

Topic 4

Spatial Interpolation

4.1 From Point Samples to Map Surfaces

Soil sampling has long been at the core of agricultural research and practice. Traditionally point-sampled data were analyzed by non-spatial statistics to identify the “typical” nutrient level throughout an entire field. Considerable effort was expended to determine the best single estimate and assess just how good the “average” estimate was in typifying a field.

However non-spatial techniques fail to make use of the geographic patterns inherent in the data to refine the estimate—the typical level is assumed everywhere the same within a field. The computed variance (or standard deviation) indicates just how good this assumption is—the larger the standard deviation the less valid is the assumption “everywhere the same.”

Surface modeling, on the other hand, utilizes the spatial patterns in a data set to generate localized estimates throughout a field.

Conceptually it “maps the variance” by using geographic position to help explain the differences in the sample values. In practice, it simply fits a continuous surface (kind of like a blanket) to the point data spikes (figure 4-1).

While the extension from non-spatial to spatial statistics is quite a theoretical leap, the practical steps are relatively easy. The left side of figure 4-1 shows 2D and 3D *point* maps of phosphorous soil samples collected throughout

the field. This highlights the primary difference from traditional soil sampling—each sample must be *geo-referenced* as it is collected. In addition, the *sampling pattern and intensity* are often different than traditional grid sampling to maximize spatial information within the data collected*.

The surface map on the right side of figure 4-1 depicts the continuous spatial distribution derived from the point data. Note that the high spikes in the western portion of the field and the relatively low measurements in the center are translated into the peaks and valleys of the surface map.

The traditional, non-spatial approach, if mapped, would be a flat plane (average phosphorous level) aligned within the bright yellow zone. Its “everywhere the same” assumption fails to recognize the patterns of larger levels (greens) and smaller levels (reds). A fertilization plan for phosphorous based on the average level (22ppm) would be ideal for the yellow zone but would likely be

inappropriate for a lot of the field as the data vary from 5 to 102ppm phosphorous.

It’s common sense that nutrient levels vary throughout a field. However, there are four important considerations before surface modeling is appropriate—

- 1) the pattern needs to be significant enough to warrant variable-rate application (*cost effective*),

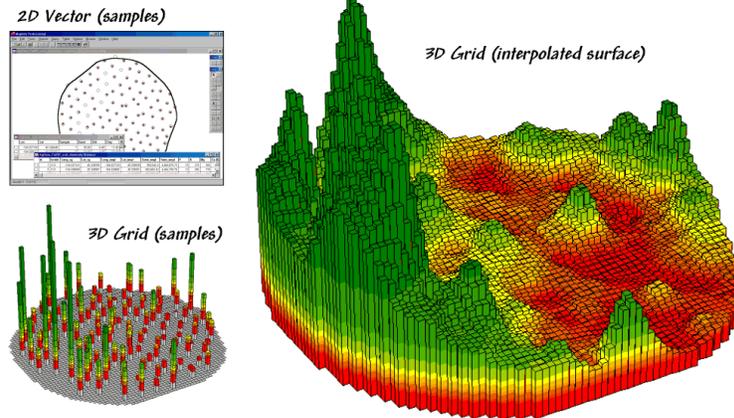
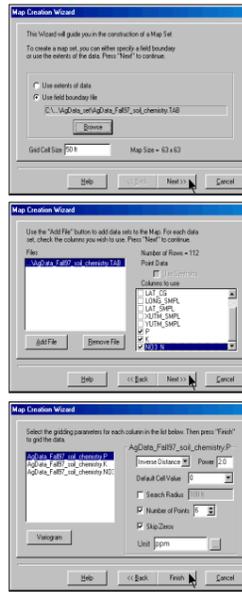


Figure 4-1. Spatial interpolation involves fitting a continuous surface to sample points.

- 2) the prescription methods must be spatially appropriate for the pattern (*decision rules*),
- 3) the derived map surface must fully reflect a real spatial pattern (*map accuracy*), and
- 4) the procedure for generating the pattern must be straightforward and easy to use (*usability*).

Let's tackle the last point on usability first. The mechanics of generating an interpolated surface involves three steps relating to *geography*, *data* and *interpolation methods* (figure 4-2).



