

Chapter 14

Uncertainty and Uncertainty Propagation in Soil Mapping and Modelling

Gerard B. M. Heuvelink

“We demand rigidly defined areas of doubt and uncertainty”!

Douglas Adams
The Hitchhiker’s Guide to the Galaxy

In previous chapters, the use of geostatistical modelling for soil mapping was addressed. We learnt that one of the advantages of kriging is that it not only produces a map of predictions but that it also quantifies the uncertainty about the predictions, through the kriging standard deviation. In this chapter we will look into this in more detail. We will also examine another way to assess the accuracy of soil prediction maps, namely, through independent validation. This approach has the advantage that it is model-free and hence makes no assumptions about the structure of the spatial variation and relationships between the target soil property and covariates. Finally, we will examine how uncertainties in soil maps propagate through environmental models and spatial analyses. Throughout this chapter we will use the Allier data set and case study, Limagne rift valley, central France, to illustrate concepts and methods. We will only consider soil properties that are measured on a continuous-numerical scale. Many of the concepts presented can also be extended to categorical soil variables, but this is more complicated and beyond the scope of this chapter.

14.1 What Is Uncertainty?

Suppose that the bulk density of the topsoil at some location in some study area equals 1.33 g/cm^3 . Suppose further that this value is unknown to us because we did not measure the bulk density at the location. All that we have is a soil property map

G.B.M. Heuvelink (✉)

Soil Geography and Landscape group, Wageningen University and ISRIC – World Soil Information, PO Box 47, 6700 AA Wageningen, The Netherlands
e-mail: gerard.heuvelink@wur.nl

that contains predictions of the bulk density for all locations in the study area, also at the location where the true bulk density is 1.33 g/cm^3 . Let the bulk density at that location according to the map be 1.45 g/cm^3 . This shows that the soil map is in error. The error equals $1.33 - 1.45 = -0.12 \text{ g/cm}^3$. Here, *error* is defined as the difference between the true and predicted value of the soil property.

In practice, we usually do not know the error, because we would need the true value to calculate it and we do not have the resources to perfectly measure the soil everywhere. In other words, we are *uncertain* about the error (and the true value). Here, *uncertainty* refers to a state of mind of a person or people that expresses a lack of confidence about reality (Heuvelink 2014). Note that uncertainty is a property of people. It is not the soil bulk density that is uncertain; it is we that are uncertain about the soil bulk density. We are uncertain because we have limited and imperfect information (i.e. only a soil map) and are aware that the information we have may be in error.

Although we are uncertain, this does not mean that we are completely ignorant. For instance, we may know that the chances are equal that the error in the bulk density prediction is positive or negative (because we used an unbiased mapping method), we may know that it is unlikely that the absolute value of the error is greater than 0.50 g/cm^3 , etc. Thus, it is not unreasonable to assume that we can come up with a large number of possible error values and attach a probability to each of these. Since the true value is the sum of the (known) prediction and the error, we can also list all possible values of the soil property and attach a probability to each. If we can do this, then we have characterised (our uncertainty about) the soil property by a probability distribution.

Now that we have characterised the soil property by a probability distribution, it has effectively become a random variable. After all, a random variable is nothing else than a variable that can take on many values, where each value has a certain probability of occurrence. Since we deal with spatially distributed variables, we must extend this concept to that of a random field. A random field is a collection of indexed random variables, where in our case the index is geographic location. We can characterise the variable at each location by a (univariate, marginal) probability distribution, but we must also characterise the (spatial) correlation between the variables at multiple locations. Geostatistics provides the methods and tools to do this (i.e. variogram estimation and kriging), and this has been explained in detail in Chaps. 9 and 10 (but see also Goovaerts 2001). However, while in previous chapters the focus was on the predictions made by kriging, in this chapter we will concentrate on the uncertainty associated with these predictions. In the next section, we will explain how geostatistics can be used to model uncertainty in mapped soil properties by means of a cokriging example.

14.2 Geostatistical Modelling of Uncertainty

14.2.1 Mapping Soil Properties for the Allier Study Area

As part of a research study in quantitative land evaluation, a crop simulation model was used to calculate potential crop yields for floodplain soils of the Allier River in the Limagne rift valley, central France. The moisture content at wilting point (Θ_{wp} , cm^3/cm^3) is an important input attribute for the crop simulation model. Because Θ_{wp} varies considerably over the area in a way that is not linked directly with soil type, it was necessary to map its variation separately to see how moisture limitations affect the calculated crop yield.

Unfortunately, because Θ_{wp} must be measured on samples in the laboratory, it is expensive and time-consuming to determine it for a sufficiently large number of data points for creating the prediction map by kriging. An alternative and cheaper way is to calculate Θ_{wp} from other soil properties which are cheaper to measure or using pedotransfer functions (see Chap. 7). Because moisture content at wilting point is often strongly correlated with moisture content at field capacity (Θ_{fc} , cm^3/cm^3) and soil porosity (Φ , cm^3/cm^3), both of which can be measured more easily and cheaply, it was decided to map these first and next derive a map of Θ_{wp} from these using multiple linear regression. We will come back to this in Sect. 14.4 and concentrate first on the kriging of Θ_{fc} and Φ .

Sixty-two measurements of Θ_{fc} and Φ were made in the field at the sites indicated in Fig. 14.1. From these data experimental variograms and an experimental cross-




Fig. 14.1 The Allier study area showing sampling points of moisture content at field capacity and porosity. Circled sites are those where in addition moisture content at wilting point was measured




Fig. 14.2 Experimental variograms (*open circles*) and fitted variogram models (*solid lines*) for Θ_{fc} (*top left*), Φ (*bottom right*) and the experimental and modelled cross-variogram (*bottom left*)

variogram were computed (see Sect. 10.5). These were then fitted using the linear model of coregionalisation (Fig. 14.2). Next both soil properties were mapped to a regular 50×50 m grid using ordinary cokriging. The cokriging yielded raster maps of means and standard deviations for both Θ_{fc} and Φ , as well as a map of the correlation of the cokriging prediction errors. Figure 14.3 displays these maps.

