

Spatial Interpolation Accuracy

Draft January 2002

Terry Kastens, Agricultural Economist, Kansas State University, tkastens@agecon.ksu.edu

Scott Staggenborg, Agronomist, Kansas State University, sstaggen@oznet.ksu.edu

prepared for:

Kansas State Precision Ag Conference, January 29-30, 2002, Great Bend, Kansas

Introduction

Many features of precision agriculture are characterized by point-based data measurements. For example, yield monitor data are crop yields assigned to specific points in time and space, as are grid-based soil samples. Even a high resolution remotely sensed image is merely a group of point values – though the points usually are quite close together. Despite the importance of point data, certain features of precision agriculture crucially depend on polygons. For example, polygons are typically used to trigger on-the-go variable rate applications of fertilizer and seed. That means data measured at specific points must be used to make inferences about expected values at other points and ultimately values assigned to whole polygons.

Using point data to make inferences about polygon values is nothing new. For example, for many decades, the recommended soil sampling procedure has been to collect soil from a series of points in a field, mix the soil together, analyze the sample, and then use the laboratory value as a measure assigned to the whole field. What has changed is that technology has lowered the cost of point data collection for some features, for example crop yields, and this has increased the demand for more point estimates of other features, for example soil samples. Yet, the cost of high resolution data measurements still prohibits collecting sample data at the management scale ultimately desired. In short, values at many points in space still must be estimated, typically using mathematical estimation algorithms along with values at sampled points.

The value of spatial information is determined by the return expected from using the information, the cost associated with information collection, and the information's accuracy. The simple fact that managers routinely use estimated data to guide decision making suggests that they are willing to trade accuracy against cost. This brief paper considers only one aspect of the return-cost-accuracy triad. Namely, we empirically examine the spatial prediction accuracy of various soil test measures using several well-accepted spatial prediction procedures. More specifically, we analyze soil pH, soil test P, and soil test K in more-or-less grid soil samples (average cell areas ranging from 0.7 to 5.9 acres) from 15 Kansas fields; soil organic matter content from 13 different fields; and soil texture (percent sand, silt, and clay) from 3 fields.

Spatial Interpolation

Perhaps because of an intrinsic assumption of spatial dependencies (values at closer points are more alike than values at more distant points), spatial interpolation has been widely used as an estimation algorithm in precision agriculture. For example, agriculturally related GIS (geographic information systems) software routinely offers point-and-click spatial interpolation features. Spatial interpolation does what its name suggests. It provides estimates of values at

unmeasured points by interpolating between the values at measured points.

Many spatial interpolation algorithms have been developed and tested over the years. However, and though there are a number of variants, generally two spatial interpolation methods especially have risen to popularity in GIS software, inverse-distance and kriging. Both methods develop an estimate for an unmeasured point by using a scheme that assigns a numerical weight to each measured point within some specified search radius of the unmeasured point. To ensure unbiased estimates, weights are required to sum to 1. To accommodate the assumption of spatial dependencies in the data, weights are usually smaller for measured points that are more distant from the point whose value is being estimated.

Inverse Distance

The inverse distance predictor is

$$V_j = \frac{\sum_{i=1}^n \left(V_i * \frac{1}{d_i^K} \right)}{\sum_{i=1}^n \frac{1}{d_i^K}}, \quad [1]$$

where V denotes value, d denotes distance, n is the number of points within the search area, K is a user-determined positive constant (usually an integer), and subscripts index spatial locations. In Eq. 1, the value at location j is estimated using the values associated with the n points within some search radius of point j . Each i th point in the search area is associated with a distance d_i , which is the distance to point j . The denominator in Eq. 1 ensures that the weights assigned to the n points will sum to 1. The K exponent determines how weights diminish across distance. For a given distance, larger K values imply smaller weights. Because d_i^0 is 1, it is easy to see that setting K equal to 0 means that the average of the n values will be used as a predictor for the value at j . In practice, K is set equal to 1 or 2 and occasionally as high as 4. Setting K equal to 4 implies that the values at the nearest points to j will have virtually all of the influence on the prediction. Thus, in mapping, using large K values leads to a “bulls eye” appearance around sampled points, whereas smaller K values leads to a much smoother map.¹

