Computers and Geosciences 112 (2018) 101–111

Contents lists available at ScienceDirect

Computers and Geosciences

journal homepage: www.elsevier.com/locate/cageo

Research paper

Phase annealing for the conditional simulation of spatial random fields

S. Hörning[a,*], A. Bárdossy[b]

[a] School of Earth and Environmental Sciences, The University of Queensland, Brisbane, Australia
[b] Institute of Modeling Hydraulic and Environmental Systems, Department of Hydrology and Geohydrology, University Stuttgart, Stuttgart, Germany

ARTICLE INFO

Keywords:
Geostatistical simulation
Simulated annealing
Phase spectrum
Singularities

ABSTRACT

Simulated annealing (SA) is a popular geostatistical simulation method as it provides great flexibility. In this paper possible problems of conditioning its realizations are discussed. A statistical test to recognize whether the observations are well embedded in their simulated neighborhood or not is developed. A new simulated annealing method, phase annealing (PA), is presented which makes it possible to avoid poor embedding of observations. PA is based on the Fourier representation of the spatial field. Instead of the individual pixel values, phases corresponding to different Fourier components are modified (i.e. shifted) in order to match prescribed statistics. The method treats neighborhoods together and thus avoids singularities at observation locations. It is faster than SA and can be used for the simulation of high resolution fields. Examples demonstrate the applicability of the method.

1. Introduction

Observed spatial fields are often the results of complicated physical, chemical and/or biological processes which are frequently not known in full detail. Results of these processes such as precipitation accumulations or groundwater levels are usually observed at a few selected locations only. Different geostatistical interpolation methods can be used to estimate values at unsampled locations. Interpolation, however, reduces the variance which leads to unrealistic, smooth fields. Using these smooth fields for subsequent non-linear modeling leads to serious biases. Instead geostatistical simulations are widely used to generate realistic fields. These reproduce measurement values, but also reflect the observed spatial variability, which makes them more appropriate for subsequent non-linear modeling and useful for uncertainty and risk assessment.

Many different geostatistical methods for the simulation in 2D and 3D are available, including methods like sequential Gaussian simulations, LU decomposition based methods and Turning bands simulations. A description and code for these methods can be found in Deutsch and Journel (1998) or Lantuejoul (2002). These methods concentrate on reproducing the observed data and the spatial variability expressed with the variogram.

An interesting technique amongst the existing methods is simulated annealing for spatial random fields (note that this specific type of simulated annealing will be abbreviated SA in the following) (Deutsch, 1992). SA is a very flexible approach as it can generate fields with desired properties without the explicit specification of a theoretical model. One of the limitations of SA is the high computational demand, specifically for conditional simulation of large fields. Some approaches exist to reduce the computational run time of SA. For example in Peredo and Ortiz (2011), the authors parallelize the algorithm. Another less known problem with SA is the often poor conditioning. Due to the nature of the algorithm observations are often not well embedded, i.e., they do not always fit into their simulated neighborhood causing "singularities".

Other interesting simulation methods were developed in the field of time series analysis. In their seminal paper, Theiler et al. (1992) described the methodology of phase randomization for investigating non-linear properties of time series. The method can easily be extended to higher dimensions (Shinozuka and Deodatis, 1991, 1996) and is frequently used, for example for the simulation of textures (Galerne et al., 2011). The link between texture synthesis and multipoint geostatistics was studied in Mariethoz and Lefebvre (2014).

The purpose of this paper is to develop a method to detect problems with conditioning and to present a new simulation methodology which combines the well-known method of phase randomization with simulated annealing to obtain conditional realizations of fields with flexible spatial properties.

After the general introduction, a brief introduction to SA is given. Problems of singularities occurring in conditional fields generated with SA are discussed. A statistical method which can be used to detect singularities at observation locations is presented. The third section

* Corresponding author.
E-mail addresses: s.hoerning@uq.edu.au (S. Hörning), a.bardossy@iws.uni-stuttgart.de (A. Bárdossy).

https://doi.org/10.1016/j.cageo.2017.12.008
Received 19 April 2017; Received in revised form 11 December 2017; Accepted 15 December 2017
Available online 21 December 2017

describes the phase randomization methodology and its modification, phase annealing to generate conditional random fields with arbitrary marginal distributions. Section four presents applications of the methodology. Different examples are given which demonstrate the flexibility of the algorithm. Fields with properties described with the help of spatial copulas are also presented. In section five, a refinement of the algorithm enabling very high dimensional simulations is described. The paper ends with discussions and conclusions.

2. Simulated annealing

Kirkpatrick et al. (1983) first introduced simulated annealing as a probabilistic approach with the goal of finding the global minimum of a given objective function. Simulated annealing represents an extension to the well-known Metropolis algorithm (Metropolis et al., 1953) which has been developed to simulate molecule behavior.

The first spatial application of simulated annealing can be found in Geman and Geman (1984), where the authors applied the approach to the restoration of degraded digital images. Deutsch (1992); Deutsch and Journel (1998) then adopted simulated annealing to the simulation of spatial random fields (abbreviated SA in the following) with the objective of reproducing observed spatial properties. In essence, the algorithm works by perturbing one or more nodes at a time, starting with an initial model (usually a random spatial distribution of values with the observations being fixed at their corresponding locations). After every perturbation, the mismatch between the current simulated statistics and the observed (target) statistics is quantified according to a predefined objective function. A perturbation is kept if it reduces the mismatch; it is rejected with a certain probability if it increases the mismatch. The acceptance probability of unfavorable perturbations is given by the predefined annealing schedule. The annealing schedule is defined by an initial temperature and a corresponding procedure to lower that temperature as the simulation progresses. A detailed description of the SA algorithm can be found in Deutsch (1992). Further theoretical background of the algorithm can be found in Hegstad et al. (1994) where its relationship to the Hastings algorithm is also discussed.