Step 1 – Geographic Specifications...
 ✓ Positioning
 ✓ Extent
 ✓ Grid configuration

Step 2 – Sample Data Specifications...
 ✓ Table to use
 ✓ Data fields

Step 3 – Interpolation Specifications...
 ✓ Method
 ✓ Parameters

Figure 4-2. A wizard interface guides a user through the necessary steps for interpolating sample data.

the default *Inverse Distance Squared (1/d²)* method was employed using the six nearest sample points. Other methods, such as Kriging, could be specified and would result in a somewhat different surface.

Figure 4-3 shows the P, K, N map surfaces that were generated by the interpolation specifications shown in figure 4-2. The entire process including entering specifications, calculation of the surfaces and display took less than a minute.

In each instance the yellow zone contains the average for the entire field. Visually comparing the

The first step establishes the geographic position, field extent and grid configuration to be used. While these specifications can be made directly, it is easiest to simply reference an existing map of the field boundary (*Positioning and Extent*) then enter the grid spacing (*Configuration—50 feet*).

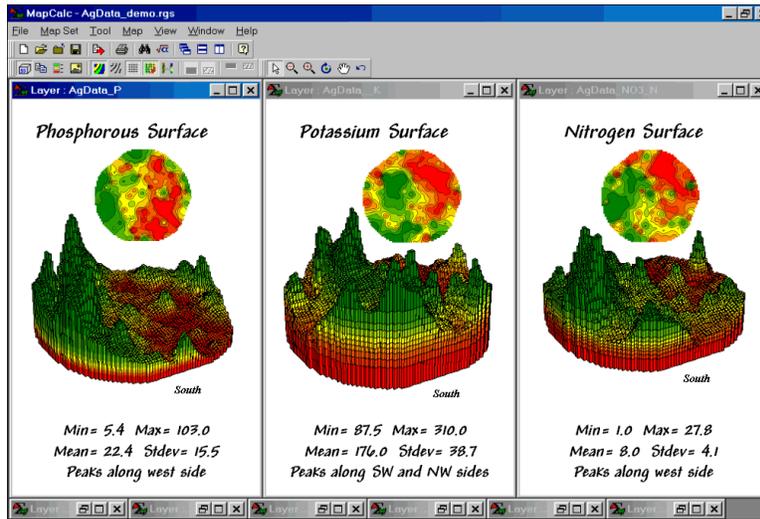


Figure 4-3. Interpolated phosphorous, potassium and nitrogen surfaces.

relative amounts and patterns of the green and red areas gives you an idea of the difference between the assumption that the average is everywhere and amount of spatial information inherent in a soil sample data set.

In the next section we'll investigate how to evaluate the accuracy of these maps—

whether they're good maps or bad maps.

4.2 The Keystone Concept

Spatial Interpolation's basic concept involves *spatial autocorrelation*, referring to the degree of similarity among neighboring points (e.g., soil nutrient samples). If they exhibit a lot similarity, or *spatial dependence*, they ought to derive a

The next step identifies the table and data fields to be used. The user navigates to the file (*Data Table*) then simply checks the maps to be interpolated (*data fields—P, K, NO3_N*).

The final step establishes the interpolation method and necessary factors. In the example,

good map. If they are spatially independent, then expect a map of pure, dense gibberish. So how can we measure whether "what happens at one location depends on what is happening around it?"

Common sense leads us to believe more similarity exists among the neighboring soil samples (lines in the left side of figure 4-4) than among sample points farther away. The *Geary* and *Moran Indices* consider the closest neighbors. They compute the differences in the values between each sample point and its closest neighbor.

If the differences in neighboring values are a lot smaller than the overall variation among all of the samples in the data set, then a high degree of positive spatial dependency is indicated. If they are about the same or if the neighbors variation is larger (a rare "checkerboard-like" condition), then the assumption of spatial dependence fails. If the dependency test fails, it means the soil nutrient map likely is just colorful gibberish. You should ask to see the Geary and Moran Indices before you accept any soil nutrient map.

These indices, however, are limited as they merely assess the closest neighbor, regardless of its distance. That's where a *variogram* comes in. It is a plot of the similarity among values based on the distance between them. Instead of simply testing whether close things are related, it shows how the degree of dependency relates to varying distances between locations. The origin of the plot at $0,0$ is a unique case. The distance between samples is zero; therefore, there is no dissimilarity (data variation = 0) as at location is exactly the same as itself.