Ordinary Kriging

Ordinary kriging, referred to here simply as kriging, is a much more complicated process than inverse-distance. Consequently, this paper provides only an introduction to the mathematical process. For more detail, readers can consult [An Introduction to Applied Geostatistics](#), by E. H. Isaaks and R.M. Srivastava (Oxford University Press, New York, 1989). It should be noted that, statistically, kriging is a theoretically more appealing process than inverse distance. However, kriging critically depends on a number of modeling assumptions required of the user. Because

¹ The distance between two points is computed as the square root of the sum of squared differences in x-coordinate values and y-coordinate values. For example, if point 1 has x,y coordinates of (x_1, y_1) and point 2 has x,y coordinates of (x_2, y_2) , then the distance between point 1 and 2 is given by $[(x_1 - x_2)^2 + (y_1 - y_2)^2]^{0.5}$. A description of the technique which calculates x,y coordinates from longitude and latitude values that results in meaningful measures of distance is beyond the scope of this paper.

the user typically must make these assumptions with limited information, it is not a foregone conclusion that kriging will result in more accurate estimates in practice than inverse distance. The best that users can hope for is a general tendency for one method to be superior based on numerous comparisons of the methods in actual data.

Kriging is based on the assumption that one half of the squared difference between point values, $G_i = 0.5*(V_i - V_j)^2$, is expected to increase with greater distance (h) between the points, until at some distance (A , referred to as the range), that value approaches the variance (squared standard deviation) of the overall data series. The idea is that nearby points are expected to have similar values and thus small values of G . However, beyond the range, the only relationship between the data points is that characterized by the natural variability in the data (no spatial dependencies beyond the range). When G is plotted against h in a two-dimensional graph it is referred to as a sample variogram (or semivariogram).

Because of certain mathematical requirements of kriging, G must continuously rise with increased distance h , faster at small h values, more slowly with larger h , ultimately approaching a plateau (referred to as the sill). Because real-world data rarely meet this stringent requirement, a sample variogram cannot directly be used in kriging. Rather, the output of a particular mathematical formula specifying G as a function of h is used instead. In practice, only a handful of functional forms are considered in kriging. Two of the most popular ones, and the ones examined in this research, are the exponential and the spherical functions. For the exponential function, the value of G for some distance h , $G(h)$, is depicted as a function of h according to

$$G(h) = \begin{cases} 0 & \text{if } h = 0 \\ C_0 + C_1 * \left(1 - \exp\left(\frac{-3 * h}{A}\right)\right) & \text{if } h > 0 \end{cases} \quad [2]$$

where C_0 , C_1 , and A are model parameters (numerical constants) whose values are determined using a nonlinear least squares algorithm which chooses those values that minimize the sum of the squared prediction errors. As used here, each field (typically involving a number of grid soil samples) and each measure of interest (e.g., soil pH) has its own values of C_0 , C_1 , and A . In the estimation data set, $1/2$ of the squared difference between point i 's value and point j 's value, or $0.5*(V_i - V_j)^2$, is matched to the distance between the two points to become a single observation of dependent and independent variable values, respectively. Because each point in a data set can be matched with every other data point, a field with 100 sample values would have $100*100$, or 10,000, observations in the data set used to estimate the parameters of Eq. 2.² Notice that, despite the fact that a search radius ultimately will be used to predict the value at some unmeasured point, all data points are used in estimating the model variogram function.

² Kriging software packages often use the average of G_i values whose distances fall within some distance range as a single G value that is paired with the average distance in that group of points (i.e., the h value). For example, all G values associated with distances ranging from 95 to 105 might be averaged, and then assigned to an h value of 100. Then, the model variogram is estimated using this series. This is necessary because of computer memory limitations. Otherwise, for example, a data set involving 40,000 yield monitor data points would require $40,000*40,000$ for each of G and h , with each vector comprising 25,600 megabytes of data at a precision requiring 16 bytes per numerical value! Clearly, this would be unworkable with today's computers.

In Eq. 2, C_0 , which is some positive value, is referred to as the nugget effect. Kriging posits a discontinuity at a distance of $h = 0$. That is, at exactly 0, G is expected to equal 0. Yet, at some infinitesimally small distance, G is expected to equal C_0 . The basic idea is that small scale variability is not 0. That is, two soil samples immediately adjacent to each other likely will not have equal values. In Eq. 2, $C_0 + C_1$ is the sill, or the plateau that is asymptotically approached at ever greater distances h . Kriging supposes the variance of the underlying data (referred to here as σ^2) and the sill to be similarly valued. In Eq. 2, A (the range) is the distance at which the variogram value is approximately $C_0 + 0.95C_1$ – which can be seen by the fact that $1 - e^{-3}$ is approximately equal to 0.95.

