Les Cahiers du GERAD ISSN: 0711–2440

Optimizing infill drilling decisions using multi-armed bandits: Application in a long-term, multi-element stockpile

R. Dirkx,
R. Dimitrakopoulos

G–2016–103

November 2016

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Optimizing infill drilling decisions using multi-armed bandits: Application in a long-term, multi-element stockpile

Rein Dirkx [a]
Roussos Dimitrakopoulos [a, b]

[a] *COSMO – Stochastic Mine Planning Laboratory, Department of Mining and Materials Engineering, McGill University, Montréal (Québec) Canada, H3A 0E8*

[b] *GERAD Montréal (Québec) Canada, H3T 2A7*

rein.dirkx@mail.mcgill.ca
roussos.dimitrakopoulos@mcgill.ca

November 2016

Les Cahiers du GERAD
G–2016–103

Abstract: Every mining operation faces a decision on additional drilling at some point during its lifetime. The two questions that always arise with this decision are if more drilling is required and if so, where the additional drill holes should be located. The method presented in this paper addresses both of these questions through an optimization in a Multi-armed Bandit (MAB) framework. The MAB optimizes for the best infill drilling pattern while taking geological uncertainty into account by using multiple conditional simulations for the deposit under consideration. MAB formulations are commonly used in many applications where decisions have to be made between different alternatives with stochastic outcomes, such as Internet advertising, clinical trials and others. In these fields MABs have proven their effectiveness in dealing with "Big Data", massive quantities of data also being faced in the mining industry.

The proposed method is applied to a multi-element, long-term stockpile, which is a part of a gold mining complex in Nevada, USA. Particular interest goes to the stockpiles in this mining complex because of difficult-to-meet blending requirements. In many periods grade targets of deleterious elements at the processing plant can only be met by using high amounts of stockpile material. Therefore, sufficient information is required to accurately assess the uncertainty of materials in the stockpiles. The conditional simulations are generated by a direct block-support sequential simulation technique using Minimum/Maximum Autocorrelation Factors to cope with multiple elements. The best pattern is defined in terms of causing the most material type changes for the blocks in the stockpile. Material type changes are the driver for changes in extraction sequence, which ultimately defines the value of a mining operation. Therefore, these are used to assess the value of additional information. The results of the proposed method demonstrate its efficacy and applicability.

Keywords: Infill drilling, value of additional information, Multi-armed Bandits, multi-element stockpiles

Acknowledgments: We would like to thank Newmont mining corporation for their collaboration in providing us with the data necessary to conduct this research. Additionally we would like to thank Dr. Cosmin Paduraru from the COSMO lab for his guidance in using MABs and his comments on our method and paper.

1 Introduction

The problem of optimizing an infill drilling pattern can be linked to a more general problem present in many different applications (e.g. stock markets, research effort direction, project development, etc.). The more common term for this problem used across these various industries is the valuation of future information (under uncertainty). Schlee (1991) proved in terms of the economic utility theory that perfect information, always has a positive value. This means that if the information were available for free, it would always be accepted. However, the requirement of a payment to retrieve this information might render the total value negative. Perfect information is defined as a piece of information that completely resolves the uncertainty regarding the topic of the information. The assessment of the value of information can be done in various ways and for various purposes. Oostenbrink et al. (2008) show how to assess the expected value of perfect information (EVPI) in an uncertain environment in a health care application. Not only the EVPI is calculated, but also the expected value of partial perfect information (EVPPI) for parts of the projects, such as to give a direction to further research in higher detail. In the mineral industry Chorn and Carr (1997) apply the theory of value of information to the purchase of additional information to guide capital investments for large oil and gas projects. The information is valued in a real options framework by option pricing techniques (OPT). Prange et al. (2008) value the additional information of a measurement on the sealing capacity of a fault in an oil field under the uncertainty of the oil price, the uncertainty inherent to the underground and operation, and the uncertainty of the measurement itself.

In all of the aforementioned work, the value of future information is usually a monetary value. The monetary value of additional information for a mineral deposit is very hard to quantify according to the authors of this work. A simple suggestion might be to look at the misclassification cost of all blocks in the deposit, as is done by Boucher et al. (2005). This work addresses a setting with a fixed cut-off grade and only one processing stream. However, this is commonly not the case in most mining operations. Menabde et al. (2004) prove that it is optimal to use a variable cut-off grade coming directly from the schedule optimization. This makes it impossible to evaluate misclassification errors without a rescheduling of the deposit. This effect is only strengthened if multiple processing streams are considered in combination with blending of material from different sources. As is the case in the application presented herein.

Value of additional information can be defined in many different ways other than a dollar value. This paper considers additional information for a mineral deposit to be valuable if the extraction schedule is influenced. Changes in extraction schedules are often caused by a swap in material types of various blocks. The application in this work does not consider scheduling, therefore value is linked to material type changes, the main driver for schedule changes. The definition of the “best” pattern is the one adding the most value to the knowledge base. According to this paper’s definition of value, the best infill drilling pattern is the one causing the most material type changes.

The related work in infill drilling is abundant and a short review of previous work is given in the paragraph below. Diehl and David (1982) look into geostatistical methods for the classification of resources and reserves. They see the kriging variance as one of the main decision factors and link this with infill drilling, which will reduce the kriging variance and thus possibly increase the classification category. Infill drilling optimization or design has been researched widely and by multiple authors in the past (Barnes, 1989; Burgess and Webster, 1980; Gershon et al., 1988; Olea, 1975). These efforts all focus on minimizing the total kriging variance in the deposit or try to locate zones of high kriging variance as prime spots for additional drilling. A flaw in using kriging variance as a measure of variability, is that it only captures the geometric part of the uncertainty for a drilled deposit (Goovaerts, 1997; Ravenscroft, 1992). The kriging variance only depends on the relative locations of the drill holes and does not take the variability of grades into account, as it does not depend on the actual data values. Ravenscroft (1992) shows that working with geological simulations is the best way to represent variability in a deposit and that estimated orebody models give a flawed representation of the true variability in a deposit. He also mentions that the only way to accurately perform a sensitivity analysis for geological uncertainty is through conditional simulations. Therefore, simulations are used in this approach to infill drilling optimization and the assessment of additional patterns under geological grade uncertainty. Englund and Heravi (1993) use conditional simulations to optimize an infill drilling design for an environmental remediation site. The main disadvantages of this approach are its brute force nature of testing many different random patterns and that the deposit is estimated rather than simulated after the additional drill hole data is added. Goria et al. (2001) also propose a method based on conditional simulations to assess the added value of additional drill holes. An important conclusion from this work is that additional drilling does not necessarily reduce the variability in the deposit. The additional drill holes only draw values from five initial simulations that represent the spread (tenth, fiftieth and ninetieth percentiles) of the initial simulations, whereas an assessment over all available simulations would have been preferable. The two aforementioned works use a similar approach as presented in this work, where the additional drill holes draw their values from simulations based on only the initial exploration data. Boucher et al. (2005) propose a method for the optimization of infill drilling that also makes use of simulated orebody models and compares the cost of drilling additional drill holes to the misclassification cost. Blocks are classified according to material types determined by the downstream processing units and their required properties. In this approach, drilling of additional holes is only justified when the cost of drilling is less than the misclassification cost of the targeted material. A similar approach to classification based on material types is taken in the following sections. A major difference, however, is the concept of misclassification cost as value deciding factor that is not applicable to this work for reasons mentioned above. Petit et al. (2012) assess the influence of additional drill holes on the uncertainty of roll front Uranium deposits based on simulations in a Gaussian framework. The assumption of a completely Gaussian simulation framework is relatively unrealistic for highly variable mineral deposit as their distributions are usually not Gaussian and the main interest is for their high grade tails, which are smoothed out in a Gaussian framework.

One major addition in this work is the use of multi-armed bandits (MAB) for the optimization. Their benefit in the proposed method is that they provide an elegant solution to the requirement of testing every single simulation multiple times in the previous work. This drastically reduces the computational intensity for an assessment of a similar set of patterns or allows for the assessment of more patterns in a similar timeframe.

The MAB framework finds its origin in the world of casinos in the early 1900s. The initial question and the name come from the problem of playing a row of slot machines (commonly referred to as an *Armed Bandit*) in a certain sequence in order to maximize the total long-term reward from playing these machines. This provides an analogy for the well-known exploitation versus exploration trade-off[1] present in many other applications (e.g. clinical trials, internet advertisements, etc.). The exploration versus exploitation trade-off in infill drilling exists in the availability of many possible locations for additional drill holes with each an unknown outcome.

The proposed method selects the best infill drilling pattern from a pre-defined set of patterns. The representation of each pattern for infill drilling will be called an *arm* in the MAB framework. These arms will be played/pulled.[2] When an arm is played, the algorithm to assess the value for the pattern associated to that arm is set in motion. The steps in this value assessment are the following (Figure 1 demonstrates the general flow of the algorithm):

- The grades in the drill holes of the selected pattern are drawn from an initial simulation, which is set to be the possibly ‘true’ representation of the deposit.
- The drill holes of the pattern linked to the pulled arm are used as additional data to the initial exploration data for the re-simulation of the deposit.
- After the re-simulation with additional data, the blocks of the re-simulation are classified according to a material classification guide.
- The reward or contribution from pulling that arm is defined as the percentage of blocks that have changed material classification in the re-simulation compared to the possibly ‘true’ deposit. Optimizing for this reward is in line with the definition of the best pattern given above.
- The solution algorithm updates the current arm and the next arm is selected.

After convergence of the algorithm, the whole procedure is repeated for other simulations of the deposit as possibly ‘true’ deposit, in order to test the sensitivity of the method to this choice. The best pattern is then selected over the results of all these possibly ‘true’ deposits.

All of the patterns in one MAB set-up are comparable in the sense that they belong to what is called the same budget class, which means that they all have the same amount of drill holes and thus represent a similar cost to be drilled. The procedure is also repeated for different budget classes of patterns where each budget class represents patterns with a different number of holes than in the other classes. A separate MAB is required to optimize within each

[1] The terms exploration and exploitation are not used in the traditional sense they are usually used for in a mining context. Exploration here refers to trying out multiple options such that you are relatively sure of what each option entails as reward. Exploitation refers to capitalizing on the best options such that you do not lose too much value by exploring inferior options.

[2] These terms will be used interchangeably throughout the rest of the text.

budget class, as the rewards of patterns with more holes cannot be directly compared to the rewards of patterns with fewer holes.

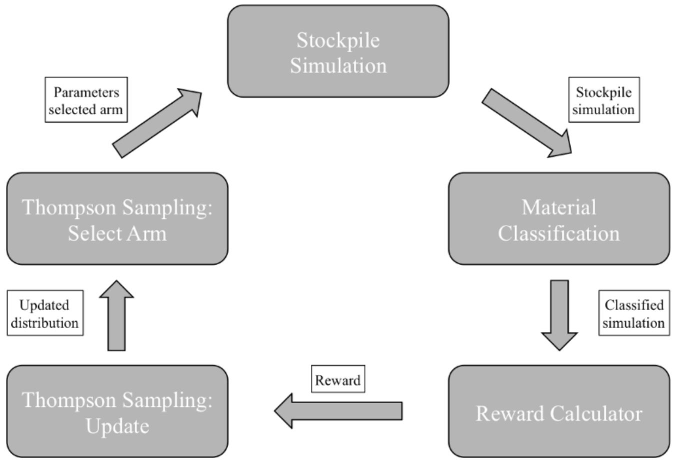

Figure 1: High-level schematic representation of the MAB optimization algorithm.

The proposed method is not only capable of selecting the best pattern within the same budget class but also to optimize between budget classes by running the algorithm multiple times in parallel. By considering multiple simulations as possibly 'true' deposit, the method is able to give an assessment of the performance of all patterns under geological uncertainty. This allows giving a maximum upside potential and downside risk for each of the patterns under geological uncertainty.

The proposed method requires simulations of a stockpile as already mentioned in the title. The stockpile in consideration for the application is a multi-element stockpile. This requires a simulation technique capable of dealing with multiple correlated elements. The simulations throughout the remainder of this work are all generated using Minimum/Maximum Autocorrelation Factors (MAF) (Switzer and Green, 1984). Desbarats and Dimitrakopoulos (2000) describe the regular MAF-technique in a mineral science context by applying it on the simulation of pore-size distribution. Boucher and Dimitrakopoulos (2009) apply it directly to the block support level; it is this direct block method that is applied in this work. The MAF simulation technique allows to co-simulate multiple correlated variables by first transforming each of them to independent factors. These factors can be simulated separately without losing any correlations. After each variable is simulated, the independent factors are back-transformed to the original data space, where the original correlations are preserved in the simulated values.

The optimization algorithm for the MAB is a heuristic algorithm called *Thompson Sampling.*[3] It is chosen because it has shown excellent performance for many problem instances (Agrawal and Goyal, 2012; Chapelle and Li, 2011; May and Leslie, 2011; Scott, 2010) and it requires weaker assumptions on the knowledge of the prior reward distributions of the arms. This is especially beneficial, because it is not known how the material type changes due to a drilling pattern are distributed. Thompson Sampling was initially described by Thompson (1933) and has become a popular heuristic in artificial intelligence and for solving MAB problems. The general idea of Thompson sampling is: *Always play the arm with the highest likelihood of being the best arm, sampled from its reward distribution.* According to this rule it will always be possible to explore arms that have not been performing as well in the previous iterations but it is most likely that the best performing arms are exploited. This gives a natural approach to the exploration vs. exploitation trade-off in MAB problems.