As the distance between points increase, subsets of the data are scrutinized for their dependency**. The shaded portion in the idealized plot shows how quickly the spatial dependency among points deteriorates with distance. The maximum range (*Max Range*)

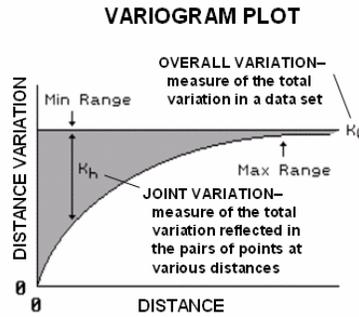
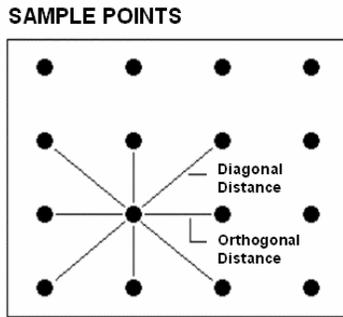


Figure 4-4. Variogram plot depicts the relationship between distance and measurement similarity (spatial autocorrelation).

position identifies the distance between points beyond which the data values are considered independent. This tells us that using data values beyond this distance for interpolation actually can mess-up the interpolation.

The minimum range (*Min Range*) position identifies the smallest distance contained in the actual data set and is determined by the sampling design used to collect the data. If a large portion of the shaded area falls below this distance, it tells you there is insufficient spatial dependency in the data set to warrant interpolation. If you proceed with the interpolation, a nifty colorful map will be generated, but likely questionable. Worse yet, if the sample data plots as a straight line or circle, no spatial dependency exists and the map will be worthless.

You should ask to see the variogram plot, as well as the Geary and Moran Indices, before you bet the farm on any soil nutrient map. If they are unavailable, then the map is just "a pig in a poke."

4.3 Benchmarking Interpolation Results

For some, the previous discussion on generating maps from soil samples might have been too simplistic—enter a few things then click on a data file and, *viola*, you have a soil nutrient surface artfully displayed in 3D with a bunch of cool colors draped all over it.

Actually, it is that easy to create one. The harder part is figuring out if the map generated makes sense and whether it is something you ought to use in analysis and important management decisions. This section discusses the relative amounts of spatial information provided in a whole-field average and site-specific map by comparing the average and two different interpolated map surfaces. The next section describes a procedure to quantitatively assess whether a particular map is a keeper.

The top-left item in Figure 4-5 shows the map of a field's average phosphorous levels. It's not very exciting and looks like a pancake but that's because there isn't any information about spatial variability in an average value—assumes 22ppm is everywhere in the field.

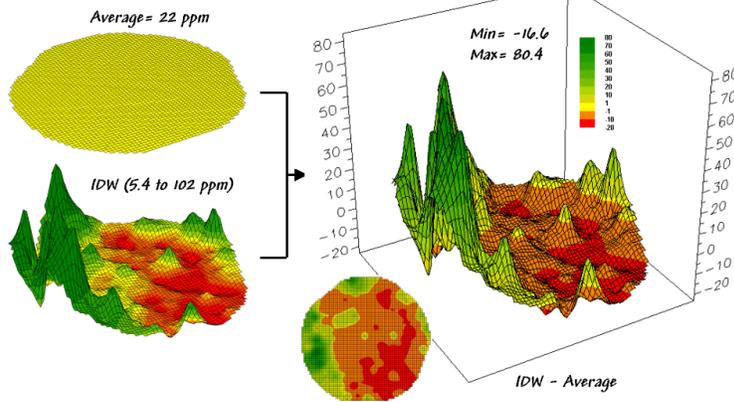


Figure 4-5. Spatial comparison of a whole-field average and an IDW interpolated map.

The non-spatial estimate simply adds up all of the sample measurements and divides by the number of samples to get 22ppm. Since the procedure didn't consider where the different samples were taken, it can't map the variations in the measurements. It thinks the average is everywhere, plus or minus the standard deviation. But there is no spatial guidance where phosphorous levels might be more, or where they might be less.

The spatially based estimates are shown in the map surface just below the pancake. As described in the last section, spatial interpolation looks at the relative positioning of the soil samples as well as their measure phosphorous levels. In this instance the big bumps were influenced by high measurements in that vicinity while the low areas responded to surrounding low values.

The map surface in the right portion of figure 4-5 objectively compares the two maps simply by

subtracting them. The colors were chosen to emphasize the differences between the whole-field average estimates and the interpolated ones.