Fig. 1 depicts an exponential model variogram for a hypothetical soil test series. Note the discontinuity shown at 0 distance between points (i.e., at $h = 0$). In Fig. 1, the nugget effect is 40 ($C_0 = 40$), the sill is 180 ($C_0 + C_1 = 180$), supposedly representing a data series with a variance of 180 (standard deviation of 13.4), and the range is 500 meters ($A = 500$). In words, this graph depicts a situation where the spatial dependencies in the data play out at a distance of around 500 meters.




Figure 1

In real world data, the model variogram estimated from the data often is not as nicely structured as that shown in Fig. 1. For example, the sill is often substantially higher than the data variance.

Further, the range is often far beyond the maximum distance observed in the data. As a demonstration of this possibility, we considered one field of 36 data points of soil test P. Those 36 data values had a mean of 25.2 ppm soil test P, a variance of 179.5 (standard deviation = 13.4 ppm), and a maximum distance between any two points of 751 meters. Thus, the hypothetical variogram of Fig. 1 might be a reasonable guess about the structure of the variogram estimated from the 36 data points. In particular, the sill in Fig. 1 is 180, which compares to the variance of 179.5 observed in the data. Further, a range of 500 depicted in Fig. 1 seems reasonable in a data set with a maximum distance of 751.

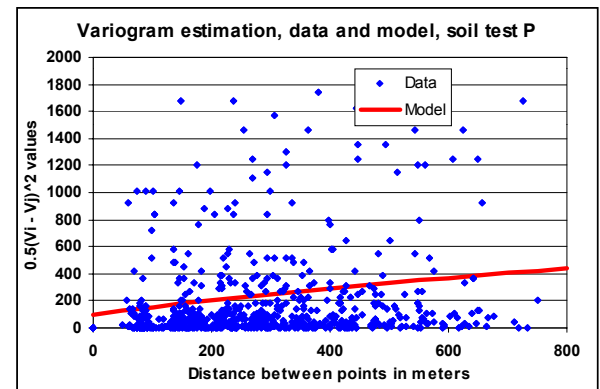


Figure 2

Using the 36 data points to estimate the model variogram meant using $36 \times 36 = 1,296$ observations of the dependent variable $0.5(V_i - V_j)^2$ and the independent variable of distance. Fig. 2 shows the

estimated variogram along with the data points used in estimation.³ Clearly, the estimated variogram of Fig. 2 is substantially different than the one in Fig. 1. The nugget effect (C_0) is 101 rather than the prior guess of 40, the sill ($C_0 + C_1$) is 898, which is 5 times as large as the variance of the data and the sill assumed in Fig. 1. The range (A) is 4,386 meters, which is nearly 6 times the maximum distance observed in the data. That Fig. 2 looks so different from the textbook example of a variogram shown in Fig. 1 is simply the reality of using grid soil sample data in a kriging exercise. However, this may not be an indicator of whether kriging is a suitable spatial interpolation technique to use in such data.

The spherical function corresponding to the exponential function of Eq. 2 is

$$G(h) = \begin{cases} 0 & \text{if } h = 0 \\ C_0 + C_1 * \left(1.5 * \frac{h}{A} - 0.5 * \left(\frac{h}{A} \right)^3 \right) & \text{if } 0 < h < A, \\ 1 & \text{if } h \geq A \end{cases} \quad [3]$$

with a similar interpretation of model parameter values as that provided for Eq. 2. Given substantial errors in estimating model variograms, as seen in the comparison of Fig. 1 to Fig. 2, real-world differences in accuracy associated with kriging using the exponential model of Eq. 2 and kriging using the spherical model of Eq. 3 will probably be small.

