

Chapter 1

Introduction

This book focuses on two methods for simulating facies and lithotypes: truncated gaussian (Matheron et al. 1987) and plurigaussian simulations (Galli et al. 1994). The first method was developed in the late 1980s for simulating the lithotypes found in oil reservoirs because it makes much more sense to simulate the geometry and internal architecture of the reservoir first, then generate suitable values of porosity and permeability once the lithotype is known. These simulations were designed for reservoirs where the lithotypes occur in a sequential order; for example, when sandstone is followed by shaly sandstone then shale. Plurigaussian simulations are a natural extension of these. They were designed to produce a much wider range of patterns and to allow for more complicated types of contacts between facies.

Figure 1.1 shows canyons along the San Juan river in Utah where phylloidal algal mounds of Pennsylvanian age are beautifully exposed. This is typical of the sedimentary oil bearing formations that can be modelled using truncated gaussian and plurigaussian methods. This series shows tabular prograding sequences of shallow carbonates and algal mounds. The carbonates can be simulated using the simpler, truncated gaussian method. The upper part of the series consisting of algal mound and intermound facies with the sandstones capping them, is more complicated. The plurigaussian method was used to simulate its structure. The details are given in Chap. 8.

Over recent years the mining industry has also started using these two approaches to simulate the facies present in orebodies. For example, the recovery during froth flotation (a mineral processing technique for concentrating sulphide ore) depends on the rock type, rather than the grades. So for both mining deposits and oil reservoirs, the message is clear: first simulate the geometry of the lithotypes or facies then simulate its properties as a function of that type. Before outlining how truncated gaussian and plurigaussian simulations work, we would like to take the time to explain why geostatistical simulations of lithofacies are important in the petroleum industry, since this motivated the development of both methods.



Fig. 1.1 Outcrops of a typical carbonate reservoirs along the San Juan River in Utah

Putting Reservoir Simulations into Context

The aim of reservoir modelling is to construct a gridded model of the reservoir containing its petrophysical properties such as porosity, permeabilities and capillary pressure, in order to simulate its behaviour during production. As the wells are generally widely spaced and as seismics only provides indirect geological information with a low resolution, the distribution of the geological heterogeneity between wells is uncertain. Geostatistical simulations are a way of quantifying the uncertainty by providing reservoir engineers with representations of the spatial distribution of reservoir heterogeneity. Fluid flow simulations can then be performed on the model or on an upscaled model to optimise the field's development.

The precise objectives vary depending on the field's stage of development. At the appraisal stage when the aim is to produce a global reservoir model rather than a detailed one, simulations are used to estimate reserves and to quantify their uncertainty. They are also used to define recovery process scenarios. At this stage only a few wells are available and so the seismics play a crucial role.

When the field has started to be developed using primary recover processes, the aim is to optimise the location of wells. At this stage more wells are available and so the geological model is constructed using the detailed well descriptions together with the seismic interpretation. Detailed reservoir characterisation studies will have been performed, and the petrophysical variables will be available from core and log analysis. Pressure measurements in wells are used to estimate the reservoir connectivity and well test information provides a way of estimating the permeability around the wells. Seismic data can also provide information about the heterogeneity

within layers. The degree of detail required in the simulations (i.e. the grid size) depends on the recovery process chosen.

During the final stage of the field's development, the initial reservoir model has to be updated in order to make the recovery process more efficient. A very detailed reservoir model is then needed, for example, to optimise the drilling of wells with a complex geometry. By this time, the well database is much larger; it contains both geological and petrophysical information. Because of its low resolution the seismic information is less important. As the production history of the field is generally available, the observed oil and water flow rates can be compared with those computed from the simulation and used to modify the model iteratively until a good fit is obtained. This process is called history matching.

The complete procedure for reservoir simulation including simulating the lithofacies on a very fine grid and transforming these into reservoir properties, is described by Matheron et al. (1987). The initial geological modelling is a critical step because it ensures the consistency of the simulation from a geological point of view. Another crucial step is upscaling the reservoir properties to a coarse grid to reduce the number of cells because the capacity of most fluid flow simulators is still limited.

Finally, all the information available on the field has to be integrated into the model. Different data have been studied at different scales and have different characteristics. For example, microscopic information such as pore throat types obtained from special core analysis has to be combined with information derived from seismic campaigns having a vertical resolution greater than 20 m. Hard data such as well logs based on physical measurements in wells, have to be combined with soft information resulting from geological interpretation of the data. So simulations have to balance hard and soft constraints.