14.2.2 Interpreting the Kriging Standard Deviation Maps

The kriging standard deviation maps shown in Fig. 14.3 are summary measures of the uncertainty about Θ_{fc} and Φ in the study area. These uncertainties are the result of interpolation errors: while we know the true values of Θ_{fc} and Φ at the 62 observation locations (assuming that measurement errors are negligibly small), we are uncertain about their true value at non-observation locations. As explained in Sect. 14.1, we are uncertain because the true value is unknown to us, and so we cannot identify a single true reality. At best we can list all possible values of the soil property and attach a probability to each of them. This is exactly what we do in kriging, because under the assumptions made (i.e. normality, stationarity, isotropy), we derived a (conditional) probability distribution of Θ_{fc} and Φ for each grid cell. In this case, the uncertainty at each grid cell is characterised by a zero-mean

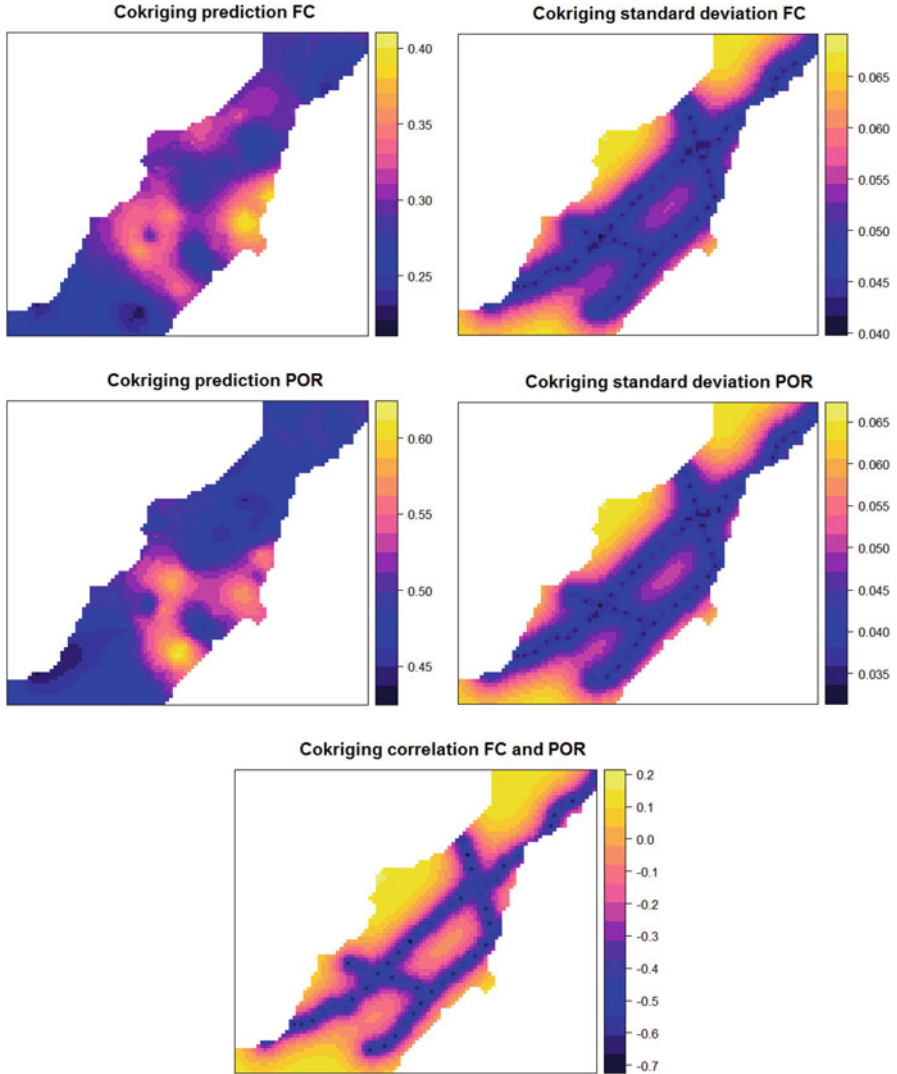


Fig. 14.3 Cokriging results for the Allier study area (50×50 m grid): conditional mean and standard deviation of Θ_{fc} (cm^3/cm^3), conditional mean and standard deviation of Φ (cm^3/cm^3) and correlation of cokriging prediction errors of Θ_{fc} and Φ

normal distribution with a standard deviation as given in Fig. 14.3. The magnitude of uncertainty is captured by the width of the probability distribution, although the shape of the distribution is important as well (see Fig. 14.4). Because of the assumption of normality, the uncertainties of the kriged Θ_{fc} and Φ all have a shape such as shown in the left-hand panels of Fig. 14.4. The width of the distribution varies in space, as is clear from Fig. 14.3.

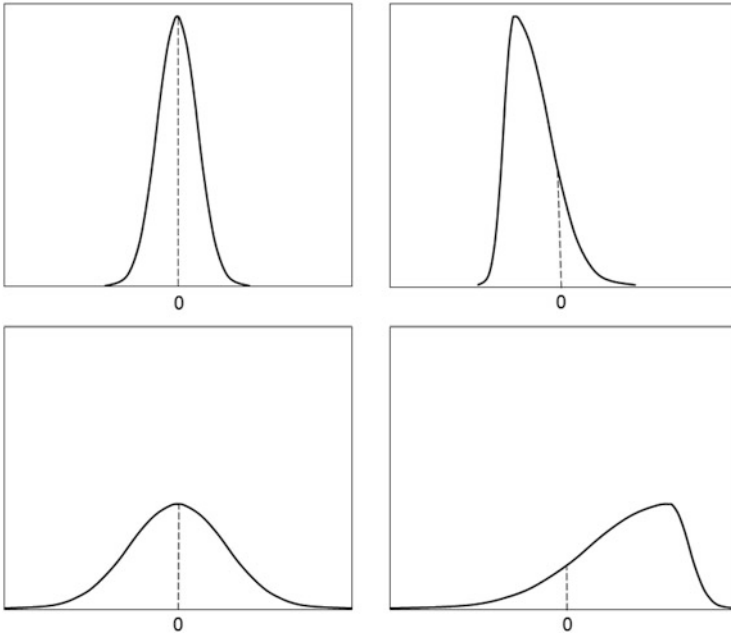


Fig. 14.4 Examples of probability distributions to characterise uncertainty. Probability distributions can be narrow (small uncertainty, *top*) or wide (large uncertainty, *bottom*). They can also be symmetric around zero (*left*) or asymmetric and biased (*right*)

The probability distributions shown in Fig. 14.4 refer to the value of a single variable. In our case study, we considered two variables, Θ_{fc} and Φ , and so each of these will have their own marginal probability distribution at any one location in the study area. But the uncertainties associated with these two variables are also correlated, because of the cross-correlation between Θ_{fc} and Φ , as characterised by the cross-variogram. It is difficult to predict how large the correlation between the prediction errors at any given location is, because the correlation between the cokriging prediction errors is not the same as that between the variables themselves. In other words, in general we have $\text{corr}(\widehat{\Theta}_{fc} - \Theta_{fc}, \widehat{\Phi} - \Phi) \neq \text{corr}(\Theta_{fc}, \Phi)$. Fortunately, cokriging provides the correlations between the cokriging errors at each location, as shown in the bottom map in Fig. 14.3. Note that there are clear spatial variations in the correlation between the cokriging errors and that these may be positive as well as negative, depending on location. A graphical illustration of the joint (bivariate) probability distribution of two uncertain variables is shown in Fig. 14.5. If the correlation between the uncertain variables were zero, then the major and minor axes of the ellipses would be along the axes of variables S_1 and S_2 , i.e. they would not be rotated. The example in the right panel of Fig. 14.5 shows a case in which there is a non-zero correlation. In this example the correlation is positive because the major axis has a positive angle. The contour lines would be circular if




Fig. 14.5 Examples of a bivariate probability distribution of two uncertain variables S_1 and S_2 . The rotated ellipses of equal probability density shown in the *right panel* indicate positive correlation

the two variables would be uncorrelated and have equal standard deviation. Figure 14.5 refers to two uncertain variables S_1 and S_2 , which could be $S_1 = \hat{\Theta}_{fc} - \Theta_{fc}$ and $S_2 = \hat{\Phi} - \Phi$, but note that it might as well refer to $S_1 = \hat{\Theta}_{fc}(x) - \Theta_{fc}(x)$ and $S_2 = \hat{\Theta}_{fc}(x') - \Theta_{fc}(x')$ for two arbitrary locations x and x' in the study area.