The predicted values of G provided by Eqs. 2 and 3 are not directly used in the kriging exercise. Rather, the parameter estimates themselves are used in functions referred to as covariance functions. The covariance function, referred to as $C(h)$, corresponding to the exponential function depicted as Eq. 2 is

$$C(h) = \begin{cases} C_0 + C_1 & \text{if } h = 0 \\ C_1 * \exp\left(\frac{-3 * h}{A}\right) & \text{if } h > 0 \end{cases}, \quad [4]$$

and the covariance function for the spherical function of Eq. 3 is

$$C(h) = \begin{cases} C_0 + C_1 & \text{if } h = 0 \\ C_1 - C_1 * \left(1.5 * \frac{h}{A} - 0.5 * \left(\frac{h}{A} \right)^3 \right) & \text{if } 0 < h < A. \\ 0 & \text{if } h \geq A \end{cases} \quad [5]$$

Continuing with the kriging description, depending on the choice of an exponential or spherical functional form, a matrix \mathbf{C} is formed as follows. An ij th element of the matrix is the model predicted value of $C(h)$, where h is the distance between points i and j . Now, only the points within the search radius are considered. Then, a row of 1's is appended to the bottom of the matrix, a column of 1's to the right side of the matrix, and finally, the lower right hand corner of the matrix is set equal to 0. Thus, in a search involving N measured data points, the matrix \mathbf{C}

³ In practice, to make Eq. 2 easier to estimate with least squares algorithms, each of the dependent variable values are first divided by σ^2 . Subsequently, right hand side predicted values are multiplied by σ^2 to derive estimates of G that correspond to the actual data.

will be a symmetric matrix dimensioned as $(N+1) \times (N+1)$. Next, a column vector \mathbf{D} is formed where the i th element is the model predicted value of $C(h)$ where h is the distance from the unmeasured point of interest to the i th measured point in the search area. Again, a 1 is appended to the bottom of the vector, resulting in \mathbf{D} being dimensioned as $(N+1) \times 1$. Then, a column weight vector \mathbf{w} is calculated by first computing the matrix inverse of \mathbf{C} , and subsequently matrix multiplying it by the vector \mathbf{D} . Thus $\mathbf{w} = \mathbf{C}^{-1}\mathbf{D}$. Next, the bottom value of \mathbf{w} is discarded, resulting in an $N \times 1$ vector, with the i th value being the weight assigned to the i th measured point in the search area. Finally, each of the N values in the search area is multiplied by its corresponding weight, and then summed to provide the kriging estimate of the unmeasured point.

Out-of-sample Testing (Jackknifing)

To derive an estimate of the expected prediction accuracy associated with using each of the interpolation methods outlined above in real-world prediction, the following jackknife procedure was used. The first data point in a field was discarded. Then, the other data points in the field were used to generate a prediction of the value at the discarded data point. Next, the first data point was returned to the data set and the second data point discarded. Again, only information in the field not associated with point 2 was allowed to be used in predicting the value at point 2. This process was continued until each point in the field was allowed to be predicted. Finally, the known values of all sample points were compared to their predicted values to derive a measure of prediction accuracy. To ensure a true out-of-sample exercise in the kriging procedures, a point being predicted generally was not allowed to assist in estimating the parameters of the covariance functions shown in Eqs. 4 and 5. Thus, a different covariance function was estimated for each data point being predicted, using only information from the other points.

Prediction Accuracy

Although a number of prediction accuracy measures might be considered, because of its appealing statistical properties, we considered only RMSE (root mean squared error). RMSE is the square root of the average squared prediction error. In a data set with 100 observations, for any given prediction procedure, there would be 100 actual values and 100 predictions. RMSE is calculated by first subtracting a predicted value from its corresponding actual value, squaring the result, averaging all 100 squared error values, and subsequently taking the square root of the result. Because we ultimately want to generalize across many different prediction procedures and a number of different soil sample measures, we normalize an RMSE value by dividing by the standard deviation of the underlying series. This value is referred to as RMSE/STD.

Using a normalized RMSE measure for comparing prediction methods is important because some soil test measures (e.g., soil test K) are scaled much differently than others (e.g., soil pH), leading to large differences in the natural variability of the data (depicted as standard deviation). Further, normalizing RMSE by STD is appealing in that the standard deviation of a series is actually the RMSE associated with using the series average as a prediction for every data point. Thus, for a particular prediction method, an RMSE/STD value of 1 (or 100% in percentage terms) indicates that the predictive accuracy for that method is no better than simply knowing the average of the data – which is essentially the same thing as having a good composite soil sample from the field. To provide some visual understanding of the accuracy associated with different

RMSE/STD values, Figs. 3 and 4 depict the relationship between actual and predicted values associated with $RMSE/STD = 100\%$, and $RMSE/STD = 70\%$, respectively. Although not shown, an $RMSE/STD$ value of 0% would imply perfect prediction, resulting in all data points in the figures falling on a straight line starting at the origin and rising at 45 degrees.

