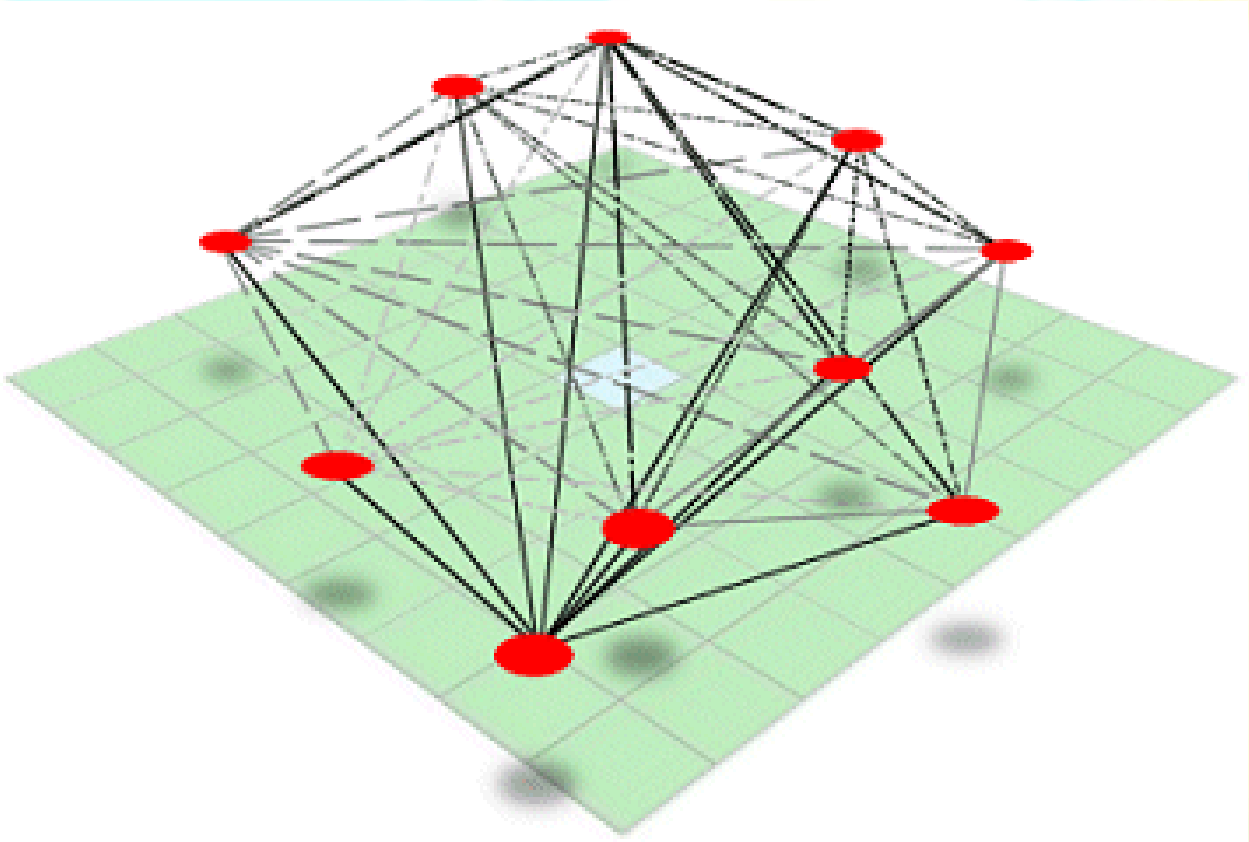

When you interpolate a surface using Kriging, the distance and direction of every point pair is quantified to provide information on the spatial autocorrelation of the sample point set. Next, a best-fit model is automatically applied to the data and the unknown values are predicted.