The next section briefly addresses some theoretical aspects of MABs, gives a detailed overview of the Thompson Sampling algorithm and shows a flow example of one iteration of the algorithm. Once the method is explained, the need for additional information for stockpiles is illustrated by the variability present in stockpile simulations and the results of the application of the method on a multi-element, long-term stockpile are presented. Finally, a conclusion and recommendations for future work are given.

[3] Sampling here is not to be confused with drilling, which is also commonly referred to as sampling. If the algorithm for solving the MAB is described, the full term Thompson Sampling will always be used. Drilling is further in this work never referred to as sampling either to avoid confusion.

2 Method

This section first gives an overview of some theoretical aspects of multi-armed bandits (MABs). An optimal solution algorithm is presented and reasons to use Thompson Sampling instead are given. After this, Thompson Sampling is explained in detail in terms of the infill drilling optimization. Finally, an overview of the algorithm for one pull of an arm is presented.

2.1 Theoretical aspects of multi-armed bandits

The classic and most simple MAB problem follows the four basic rules formulated below as in Mahajan and Teneketzis (2008). The state for an arm in a simple MAB is the internal time, which is equal to the number of times the arm has been played so far.

- Only one arm is played at each time-step, only this arm will change state and gives a reward. The reward is uncontrolled, meaning that the operator only controls which arm is played not how it is played.
- Not played arms remain *frozen*, meaning they do not change state.
- A *frozen* arm will not give a reward.
- All arms are independent.

There are many variants of MAB problems where one or multiple of these rules are violated or altered, which makes them considerably more complicated to understand and to solve. Therefore, the focus of the following sections will only be on these simple bandits, as these suffice to solve the infill drilling problem.

The mathematical objective function of the MAB problem, with *k* arms, is given in Equation (1). The same formulation as in Mahajan and Teneketzis (2008) is used.

$$J^{\gamma} = \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \sum_{i=1}^{k} R_i\left(X_i\big(N_i(t)\big), U_i(t)\right) | Z(0)\right] \tag{1}$$

With:

- J^{γ}, the objective function value under policy γ, defined by $U_i(t)$ below.
- β^t, the discount factor with time *t*, necessary to yield a finite total reward even when the time scale is infinite. The discounting also reflects that earlier rewards are higher valued than later rewards, much like in an NPV calculation. The time *t* is a measure for the number of iterations/plays.
- $R_i\left(X_i\big(N_i(t)\big), U_i(t)\right)$, the reward of arm *i* depending on:
- $X_i\big(N_i(t)\big)$, the state of arm *i* at time *t*, which in term depends on:
- $N_i(t)$, the local time of arm *i* at time *t*; the number of times arm *i* has been played before time *t*.
- $U_i(t)$, the policy decision, which determines whether or not arm *i* is played at time *t*.
- $Z(0)$, the initial state of all arms *i*.

The total objective function value is then found by taking the expected value over the summations over time and the number of arms conditional to the initial states of all arms. The expected value is required because of the stochastic rewards of the arms. The goal is to maximize the objective function value by choosing the scheduling policy γ that determines the values of $U_i(t)$ for every arm *i* and time *t*:

$$U_i(t) = \begin{cases} 1 & \text{if arm } i \text{ is played at time } t \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

An optimal policy γ will determine $U_i(t)$ for every arm *i* and time *t* in such a way that the total expected reward is maximized. If an arm is not played ($U_i(t) = 0$) the reward will be zero (rule 3). A very important remark regarding the solution of the MAB problem is that it can also be used to find the best performing arm, instead of a schedule of when each arm should be played. In that set-up the optimal schedule will converge to only playing one certain arm after some initial iterations and it can be concluded that this arm is the optimal arm. The approach to find an optimal arm rather than a schedule is used in this work. The local time $N_i(t)$ of arm *i*, changes only when arm *i* is played at time *t* (rule 2):

$$N_i(t+1) = \begin{cases} N_i(t) + 1 & \text{if } U_i(t) = 1 \\ N_i(t) & \text{if } U_i(t) = 0 \end{cases} \tag{3}$$

Gittins (1979) showed that there is an optimal solution to the simple bandit problem in the from of *dynamic allocation indices*, later coined *Gittins Indices* after the inventor. The solution method of Gittins is based on forward induction, which is only possible because decisions are irrevocable. For example, the decision not to play a given arm *i* at time *t*, while it is in state $X_i\big(N_i(t)\big)$, does not make this play unavailable in the future. The state remains frozen while the arm is not played during a certain period Δt: $X_i\big(N_i(t+\Delta t)\big) = X_i\big(N_i(t)\big)$. Therefore this arm can be played in the same state at time $t + \Delta t$ as it was at time *t*.

Gittins index solutions are only possible because of the four basic rules of the simple MAB mentioned above and in a Bayesian setting with known prior reward distributions. May et al. (2012) comment on the Gittins index solution method that the quality of its results relies heavily on the selected priors, as these priors are unknown in our case it is hard to justify the use of this solution scheme. Other shortfalls of the Gittins index solution method according to Scott (2010) are that it is computationally difficult, it is hard to generalise for all reward distributions and that it relies on artificial discounting terms. The combination of these arguments against the Gittins index solution and the aforementioned excellent performance of the Thompson Sampling algorithm for MAB problems, favours the use of Thompson Sampling to solve the infill drilling MAB.

2.2 Thompson sampling in terms infill drilling

Thompson Sampling requires an assumption for the reward distributions for each arm. As already mentioned above the true reward distributions are unknown. This is common in many applications of MABs and is overcome in Thompson Sampling by assuming a uniform prior over [0,1] for each arm. The prior is represented as a Beta-distribution with parameters $\alpha = 1$ and $\beta = 1, B(1,1)$. This will lead to a pure exploration phase at the start because all distributions are identical. The algorithm will gradually evolve to a pure exploitation phase of the highest reward-giving arm at convergence. The Beta-distribution is chosen to represent the reward distribution of the arms, because they are easy to interpret for the rewards and they are easy to update during the Thompson Sampling algorithm.

The selection of the arm to be played at each iteration is based on a random sampling of each arm's posterior distribution. The arm with the highest sample value is played. All other distributions remain unchanged, as not played arms are frozen. This procedure is repeated until convergence.

After the selection of which arm to play, the deposit gets re-simulated with the additional data from the infill drilling pattern associated to that arm and the initial exploration data. The reward is, as mentioned above, the number of blocks that have changed their material type. This is determined by comparing every block in the re-simulation to the same block in the possibly 'true' deposit. The number of blocks that have changed their material type is expressed as a percentage over the total number of blocks and this is the reward of that play. The material type classification guide is provided by the mining operation, which takes the grade of all relevant elements into account. The reward is fed back to the Thompson Sampling procedure and the distribution of that arm is updated.

The distribution update depends on the reward of the pull. Traditionally the distribution is updated according to the fact whether or not a pull is a success (reward is 1) or a failure (reward is 0). If it is a success the posterior of arm i is updated to $B_i(\alpha_t + 1, \beta_t)$ in iteration *t*, otherwise it is updated to $B_i(\alpha_t, \beta_t + 1)$ in iteration *t*. With α_t and β_t the parameters of the beta-distribution in iteration *t* updated for all previous rewards. Because the reward is defined as the percentage of blocks that change their material, it is not binary but it is defined on [0,1]. This adds a slight complication to the updating of the distribution. This complication is resolved by first performing a Bernoulli trial with the reward defined on [0,1] as the success rate. The binary result of the Bernoulli trial is then used to update the distribution in the same way as expressed above. Pseudo-code for the Thompson Sampling algorithm with Bernoulli trial updates is provided in the box below. The parameters of the beta-distribution follow the same definition as above. The parameter T in the algorithm is a generic parameter to denote the stopping time of the Thompson Sampling algorithm.

Algorithm 1: Thompson Sampling

for t = 1,2,…,T **do**
  for i = 1, 2, …, numberOfPatterns **do**
    sample(i) = sample reward distribution of pattern i in iteration t from $B_i(\alpha_t, \beta_t)$
  End for
  Play arm j = argmax$_i$ [sample (i)] and observe reward r
  success = outcome of Bernoulli trial with success rate r
  if success == 1 **then**
    $B_i(\alpha_t + 1, \beta_t)$
  else
    $B_i(\alpha_t, \beta_t + 1)$
  End if
End for

The Thompson Sampling algorithm is proven to converge for the MAB problem by Agrawal and Goyal (2012) and May et al. (2012), but it has no pre-defined termination moment. Therefore, it is required to define a convergence criterion that works for the application in consideration. In this case a simple criterion is defined, which has shown satisfactory results. The applied criterion is: *If in the last 20% of iterations one arm is selected to be played more than 90% of the iterations, this is the best arm and the algorithm has converged.* Important to note here is that this criterion is not checked for the first iterations (e.g. 100) to allow for some initial exploration. However, in some cases the convergence criterion is only met after a very large number of iterations. Therefore, an extra measure is imposed in the form of a maximum number of iterations (e.g. 2000 or 5000). If the maximum number of iterations criterion terminates the algorithm, the best arm is the one that was played most and has the highest average reward over all plays. A case like this can occur when two patterns are performing very similarly and no statistically significant distinction can be made between the rewards of these two patterns.

2.3 Detailed flow of one pull of an arm

The flow through the algorithm is illustrated with a series of images (Figure 2 to Figure 7) that show all steps in detail. Figure 2 shows the starting point of the algorithm: the deposit in consideration has been explored by an initial exploration drilling campaign and this has revealed some values in the deposit. The deposit in this example has a total of 75 blocks, as each square in the figure represents one block. Figure 3 shows one simulation of the deposit based only on the data of the initial exploration drill holes. This simulation is selected as the possibly 'true' deposit for this run of the algorithm. A new possibly 'true' deposit is selected when the algorithm has converged. The whole procedure is repeated with this new possibly 'true' deposit. For the example in this section a two-armed bandit is modelled with two infill drilling patterns demonstrated in Figure 4 and Figure 5. Each of these patterns corresponds to the same budget, as they both have three additional drill holes represented by the yellow lines. For the sake of example assume that Thompson Sampling selects arm 1, linked to pattern 1 in Figure 4, to be played in this iteration. The data used for the re-simulation of the deposit is illustrated in Figure 6. The used data are the initial exploration data and the values in the additional drill holes from pattern 1, drawn from the possibly 'true' deposit. Figure 7 illustrates the result of the re-simulation. The re-simulation is now the key measure for the reward definition in this iteration. The reward is calculated by comparing the re-simulation (with additional data) in Figure 7 with the possibly 'true' deposit in Figure 2. As can be seen thirteen blocks changed material type and the reward is 17.3%. After the update the next arm to be played is selected and all the above steps are repeated until the algorithm has converged.

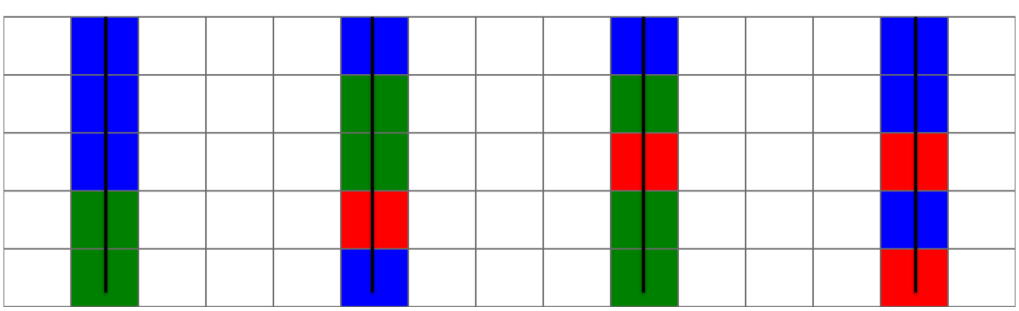

Figure 2: The original deposit with exploration drill holes. Blue is low, green is medium, and red is high value. The lines in black represent the initial exploration drill holes.

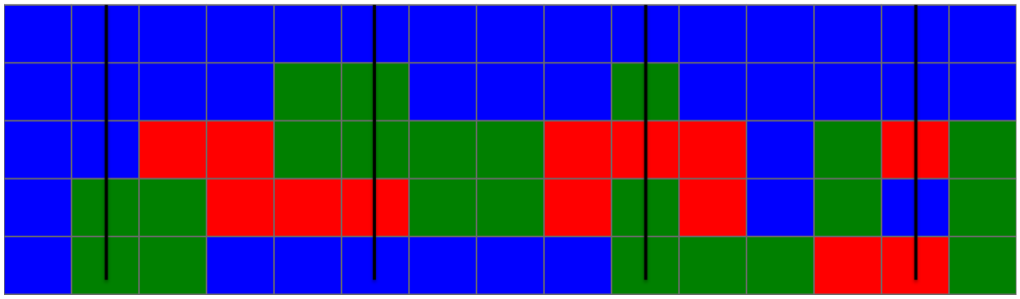

Figure 3: The simulation based on the initial exploration drill hole data, which is set to be the possibly 'true' deposit for this run of the algorithm.