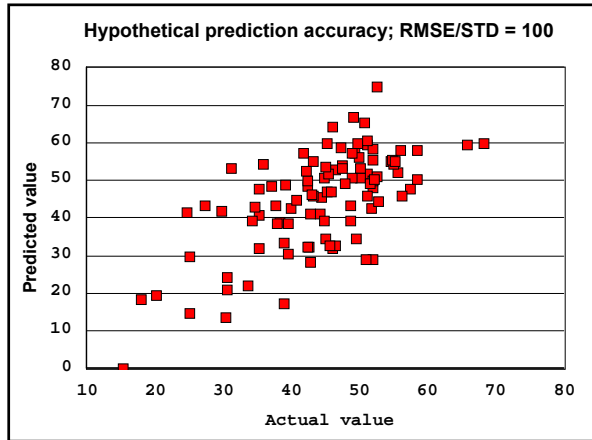


Figure 3




Figure 4

Prediction Methods Considered

Generally, in our analysis, two kriging methods were used – *Skrig*, based on a spherical variogram and *Ekrig*, based on an exponential variogram. Five inverse distance methods were used, *Inv0*, *Inv1*, *Inv2*, *Inv3*, *Inv4*, corresponding to different powers (the K values in Eq. 1). For each of the kriging and inverse difference methods, search radii were considered in 50 meter increments, starting with the first 50-meter value that contained at least one data point, and ending with the 50-meter value that contained all data points in the field.

Our prediction process considered other methods and features besides those already noted. For example, we considered an untypical method where the soil test measures were regressed directly on the x-coordinate and the y-coordinate of the location. Regressions were based alternatively on all of the data in a field (i.e., all but the point being predicted), or only the data in the search area. It was believed that these methods might prove fruitful if strong spatial trends were present in the data. Additionally, we considered kriging procedures where all of the data in a field were used to estimate the model variogram, and where the point being forecasted was excluded during model variogram estimation. The purpose for this distinction was to see how much the standard practice of using all data for model variogram estimation might misrepresent the prediction accuracy in real-world data where the point being predicted truly was missing from the data set.

Generalizing the Results

For any given soil measure, for example soil pH, each field had a unique $RMSE/STD$ measure for each of the 2 kriging methods (spherical variogram and exponential variogram), for variogram estimation with and without using the point being forecasted, for each of the 5 inverse distance methods (the different K powers, 0-5), and for each of the different search radii. The

end result is numerous measures of prediction accuracy that must be summarized into meaningful generalizations about the procedures. To accomplish that, we used a linear regression model, where RMSE/STD was considered to be a function of binary variables associated with prediction method, the average distance between data points in a field, a measure of the search radius, and a number of interactions between method variables and the measure of search radius. Ultimately, the generalizing regression model was used to predict the point-prediction accuracy expected from using certain methods of spatial interpolation. That is, certain variables in the model were held constant while varying other variables over meaningful ranges of interest – to present graphical generalizations of our results.

To provide a useful measure of search radius for managers ultimately wanting to use our results in the future, we considered “percent of total points in the field” as our measure of interest. That is, we characterized search radius information by the percent of total field points in the search area rather than by some distance measure such as feet or meters.

Research Results

Table 1 reports the average (across fields and across the search radii where relevant) predictive accuracy measures associated with the two ordinary least squares regression prediction methods, OLS1 (all data used in regression on coordinates) and OLS2 (only data from search area used in regressions), the spherical kriging methods where the model variogram was estimated using all data (Skr0) and where the point being predicted was excluded during the model variogram estimation (Skr1), and the two exponential kriging counterparts (Ekr0 and Ekr1).