The kriging standard deviations shown in Fig. 14.3 vary spatially and tend to be small in the neighbourhood of observation locations and are large further away from these, particularly at the boundary of the study area. This is as one would expect intuitively, because the magnitude of the interpolation error depends on the closeness of observations and their local density and because (spatial) extrapolation is more error prone than interpolation. It can also be inferred from the kriging variance equation (see Sect. 10.3):

$$\sigma_K^2(x_0) = E \left[\left(\hat{S}(x_0) - S(x_0) \right)^2 \right] = \sum_{i=1}^n \lambda_i \cdot \gamma(|x_i - x_0|) + \varphi \quad (14.1)$$

where n is the number of observations used in kriging, the λ_i are kriging weights, the x_i are observation locations and x_0 is the prediction location, γ is the variogram model and φ is a Lagrange parameter. In most practical cases, the latter is relatively small so that we can concentrate on the summation part. Inspection shows that this part will be small when the distances between the x_i and x_0 are small, hence when observation locations are close to the prediction location. Of course the exact result also depends on the shape of the variogram model. For instance, in case of a pure nugget variogram, the kriging variance (and hence its square root, the kriging standard deviation) will be constant: if there is no spatial correlation, then interpolation cannot benefit from nearby observations, and the interpolation error (variance) will be equal everywhere. Note that the kriging variance can never be smaller than the nugget variance, except when we interpolate to an observation location. For Φ , which has a nugget variance of 0.0008 (see Fig. 14.2), this means

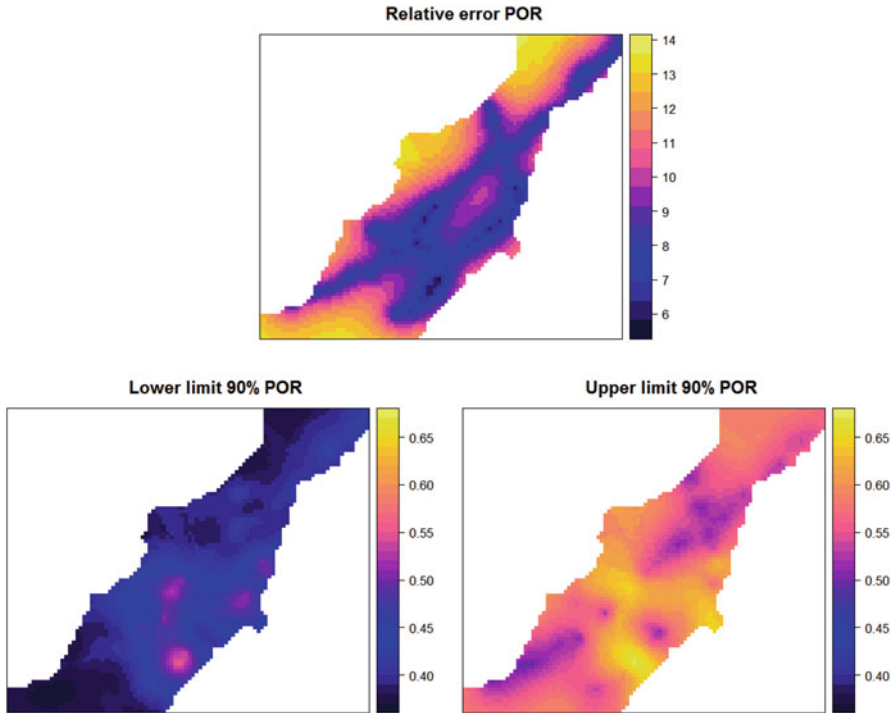


Fig. 14.6 Relative error (*top*) and lower and upper limits of the central 90% prediction interval of the topsoil porosity for the Allier study area

that the uncertainty about the predicted Φ at any non-observation location in the Allier study area will be at least $0.028 \text{ cm}^3/\text{cm}^3$ (as confirmed by Fig. 14.3). Note also that in the Allier example, the kriging variance will be calculated in a slightly different way than using Eq. 14.1, because in that case predictions were made using cokriging instead of kriging (Wackernagel 2003).

For users with little background in (geo)statistics, it may not be that easy to interpret a standard deviation map. More appealing for uncertainty communication are maps of the relative error (computed as the ratio of the kriging standard deviation and kriging prediction maps, multiplied with 100%) and maps of the lower and upper limits of a central 90% prediction interval, which are derived by subtracting and adding 1.64 times the kriging standard deviation map from the kriging prediction map, respectively. Note that here we assumed that the kriging error is normally distributed. These maps are shown for the soil porosity in the Allier study area in Fig. 14.6. The relative error is nowhere greater than 15%, indicating that the uncertainty about soil porosity is small compared to its predicted value. Nonetheless, the differences between the lower and upper limits of the 90% prediction interval maps are large, indicating that the kriging interpolation error is far from negligible.

14.2.3 Spatial Stochastic Simulation

Kriging makes predictions, such that the expected squared prediction error is minimised. This is attractive because it means that the predicted value is on average closest to the true (unknown) value. As explained in Chap. 9, *spatial stochastic simulation* has an entirely different objective. Here, the goal is to generate ‘possible realities’ from the probability distribution of the uncertain variable. This is done by sampling from the probability distribution using a *pseudo-random number generator*. The result of a spatial stochastic simulation exercise is not unique, because there are an infinite number of possible realities, from which just one or several are taken. To illustrate the difference between optimal prediction and stochastic simulation, take the example of the outcome of a throw of a fair die. Optimal prediction would produce the value of 3.5, because on average this is the number closest to any of the outcomes 1–6. However, stochastic simulation would randomly take one of the numbers 1–6, where each of the six values would have equal chance of being selected. Figure 14.7 shows four realisations (‘possible realities’) from the kriging probability distribution of topsoil porosity in the Allier study area. These were created using *conditional* simulation, meaning that the