Kriging

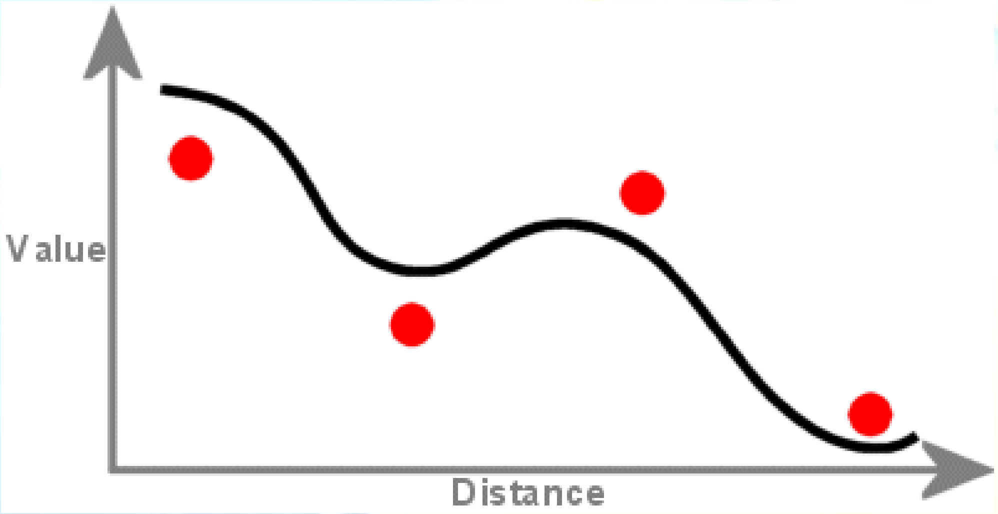

A surface created with kriging can exceed the value range of the sample points, but will not pass through the points.

Kriging

- Kriging aficionados consider the initial kriged surface a first draft—a test surface against which they compare future iterations as they search for the perfect surface. Directional influences, such as prevailing winds and random error, can be accounted for using Kriging, but you will need a statistical tool such as ArcGIS™ Geostatistical Analyst to visualize these trends.

Kriging Methods

- Two general and widely used Kriging methods are Ordinary and Universal Kriging.
- **Universal Kriging** assumes that there is an overriding trend in the data. For example, you may know that there is a prevailing wind or a gently sloping hillside across your study area.
- **Ordinary Kriging** assumes there is no trend in the data, which should be your standard operating assumption.