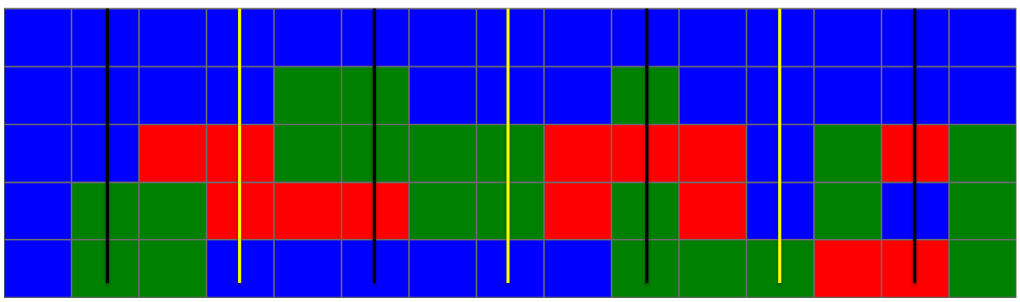

Figure 4: The yellow lines are the additional drill holes associated with pattern 1. Shown with the possibly 'true' deposit as values in the drill holes.

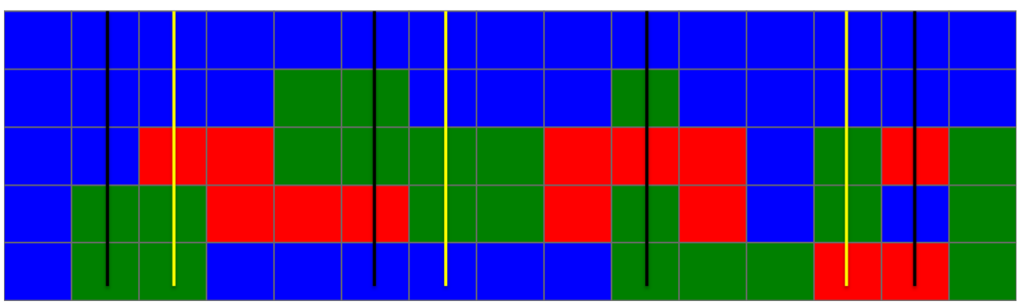

Figure 5: The yellow lines are the additional drill holes associated with pattern 2. Shown with the possibly 'true' deposit as values in the drill holes.

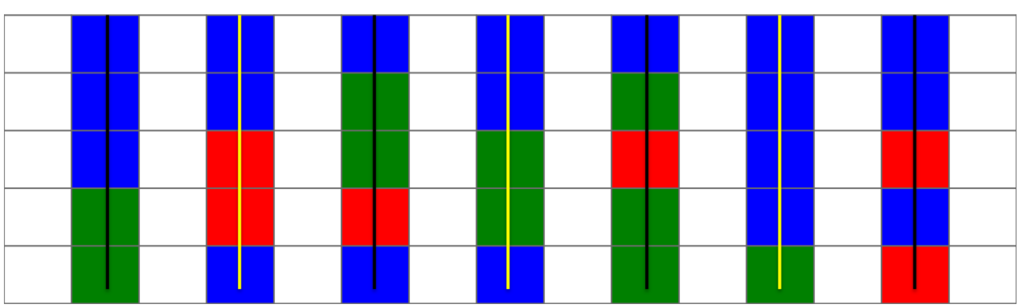

Figure 6: The data used for the re-simulation of the deposit after the selection of arm 1. The yellow lines are the drill holes from the selected infill drilling pattern and the black lines are the initial exploration drill holes.

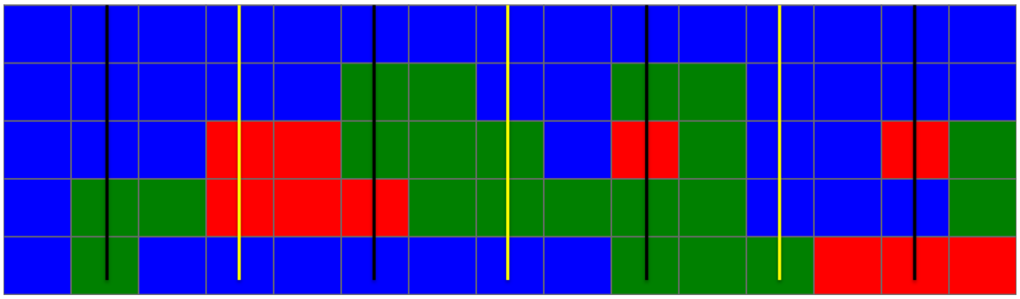

Figure 7: The re-simulation for the deposit based on the data shown in Figure 6 above.

3 Case study – A multi-element, long-term stockpile

The case study presented below illustrates the optimization of the infill drilling decision for a multi-element, long-term stockpile of a gold mining complex in Nevada, USA. This complex consists of two open pit mines, an underground mine and several stockpiles. The downstream processing includes an autoclave, an oxide mill and multiple heap leaches. Deleterious elements are co-simulated with the gold grade because of strict grade requirements on the deleterious elements to guarantee efficient processing recoveries. Knowledge of these elements can assure blending of ore from different sources to meet the constraints at the processing facilities. The elements begin considered are gold, sulphide sulphur, organic carbon and carbonate.

For this case study three classes of patterns are considered, each corresponding to a different budget. The first budget class has patterns with five holes, the next class has patterns with ten holes and the last class has patterns with fifteen holes. Eight patterns are considered for the optimization in each class. For every class the procedure is repeated for twenty simulations as possibly 'true' stockpiles to guarantee that the results are independent from the selected possibly 'true' stockpile. Optimizing over multiple simulations also allows for an indication of how these patterns perform under geological uncertainty. The performance under geological uncertainty is highlighted in the final part of this section through a validation-step that tests all patterns for a set of simulations not used in the optimization procedure.

3.1 Stockpile simulations

As mentioned above it is required to use a technique capable of co-simulating multiple correlated variables for the stockpile simulations because there are four correlated elements of interest. The first element is gold (Au), which is the produced metal. Sulphide sulphur (SS), organic carbon (OC) and carbonate (CO_3) are of importance for the working of the downstream processing facilities. The main constraints at the autoclave, where the sulphide material from the stockpile is processed, are a maximum organic carbon content and an upper- and lower bound on the sulphide sulphur – carbonate ratio.

Figure 27 to Figure 30 in the appendix show histogram comparisons for each simulated variable between the data and the simulations. The two point statistics in the form of variograms and cross-variograms are presented for all variables for the data and the simulations in Figure 31 to Figure 40 in the appendix. All these figures show a good representation of the data in the simulations.

Figure 8 below shows the shape of the stockpile with the initial exploration drill holes. The exploration drilling is done on a regular grid with a 40 feet spacing. In total there are 104 drill holes with one sample each. The height of the stockpile is 20 feet. The size is ± 200,000 m^3, which corresponds to ± 400 *kton*.

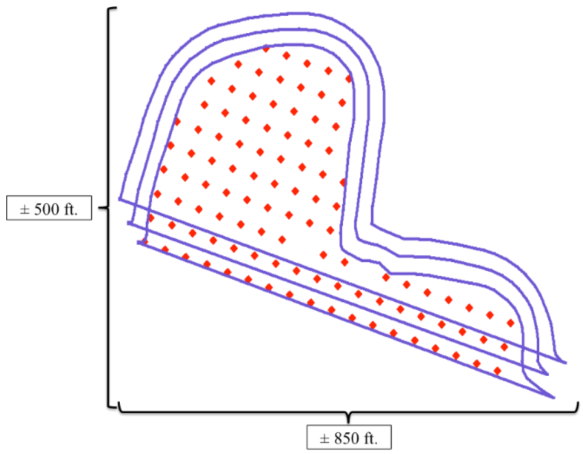

Figure 8: Planar view of the stockpile under consideration. The red diamonds are the exploration drill holes. The blue lines represent the outline of the stockpile on different height levels 10 feet apart.

Twenty simulations are generated each containing 204 blocks within the shape of the stockpile. The block size is 30x30x20 feet. This is the same as in the mines feeding the stockpile, because the same mining selectivity is used for both the stockpiles and the mines. The point discretization for the stockpile simulations is 6x6x2, meaning that each block contains 50 points.

Figure 9 below shows the material variability over all twenty stockpile simulations at the block level. The material is classified based on the material classification guide provided by the mining complex. The distinction between the different sulphide materials (e.g. high sulphide versus medium sulphide) is made based on gold grade. The difference between oxide versus sulphide material comes not only from the gold grade but also form the other elements. The mining complex flags the stockpile under consideration as medium sulphide material, but Figure 9 shows that in all simulations there is less than 30% medium sulphide material present. High and low sulphide material usually dominate the material distribution in the stockpile. These two material types also have a much higher variability in quantity over all simulations than the medium sulphide material. This can be in part explained by the fact that the gold grade criteria for the medium sulphide material are tighter than for the low sulphide or the high sulphide, which does not have a cap for gold grade. Similar results are seen in Figure 10, which demonstrates the spatial material variability for three simulations on the block level. It is usually the high and low sulphide materials that dominate the material types in the stockpile.

Figure 10 shows that there is a large amount of spatial material variability through the stockpiles, within one simulation but also from one to another simulation. This observation is one of the main motivations for the application of the proposed method to stockpiles. It shows that more information is required to assess the local scale variability in stockpiles. This is especially true for the stockpile in this mining complex. Because in some periods it is the main contributor to the processed blend at the autoclave, which has tight constraints on the deleterious elements.

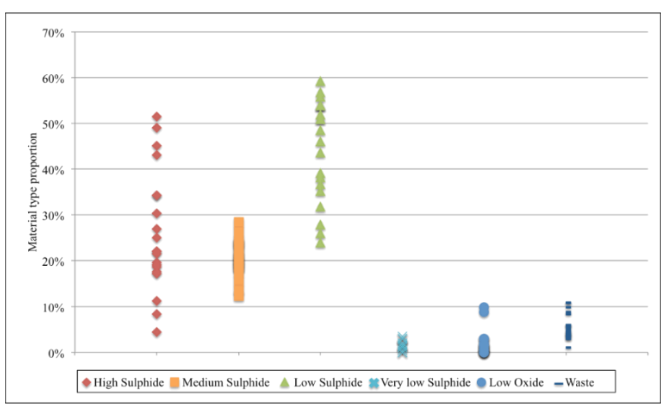

Figure 9: Material variability at the block level over all twenty stockpile simulations.

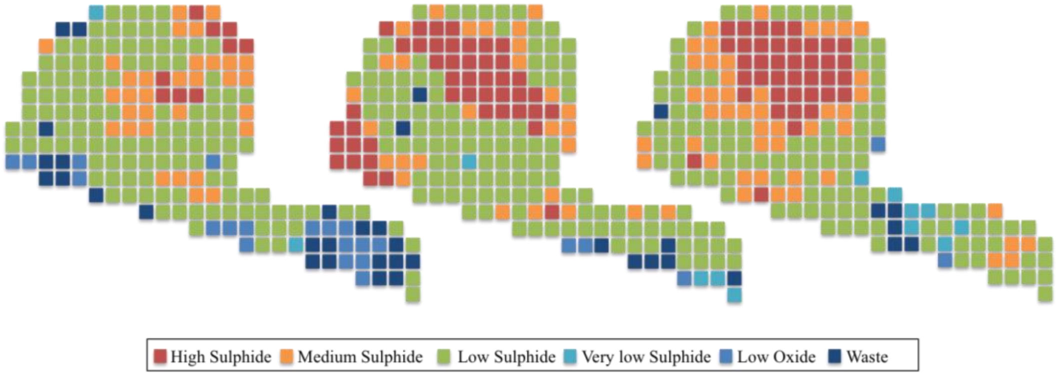

Figure 10: Three simulations that show the spatial material variability of the stockpile at the block level. Block size: 30x30x20 feet.

Figure 11 to Figure 14 below show the spatial material variability in two simulations for the four elements of interest. All simulations show that the same main features are respected between two simulations but that there is local scale variability for the grade within each simulation. The gold simulations, in Figure 11, show the presence of high grades in the centre and toward the top of the ball shape on the east of the stockpile. The high sulphide material zones shown in Figure 10 are a result of these high gold grades.

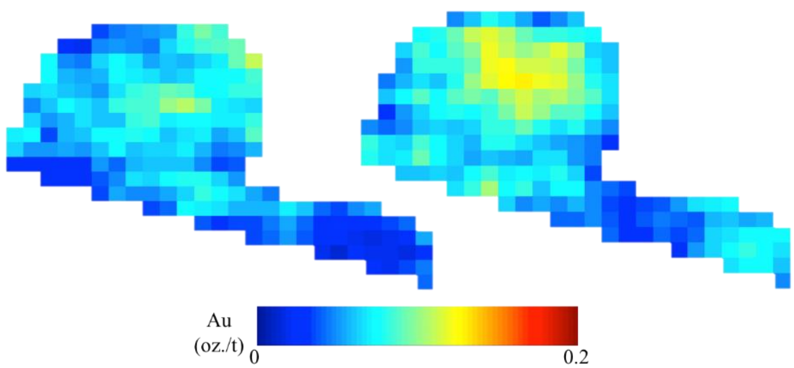

Figure 11: Gold grade variability in two stockpile simulations at the block level. Block size: 30x30x20 feet.