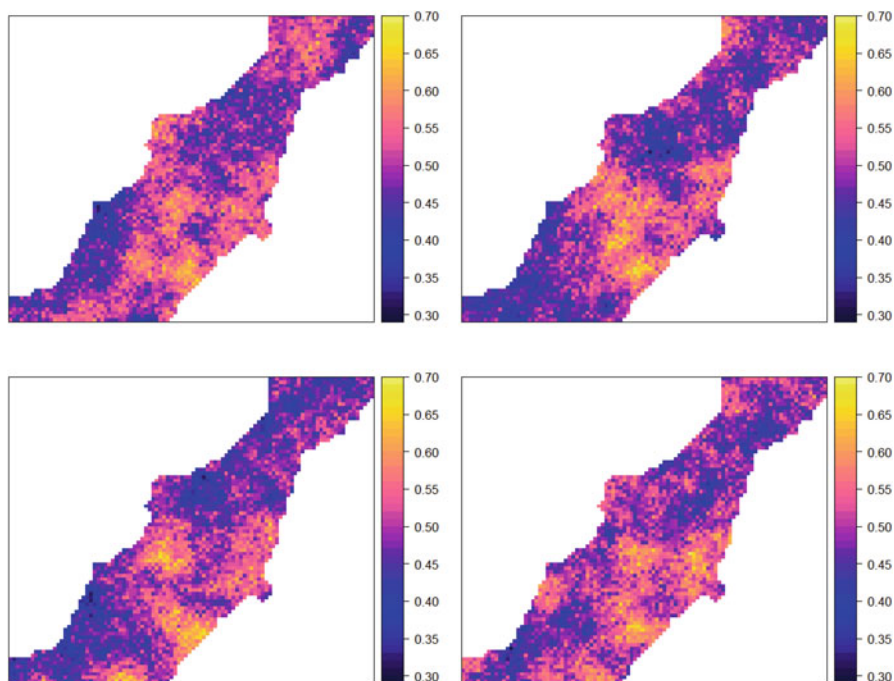


Fig. 14.7 Four possible realities of the topsoil porosity for the Allier study area generated with conditional spatial stochastic simulation

62 observations of soil porosity were used as conditioning data. The differences between the simulated maps convey the uncertainty about the true porosity, and when shown in animation mode, they are an attractive means to communicate the uncertainty of an interpolated map to non-experts. We will also make use of simulated maps in Sect. 14.4, when we explain the Monte Carlo method for analysing uncertainty propagation through environmental models.

If we would generate many more than four realisations of topsoil porosity, then their average would equal the kriging prediction map shown in Fig. 14.3, while their standard deviation would equal the kriging standard deviation map, also shown in Fig. 14.3. Thus, although perhaps not easily noticeable from Fig. 14.7, the differences between the realisations are greater far from observation locations than close to observation locations.

14.2.4 Change of Support

Often users do not want to predict soil properties at points but instead are interested in the average value over a larger piece of land. For instance, perhaps for a farmer it is not that relevant to know Θ_{fc} and Φ at point locations within the Allier study area, but instead the real interest is in parcel averages. Such averages over spatial units or ‘blocks’ can be predicted using *block kriging*, as explained in Sect. 10.3. The blocks need not be rectangular or square but may take irregular shapes as well. They may even be as large as the entire study area. When the blocks are relatively small, then block kriging produces similar predictions as point kriging, but the associated kriging standard deviation is usually much smaller, especially when the nugget variance is large. Figure 14.8 shows the cokriging standard deviation maps of Θ_{fc} and Φ for the case where the blocks are equal to the grid cells. Note that the standard deviations are substantially smaller compared to those of point kriging shown in Fig. 14.3. The explanation is that within-block spatial variation averages

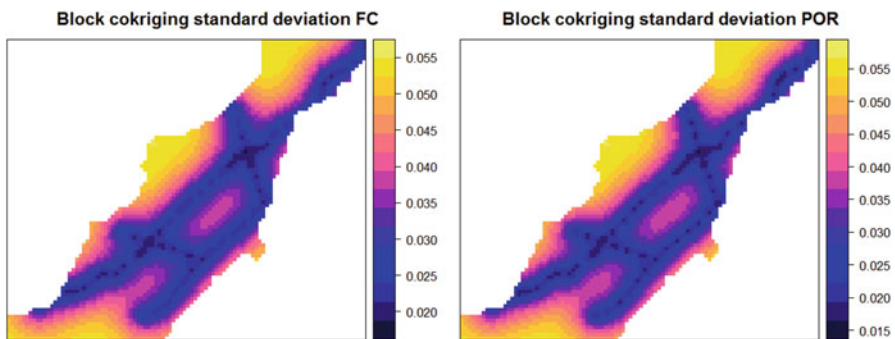


Fig. 14.8 Block cokriging standard deviation maps of Θ_{fc} and Φ using blocks of 50×50 m

out when the block mean is predicted. In other words, in case of large short-distance spatial variation (i.e. high nugget), predictions of block averages are much more accurate than point predictions. This tells us that it is crucially important to choose the right support when addressing uncertainty in interpolated soil maps.

Block kriging is used when *spatial aggregation* is the objective, in other words it is used in a case where the observations have a smaller support than the predictions. The opposite, i.e. making predictions at a smaller support than the observations, is known as *area-to-point kriging*. It will be no surprise that in this case uncertainty increases instead of decreases. There are not many examples of area-to-point kriging in soil science, because usually the starting point is observations at point support, but an exception is vertical spatial interpolation. In this case observations are often averages over soil horizons or layers, while predictions may be required for smaller or different depth intervals (Orton et al. 2016).

14.2.5 Extension to Kriging with External Drift

So far we discussed the uncertainty resulting from ordinary (co)kriging. But in recent years ordinary kriging is used less frequently and often replaced with kriging with external drift (KED), also termed universal kriging and regression kriging (Odeh et al. 1995; Hengl et al. 2004). Chapter 9 provides the details. This is because we rarely have only soil point observations as a source of information, but in addition we may have a large suit of covariate maps that provide valuable information about the soil property of interest. The additional information may then be used to improve the mapping and reduce uncertainty. The mathematical expression for the KED variance, which quantifies the uncertainty in the resulting map, is more complicated than that of the ordinary kriging variance given in Eq. 14.1. It is the sum of the trend estimation variance and the kriging variance. Even though the trend estimation variance is added, in practice the KED variance will usually be smaller than the ordinary kriging variance. This is because the KED variance is based on the variogram of the residual (defined as the difference between the soil property and the trend), which typically is much smaller than the variogram of the soil property itself, and hence the kriging component of the KED variance will decrease. See Sect. 10.3 for a more detailed discussion of KED and comparison with ordinary kriging.

Recall from Sect. 14.2.2 that the kriging standard deviation tends to be small near observation locations and large further away from them. As noted, this can be explained from a closer look at Eq. 14.1, which shows that the kriging variance will be larger if the distance between observation and prediction locations is large. Note also from Eq. 14.1 that the kriging variance does not depend on the observations themselves but only on the variogram and configuration of the observation locations. This allows optimisation of spatial sampling designs that minimise the spatially averaged kriging variance, as explained in Sect. 11.6. Sampling design optimisation under the ordinary kriging model typically leads to a fairly uniform distribution of

the sampling locations, with a slightly higher concentration of sampling locations near the study area boundary. In case of KED another aspect is also included in the sampling design optimisation. In that case the trend estimation error variance needs also to be minimised, which calls for a joint optimisation in geographic and ‘feature’ space. In other words, we must make sure that the observations also cover the covariate space well (Minasny and McBratney 2006; Brus and Heuvelink 2007).