Table 1. Average RMSE/STD (as a percent) values associated with different prediction methods.

soil test measure	Interpolation Method					
	OLS1	OLS2	Skr0	Skr1	Ekr0	Ekr1
soil pH	96.9	110.6	86.6	87.2	86.6	87.3
soil test P	104.7	109.9	96.6	86.9	94.5	94.6
soil test K	93.9	110.0	87.3	88.2	87.3	88.8
soil organic matter %	101.9	131.8	92.7	95.5	92.7	95.1
sand %	103.4	104.3	107.3	108.4	103.4	103.4
silt %	107.3	102.2	106.9	104.3	100.6	101.0
clay %	100.2	95.4	97.2	96.0	87.5	91.1
Overall average	101.2	109.2	96.4	96.6	93.2	94.5

Notes: soil pH, P, and K from 15 fields, organic matter from 13 fields, texture variables from 3 fields

Several features stand out in Table 1. First, based on the averages reported, none of the interpolation methods is especially accurate. The lowest (most accurate) value in the table, 86.6, pertains to soil pH and the two kriging methods, Skrg0 and Ekr0. On average, these methods have a prediction RMSE that is 86.6% of the standard deviation of soil pH for the field in question. On the other hand, we have little prior information about the level of accuracy to

expect for interpolation methods. Nor do we know the economic impacts associated with interpolating at different levels of accuracy. Second, the two “regressing on x,y coordinates” methods, OLS1 and OLS2, are highly inaccurate. In particular, a good composite soil test for a field would predict data points more accurately than these interpolation methods. Third, whether the point being predicted is excluded from variogram estimation is largely irrelevant regarding predictive accuracy. In particular, the “include all data points” columns of Skrg0 and Ekrg0 are only slightly lower (more accurate) than their “skip the point being predicted” counterparts, Skrg1 and Ekrg1. This is useful to know in that an analyst wanting to examine kriging’s accuracy may not have to spend the additional programming and computer time required of estimating a separate variogram for each data point.

Given the findings of Table 1, the results of our information generalizing regression focus only on the Skrg1 and Ekrg1 kriging methods and certain of the inverse-distance methods – all across different search areas. In the simulations, the average distance between points was held constant at its overall average. These results are presented graphically as Figs. 5, 6, 7, 8, and 9, covering soil pH, soil test P, soil test K, soil organic matter percent, and clay percent, respectively.

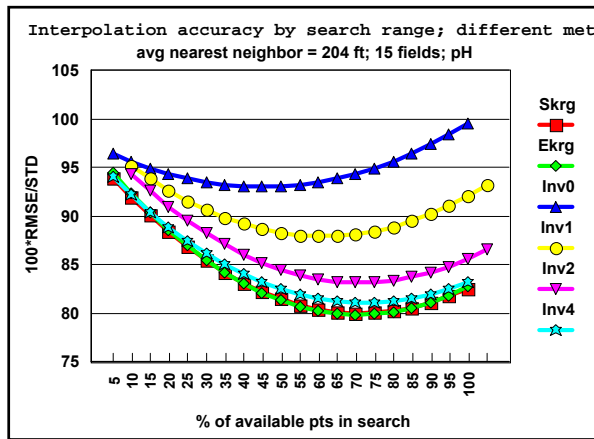


Figure 5




Figure 6




Figure 7




Figure 8

Generally, the figures speak for themselves. Nonetheless, some findings are noteworthy. First, among the methods considered in the figures, the *Inv0* method appears to perform the worst.

This should not be surprising in that the *Inv0* method simply averages the values in the search area to provide the necessary prediction. In short, the only spatial dependency feature of the *Inv0* method is that different search radii are considered – more distant points do not have smaller weights in prediction. Second, the *Inv2* method is nearly as accurate as the two kriging methods. This confirms Isaaks and Srivastava’s assertion that, practically speaking, unless data are clustered, kriging procedures have little benefit over the much simpler inverse distance methods. The data sets examined here generally involved regularly spaced grid soil samples.

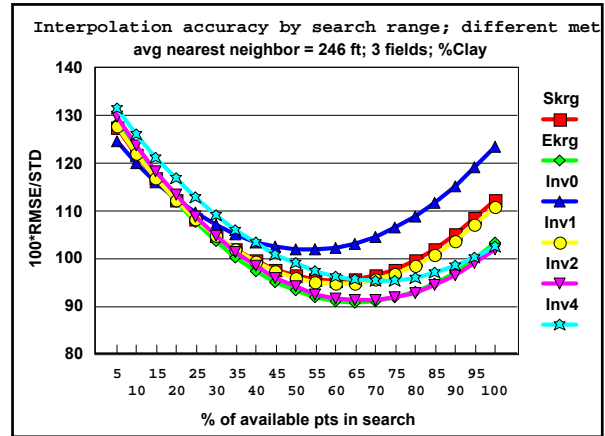


Figure 9

Consequently, it should not be surprising that kriging was found to have predictive accuracies similar to those of inverse distance methods. Of course, this presumes that the analyst would know to use *Inv2* as opposed to other inverse distance procedures. Third, for the type of data analyzed here (more-or-less grid soil samples), the optimal search radii should be rather large, on average engulfing a number of data points approximately equal to 65% of the total number of available data points in a field.