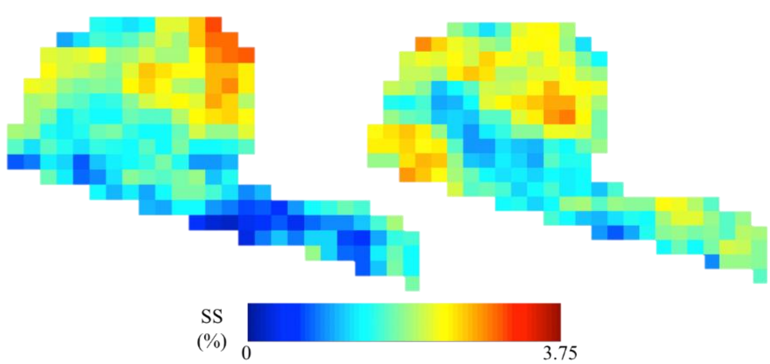

Figure 12: Sulphide sulphur content variability in two stockpile simulations at the block level. Block size: 30x30x20 feet.

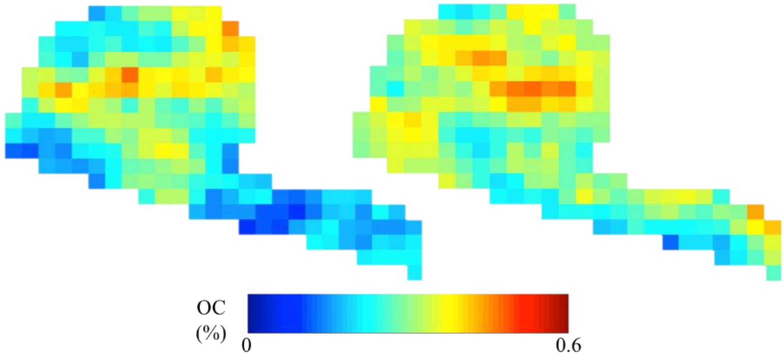

Figure 13: Organic carbon content variability in two stockpile simulations at the block level. Block size: 30x30x20 feet.

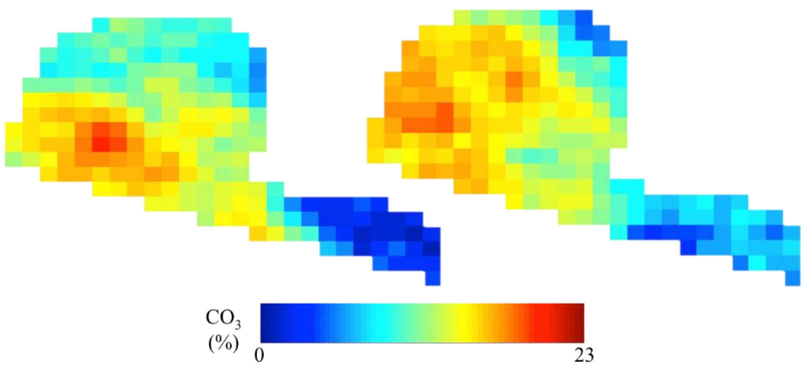

Figure 14: Carbonate content variability in two stockpile simulations at the block level. Block size: 30x30x20 feet.

3.2 Tested patterns

Figure 15 to Figure 18 below show the eight patterns that are considered in the optimization. The figures always show all fifteen additional drill holes for each pattern on top of the initial exploration drill holes. However, if only the red squares (“First 5 Holes”) are considered the patterns represent the budget class with five holes per pattern. If the red squares and the green triangles (“Second 5 Holes”) are considered the patterns represent the budget class with ten holes per pattern. The patterns have the same numbering over the budget classes because they are always based on the same idea for the location of the holes.

Pattern 1 is constructed as an evenly spread pattern through the whole stockpile. Pattern 2 is constructed as a pattern that focuses on a high gold grade zone in the top of the stockpile. Pattern 3 is constructed to mimic the delineation of the stockpile on the road side. Pattern 4 is constructed to target some specific zones of high grade spread over the stockpile. Pattern 5 is constructed to delineate the borders of the stockpile more precisely. Pattern 6 is constructed based on the blocks that have the highest coefficient of variation in gold grade over the twenty stockpile simulations. The blocks considered here are the blocks on which the stockpile was simulated with size 30x30x20 feet. Pattern 7 is constructed by randomly selecting drill hole locations. Pattern 8 is constructed as a mix of points and blocks with a size of 10x10x10 feet with the highest coefficient of variation over all twenty simulations. It was required to mix these two scales to look at the highest coefficients of variation because all the fifteen points with the highest coefficients of variation are located right next to another. A pattern like that would not make sense. To add some diversity in the locations of the additional drill holes the blocks of size 10x10x10 feet with the highest coefficient of variation over all twenty simulations are added.

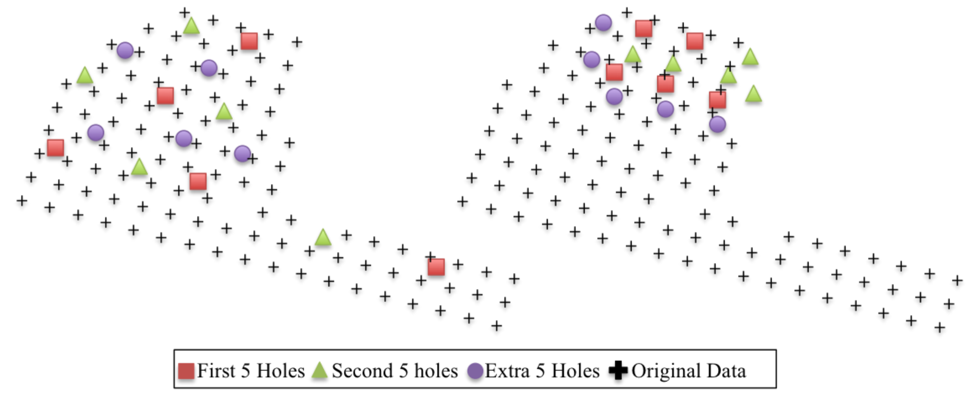

Figure 15: Pattern 1 (left) and pattern 2 (right) with fifteen additional drill holes. The black crosses are on the same locations as the red diamonds in Figure 8.

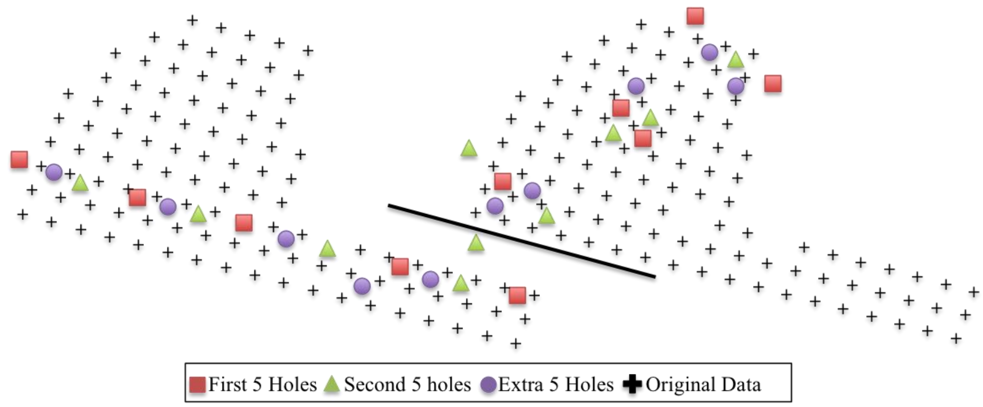

Figure 16: Pattern 3 (left) and pattern 4 (right) with fifteen additional drill holes. The black crosses are on the same locations as the red diamonds in Figure 8. The black line indicates that the green triangle below it belongs to pattern 3.

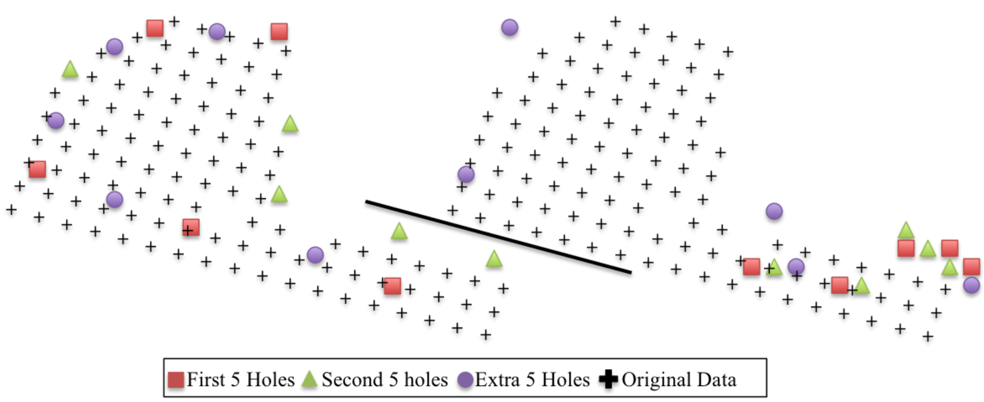

Figure 17: Pattern 5 (left) and pattern 6 (right) with fifteen additional drill holes. The black crosses are on the same locations as the red diamonds in Figure 8. The black line indicates that the green triangles below it belong to pattern 5.

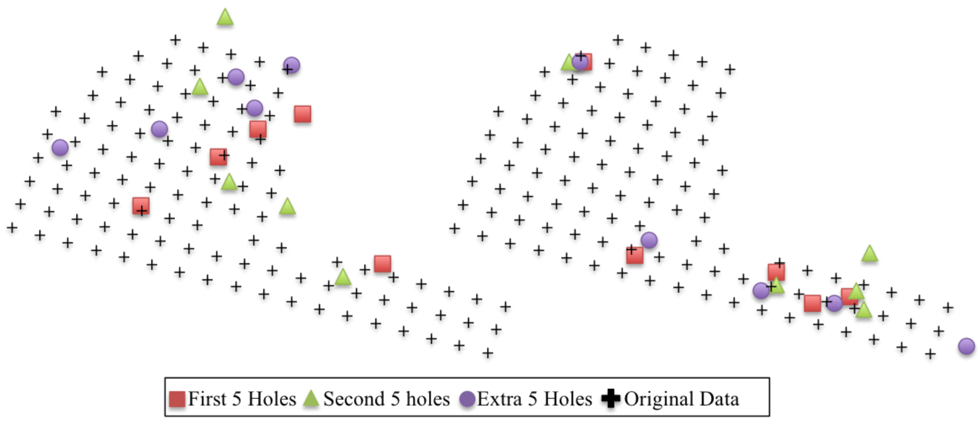

Figure 18: Pattern 7 (left) and pattern 8 (right) with fifteen additional drill holes. The black crosses are on the same locations as the red diamonds in Figure 8.

3.3 Case study results

Figure 19 shows the convergence results for the algorithm over all simulations for every pattern in each budget class. Convergence results show how often the algorithm has converged on that specific pattern in that budget class. This is divided by the total number of times the algorithm is run to give the percentages in the graph. In this case study the algorithm is run twenty times, once for each stockpile simulation as the possibly 'true' stockpile. The pattern that has been converged upon the most is called the 'winner' of that budget class. For the ten- and fifteen-hole budget classes there is always a clear winner in pattern 1. However, for both budget classes it has to be noted that pattern 7 is also a strong performing pattern, especially for the fifteen-hole budget class, where the difference between pattern 1 and 7 is only 10%. The observation on a winner for the five-hole budget class is much less clear. Pattern 2 has 30% convergence but pattern 1, 4 and 5 have each 20% too. This is only a small difference and no strong, clear decision on the winner can be made from this.

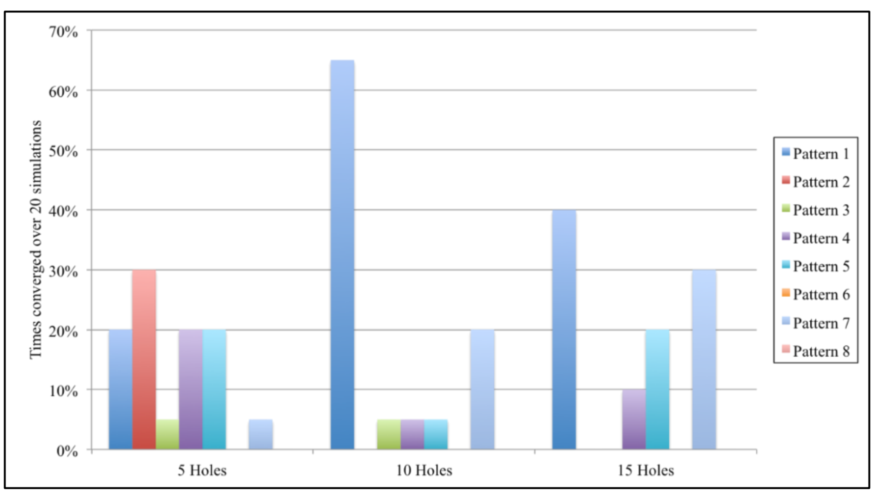

Figure 19: Convergence results over all twenty simulations for all patterns and every budget class.