The yellow band indicates the average level while the progression of green tones locates areas where the interpolated map thought there was more phosphorous than the average. The progression of red tones identifies the opposite with the average estimate being more than the interpolated ones.

The information in figure 4-6 shows that the difference between the two maps ranges from -20 to +80ppm. If one assumes that +/- 10ppm won't significantly alter a fertilization recommendation, then about two-thirds (47 +5.1 +12= 64.1 percent of the field) is adequately covered by the whole-field average. But that leaves about a third of the field that is receiving too much (6.6 + 3+ 2.5 +2.1 +1.3 +.79 +.18= 16.5 percent) or too little (19 percent)

phosphorous compared to a fertilization program based on the whole-field average.

Now turn your attention to figure 4-7 that compares maps derived by two different interpolation techniques—

IDW (inverse distance-weighted) and Krig. Note the similarity in the peaks and valleys of the two surfaces. While subtle differences are visible the general trends in the spatial distribution of the data are identical.

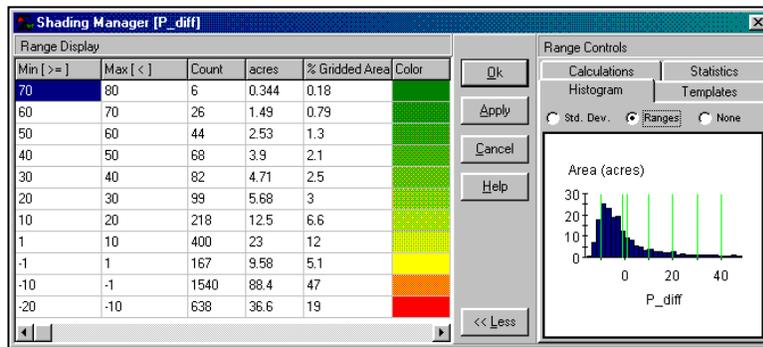


Figure 4-6. Statistics summarizing the difference between the maps in figure 4-5.

The difference map on the right confirms the coincident trends. The broad band of yellow identifies areas that are +/- 1 ppm. The brown color identifies areas that are within 10 ppm with the IDW surface estimates a bit more than the Krig ones. Applying the same assumption about +/- 10 ppm difference being negligible in a fertilization program the maps are effectively identical.

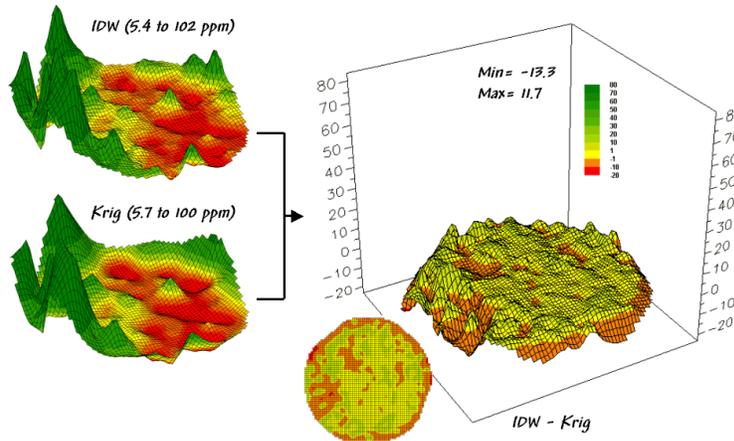


Figure 4-7. Spatial comparison of IDW and Krig interpolated maps.

field. While there was considerable differences between the average and the derived surfaces (from -20 to +80ppm), there was relatively little difference between the two surfaces (+/- 10ppm).

But which surface best characterizes the spatial distribution of the sampled data? The answer to this question lies in **Residual Analysis**—a technique that investigates the differences between

estimated and measured values throughout a field. It's common sense that one shouldn't simply accept a soil nutrient map without checking out its accuracy. Cool graphics just aren't enough.

Ideally, one designs an appropriate sampling pattern and then randomly locates a number of

So what's the bottom line? That there are substantial differences between a whole field average and interpolated surfaces—at least for this data set. It suggests that quibbling about the best interpolation technique isn't as important as using an interpolated surface (any surface) over the whole field average. However, what needs to be addressed is whether an interpolated surface (any surface) actually reflects the real spatial distribution. That weighty question is the focus of the next section.