Kriging Settings

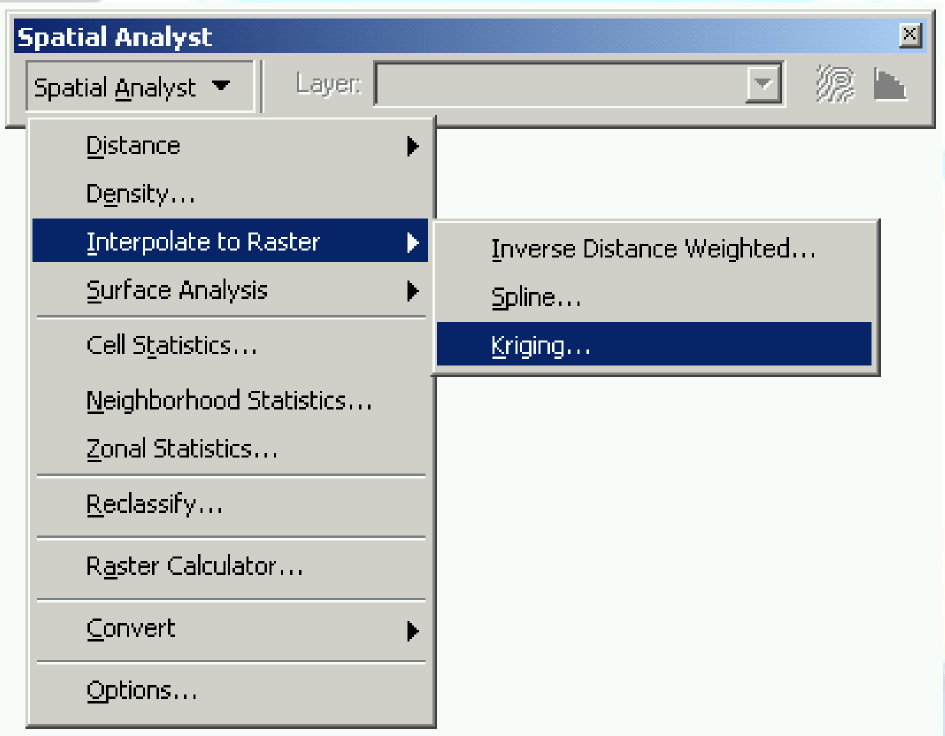

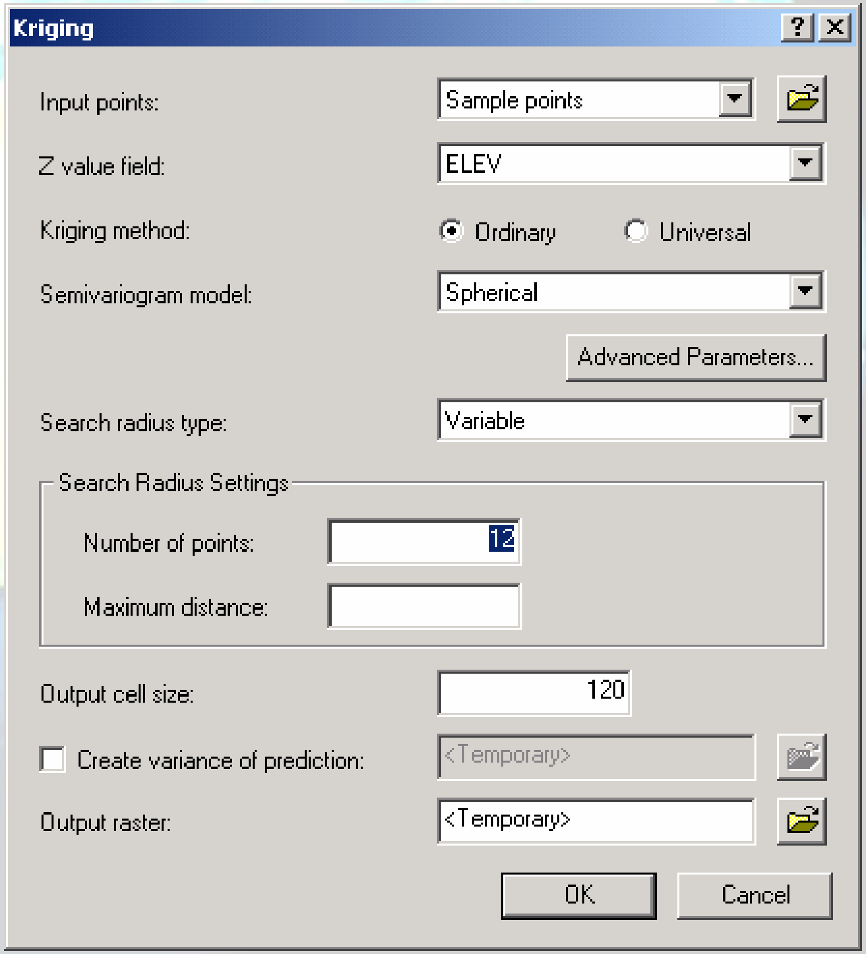

Kriging Methods

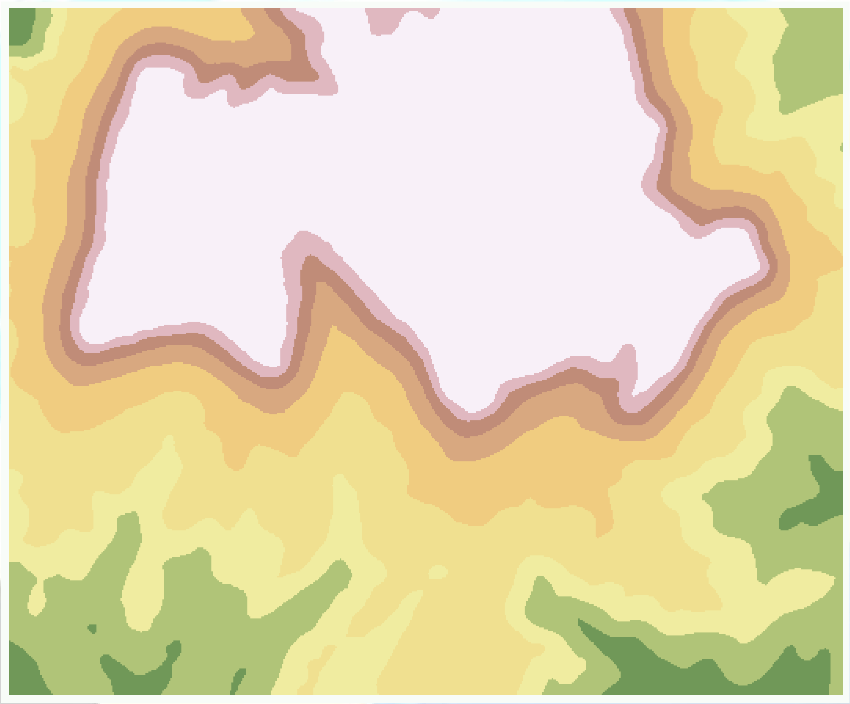

Ordinary Kriging

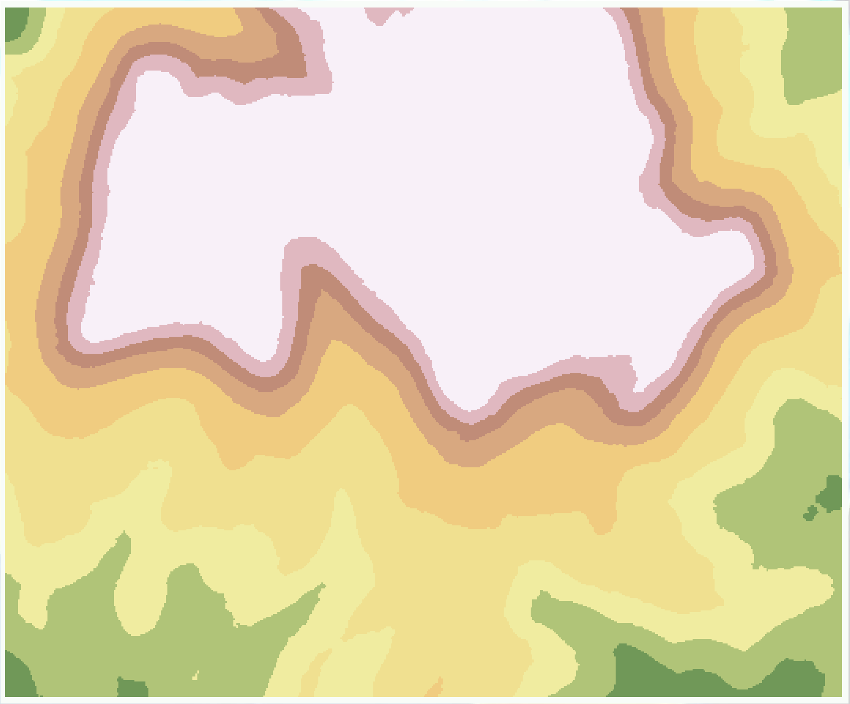

Universal Kriging

Which Interpolation Method?

- The type of interpolation method you use will depend on many factors. Rather than assume one interpolation method is better than another, you should try different interpolation methods and compare the results to determine the best interpolation method for a given project.
- Your real-world knowledge of the subject matter will initially affect which interpolation method you use. If you know that some of the features in your surface exceed the z value, for example, and that IDW will result in a surface that does not exceed the highest or lowest z value in the sample point set, you might choose the Spline method.
- If you know that the splined surface might end up with features that you know don't exist because Spline interpolation doesn't work well with sample points that are close together and have extreme differences in value, you might decide to try IDW.
- The quality of your sample point set can affect your choice of interpolation method as well. If the sample points are poorly distributed or there are few of them, the surface might not represent the actual terrain very well. If you have too few sample points, you might experiment with adding more sample points in areas where the terrain changes abruptly or frequently, then try using Kriging.

Which Interpolation Method?