Figure 20 below shows the average reward over all twenty simulations for every pattern in each budget class. This graph helps to better formulate a conclusion on which pattern is the best performer in every budget class. For the ten- and fifteen-hole budget classes it is pattern 1 that has the highest average reward as is expected from the convergence results in Figure 19. The link between these figures is apparent. For the fifteen-hole patterns it is also pattern 7 that has an average reward that is close to the highest average reward of pattern 1. However, the results for the five-hole budget class are not as one would expect in the sense that pattern 2 does not have the highest average reward. Pattern 1 has the highest average reward for the five-hole budget class. This can be explained by having pattern 2 strongly outperform pattern 1 for some simulations as the possibly 'true' stockpile, while in others pattern 2's performance is much worse than pattern 1's. Therefore, it is concluded that pattern 1 is more robust against geological uncertainty than pattern 2. Figure 20 strengthens the observation of no clear winner for the five-hole budget class by demonstrating that the average rewards for all patterns are very close to each other. It is hard for the algorithm to distinguish patterns that yield such similar results. This is the reason why the convergence results for the five-holes budget class in Figure 19 are less clear than for the two other budget classes.

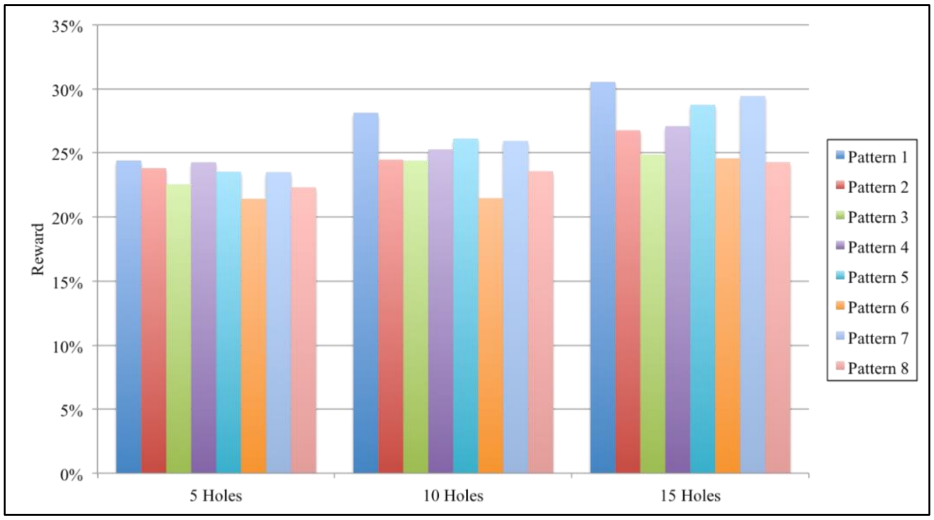

Figure 20: Average reward over all twenty simulations for all patterns and every budget class.

The differences between the running times observed in Table 1 can also be explained by the observations in Figure 19 and Figure 20 mentioned above. Because the five-hole patterns yield similar average rewards, the algorithm has a hard time finding the winner and therefore it requires more iterations, and thus time, to evaluate the patterns over all simulations. The five-hole budget class requires almost three times more time than the fifteen-hole budget class and almost two times more time than the ten-hole budget class to be optimized. A special note (denoted by the asterisk*) has to be made on the MAB set-up of the algorithm for the five-hole budget class. This set-up allows for 5000 instead of 2000 iterations as the maximum iterations without convergence. This is because initial runs with only maximum 2000 iterations had ten simulations that did not converge. Another addition to this set-up is a forced exploration-phase at the start of the algorithm. This phase forces that each pattern is played twenty times before the algorithm starts doing its normal steps. Again this is done to overcome convergence issues observed in initial tests without the forced exploration, as there were seven simulations that did not converge.

Table 1 also shows the number of times the algorithm did not converge for every budget class. For the ten- and fifteen-hole budget classes, this is after 2000 iterations. For the five-hole budget class it is after 5000 iterations. The number of non-converging simulations is low for every budget class. To get it to zero, it might require unnecessary large amounts of iterations because it is always possible that two patterns perform very similar for a certain simulation. Therefore, it is not really seen as a problem because the algorithm still identifies the two patterns that perform equally well in these cases. As a tiebreaker the average rewards are compared and the pattern with the highest average reward is attributed the convergence. A summary of which patterns and simulations caused the non-convergences for each budget class is presented in Table 2. From these results it is clear that the non-convergence is concentrated on six different simulations over the three budget classes. One of the tied patterns is in most cases pattern 1.

Table 1: Overview of the statistics of the algorithm for each budget class. *Indicates a slightly different set-up for the five-hole budget class, more on this in the text.

Budget Class	Running time	Number of iterations	Number of times not converged
5 Holes*	24 h 24 min 46 s	25,465	3
10 Holes	13 h 9 min 31 s	13,085	4
15 Holes	8 h 41 min 31 s	9,378	2

Table 2: Detailed overview of the non-converging simulations and patterns. The pattern that has the highest average reward in case of a tie is marked in bold and is attribute the convergence.

Budget Class	Non-converging simulations	Tied patterns
5 Holes	3	**2** vs. 6
	6	1 vs. **2**
	17	1 vs. **4**
10 Holes	3	1 vs. **3**
	5	**1** vs. 2
	7	**1** vs. 2
	17	**1** vs. 4
15 Holes	5	**1** vs. 2
	8	3 vs. **4**

Figure 21 below shows the convergence results if the ties in the non-convergence case of Table 2 are broken in the opposite way. For the five-hole budget class, pattern 1 gains 10% and pattern 2 looses 10%, making pattern 1 the most converged upon as indicated by the highest average reward in Figure 20. Pattern 4 looses 5% and pattern 6 gains 5%, but this does not influence any conclusions made above. For the ten-hole budget class, pattern 1 looses 15% but also gains 5%, this means a net-loss of 10%. Pattern 2 gains 10%, pattern 4 gains 5% and pattern 3 looses 5%. All these changes do not alter the conclusions for the ten-hole budget class; it is still pattern 1 that performs the best in average reward and convergence percentage. For the fifteen-hole budget class, pattern 1 and 4 each loose 5% and patterns 2 and 3 each gain 5%. This does not alter the conclusions either. However, the difference between the convergence percentage of patterns 1 and 7 is only 5% in this case. Making pattern 1 a less clear winner in terms of convergence percentage. The difference in average reward is still sufficient to conclude that pattern 1 is better than pattern 7 for the fifteen-hole budget class.

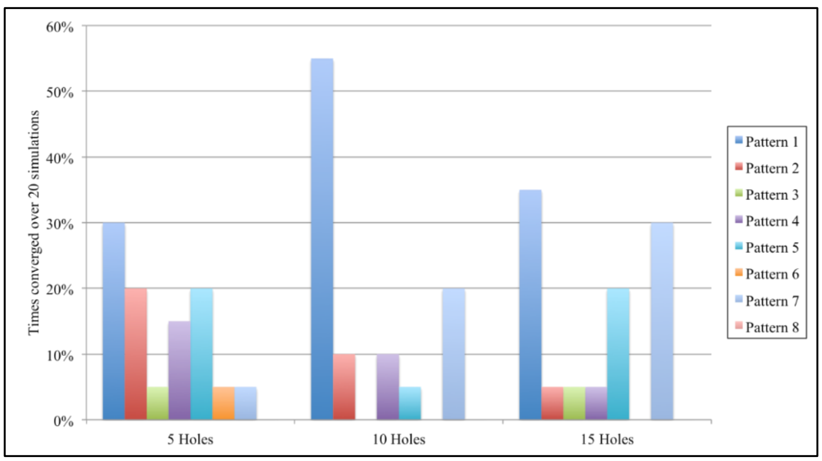

Figure 21: Convergence results over all twenty simulations for all patterns and every budget class when the ties presented in Table 2 are broken in the opposite way.

As an intermediate conclusion from the graphs in Figure 19 and Figure 20, and the discussion on the non-convergence above (Figure 21), pattern 1 is the best performing pattern in terms of average reward for each budget class. In addition to this it is important to remark that the five-hole patterns might not be sufficient for this stockpile. This is concluded from the fact their average rewards are all very similar no matter where the exact locations of the drill holes are.

To reach a conclusion concerning what pattern between the budget classes is preferred, an extra graph is presented in Figure 22. This graph demonstrates the spread on the best pattern's average reward compared to the average of the rewards of all patterns over all simulations. The tenth percentile (P10) for the downside risk and the ninetieth percentile (P90) for the upside potential under geological uncertainty are the measures for the spread on the best pattern's performance.

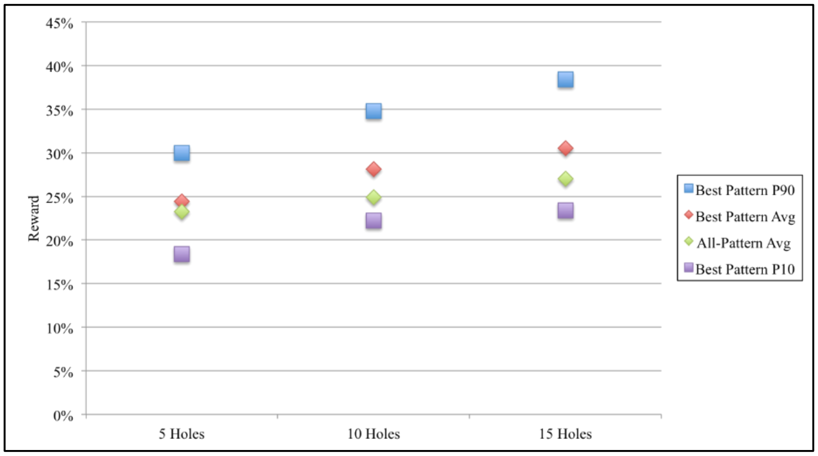

Figure 22: Profile of the best pattern for every budget class versus the all-pattern average for every budget class.

Figure 22 shows that the five-hole budget class performs worse than the other two budget classes. The five-hole best pattern average is lower than the all-pattern average for both other budget classes. Also the downside risk is much higher and the upside potential is much lower. Therefore, in combination with the reasons mentioned above, it is not recommended by the author to go with a five-hole pattern for infill drilling in this stockpile. Purely based on performance pattern 1 with fifteen holes is better than pattern 1 with ten holes because it has a higher best pattern average and a higher upside potential. However, the extra five holes also cost more to drill, this could be an extra consideration to take into account and choose for pattern 1 with ten holes. Another observation is that the downside risk for the ten- and fifteen-hole pattern 1 is almost equal, this means that they both guarantee similar results in the worst-case scenario. This is an extra argument in favour of pattern 1 with ten holes, especially if there are constraints on the budget that are more important than having the best possible knowledge.

3.4 Results validation

As a validation step the performance of the patterns is tested on an alternate set of twenty simulations, which are not used in the optimization process. Every pattern is tested twenty times for each simulation in the alternate set. One test is the same as the evaluation procedure in the MAB described above; the stockpile is re-simulated based on the pattern data from the possibly 'true' stockpile (in this case a simulation from the alternate set of simulations) plus the original data and the reward is calculated. Twenty tests for every simulation result in 400 rewards for every pattern in each budget class. The rewards are summarized by their average and profile based on the ninetieth and tenth percentiles as demonstrated in Figure 23 to Figure 26 below.

Figure 23 to Figure 25 are the direct counterparts of Figure 20 above. The comparison made in the three figures below shows that the results on the average reward for each pattern are not strongly dependent on the simulations used in the optimization process. For every budget class it is still pattern 1 that has the highest reward, which is more pronounced for the ten-hole and fifteen-hole classes. The five-hole class shows the same behaviour in regards of very similar average rewards for all patterns. Therefore, the same conclusions are made based on these graphs as on the graph in Figure 20.

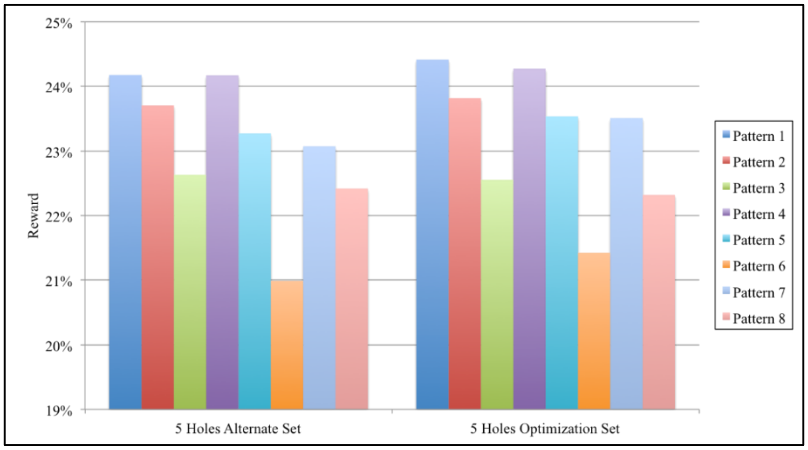

Figure 23: Average reward over all twenty simulations for all patterns in the five-hole budget class. The bars on the left are based on the reward from the alternate set of simulations, the bars on the right on the reward from the simulations used in the optimization.