Idea Behind Truncated and Plurigaussian Simulations

The basic idea is to start out by simulating one or more gaussian variables (i.e. with a $N(0,1)$ distribution) at every point in the study area and then use the *rock type rule* to convert these values back into lithotypes. Figure 1.2 summarises the procedure when only one gaussian is used. This is called the truncated gaussian approach. The greytone image (left) represents the gaussian value at each point in space. In the image on the right, values below -0.6 have been coloured in dark grey, indicating one facies. Values above 0.5 have been shaded white and intermediate values have been coloured light grey. These interval definitions constitute the rock type rule.

Looking at Fig. 1.2b we see an important feature of the truncated gaussian approach: the grey facies can touch the other two facies, but black and white never touch. Or if they do, it is because the pixel size is too large. If four or five facies were simulated in this way, they would occur in a fixed order, which is defined by the sequence stratigraphy.




Fig. 1.2 (a) Simulated greytone image. Values have a $N(0,1)$ distribution. (b) Same image after being truncated at the cutoffs -0.6 and 0.5 . Values below -0.6 have been *shaded dark grey*, those between -0.6 and 0.5 are *coloured light grey* while values above 0.5 are shown in *white*

Fig. 1.3 Histogram of a standard normal distribution [i.e. $N(0,1)$] showing the two cut-offs, -0.6 and 0.5




Figure 1.3 shows the two thresholds, -0.6 and 0.5 , on a standard normal $N(0,1)$ distribution. It is easy to calculate the areas under the three parts of the curve (25, 45 and 30%) and hence deduce the proportion of space occupied by each facies. In fact there is a one-to-one relationship between the proportions and the thresholds, once the order has been established. In practice we compute the proportions experimentally and use these to deduce the thresholds.

In many cases, the truncated gaussian approach proves to be too restrictive; for example, if there is no natural sequence in the facies or if certain facies can be in contact with more than two facies. So it has been extended to two or more gaussians. Figure 1.4 illustrates this *plurigaussian* procedure for the case of two gaussians. At the top we see two gaussian images (simulations obtained using a gaussian random function). The one on the left has its long range in the NS direction while the other one has its long range in the EW direction. The square shown in the bottom right summarises the *rock type rule*, the rule that is used to assign each point to a facies. Values of the first gaussian, Y_1 , can range from $-\infty$ to $+\infty$. They are plotted along the horizontal axis. Similarly for those of Y_2 along the vertical axis. This square is divided into three regions corresponding to different lithotypes. If $Y_2 < 0$, the rock is coded as dark grey; if $Y_2 > 0$ and $Y_1 < 0$, the rock is classified




Fig. 1.4 How plurigaussian simulations work. The two images at the top are realisations of gaussian [i.e. $N(0,1)$] random functions with anisotropies in the NS and EW directions. The rock-type rule shown bottom *right* has been used to truncate these two images to obtain the facies or lithotypes shown bottom *left*. Note the contrast between the anisotropies in the simulated facies

as being white whereas if $Y_2 < 0$ and $Y_1 > 0$, the rock is light grey. The resulting three facies are shown in the lower left square.

The underlying idea in both truncated gaussian and plurigaussian simulations is to set up one or more simulations of standard normal random functions in the area of interest and to attribute the lithotype or facies depending on the simulated values at each point. This is done by truncating. When only one gaussian is used, the truncation is effectively defined by the values of thresholds. When two or more gaussians are used, the situation is more complex. It is represented graphically via the rock type rule.

In this book we will almost always use rectangles to divide the rock type rule into lithofacies. We do this because it makes it easier to work out the thresholds from the experimental proportions.

In most plurigaussian applications, the two gaussian random functions are independent of each other but it is also possible to use correlated gaussian RFs. When this is done the facies tend to “wrap around” each other. This makes it more difficult to compute the thresholds from the experimental proportions. It is also

possible to use more than two gaussians but this makes it more difficult to compute the variograms, to estimate the parameters and to carry out the simulations.

As well as being able to use independent or correlated RFs, more advanced models can be obtained by using derivatives, or translated RFs, or any linear transform of RFs, even random ones. Even more complicated models can be obtained by combining these.

Key Steps in a Plurigaussian Simulation

Having explained the broad principles behind plurigaussian simulations we need to go into more detail. Here we assume that readers are familiar with basic geostatistics (variograms and kriging) and with simulating gaussian random functions. The four main steps in a plurigaussian simulation are choosing the appropriate type of model, estimating the values of its parameters, generating gaussian values corresponding to the lithotypes at sample points and lastly running the conditional simulation, using the gaussian values generated in the previous step.