“test points” to assess interpolation performance. Since a lot is riding on the accuracy of the interpolated maps, it makes sense to invest a little bit more just to see how well things are going.

...Residual Analysis is used to evaluate interpolation performance (IDW at .18 appears best)

	A	B	C	D	E	F	G	H
1	Test Set (randomly selected)							
2	Id	P_actual	Average	Avg-Actual	P_IDW	IDW-Actual	P_Krig	Krig-Actual
3	59	20.0	21.6	1.6	18.80	-1.20	17.80	-2.20
4	2	15.0	21.6	6.6	11.70	-3.30	13.50	-1.50
5	25	5.0	21.6	16.6	5.39	0.39	5.69	0.69
6	63	27.0	21.6	-5.4	28.80	1.80	27.40	0.40
7	18	9.0	21.6	12.6	9.68	0.68	7.90	-1.10
8	34	14.0	21.6	7.6	12.80	-1.20	13.30	-0.70
9	22	10.0	21.6	11.6	10.00	0.00	10.20	0.20
10	91	20.0	21.6	1.6	31.60	11.60	42.30	22.30
11	109	103.0	21.6	-81.4	91.10	-11.90	90.70	-12.30
12	46	12.0	21.6	9.6	16.80	4.80	15.80	3.80
13	89	93.0	21.6	-71.4	67.60	-25.40	68.60	-24.40
14	50	24.0	21.6	-2.4	24.60	0.60	27.40	3.40
15								
16	Average values=	29.3	21.6		27.41		28.38	
17	Residual sum=			-92.8		-23.1		-11.4
18	Average error=			19.0		5.24		6.08
19	Normalized error=			0.65		0.18		0.21

Figure 4-8. A residual analysis table identifies the relative performance of average, IDW and Krig estimates.

So which surface, IDW or Krig, did a better job in

4.4 Assessing Interpolation Results

Last section's discussion compared the assumption of the field average with map surfaces generated by two different interpolation techniques for phosphorous levels throughout a

estimating the measured phosphorous levels in the test set? The table in figure 4-8 reports the results for twelve randomly positioned test samples. The first column identifies the sample ID and the second column reports the actual measured value for that location.

Column C simply depicts estimating the whole-field average (21.6) at each of the test locations. Column D computes the difference of the “estimate minus actual”—formally termed the *residual*. For example, the first test point (ID#59) estimated the average of 21.6 but was actually measured as 20.0 so the residual is 1.6 (21.6-20.0= 1.6ppm)... not bad. However, test point #109 is way off (21.6-103.0= -81.4ppm)... nearly 400% under-estimate error.

The residuals for the IDW and Krig maps are similarly calculated to form columns F and H, respectively. First note that the residuals for the whole-field average are generally larger than either those for the IDW or Krig estimates. Next note that the residual patterns between the IDW and Krig are very similar—when one is way off, so is the other and usually by about the same amount. A notable exception is for test point #91 where Krig dramatically over-estimates.

The rows at the bottom of the table summarize the residual analysis results. The *Residual sum* row characterizes any bias in the estimates—a negative value indicates a tendency to underestimate with the magnitude of the value indicating how much. The -92.8 value for the whole-field average indicates a relatively strong bias to underestimate.

The *Average error* row reports how typically far off the estimates were. The 19.0ppm figure for the whole-field average is three times worse than Krig’s estimated error (6.08) and nearly four times worse than IDW’s (5.24).

Comparing the figures to the assumption that a +/-10ppm is negligible in a fertilization program it is readily apparent that the whole-field estimate is inappropriate to use and that the accuracy differences between IDW and Krig are minor.

The *Normalized error* row simply calculates the average error as a proportion of the average value for the test set of samples (5.24/29.3= .18 for IDW). This index is the most useful as it allows you to compare the relative map accuracies between different maps. Generally speaking, maps with normalized errors of more than .30 are suspect and one might not want to make important decisions using them.

So what’s the bottom line? That Residual Analysis is an important component of precision agriculture. Without an understanding of the relative accuracy and interpolation error of the base maps, one can’t be sure of the recommendations derived from the data. The investment in a few extra sampling points for testing and residual analysis of these data provides a sound foundation for site-specific management. Without it, the process can become one of blind faith and wishful thinking.

* See www.innovativegis.com/basis/pfprimer/, an online book entitled *Precision Farming Primer*, Topic 2 Point Sampling for further discussion on point sampling.