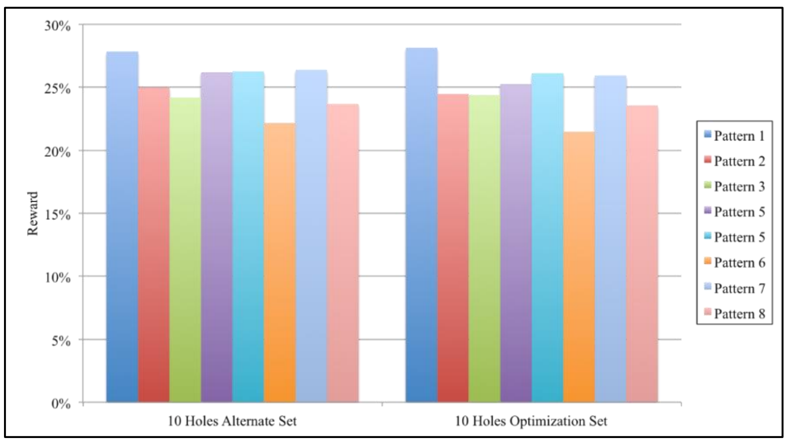

Figure 24: Average reward over all twenty simulations for all patterns in the ten-hole budget class. The bars on the left are based on the reward from the alternate set of simulations, the bars on the right on the reward from the simulations used in the optimization.

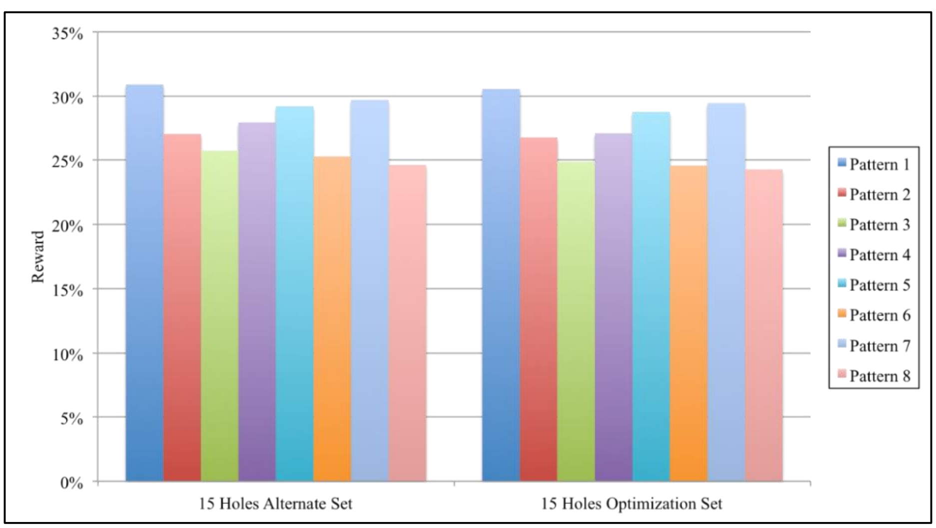

Figure 25: Average reward over all twenty simulations for all patterns in the fifteen-hole budget class. The bars on the left are based on the reward from the alternate set of simulations, the bars on the right on the reward from the simulations used in the optimization.

Figure 26 is the counterpart of Figure 22 and when they are compared, not many differences in general trends are seen. For ease of comparison the results presented in Figure 22 are summarised with the black lines in Figure 26. The black lines are always very close to their counterpart from the tests on the alternate set of simulations. This leads to the same conclusion as made for Figure 22 above.

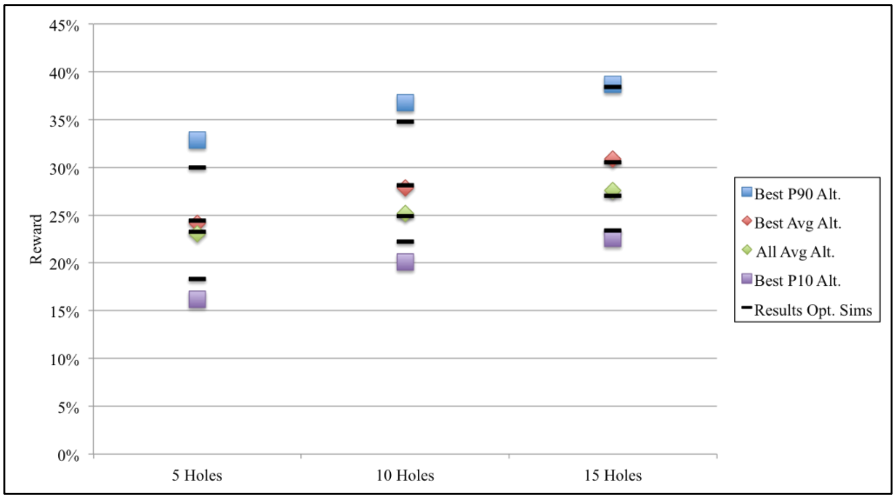

Figure 26: Profile of the best pattern for every budget class versus the all-pattern average for every budget class based on the reward from the alternate set of simulations presented by the squares and diamonds. The black lines represent the same metrics for the reward based on the simulations used in the optimization as shown in Figure 22.

All comparisons made in this section show that the patterns have similar performance on a set of simulations, which are not used in the optimization process as for the simulations used in the optimization. This demonstrates that the proposed optimization method produces robust results that hold true in situations not directly optimized for.

4 Conclusions and future work

The case study presented above shows that the proposed MAB algorithm works as intended. It is able to select the best pattern within a set of patterns of the same budget and provides a decision tool to select the best pattern between sets of a different budget. It is also demonstrated how the method can successfully quantify the influence of geological uncertainty on the performance of an infill drilling pattern.

Additional to the performance of the proposed infill drilling optimization algorithm it is demonstrated that stockpiles can be effectively simulated and that it can be required to look at their local scale variability as they can be much more variable than is expected.

For future work it would be interesting to look into the updating the original simulation after every play according the infill drilling data instead of performing a re-simulation at every iteration. Two methods that are suitable for this are conditional simulation by successive residuals (Jewbali and Dimitrakopoulos, 2011) and ensemble Kalman filters (Benndorf, 2015). Another improvement to the algorithm could be to let it choose the locations of the additional drill holes independently rather than to constrain it to a set of predefined patterns. A different focus for future research is the application of the MAB framework to other applications in the mining industry.

Appendix

The figures in this appendix are a comparison between the data and the simulations at point support level for the histograms, variograms and cross-variograms for all elements. The thicker red line in all figures represents the data from the exploration drill holes. The twenty thinner black lines represent the results from the simulations. Figure 27 to Figure 30 show the histograms, Figure 31 to Figure 34 show the variograms and Figure 35 to Figure 40 show the cross-variograms. The variography is only presented in one direction, azimuth 107.5° and dip 0°. For the simulation the perpendicular direction (azimuth 17.5° and dip 0°) is also used, these graphs are not shown in this appendix. The lag in all variograms and cross-variograms is 40 feet.

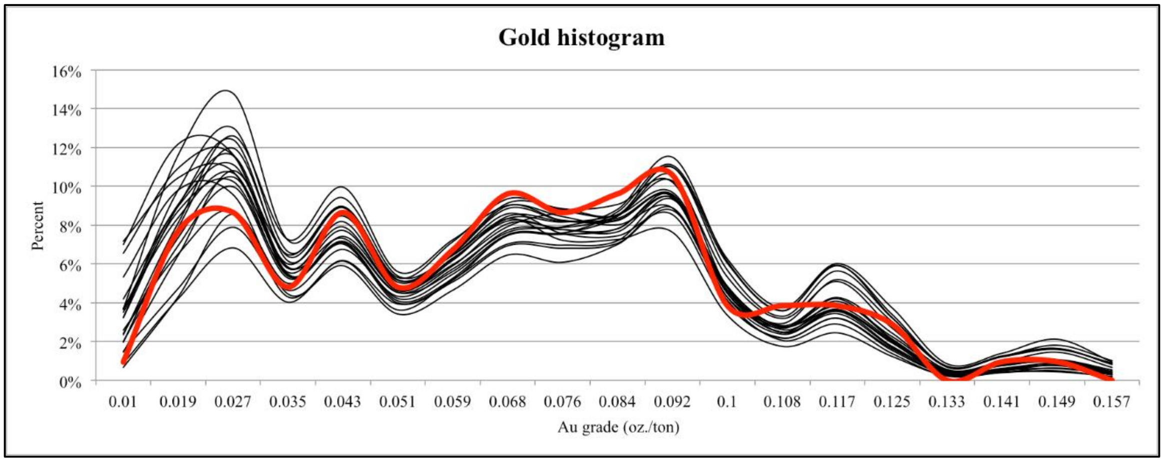

Figure 27: Comparison of the histograms of the drill hole data for gold (red line) and the twenty simulations at point support level (black lines).

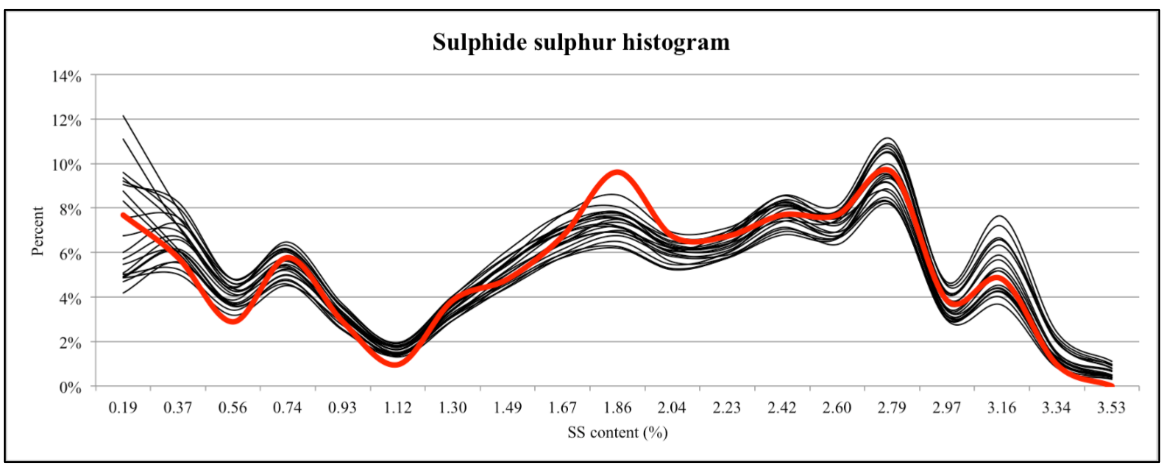

Figure 28: Comparison of the histograms of the drill hole data for sulphide sulphur (red line) and the twenty simulations at point support level (black lines).

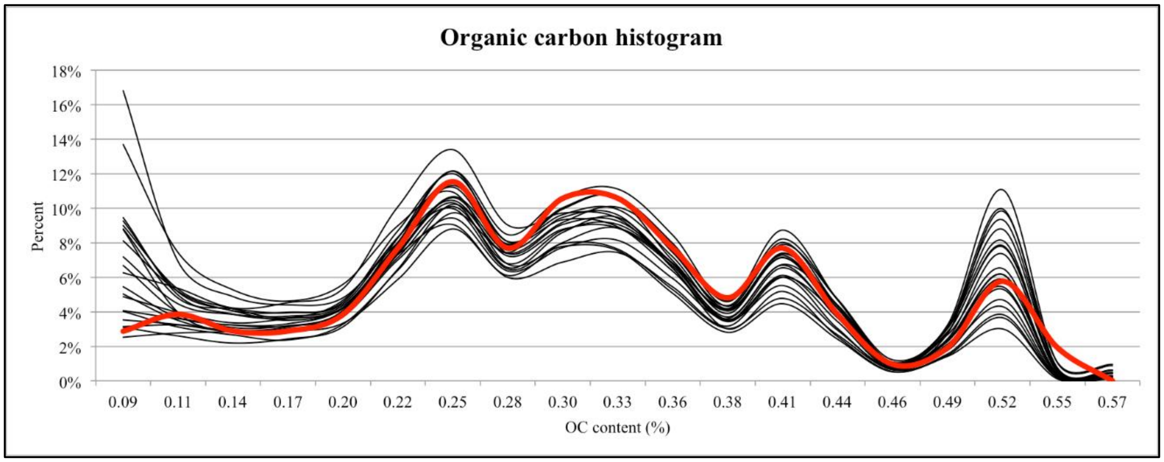

Figure 29: Comparison of the histograms of the drill hole data for organic carbon (red line) and the twenty simulations at point support level (black lines).