Step 1: Choosing the Model Type

Plurigaussian simulations can be divided into several broad families, depending on the types of relations between the facies or lithotypes. For example, in some cases there is a natural sequential order among the facies. In fluvial channel reservoirs, the lithotypes – sandstone, shaly sandstone and shale – generally occur in that order because of depositional conditions. To help the reader choose the type of model that is best suited to his/her data, we provided a catalogue of examples with the freeware that is available online (see Chap. 9). When there is a clear sequential order in the lithotypes, a single gaussian usually suffices; otherwise two or more gaussians can be used. In some cases, one of the facies may appear to be a « shifted » version of another facies. Increments or derivatives of gaussians can also be used to obtain special effects.

Step 2: Estimating the Parameter Values

Two key factors control plurigaussian simulations: the thresholds at which the different gaussians are truncated and the variogram model of the underlying gaussian variable. The proportion of each facies, the “rock type” rule and the correlation between the underlying gaussian random functions determine the

thresholds. Knowing the mathematical relation between the indicator variograms and the variograms of the underlying gaussian variable(s), we can find a suitable model for the underlying variogram and estimate the values of its parameters.

Step 3: Generating Gaussian Values at Wells/Drill-Holes

The facies or lithotype is known for each sample at wells (or drill-holes) but this does not tell us what the corresponding gaussian values are. We merely know that they must fall in certain domains which reduce to intervals when a rectangular partition is used. So the third step consists of generating gaussian values in the appropriate intervals and with the right properties (e.g. so as to respect the variogram model). Obviously these values are not unique. A special statistical method called a Gibbs sampler is used to generate these values.

Step 4: Simulating Values at Grid Nodes Given Values at Wells

Once the gaussian values corresponding to the facies/lithotypes have been generated at sample locations, the rest of the simulation procedure is quite straightforward. Any algorithm can be used for conditionally simulating the gaussian values at grid nodes. The last step is to convert the gaussian values at grid nodes back into facies using the rock type rule.

Recent Developments

A few theoretical advances have been made in the past 8 years since the first edition was published: Xu et al. (2006a, b) developed an approach for handling many underlying gaussians; Emery and Gonzalez (2007a, b) developed an approach for incorporating uncertainty on geological boundaries and Emery and da Silva (2009) developed a hybrid method for conditionally co-simulating continuous and categorical variables. But the main feature has been a wide range of applications in disciplines throughout the earth sciences: petroleum, mining, hydrology and environmental science. The technique has come of age in the oil industry and is now widely used for building the underlying geological model. For example, de Galard et al. (2005) used plurigaussian simulations in a study of a giant highly fractured carbonate reservoir in Iran in order to improve production; Le Maux et al. (2005) used them to study the impact of fracture modelling on reservoir performance in carbonate reservoirs; Albertao et al. (2005) used plurigaussian simulations to model an offshore reservoir in a cretaceous turbidite environment; Mubarak et al. (2009)

used them when comparing a double porosity/double permeability model to a conventional single porosity/single permeability model in fractured carbonate oil reservoirs and Pontiggia et al. (2010) used them in a study of the diagenesis of an Egyptian siliclastic reservoir. Gunning et al. (2007) used the truncated gaussian method to generate a geological model for seismic inversion.

Over the past 10 years, plurigaussian simulations have started to be used for modelling deposits in the mining industry. One of the first applications was to a granite-hosted uranium deposit (Skvortsova et al. 2000, 2002). Since then, they have been applied to a porphyry copper deposit in Chile (Carrasco et al. 2007), a roll front uranium deposit in Kazakhstan (Fontaine and Beucher 2006), a nickel laterite deposit (Rondon 2009) and to the upper portion of a diamond pipe in Botswana which is filled with sedimentary crater facies (Deraisme and Field 2006). Xu and Dowd (2008) used them when simulating rock fractures.

Applications of plurigaussian simulations are beginning in hydrogeology and environmental science. Mariethoz et al. (2006, 2009) used them to characterize aquifer heterogeneity; Cherubini et al. (2009a, b) applied them when modelling aquifer contamination; Babish (2006) included the method in a textbook designed for environmental scientists.

Probably the most significant set of new applications of plurigaussian simulations has been to construct the geological model in history matching using either an ensemble Kalman filter or gradual deformation methods, both of which require an underlying gaussian framework.