3. Singularities

Conditional simulations are used to restrict the possible realizations by generating fields which honor observed values at the observation locations. Observed values should be embedded into their neighborhood so that they are indistinguishable from points of the simulated field. If observations are visually identifiable in a simulated realization, then conditioning was not appropriate. In general, values differing significantly from their neighbors will subsequently be called singularities. SA unfortunately can produce realizations with singularities at observation locations. The reason for this is that objective functions used for SA are defined over the whole field, and when few observations are available relative to the dimensions of the field, the objective functions may thus not penalize local unusual behavior at observation locations.

The following example illustrates the problem. Consider two independently generated realizations, both have the same variogram and an exponential marginal distribution. The fields under consideration are defined on a regular 128×128 grid. $N = 95$ random locations are selected and considered as observation points. The values of the first field at those observation points are replaced by the values of the second field at those N observation points. Then the variogram of the field with the changed values is calculated again. Fig. 1 shows parts of the original and the modified field. Both fields exhibit pixels which deviate from their neighbors, but the modified field shows a few more deviations. Looking

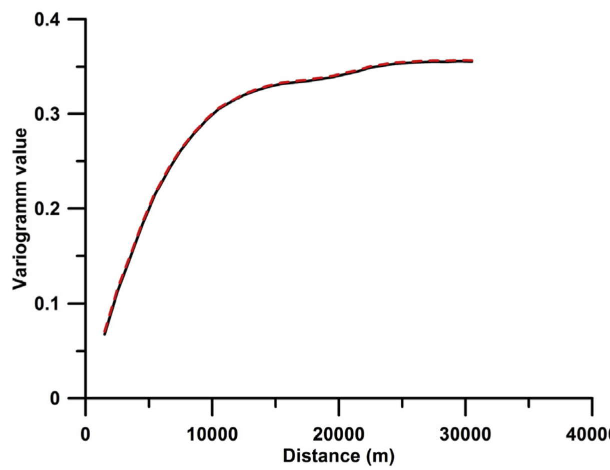

Fig. 2. Variogram of the simulated data (solid line) and of the field modified at $N = 95$ locations.

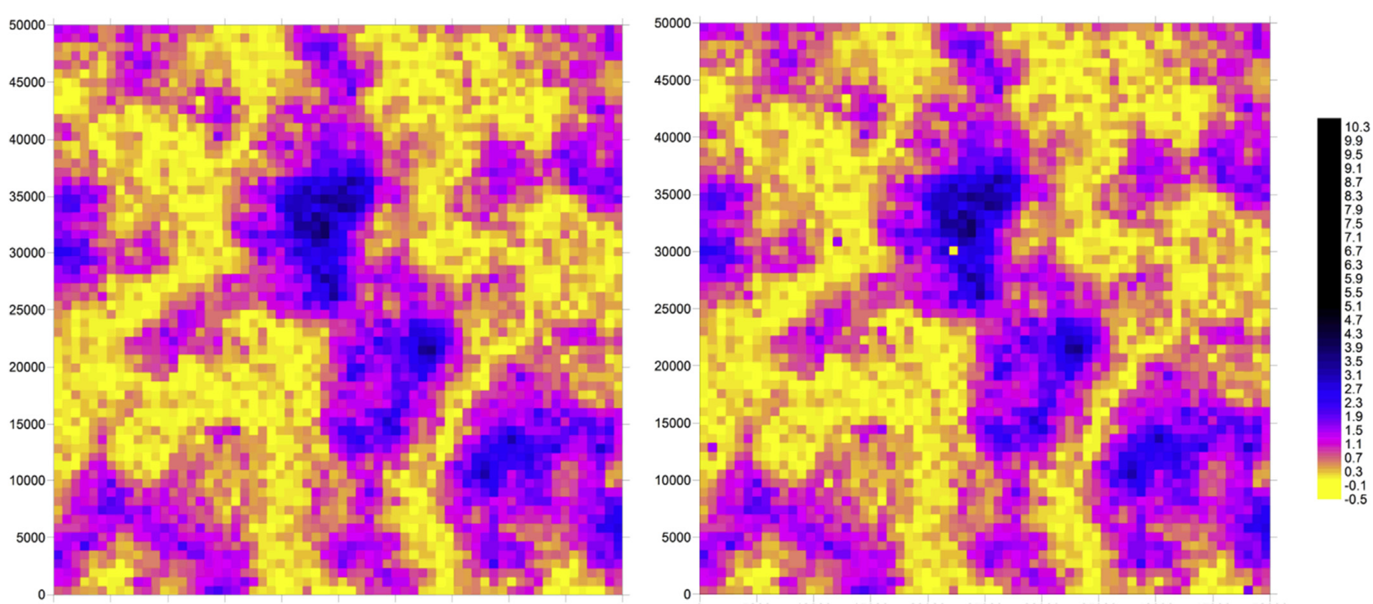

Fig. 1. Simulated field (left) and simulated field with exchanged values at observation points (right).