Other Prediction Methods for Soil Samples

Relative to using the average soil sample as a predictor everywhere, the spatial interpolation methods examined here often do not appear to offer substantial improvement in accuracy. This suggests that other methods, relying less on spatial dependency and more on causal dependency, may need to be considered when predictions of unmeasured points are needed. Such methods would bring in other, perhaps less expensive, information that is available at points requiring a prediction. For example, a causal yield model might be available that mathematically describes crop yield as a function of fertilizer rates and soil test measures, and possibly also as a function of other more spatially dense data such as electrical conductivity or remote sensing information. Then, when a prediction of some soil test measure, say soil test P, is needed for a point that has not been measured for soil test P, the available information at that point (yield, fertilizer rates, electrical conductivity, remote sensing) could be brought into a framework designed to predict soil test P using those variables.

Although little is known about the accuracy associated with non-interpolative predictive frameworks such as those briefly described above, preliminary work by the authors suggests that they could offer accuracy comparable with the standard spatial interpolation methods. Even if they are less accurate than traditional spatial interpolation methods, they likely should be considered on the basis of cost alone. That is, the predictive methods considered in this research each relied on more-or-less grid-based soil test data, which means such data must first be collected. Yet, the cost of such data collection is often considered too high for general use – currently, only a small percentage of fields have actually been grid soil sampled. More to the point, those managers wishing to variably apply fertilizer and other crop inputs may wish to proceed without gathering grid soil samples. That is, they may be willing to proceed on the basis

of less accurate but also less expensive predictions of soil test measures required of their variable rate prescriptions.

An Aside: Interpolating Dense Data

This paper principally deals with using spatial interpolation methods for sparse data such as grid soil samples. However, as an aside, it should be useful to include an interpolation example involving more dense data for contrast. Our example is elevation and yield data taken from a combine yield monitor. The first case involved a corn field (nearly 50,000 data points) that was harvested using a yield monitor that also collected elevation data. The second case simulated “2 combines, 1 without a yield monitor” for the same field (points within a random distance below 400 feet were excluded from helping make predictions, which assumes a maximum unmeasured distance from any point of 400 feet). Because of the large data size, we used average results from three random draws of 300 data points each.

