

Document downloaded from:

<http://hdl.handle.net/10251/74232>

This paper must be cited as:

Gómez-Hernández, JJ. (2015). Multiple-point Geostatistics: Stochastic Modeling with Training Images by G. Mariethoz and J. Caers, December 2014. Wiley-Blackwell, United States (2014). Computers and Geosciences. 83:231-235. doi:10.1016/j.cageo.2015.06.010.



The final publication is available at

<http://dx.doi.org/10.1016/j.cageo.2015.06.010>

Copyright Elsevier

Additional Information

This is a book review

Elsevier Editorial System(tm) for Computers & Geosciences
Manuscript Draft

Manuscript Number:

Title: Multiple-point geostatistics: Stochastic imaging with training images by G. Mariethoz and J. Caers

Article Type: Book or Software Review

Keywords: This is a book review, I do not think there is a need for keywords

Corresponding Author: Prof. J. Jaime Gómez-Hernández,

Corresponding Author's Institution: Universitat Politècnica de València,

First Author: J. Jaime Gómez-Hernández

Order of Authors: J. Jaime Gómez-Hernández

1 Multiple-point geostatistics: Stochastic imaging with training images by G.
2 Mariethoz and J. Caers

3

4 J. Jaime Gómez-Hernández
5 Research Institute of Water and Environmental Engineering
6 Universitat Politècnica de València, Spain
7 Tel.: +34 96 387 9614, email: jgomez@upv.es

8

9 I was especially pleased when I was asked to review this book, because I was a
10 graduate student at Stanford when multiple-point geostatistics (MPS) was born
11 back in the late 80's of the past century. Indicator geostatistics had been a major
12 breakthrough, but André Journel felt the need to go beyond the bivariate
13 characterization offered by indicator variograms. He came up with the idea of
14 computing local conditional probability distributions for a binary variable
15 accounting for more than two-point variograms by scanning for high-order
16 moments through a training image. Although André is never credited when
17 referring to the history of MPS, he was truly its father. The initial idea became the
18 subject of the master thesis of Felipe Guardiano, with Mohan Srivastava in charge
19 of all the implementation subtleties. That was 25 years ago. Much action has
20 taken place since then, multiple-point geostatistics has matured into a discipline
21 of its own, and today we welcome a comprehensive book on the subject that has
22 been put together by Gregoire Mariethoz and Jef Caers.

23

24 The reader is confronted with a 378-page book and an accompanying website
25 with complementary material. A book that is oriented to the practical aspects of
26 MPS, not to create any new theory, a book about practice and the solution of real
27 problems. Beware theoreticians, welcome engineers! And, although the
28 generality of the MPS framework is recognized, the book centers in applications
29 to the physical sciences, not to other areas in which spatial statistics are also
30 used, and for which MPS could be applicable.

31

32 The book is divided in three major parts: concepts, methods and applications. In
33 the first part a case study, unrelated to the physical sciences that will underlie
34 the rest of the book, is used to motivate the reader on why the need for MPS, the
35 fundamental concepts are then introduced, and the case study is revisited to
36 show, indeed, the need for MPS when dealing with complex spatial patterns. In
37 the second part a thorough review of all methods, algorithms and tools needed to
38 perform MPS in a variety of contexts are described. And in the third part, three
39 real applications are presented, in the fields of petroleum engineering, mining
40 engineering and climate modeling.

41

42 The book is written with a very direct prose, trying to explain the practical
43 implications of the different decisions that a modeler has to take in his quest for
44 producing a representation of a partially known reality, whether you call it
45 estimation or simulation. In this respect, the book is truffled with tips and
46 recommendations, as well as lengthy explanations why many modeling decisions
47 are taken. The book cannot be read on its own, the reader must have previous
48 knowledge about geostatistics, uncertainty analysis and stochastic simulation to
49 fully understand the benefits of MPS. The authors make numerous reference of

50 the standard geostatistical approach based on two-point statistics; therefore,
51 previous knowledge on classic geostatistics is a must.

52

53 Part I, on concepts, sets the stage.

54

55 Chapter I.1, with practicality in mind, poses a problem (somehow unrelated to
56 the geosciences, not statistical at all), in which the goal is to estimate, from
57 limited information, the topography of a certain area through which hikers will
58 cross to provide them with sufficient (but not in excess) food supplies. This
59 problem will be revisited at the end of this part to support the need to use MPS
60 to model complex spatial patterns.