14.2.6 Uncertainty Quantification of a Given Soil Map

The geostatistical uncertainty quantification approach works well in situations where one starts from scratch and where it is feasible to build a geostatistical model of reality. However, what to do in situations in which a soil property map has already been derived, without using geostatistical models? These could be maps made using deterministic algorithms, such as inverse distance or nearest neighbour interpolation. Alternatively, soil property maps may have been derived from an existing soil class map, by assuming that soil properties within each soil class are constant and assign these map-unit mean values using expert judgement or data from ‘representative’ profiles. It may also be that a soil property map is provided without additional information about how the map was made and without quantification of the uncertainty.

In such situations the map uncertainty may still be modelled geostatistically if there are sufficient independent observations of the soil property. This boils down to building a geostatistical model of the differences between the soil property map and the independent observations and kriging these errors (Heuvelink 2014). Here, it is essential that the observations are truly independent, i.e. have not been used for map making, because otherwise it might result in a severe underestimation of the map uncertainty. In practice, truly independent data are rarely available unless a new sampling campaign is initiated after the map was made. If uncertainty quantification is important, it is worthwhile to spend extra budget on collecting new data and quantifying the map uncertainty as described above. Note that this will not only quantify the map uncertainty but will also improve the map accuracy, because the existing map could be adjusted by adding the interpolated error to it. In a way, this approach comes close to the KED approach described in Sect. 14.2.5, but now using a single external covariate that is an existing map of the target soil property.

When soil property maps are derived without an underlying geostatistical model and there are no independent observations to build a geostatistical model of the map error, then the only resort is to base the uncertainty model of the map on expert judgement (Truong and Heuvelink 2013). This introduces subjectivity because different experts tend to have different opinions. Also, it will prove to be practically impossible to extract from experts a full probabilistic uncertainty model that also includes spatial and cross-correlations. Expert elicitation procedures are cumbersome and often limited to estimation of quantiles of the (marginal) distribution (O’Hagan et al. 2006).

14.3 Uncertainty Assessment Through Statistical Validation

Uncertainty quantification as described in Sect. 14.2 takes a model-based approach, by defining a geostatistical model of the soil property of interest and deriving an interpolated map and the associated uncertainty from that or by constructing a geostatistical model of the error in an existing map. The approach yields a complete probabilistic characterisation of the map uncertainty, but such characterisation is only valid under the assumptions made. Perhaps the stationarity assumption of ordinary kriging can be relaxed by using a more elaborate geostatistical model, such as that underlying kriging with external drift, but such a model typically needs more data, and in the end no modelling approach is free of assumptions. Therefore it is worthwhile to discuss an alternative, model-free approach to assess the accuracy of soil property maps. This is achieved through (statistical) *validation*.

Validation is defined here as an activity in which the soil map predictions are compared with independent observations. Unlike in Sect. 14.2.6, the outcomes are not used to build a geostatistical model of the map error, but instead summary measures of the observed errors are computed and reported. Common summary measures are the mean error and the root mean squared error. As before, it is essential that the validation observations are independent and have not been used in map making. The safest way to ensure this is to collect validation data after the map was made.

In practice, often we are not that much interested in how well the map predicts the soil property at the limited set of validation locations, but instead we want to know how well the map performs for the entire study area. Summary measures of the entire area cannot be computed but only estimated, because we cannot afford to collect validation observations everywhere. It is then strongly advised to select the validation locations using probability sampling (Brus et al. 2011). The important advantages are that in such case unbiased estimation of summary measures can be ensured and that confidence intervals around the estimated summary measures can be calculated, which is also a prerequisite for significance testing (e.g. to compare whether map A is more accurate than map B). The simplest probability sampling design is *simple random sampling*, but efficiency can be improved by using more elaborate designs. In practice, *stratified simple random sampling* is often used. Model-free estimation of map accuracy has the important advantage that it makes no assumptions, but the disadvantages are that a probability sample is required and that the method can only produce summary measures of the map accuracy.

Validation is based on a comparison of map predictions with independent observations. Typically the observed differences are attributed to map errors. However, it is important to recognise that part of the differences may also be caused by errors in the observations. It is not difficult to incorporate this if the observation error is known in statistical terms (i.e. bias and variance). If observation error is negligibly small compared to map error, as may be the case when a poor map is validated with observations analysed in a high-quality lab, then the influence of observation error on validation statistics may be ignored.

Although it is advised to collect an independent validation data set using probability sampling, this does not mean that summary measures of map accuracy cannot be calculated in case of a non-probability sample, such as a convenience or purposive sample. But in such case, it is important to be aware that there is a risk that the measures may not represent the overall map accuracy very well, such as when the validation data are from specific parts of the study area that have a different map accuracy as other parts.

14.3.1 Cross-Validation

Summary accuracy measures may also be derived using *cross-validation* (see Sect. 11.5.2). In the case of leave-one-out cross-validation, all observations are put aside one by one and the remaining data are used to calibrate the soil mapping model and predict at the location that was put aside. Validation measures are then computed by comparing the predictions with the put-aside observations for all observation locations.

Table 14.1 shows the accuracy measures for Θ_{fc} and Φ as obtained using leave-one-out cross-validation. The *mean error* is close to zero for both properties, indicating that cokriging is unbiased. The *root mean squared error* is $0.050 \text{ cm}^3/\text{cm}^3$ for Θ_{fc} and $0.044 \text{ cm}^3/\text{cm}^3$ for Φ . These values are not much smaller than the spatial variation of these properties, which can be gleaned from comparison with the square root of the variogram sills shown in Fig. 14.2 (because the sill of the variogram is approximately equal to the variance of the variable of interest). Poor prediction performance is also evidenced by the low values for the *amount of variance explained*, which is defined as one minus the ratio of the mean squared error and the variance. Apparently the sampling density is insufficient to capture a large part of the spatial variation. In fact this is already foretold by the variograms in Fig. 14.2, which have fairly high nugget variances and small ranges. The poor prediction accuracy is also confirmed by the scatter plots of cross-validation predictions against observations (Fig. 14.9). Note also that the cross-validation accuracy measures might still be somewhat too optimistic about the overall map accuracy, since all observations are on transects, and hence any cross-validation location always has at least a few nearby observations. The last column of Table 14.1 shows the

Table 14.1 Accuracy measures of cokriging predictions of Θ_{fc} and Φ as obtained with leave-one-out cross-validation

	Mean error (cm^3/cm^3)	Root mean squared error (cm^3/cm^3)	Amount of variance explained (—)	Standardised root mean squared error (—)
Water content at field capacity (Θ_{fc})	-0.002	0.050	0.30	1.05
Porosity (Φ)	-0.001	0.044	0.37	1.05




Fig. 14.9 Scatter plots of predicted against observed Θ_{fc} (left) and Φ (right) as obtained using leave-one-out cross-validation




Fig. 14.10 Spatial plot of leave-one-out cross-validation errors for Θ_{fc} (left) and Φ (right)

standardised root mean squared error (SRMSE), which is obtained by taking the square root of the average squared z_{score} , where z_{score} is defined as the difference between the observed and predicted soil property, divided by the cokriging standard deviation. If the cokriging standard deviation is a proper measure of the map prediction error, then SRMSE should be close to one. The obtained values are fairly close to one and do not indicate a significant over- or underestimation of the uncertainty. Figure 14.10 shows bubble plots of the spatial distribution of the cross-validation errors.