History Matching

Aanonsen et al. (2009) provide an excellent review of applications of the ensemble Kalman filter (EnKF) in reservoir engineering. The authors point out that it can be considered either as an improvement on the extended Kalman filter designed to handle more complicated data or alternatively as a sequential Bayesian inversion method. As the Kalman filter was originally developed for gaussian variables, the EnKF works best when the data are approximately gaussian.

The first step in history matching is to construct a geological model of the reservoir as input into a fluid flow simulator. As the variables in the geological model (permeability, porosity etc) depend on the lithofacies, the lithofacies are simulated first, then the porosity and permeability are simulated conditional on the lithotype. Different methods are available for simulating the lithofacies: pixel based methods such as the truncated gaussian or the plurigaussian, indicator simulations or object based methods. The advantage of using plurigaussian simulations is that they provide an underlying gaussian process on which EnKF can be applied.

Liu and Oliver (2003a, b, 2005a, b) seem to have been the first to use plurigaussian simulations to model the geology of the reservoir in history matching problems, firstly using gradient methods and then in 2005 with EnKF. Agbalaka and Oliver (2008) and Zhao et al. (2008) considered more realistic 3D reservoirs with

production data. Gu and Oliver (2007) considered multiphase fluid flow; Chen and Oliver (2010) focused on improving the mass balance; Haugen et al. (2006) and Evensen et al. (2007) tested EnKF on real North Sea reservoirs.

It is important to know how this approach (plurigaussian simulations plus EnKF) compares to others and in particular whether it gives better results than the existing method, manual history matching. Aanonsen et al. (2009) commented that: Applications of EnKF to both pseudo field cases (Naevdal et al. 2005; Gu and Oliver 2005; Lorentzen et al. 2005; Gao et al. 2006) and real field cases (Skjervheim et al. 2007; Haugen et al. 2008; Evensen et al. 2007; Bianco et al. 2007) have been very encouraging. Perhaps the most important observation is that a significantly better match of production data was obtained with EnKF than using manual history matching.

Gradual deformation methods (Hu 2000) are another approach that is used in history matching. Plurigaussian simulations have been used to set up the underlying geological model. See for example, Le Ravalec-Dupin et al. (2004), Thomas et al. (2006), Gervais et al. (2009) and da Veiga and Le Ravalec (2010).

Layout of the Book

This guidebook to plurigaussian simulations is divided into nine chapters. The structure of the second edition of the book has been changed slightly compared to the first one. A new theory chapter (Chap. 2) has been included to present the method more generally from a mathematical point of view. The Gibbs sampler is used to generate gaussian values corresponding to the lithotypes/facies at the sample points. The more mathematical aspects of its convergence which were in Chap. 6 before are now at the end of Chap. 2, together with an Appendix that provides reminders on conditional gaussian distributions.

Throughout the rest of the book, a “maths-lite” approach has been used. Having said, care is required because indicator variables are deceptively simple. People have the mistaken impression that they can be treated like ordinary variables – like grades, or porosity, or the depth to a horizon. Chapter 3 reviews their properties and those of indicator variograms. Chapter 4 is on proportions. It shows how to calculate vertical proportion curves which summarise the vertical variability in the lithotypes found in oil reservoirs, and generalises this for the case where there is horizontal non-stationarity as well. Once the proportions have been modelled, the thresholds separating the facies can be calculated. Chapter 5 presents this step. The next step (Chap. 6) is to calculate the experimental indicator variograms and fit models to the variograms of the underlying gaussian variables. In Chap. 7 we illustrate how the Gibbs sampler works via several examples.

Chapter 8 focuses on applications of truncated gaussian and plurigaussian simulations, starting with several petroleum case studies. The first is on algal mounds in the Paradox Basin in Utah. The second is a synthetic case study of a reef reservoir. Two shorter studies consider prograding patterns and fracturing in

one of the facies. Next we address the sensitivity of the variogram parameters and the rock-type rule on the simulations. This is followed by three new case studies: one on how to handle heterotopic data and then two mining cases. The first is on a roll-front uranium deposit in Kazakstan (i.e. a sedimentary deposit). The second is on a porphyry copper deposit in Chile. This is particularly interesting because it is not of sedimentary origin.

Chapter 9 presents three freeware programs (*PluriDemoSimu*, *PluriDemoVario* & *PluriDemoSet*) which were designed to allow readers to test plurigaussian simulations for themselves. They can be down-loaded from the web.