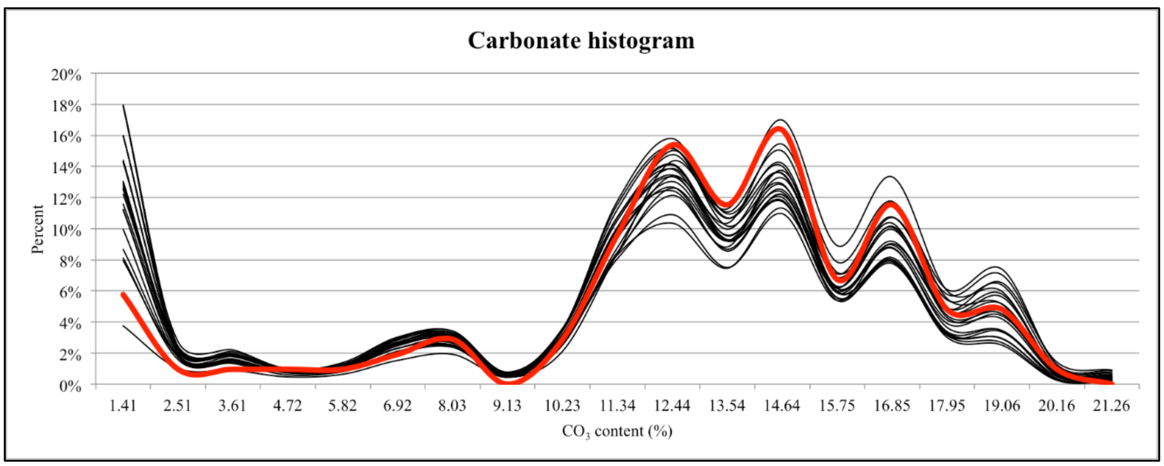

Figure 30: Comparison of the histograms of the drill hole data for carbonate (red line) and the twenty simulations at point support level (black lines).

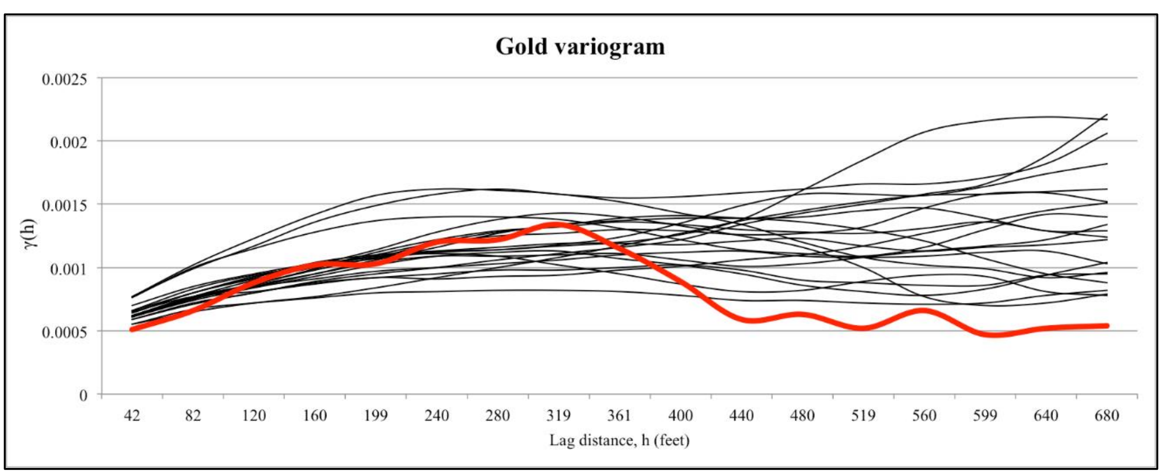

Figure 31: Comparison of the variograms of the drill hole data for gold (red line) and the twenty simulations at point support level (black lines). Azimuth 107.5°, dip 0° and lag 40 feet.

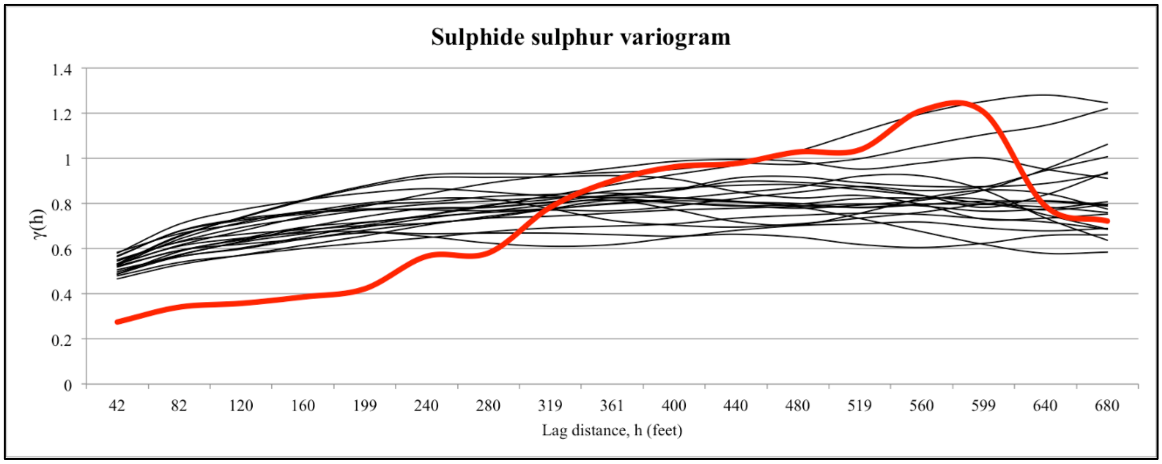

Figure 32: Comparison of the variograms of the drill hole data for sulphide sulphur (red line) and the twenty simulations at point support level (black lines). Azimuth 107.5°, dip 0° and lag 40 feet.

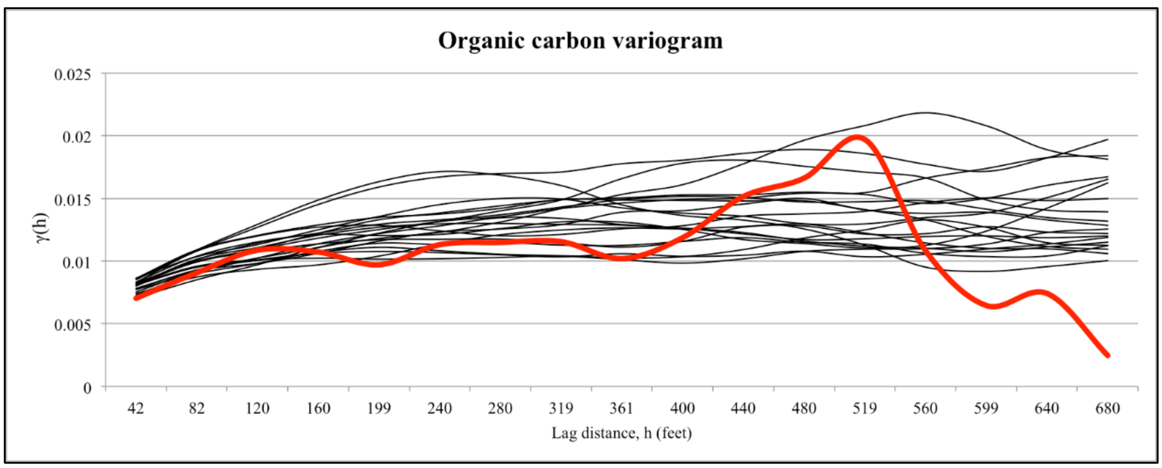

Figure 33: Comparison of the variograms of the drill hole data for organic carbon (red line) and the twenty simulations at point support level (black lines). Azimuth 107.5°, dip 0° and lag 40 feet.

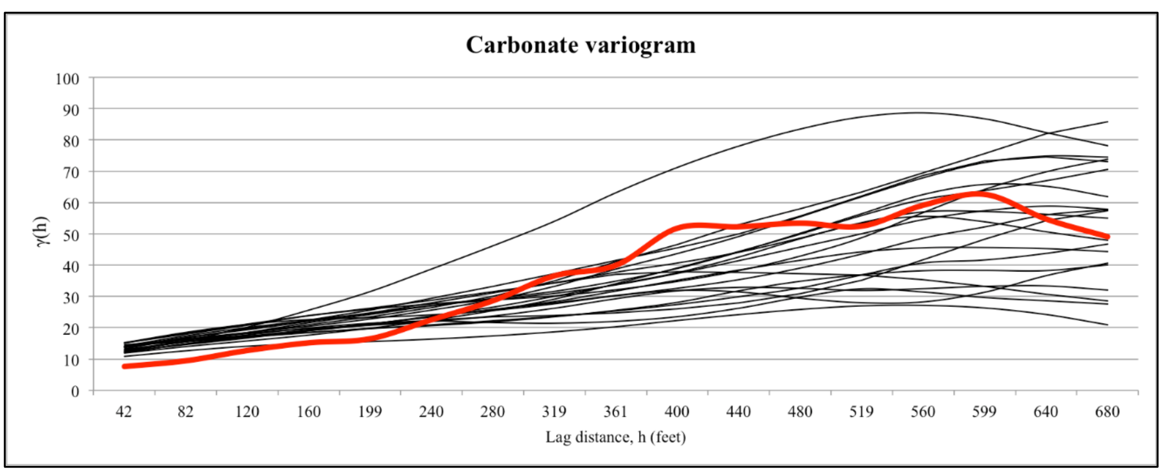

Figure 34: Comparison of the variograms of the drill hole data for carbonate (red line) and the twenty simulations at point support level (black lines). Azimuth 107.5°, dip 0° and lag 40 feet.

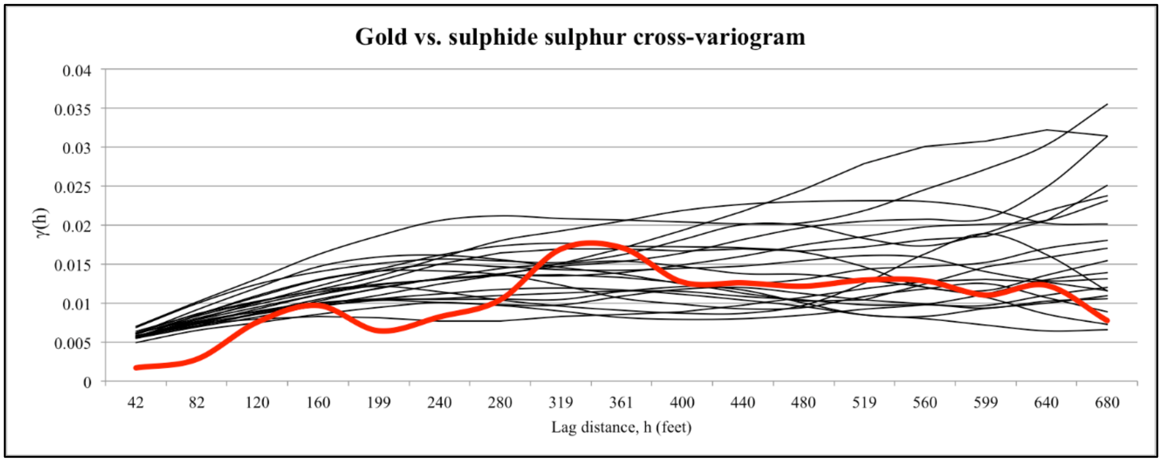

Figure 35: Comparison of the cross-variograms for gold and sulphide sulphur of the drill hole data (red line) and the twenty simulations at point support level (black lines). Azimuth 107.5°, dip 0° and lag 40 feet.

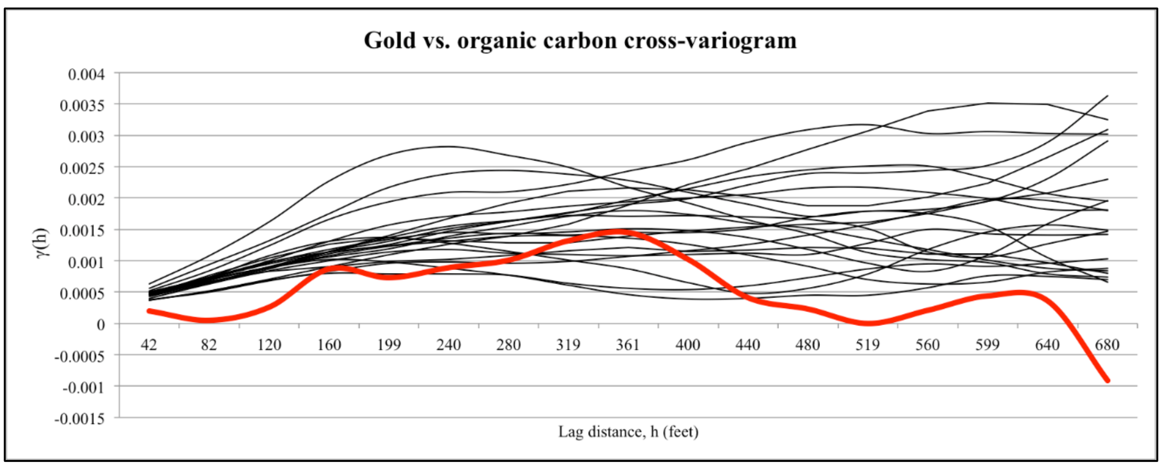

Figure 36: Comparison of the cross-variograms for gold and organic carbon of the drill hole data (red line) and the twenty simulations at point support level (black lines). Azimuth 107.5°, dip 0° and lag 40 feet.