14.4 Uncertainty Propagation

Previous sections of this chapter explained that uncertainty about soil properties can be conveniently represented using probability distributions. Specific attention was paid to quantification of spatial interpolation errors using geostatistics. The methods were illustrated using a case study on mapping water content at field capacity (Θ_{fc}) and soil porosity (Φ) in the Allier study area, France. This chapter takes the analysis a step further by analysing how uncertainties in soil properties propagate through an environmental model that uses these soil properties as input (Heuvelink 1998). More specifically, we will analyse how uncertainties in Θ_{fc} and Φ propagate through a multiple linear regression model that predicts the soil water content at wilting point (Θ_{wp}) from Θ_{fc} and Φ . Before we do this, we first present the statistical uncertainty propagation methodology.

The uncertainty propagation analysis can be formulated mathematically as follows. Let U be the output of an environmental model g on m input variables S_i :

$$U = g(S_1, S_2, \dots, S_m) \quad (14.2)$$

The model g may be of various types, ranging from a simple pedotransfer function to a complex soil erosion or crop yield model. The objective of the uncertainty propagation analysis is to determine the uncertainty in the output U , given the operation g and the inputs S_i and their associated uncertainties. Let us denote the means and variances of the S_i by μ_i and σ_i^2 , respectively. Since the inputs are random variables or random fields, the output will be a random variable or random field as well. Important parameters of U are its mean ξ and variance τ^2 . From an uncertainty propagation perspective, the main interest is in the uncertainty of U , as contained in its variance τ^2 .

It must first be observed that the uncertainty propagation problem is relatively easy when g is a linear function of its inputs S_i . In that case the mean and variance of U can be directly and analytically derived. In case of non-linear models, analytically driven methods exist only in a few cases, and one must nearly always rely on approximation methods for a complete evaluation. Two of these methods will now be discussed.

14.4.1 Taylor Series Method

The idea of the Taylor series method is to approximate g by a truncated Taylor series centred at the means μ_i . In case of the first-order Taylor method, g is linearised by taking the tangent of g in μ_i . Fig. 14.11 illustrates this for the one-dimensional case ($m = 1$). The linearisation greatly simplifies the uncertainty analysis, but only at the expense of introducing an approximation error.




Fig. 14.11 Graphical illustration of the first-order Taylor method for the case where the model has a single input. The model (*red line*) is approximated by a linear function (*green line*) that has a small approximation error near the centre of the input probability distribution. Two cases are depicted: model is sensitive to changes in input (*blue*), and model is insensitive to changes in input (*purple*)

Using the first-order Taylor series method, the variance τ^2 of the output U is given by (Heuvelink 1998):

$$\tau^2 \cong \sum_{i=1}^m \sum_{j=1}^m \rho_{ij} \sigma_i \sigma_j g'_i g'_j \tag{14.3}$$

where ρ_{ij} is the correlation coefficient between the uncertainties in S_i and S_j and g'_i is the first derivative of g with respect to S_i , which is evaluated in the means $\mu_i, i = 1, \dots, m$. Equation 14.3 shows that the variance of U is the sum of various terms, which contain the correlations and standard deviations of the S_i and the first derivatives of g . These derivatives reflect the sensitivity of U to changes in the inputs (see Fig. 14.11 for a graphical illustration). From Eq. 14.3 it also appears that the correlations of the input uncertainties can have a marked effect on the variance of U .

To decrease the approximation error invoked by the first-order Taylor method, one option is to extend the Taylor series of g to include a second-order term as well. This is particularly useful when g is a quadratic function, in which case the second-order Taylor method is free of approximations and the first-order Taylor method is not. The application to the Allier case study discussed later in this section gives an example. However, it should be noted that in other cases including a second-order term might worsen the results. For instance, this might happen if the variance of the input is large and the approximation by a quadratic function, which is more accurate locally, is less accurate than the linear approximation at a greater distance from the approximation point.

14.4.2 Monte Carlo Method

The Monte Carlo method uses an entirely different approach to analyse the propagation of uncertainty. The idea of the method is to compute the result of the model repeatedly, with input values s_i that are randomly sampled from their joint distribution. The model results form a random sample from the distribution of U , so that parameters of the distribution, such as the mean ξ and variance τ^2 , can be estimated from the sample.

The method thus consists of the following steps:

1. Repeat N times:
 - (a) Generate a set of realisations $s_i, i = 1, \dots, m$.
 - (b) For this set of realisations s_i , compute and store the output $u = g(s_1, \dots, s_m)$.
2. Compute and store sample statistics from the N outputs u .

A random sample from the m inputs S_i can be obtained using an appropriate pseudorandom number generator (Lewis and Orav 1989; Ross 1990). Note that a conditioning step will have to be included when the S_i are correlated. In case of spatial inputs, these may be sampled using spatial stochastic simulation as explained in Sect. 14.2.3.

The accuracy of the Monte Carlo method is inversely related to the square root of the number of runs N . This means that to double the accuracy, four times as many runs are needed. The accuracy thus slowly progresses as N increases.

14.4.3 Evaluation and Comparison of Uncertainty Propagation Techniques

The main problems of the Taylor method are that it only works with models that are continuously differentiable with respect to their uncertain inputs, that it only provides estimates of the mean and variance of the model output and that the results are approximate only. It will not always be easy to determine whether the approximations involved using this method are acceptable. The Monte Carlo method also involves approximation errors, but these can be made arbitrarily small by increasing the number of Monte Carlo runs.

The Monte Carlo method brings along other problems, though. High accuracies are reached only when the number of runs is sufficiently large, which may cause the method to become extremely time-consuming. This will remain a problem even when variance reduction techniques such as Latin hypercube sampling are employed. Another disadvantage of the Monte Carlo method is that the results do not come in an analytical form.

As a general rule, it seems that the Taylor method may be used to obtain crude preliminary answers for simple models. These should provide sufficient detail to be able to obtain an indication of the quality of the model output. When exact values or

quantiles and/or percentiles are needed, the Monte Carlo method may be used. The Monte Carlo method will probably also be preferred when uncertainty propagation with complex models is studied, because the method is easily implemented and generally applicable. It is no more than an extra loop around an existing model. This, and the fact that computer power is ever increasing, means that nowadays the majority of uncertainty propagation studies uses the Monte Carlo method. Some examples from soil science are Brown and Heuvelink (2005), Bishop et al. (2006), Hastings et al. (2010), Kros et al. (2012), Van Den Berg et al. (2012), Brodsky et al. (2013), Poggio and Gimona (2014), Malone et al. (2015) and Xiong et al. (2015).