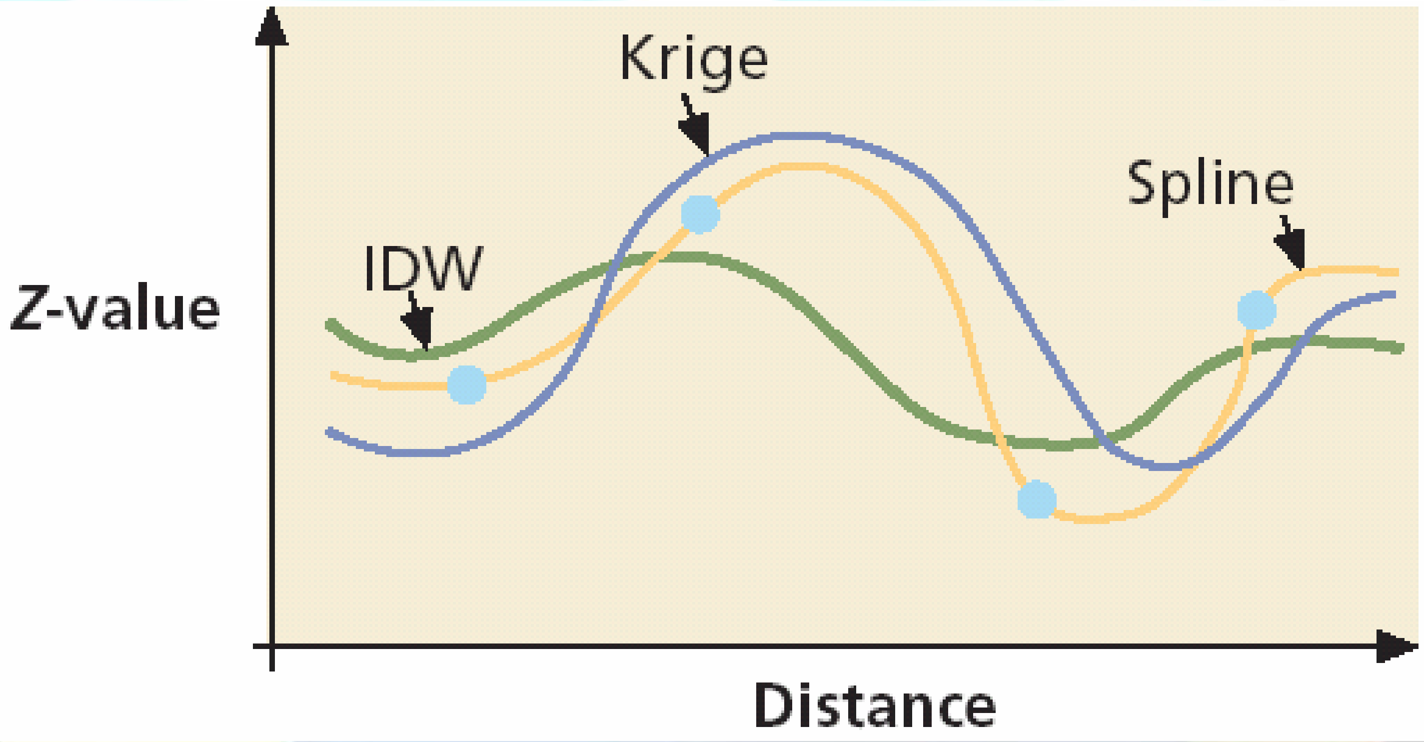

Summary

- When you measure elevation, the depth of a well, or the level of noise, you make that measurement at a precise location. A point layer can represent a set of measurements. The location of the points and the point values form the basis for interpolation. Interpolation is a method of estimating unknown values based on known values.
- There are different methods of interpolation. Choosing an interpolation method is influenced by your knowledge of the surface you are modeling. Each method works differently, but most utilize the concept of spatial autocorrelation; near things are more alike than things far away.

Summary

- IDW interprets spatial autocorrelation in a literal fashion. A surface created with IDW will not exceed the known value range or pass through any of the sample points. IDW is a good interpolator for phenomena whose distribution is strongly correlated with distance, such as noise. In some cases, the accuracy of an IDW surface can be improved by using line layers as barriers.
- The Spline interpolation method incorporates a curvilinear model as part of the calculation. A surface created with Spline can exceed the known value range, but must pass through all of the sample points.
- Kriging is one of the most complex interpolators. It measures the relationships between all of the sample points and then predicts the cell value. A surface created with Kriging can exceed the known value range, but does not pass through any of the sample points.