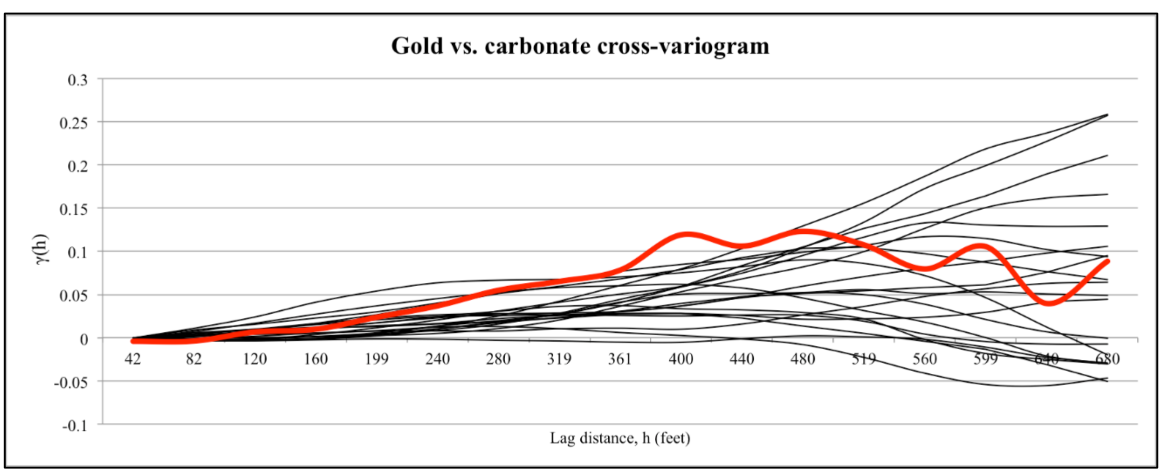

Figure 37: Comparison of the cross-variograms for gold and carbonate of the drill hole data (red line) and the twenty simulations at point support level (black lines). Azimuth 107.5°, dip 0° and lag 40 feet.

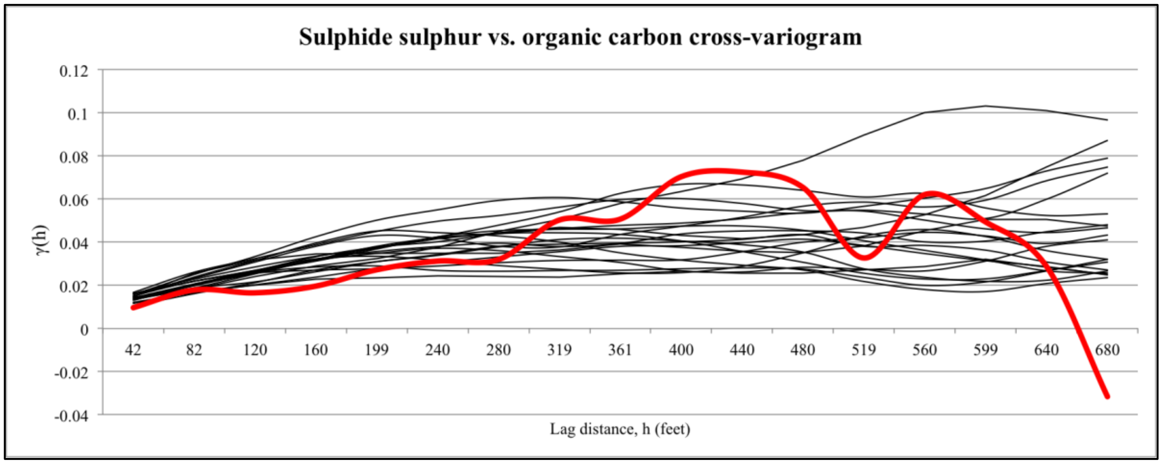

Figure 38: Comparison of the cross-variograms for sulphide sulphur and organic carbon of the drill hole data (red line) and the twenty simulations at point support level (black lines). Azimuth 107.5°, dip 0° and lag 40 feet.

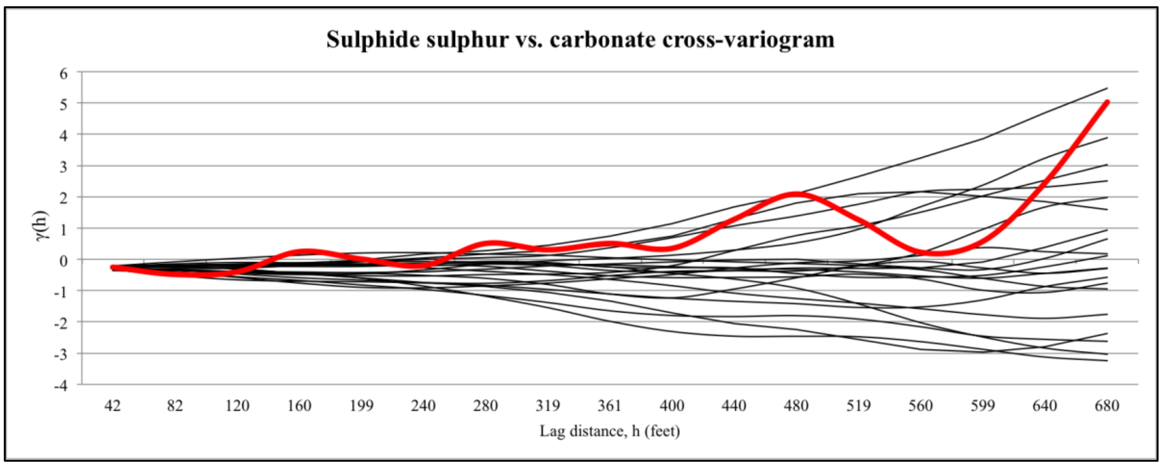

Figure 39: Comparison of the cross-variograms for sulphide sulphur and carbonate of the drill hole data (red line) and the twenty simulations at point support level (black lines). Azimuth 107.5°, dip 0° and lag 40 feet.

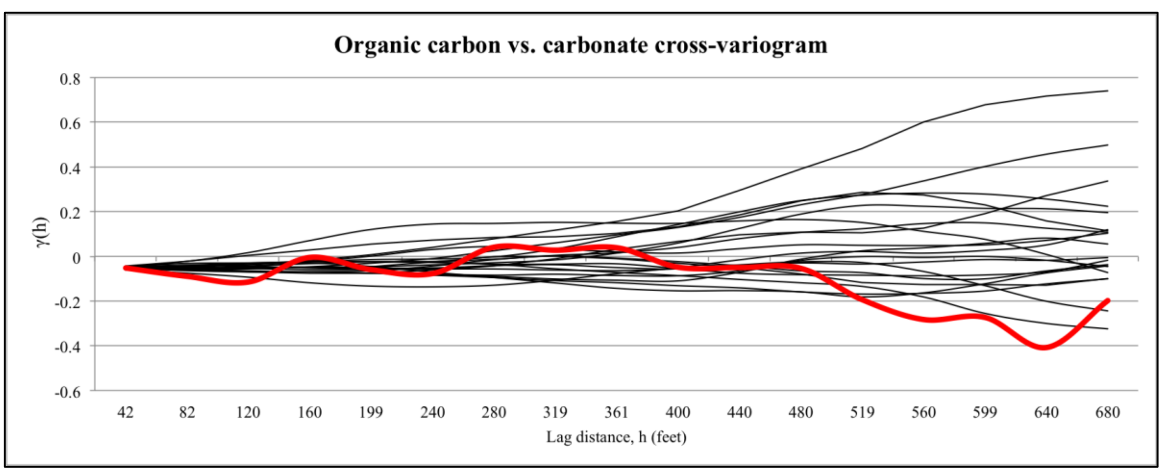

Figure 40: Comparison of the cross-variograms for organic carbon and carbonate of the drill hole data (red line) and the twenty simulations at point support level (black lines). Azimuth 107.5°, dip 0° and lag 40 feet.

References

Agrawal, S., and Goyal, N. (2012). Analysis of Thompson sampling for the multi-armed bandit problem. Journal of Machine Learning Research, 23(Workshop and Conference Proceedings).

Barnes, R. J. (1989). Sample Design for Geologic Site Characterization. In M. Armstrong (Ed.), Geostatistics: Springer Netherlands, 4, 809–822.

Benndorf, J. (2015). Making Use of Online Production Data: Sequential Updating of Mineral Resource Models. Mathematical Geosciences, 47(5), 547–563.

Boucher, A., and Dimitrakopoulos, R. (2009). Block Simulation of Multiple Correlated Variables. Mathematical Geosciences, 41(2), 215–237.

Boucher, A., Dimitrakopoulos, R., and Vargas-Guzmán, J. A. (2005). Joint Simulations, Optimal Drillhole Spacing and the Role of the Stockpile. In O. Leuangthong and C. Deutsch (Eds.), Geostatistics Banff 2004: Springer Netherlands, 14, 35–44.

Burgess, T. M., and Webster, R. (1980). Optimal Interpolation and Isarithmic Mapping of Soil Properties. Journal of Soil Science, 31(2), 333–341.

Chapelle, O., and Li, L. (2011). An empirical evaluation of thompson sampling. Paper presented at the Advances in neural information processing systems.

Chorn, L. G., and Carr, P. P. (1997). The Value of Purchasing Information to Reduce Risk in Capital Investment Projects.

Desbarats, A. J., and Dimitrakopoulos, R. (2000). Geostatistical Simulation of Regionalized Pore-Size Distributions Using Min/Max Autocorrelation Factors. Mathematical Geology, 32(8), 919–942.

Diehl, P., and David, M. (1982). Classification of Ore Reserves/Resources based on Geostatistical Methods. CIM Bulletin, 75(838), 127–136.

Englund, E., and Heravi, N. (1993). Conditional Simulation: Practical Application for Sampling Design Optimization. In A. Soares (Ed.), Geostatistics Tróia '92: Springer Netherlands, 5, 613–624.

Gershon, M., Allen, L. E., and Manley, G. (1988). Application of a new approach for drillholes location optimization. International Journal of Surface Mining, Reclamation and Environment, 27–31.

Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. Journal of the Royal Statistical Society. Series B (Methodological), 148–177.

Goovaerts, P. (1997). Geostatistics for Natural Resources Evaluation. New York: Oxford University Press, Inc.

Goria, S., Armstrong, M., and Galli, A. (2001). Quantifying the impact of additional drilling on an openpit gold project. Paper presented at the 2001 Annual Conference of the IAMG, Cancun, Mexico.

Jewbali, A., and Dimitrakopoulos, R. (2011). Implementation of conditional simulation by successive residuals. Computers & Geosciences, 37(2), 129–142.

Mahajan, A., and Teneketzis, D. (2008). Multi-Armed Bandit Problems. In A. Hero, III, D. Castañón, D. Cochran, and K. Kastella (Eds.), Foundations and Applications of Sensor Management: Springer US, 121–151.

May, B. C., Korda, N., Lee, A., and Leslie, D. S. (2012). Optimistic Bayesian sampling in contextual-bandit problems. The Journal of Machine Learning Research, 13(1), 2069–2106.

May, B. C., and Leslie, D. S. (2011). Simulation Studies in Optimistic Bayesian Sampling: Statistics Group, Department of Mathematics, University of Bristol.

Menabde, M., Froyland, G., Stone, P., and Yeates, G. (2004). Mining schedule optimisation for conditionally simulated orebodies. Paper presented at the Proceedings of the international symposium on orebody modelling and strategic mine planning: uncertainty and risk management.

Olea, R. (1975). Optimum Mapping Techniques Using Regionalized Variable Theory. Lawrence: Kansas Geological Survey.

Oostenbrink, J. B., Al, M. J., Oppe, M., and Rutten-van Mölken, M. P. M. H. (2008). Expected Value of Perfect Information: An Empirical Example of Reducing Decision Uncertainty by Conducting Additional Research. Value in Health, 11(7), 1070–1080.

Petit, G., De Boissezon, H., Langlais, V., Rumbach, G., Khairuldin, A., Oppeneau, T., and Fiet, N. (2012). Application of Stochastic Simulations and Quantifying Uncertainties in the Drilling of Roll Front Uranium Deposits. In P. Abrahamsen, R. Hauge, and O. Kolbjørnsen (Eds.), Geostatistics Oslo 2012: Springer Netherlands, 17, 321–332.

Prange, M., Bailey, W. J., Couët, B., Djikpesse, H., Armstrong, M., Galli, A., and Wilkinson, D. (2008). Valuing Future Information Under Uncertainty Using Polynomial Chaos. Decision Analysis, 5(3), 140–156.

Ravenscroft, P. J. (1992). Risk analysis for mine scheduling by conditional simulation Transaction of the institution of mining and metallurgy Section A, 101, A104–A108.

Schlee, E. (1991). The value of perfect information in nonlinear utility theory. Theory and Decision, 30(2), 127–131.

Scott, S. L. (2010). A modern Bayesian look at the multi-armed bandit. Applied Stochastic Models in Business and Industry, 26(6), 639–658.

Switzer, P., and Green, A. A. (1984). Min/Max Autocorrelation Factors for Multivariate Spatial Imagery: Stanford Univeristy, Department of Statistics.

Thompson, W. R. (1933). On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. Biometrika, 25(3/4), 285–294.