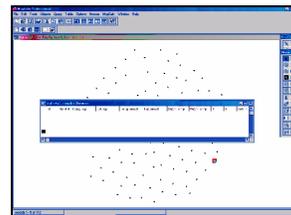
** See same link, Appendix A, Part 3 “More on Spatial Dependency” for further discussion on spatial interpolation and dependency.

4.5 Exercises

Access *MapCalc Learner* without specifying an existing database by selecting **Start** → **Programs** → **MapCalc Learner** → **MapCalc Learner** → **Create a new map set**.

 If MapCalc is actively using another database, select **File** → **Close** then press the **Create a new file** button on the main menu bar to pop-up the first dialog box in the *Map Creation Wizard*.

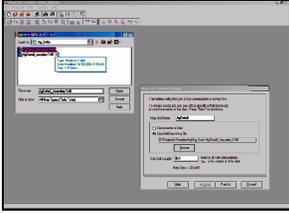
4.5.1 Interpolating Point Data (IDW)



Point Samples (MapInfo)

Note: point data for interpolation must be in either ESRI (.shp) or MapInfo (.tab) format.

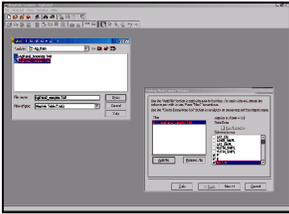
The first step in the *Wizard* identifies the *Name* of the new data set, the *Extent* (boundary) and the *Gridding Resolution* (cell size).



Wizard Step 1

- ✓ Enter the new map set name as **AgData2**
- ✓ Select **Use Field Boundary File**, browse to the **PA_MapCalc Data** folder and specify **AgData2_boundary.tab**
- ✓ Specify the **Grid Cell Size** as **50 feet**
- ✓ Press the **Next** button.

The second step identifies the point sampled data you want to map (**P** <phosphorous>, **K** <potassium>, **NO3_N** <nitrogen>). These specifications identify the “data fields” (columns) in the sample data table that will be used—the X,Y (Longitude, Latitude) positions each data point and the measured value reports the results of chemical analysis for each soil sample.

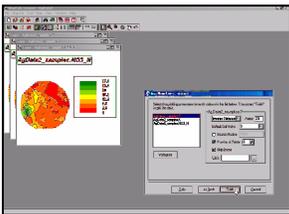


Wizard Step 2

Note: The data must be in lat/Lon WGS84 projection for the MapCalc Learner and Academic versions. MapCalc Professional contains on-the-fly projection conversion.

- ✓ Press the **Add File** button and browse to the **AgData2_samples.tab** table containing the sample data
- ✓ Check the **P**, **K** and **NO3_N** as the *Columns to use* at the bottom of the listing
- ✓ Press the **Next** button.

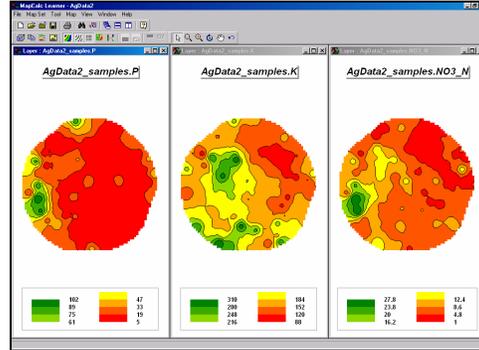
The final step specifies the interpolation technique and appropriate parameters to use.



Wizard Step 3

- ✓ Press the **Finish** button to interpolate the three maps using the default *Inverse Distance Square* technique and parameters.

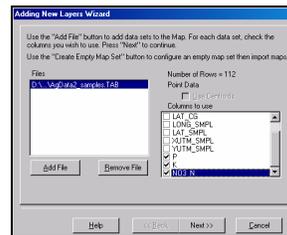
 Use the **Tile Vertically** button to generate side-by-side displays of all three maps.



4.5.2 Interpolating Point Data (Kriging)

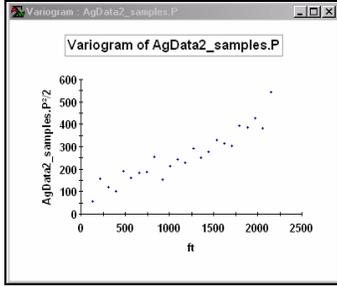
 Click on the *View* button, select the **AgData2_samples.K** map and press the **Rename** button. Rename the map to **AgData2_samples.K_IDW** to identify the interpolation technique that was used. Repeat to add the **_IDW** extension to the **P** and **NO3_N** map names then **Close** the window for managing map layers.