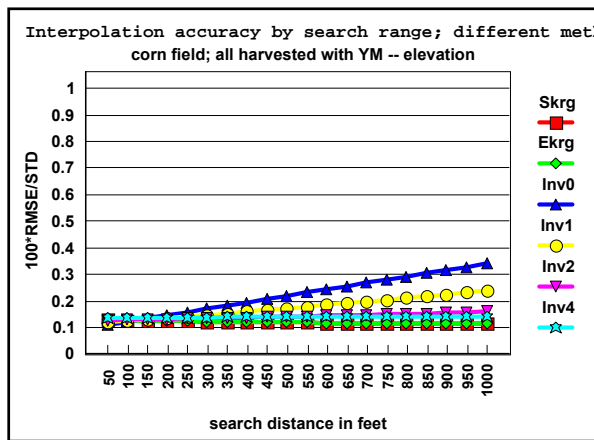


Figure 10

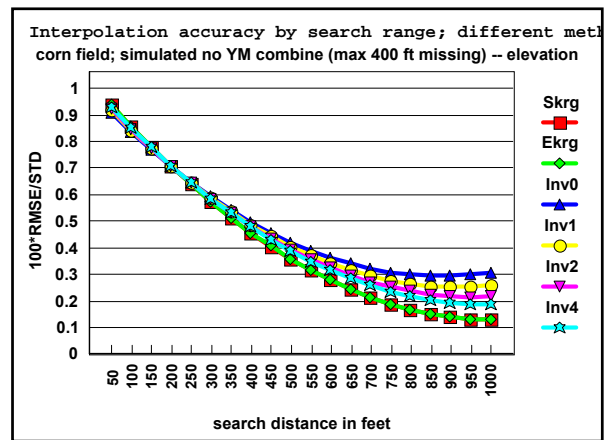


Figure 11

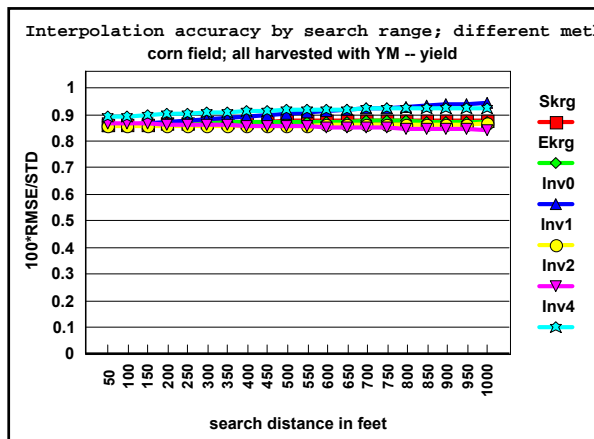


Figure 12

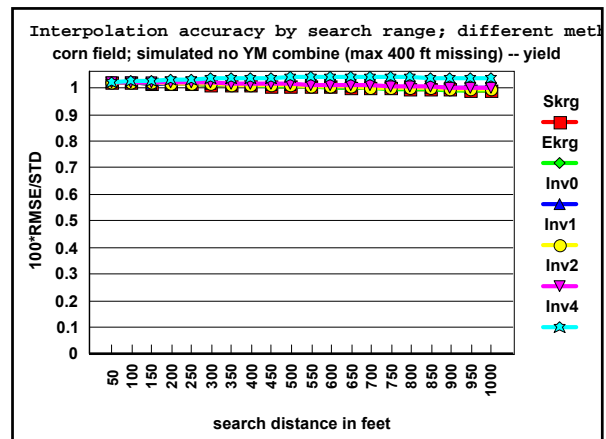


Figure 13

Fi

g. 10 shows that, with dense and highly spatially dependent data such as elevation, most interpolation methods are quite successful (very low RMSE/STD relative to soil sample data), but the best method is *Inv0* at 50 feet (a simple average of values within 50 feet of unmeasured point). Fig. 11 shows that virtually the same accuracy might be achieved when one combine has

no yield monitor – as long as the search radius is sufficiently large. On the other hand, with dense data that are more spatially independent, such as yield, Fig. 12 reveals predictive accuracy more like that of the soil sample data discussed earlier. Further, when large areas may be unmeasured, Fig. 13 shows no benefit to interpolation over simply using the field average.

Summary

In a comprehensive analysis involving 15 more-or-less grid sampled Kansas fields, several soil test measures, and several spatial interpolation techniques, the following generalizations emerge.

1. Using a search radius that is large enough, on average, to capture a number of points equal to about 65% of the total number of points in a field is probably optimal. Smaller and larger search radii diminish prediction accuracy.
2. Among the inverse distance methods, the most consistently best method is probably inverse-distance squared (*Inv2* in the paper). *Inv0* and *Inv1* are generally less accurate and high powers, such as *Inv4*, are probably inconsistent in terms of prediction accuracy.
3. In using actual data for testing the expected accuracy associated with kriging, it likely is irrelevant whether the model variogram is estimated from all data points or whether the point being predicted is excluded from the estimation.
4. For kriging, the choice between a spherical and an exponential variogram probably is inconsequential. Both behave similarly in terms of predictive accuracy.
5. For more-or-less grid soil samples, using inverse-distance squared (*Inv2*) is generally as accurate as the more sophisticated kriging procedures.
6. Relative to using a field's composite soil sample (average) value as a prediction everywhere, the various spatial interpolation procedures examined here often do not greatly improve accuracy. That is, none of the interpolation methods are particularly accurate for the coarse soil sample data examined. This suggests that non-interpolative methods, such as those that would bring in other available information for the point being predicted, should be considered as competitive prediction methods.
7. As long as the search radius is kept small, highly dense and highly spatially dependent data (e.g., elevation data taken from a yield monitor) may require little to no interpolation – simply average the points in the search area. A substantially larger search radius can help with such interpolations when large areas have unmeasured values. On the other hand, for more spatially independent data, such as yield, large areas of unmeasured values implies virtually no benefit to spatial interpolation over simply using the field average.