14.4.4 Sources of Uncertainty Contributions: The Balance of Errors

When the uncertainty propagation analysis reveals that the output of g contains too large an error, then measures will have to be taken to improve accuracy. When there is a single input to g , then there is no doubt where the improvement must be sought, but what if there are multiple inputs? Also, how much should the uncertainty of a particular input be reduced in order to reduce the output uncertainty by a given factor? It is useful to explore these questions briefly.

To obtain answers to the questions above, consider Eq. 14.3 again, which gives the variance of the output U using the first-order Taylor method. When the inputs are uncorrelated, this reduces to:

$$\tau^2 \cong \sum_{i=1}^m \sigma_i^2 g_i'^2 \quad (14.4)$$

This shows that the variance of U is a sum of parts, each to be attributed to one of the inputs S_i . This *partitioning property* allows one to analyse how much each input contributes to the output variance. Thus from Eq. 14.4, it can directly be seen how much τ^2 will reduce from a reduction of σ_i^2 . Clearly the output will mainly improve from a reduction in the variance of the input that has the largest contribution to τ^2 . Note that this need not necessarily be the input with the largest error variance, because the sensitivity of the model g to changes in the input is also important. Figure 14.11 shows an example where in the purple case the input uncertainty is greater than in the blue case, but where the output uncertainty still is the greatest in the blue case. This is because in the blue case, the model is more sensitive to changes in the input. Note also that Eq. 14.4 is derived under rather strong assumptions. When these assumptions are not realistic, it may be advisable to derive the uncertainty source contributions using a modified Monte Carlo approach.

14.4.5 Application of Uncertainty Propagation to the Allier Case Study

Recall from Sect. 14.2.1 that our aim is to map the soil moisture content at wilting point Θ_{wp} from maps of the moisture content at field capacity Θ_{fc} and soil porosity Φ . We obtained maps of both input soil properties and their associated uncertainties using cokriging in Sect. 14.2. We now need to discuss how these can be used to derive a map of Θ_{wp} and its associated uncertainty. Recall that we use a multiple linear regression model to predict Θ_{wp} from Θ_{fc} and Φ . The model is very simple, and hence we can use the Taylor series method to analyse the uncertainty propagation.

Figure 14.1 shows 12 circled sites where all three properties Θ_{wp} , Θ_{fc} and Φ were determined in the laboratory. These measurements were used to set up a pedotransfer function, relating Θ_{wp} , Θ_{fc} and Φ , which took the form of a multiple linear regression:

$$\Theta_{wp} = \beta_0 + \beta_1 \Theta_{fc} + \beta_2 \Phi + \varepsilon \quad (14.5)$$

The coefficients β_0 , β_1 and β_2 were estimated using standard ordinary least squares regression. The estimated values for the regression coefficients and their respective standard deviations were $\beta_0 = -0.263 \pm 0.031$, $\beta_1 = 0.408 \pm 0.096$ and $\beta_2 = 0.491 \pm 0.078$. The standard deviation of the stochastic residual ε was estimated as 0.0114. The correlation coefficients of the estimation errors of the regression coefficients were $\rho_{01} = -0.221$, $\rho_{02} = -0.587$ and $\rho_{12} = -0.655$. The regression model explains 94.8% of the variance of the observed Θ_{wp} , indicating that the model is satisfactory. Note that presence of spatial correlations between the observations at the 12 locations was ignored in the regression analysis.

The maps of Θ_{fc} and Φ as derived in Sect. 14.2.1 were substituted in the regression (Eq. 14.5) yielding a map of Θ_{wp} . The associated uncertainty was computed using the Taylor series method. Note that Eq. 14.5 is a quadratic function of six uncertain inputs. To avoid approximation errors, it was therefore decided to use the second-order Taylor method, which is a logical extension of the first-order Taylor method. Because the model coefficients and the field measurements were determined independently, the correlation between the β_i and cokriging prediction errors was taken to be zero. Also, the stochastic residual ε is uncorrelated with all other uncertain inputs.

The results of the uncertainty propagation are given in Fig. 14.12. The accuracy of the map of Θ_{wp} is reasonable: the standard deviation of Θ_{wp} rarely exceeds 50% of the predicted value. The uncertainty is much larger in those parts of the study area where there are no observations. This suggests that uncertainty in the maps of Θ_{fc} and Φ are the main source of uncertainty because these uncertainty maps had similar spatial patterns. Indeed Fig. 14.13 shows that the contribution of the regression model uncertainty is small. Improvement of the Θ_{wp} map can thus best be done by improving the maps of Θ_{fc} and Φ , for instance, by taking more measurements over

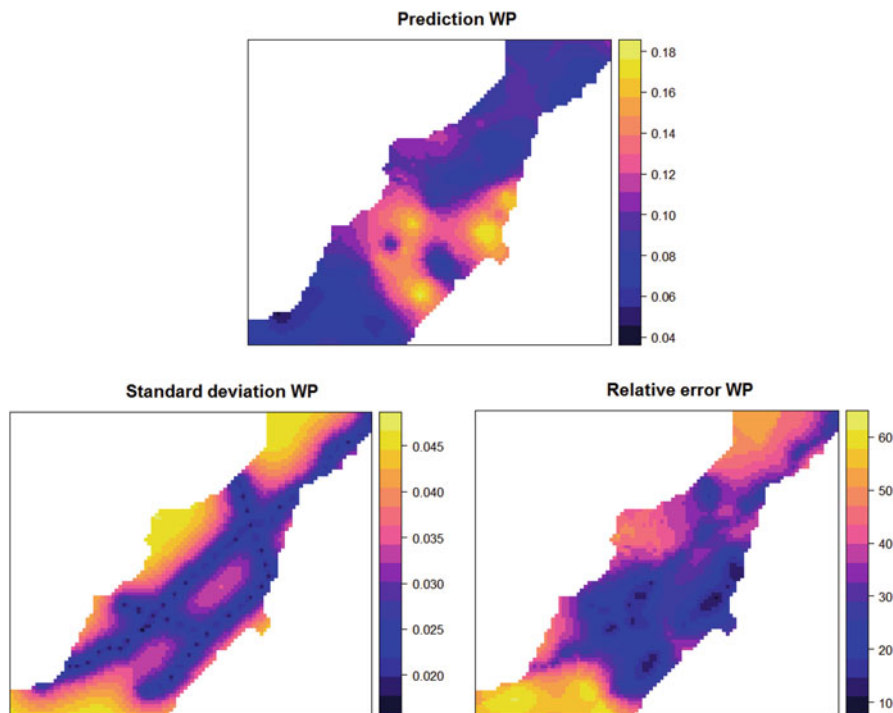


Fig. 14.12 Prediction map (*top*) of soil moisture at wilting point computed from cokriging maps of soil moisture at field capacity and porosity using a pedotransfer function, associated prediction error standard deviation map (*bottom left*) derived using the second-order Taylor method and relative error (*bottom right*)