Select **Map Set** → **Add New Layers** to re-access the *Interpolation Wizard*.



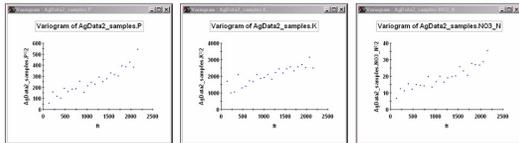
Click on the **Add File** button and specify **AgData2_samples.tab** table and **P**, **K** and **NO3_N** columns as before. Press **Next**.

 Highlight the **AgData2_sample.P** set of samples and press the **Variogram** button to generate a plot of the data’s spatial autocorrelation.

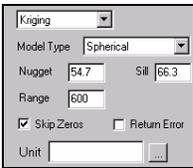


Generally speaking, a horizontal (flat) variogram indicates minimal information for interpolation and the map surface generated may be of negligible value. A variogram exhibiting a tightly-clustered linear pattern suggests better map surfacing results.

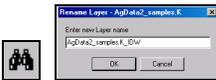
Repeat to generate *Variograms* of the other two variables.



Note that the samples for all three variables exhibit an upward-tilted linear trend that suggests that the data likely contains sufficient spatial autocorrelation for interpolation.

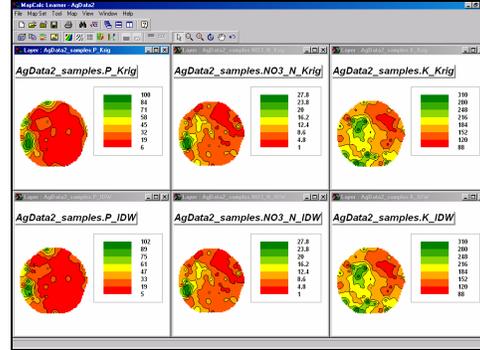


Change the technique from *IDW* to **Kriging** using the default set of parameters. Press the **Finish** button to generate another set of interpolated maps.



Rename the new set of maps with the extension **_KRIG**.

Close the windows containing the *Variogram* plots and press the **Tile Vertically** button to display all six maps at the same time. Move the map windows into a sensible order by clicking along their top edge and dragging them.

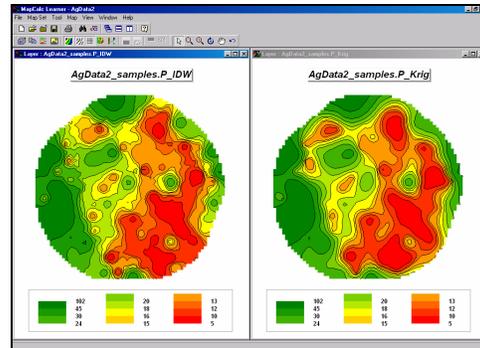


4.5.3 Comparing Interpolation Results

Click on the **Shading Manager** to generate a display of the **AgData2_samples.p_IDW** with the following formatting:

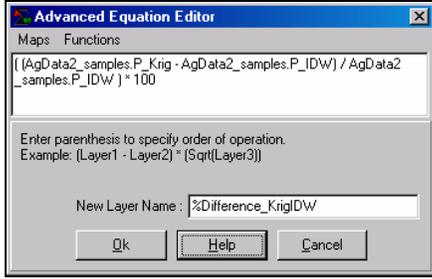
- ✓ 2D Continuous lattice
- ✓ Equal Count
- ✓ 11 Ranges
- ✓ Color Ramp with the lowest range set to **red**, highest range set to **green** and a **yellow** color inflection point at the mid range (6th interval).

Use the *Templates* tab to save the color ramp then display the **AgData2_samples.p_Krig** map and apply the template. Generate side-by-side displays of both maps as shown below.

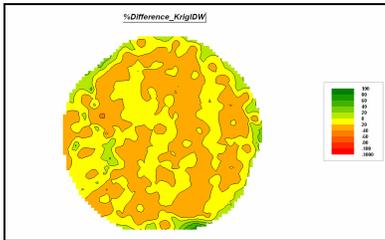
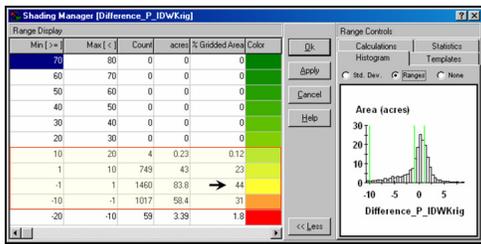


IDW and Krig interpolated maps ...can you see a difference?

Press the **Map Analysis** button, select **Overlay** → **Calculate** and complete the following dialog box that generates a percent difference map between the *IDW* and *Krig* interpolated surfaces.



Click on the **Shading Manager** and use the *Statistics* and *Histogram* tabs to get an idea of the *numeric distribution* of the map data— Min, Max, Mean, Std. Dev., etc.— of the map you just created. Enter the settings for display as shown below to get an idea of the *spatial distribution* of the data.



Note that 44% of the field estimates differ by less than plus or minus 1ppm (98.2% differs by less than +/-10ppm)—the predicted phosphorous surfaces for this data using default parameters for IDW and Krig exhibit minimal differences.

On your own, calculate percent difference maps comparing the average phosphorous level for the field (20.9ppm) with the .P_IDW and .P_Krig interpolated maps. What is the range of percent differences for both calculations? What is the average percent difference? What is the standard deviation of percent differences?

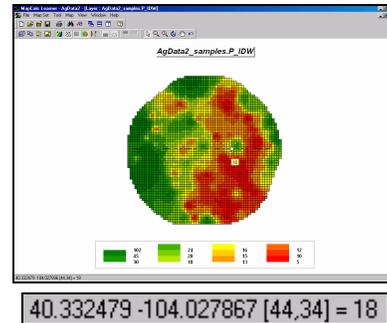
Hint: use the percent difference equation—
 $\%Diff = ((estimated - actual) / actual) * 100$
 ...substituting the average value for the *actual* term and the interpolated map for the *estimated* term.

4.5.4 Evaluating Interpolation Performance

The phosphorous levels for ten randomly selected “test” locations were measured as follows:

Column	Row	Value
44	34	20.1
24	27	13.5
17	47	19.2
28	61	50.8
43	6	9.0
5	27	63.1
56	46	17.8
33	31	18.3
41	58	17.6
5	38	49.9

Display the **AgData2_samples.P_IDW** map in *2D continuous grid* format with the *Layer Mesh* turned on.



As you move the cursor about the map notice that the Lat/Lon and Column/Row position is reported in the lower-left corner of the display. Move the cursor until you are over the first test sample location at *Column=44, Row= 34* and note the *.P_IDW estimated value= 18.0*. On a sheet of paper or an Excel spreadsheet enter these values in the *.P_IDW* column as shown in the table below.

	A	B	C	D	E	F	G	H	I
1	Column	Row	Value	.P_IDW	IDW_diff	.P_Krig	Krig_diff	Average	Avg_diff
2	44	34	20.1	18.0	-2.1	17.7	-2.4	27.9	-7.8
3	24	27	13.5					27.9	14.4
4	17	47	19.2					27.9	8.7
5	28	61	50.8					27.9	-22.9
6	43	6	9.0					27.9	18.9
7	5	27	63.1					27.9	-35.1
8	56	46	17.8					27.9	10.1
9	33	31	18.3					27.9	9.6
10	41	58	17.6					27.9	10.3
11	5	38	49.9					27.9	-22.0
12		Average=	27.9						0.0
13		StDev=							19.0

Once you have entered the other nine estimated values for the “test” location calculate the difference between the estimated value and the actual value (IDW_diff= 18.0 – 20.1= -2.1).

Complete the calculated differences for column **IDW_diff** then compute the column average and standard deviation. Complete the same processing for the **.P_Krig** and **Krig_diff** columns in the table.

The best technique (whole-field *Average*, *IDW* or *Krig*) is the one that has the smallest **Average difference** and **Standard Deviation**. So which technique is best? How much better is it than the other two?

You can exit the program by selecting **File → Exit** or by clicking on the “X” in the upper-right corner of the MapCalc program window. If you want to save your work, specify a new file name, such as *AgData_Topic4_exercises.rgs*. Each exercise set assumes you will start with the basic *AgData.rgs* data set and this database will become cluttered with exercise maps if you save your results to it each time.