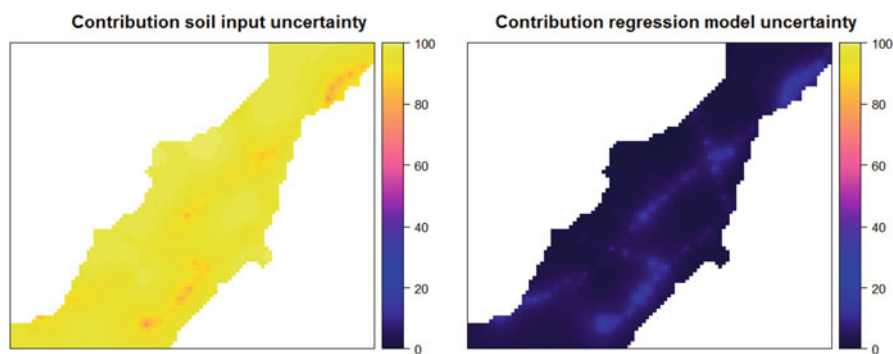


Fig. 14.13 Contributions in percentages to the overall error variance of the soil moisture at wilting point predictions as caused by cokriging errors in soil input maps (*left*) and by uncertainty in the multiple linear regression model (*right*)

the study area. The variograms and cross-variogram of Θ_{fc} and Φ could be used to assist in optimising sampling. This technique would allow judging in advance how much improvement is to be expected from the extra sampling effort.

14.5 Conclusions

No soil map is perfect. It is important to quantify the errors and uncertainties associated with soil maps because this determines whether a map is usable for an intended purpose. Any end user of soil maps should therefore require that the maps are accompanied by accuracy measures. Such measures can be computed from comparison of map predictions with independent validation data, but for spatially explicit uncertainty measures, a geostatistical approach that quantifies the map accuracy through the kriging standard deviation is recommended. Geostatistics also provides the tools to generate ‘possible realities’ by sampling from the conditional spatial probability distribution of the uncertain soil property. These possible realities may be used to communicate uncertainty and are also useful in Monte Carlo uncertainty propagation analyses.

Uncertainty propagation analysis is used to analyse how uncertainty in input (soil) maps propagates through spatial analyses and environmental models. This not only quantifies the uncertainty in the model output but can also tell us which are the main sources of uncertainty, which is essential information for taking informed decisions about how to improve the quality of maps and model results.

This chapter was limited to uncertainty quantification and uncertainty propagation of continuous-numerical soil properties and variables, but generalisation to categorical variables can be made, although it is more complicated.

This chapter concentrated on errors and uncertainties that arise from spatial interpolation and from fitting and applying linear regression models. There are many more sources of uncertainty, such as field and lab measurement error, positional error, classification error and model parameter and structural errors. These can be handled in similar ways, but the main challenge often is to characterise the error sources with realistic probability distributions. Once this is done, the uncertainty propagation analysis itself is not difficult, although it might be computationally demanding.

References

- Bishop TFA, Minasny B, McBratney AB (2006) Uncertainty analysis for soil-terrain models. *Int J Geogr Inf Sci* 20:117–134
- Brodsky L, Vasat R, Klement A, Zadorova T, Jaksik O (2013) Uncertainty propagation in VNIR reflectance spectroscopy soil organic carbon mapping. *Geoderma* 199:54–63
- Brown JD, Heuvelink GBM (2005) Assessing uncertainty propagation through physically based models of soil water flow and solute transport. In: Anderson MG et al (eds) *Encyclopaedia of hydrological sciences*. Wiley, Chichester, pp 1181–1195

- Brus DJ, Heuvelink GBM (2007) Optimization of sample patterns for universal kriging of environmental variables. *Geoderma* 138:86–95
- Brus DJ, Kempen B, Heuvelink GBM (2011) Sampling for validation of digital soil maps. *Eur J Soil Sci* 62:394–407
- Goovaerts P (2001) Geostatistical modelling of uncertainty in soil science. *Geoderma* 103:3–26
- Hastings AF, Wattenbach M, Eugster W, Li C, Buchmann N, Smith P (2010) Uncertainty propagation in soil greenhouse gas emission models: an experiment using the DNDC model and at the Oensingen cropland site. *Agric Ecosyst Environ* 136:97–110
- Hengl T, Heuvelink GBM, Stein A (2004) A generic framework for spatial prediction of soil properties based on regression-kriging. *Geoderma* 120:75–93
- Heuvelink GBM (1998) Error propagation in environmental modelling with GIS. Taylor & Francis, London. 127 pp
- Heuvelink GBM (2014) Uncertainty quantification of GlobalSoilMap products. In: Arruauys D, McKenzie N, Hempel J, Richer de Forges A, McBratney A (eds) *GlobalSoilMap. Basis of the global spatial soil information system*. CRC Press, Boca Raton, pp 335–340
- Kros J, Heuvelink GBM, Reinds GJ, Lesschen JP, Ioannidi V, De Vries W (2012) Uncertainties in model predictions of nitrogen fluxes from agro-ecosystems in Europe. *Biogeosciences* 9: 4573–4588
- Lewis PAW, Orav EJ (1989) *Simulation methodology for statisticians, operations analysts, and engineers*, vol. 1. Wadsworth and Brooks/Cole, Pacific Grove
- Malone BP, Kidd DB, Minasny B, McBratney AB (2015) Taking account of uncertainties in digital land suitability assessment. *Peer J* 3:e1366
- Minasny B, McBratney AB (2006) A conditioned latin hypercube method for sampling in the presence of ancillary information. *Comput Geosci* 32:1378–1388
- O'Hagan A, Buck C, Daneshkhah A, Eiser L, Garthwaite P, Jenkinson D, Oakley J, Rakow T (2006) *Uncertain judgements: eliciting experts' probabilities*. Wiley, Chichester
- Odeh IOA, McBratney AB, Chittleborough DJ (1995) Further results on prediction of soil properties from terrain attributes: heterotopic cokriging and regression-kriging. *Geoderma* 67:215–226
- Orton TG, Pringle MJ, Bishop TFA (2016) A one-step approach for modelling and mapping soil properties based on profile data sampled over varying depth intervals. *Geoderma* 262:174–186
- Poggio L, Gimona A (2014) National scale 3D modelling of soil organic carbon stocks with uncertainty propagation – an example from Scotland. *Geoderma* 232:284–299
- Ross SM (1990) *A course in simulation*. Macmillan, New York
- Truong PN, Heuvelink GBM (2013) Uncertainty quantification of soil property maps with statistical expert elicitation. *Geoderma* 202–203:142–152
- Van den Berg F, Tiktak A, Heuvelink GBM, Burgers SLGE, Brus DJ, de Vries F, Stolte J, Kroes JG (2012) Propagation of uncertainties in soil and pesticide properties to pesticide leaching. *J Environ Qual* 41:253–261
- Wackernagel H (2003) *Multivariate geostatistics. An introduction with applications*. Springer, Berlin
- Xiong X, Grunwald S, Brenton MD, Kim J, Harris WG, Bliznyuk N (2015) Assessing uncertainty in soil organic carbon modeling across a highly heterogeneous landscape. *Geoderma* 251:105–116