61

62 Chapter I.2 introduces standard estimation based on random function theory,
63 that is, kriging. This chapter is representative of the type of discussions that are
64 included throughout in the book. It begins with a very lengthy and appropriate
65 discussion about the assumptions of stationarity and ergodicity, the fact that
66 there is a unique reality, and that to construct a model many assumptions must
67 be made that cannot be corroborated with data (the authors make clear that
68 stationarity and ergodicity are assumptions, never hypothesis). The chapter goes
69 on presenting the concepts of stationarity, unbiasedness, loss functions, local
70 stationarity, drifts, trends; it develops the equations of kriging, with some
71 discussion about how to infer the covariances and the variograms from data, and
72 the dangers of estimating these from the residuals obtained after a least-square
73 fitting of a trend. The chapter ends with a warning that no statistical model is
74 good for universal kriging.

75

76 Chapter I.3 builds on the conclusion from the previous chapter to provide a
77 framework to perform universal kriging without resorting to a random function
78 model, much as Journel did with his “deterministic geostatistics”. The concept of
79 training image (TI), although mentioned several times before, is presented for
80 the first time in a formal way: a TI is an image, which provides an analog to the
81 field of study, and which is “deemed representative” of the spatial variability
82 within the study area. Now, the big model decision is not stationarity or
83 ergodicity, but rather to decide whether a TI is deemed representative.

84

85 One of the problems when presenting a new framework is nomenclature. This
86 chapter introduces a large number of new elements. The authors have made a
87 tremendous effort to be consistent throughout the entire book, and they have
88 succeeded; however, they have introduced so many new terms and concepts that
89 the reader may become overwhelmed with all the new symbols, acronyms and
90 abbreviations. This wealth of new terminology may preclude some readers to
91 use the book to address a specific problem that can be solved with the tools
92 presented in a specific chapter, because they may stumble with symbols and
93 acronyms for which no definition is readily available. The authors are urged to
94 include, in the next printing of the book, a nomenclature section with all the
95 symbols used in the book and their meaning, and possibly, the page in which
96 they are first introduced.

97

98 The authors continue presenting a version of ordinary kriging with TI, in which
99 variogram estimation/evaluation is replaced by the calculation of sums of
100 products from the TI. Then, they present the simple kriging version, and move on
101 to analyze the impact of the choice of search neighborhood, the size of the
102 training image, the nature of the training image and other aspects. An interesting
103 result is that since the products used are two-point products, two TIs with
104 similar histograms and variograms but with very different textures will result in
105 very similar estimates. Then, the authors move into the introduction of non-
106 stationary models, stating that, for the model to be non-stationary, the TIs must
107 be locally representative; one way to achieve this is by the introduction of
108 auxiliary variables. Some examples are given, but unlike for the stationary case,
109 the description is not as exhaustive and the reader is referred to later chapters to
110 fully understand how to work with non-stationary MPS models.

111

112 It is important to note that in universal kriging with TI, conditioning is separated
113 from model specification (that is, data are not used to build the model, in this
114 case, the training image) and therefore there could be some inconsistency
115 between data and training image if lots of data are present. An issue that will be
116 discussed in depth later in the book.

117

118 Chapter 1.4 deals with stochastic simulation based on random function theory. It
119 starts by explaining in which context one would need stochastic realizations
120 instead of a single estimation map, and then carry on to present several
121 stochastic simulation techniques: first, sequential Gaussian simulation, then
122 direct sequential simulation, and finally pluriGaussian simulation.

123

124 Chapter 1.5 presents stochastic simulation without random function theory. It
125 starts by presenting the direct sampling approach in the context of producing a
126 paragraph that may look like it is written in French but that, in fact, is a random
127 realization of words and punctuation signs that resemble true French. The new
128 paragraph simply replicates the patterns observed in a text by Flaubert. Then, it
129 shows how direct sampling would perform in the Walker Lake data set. These
130 two demonstrations show how stochastic simulation can be performed using a
131 TI instead of a formal random function model. The algorithm itself will be
132 explained in detail later in the book.

133

134 Next, the extended normal equation is formulated in detail. This equation is the
135 foundation of all MPS algorithms. The extended normal equation, at the end,
136 reduces to counting how many patterns are found in the TI close to a given data
137 event (a specific spatial pattern of point values). From an implementation point
138 of view, there are two ways to address this counting, one by scanning the image
139 each time that a count is needed, and another by precomputing all potential
140 counts a priori. The first approach needs CPU-time; the second one needs
141 random access memory (RAM).

142

143 There is an interesting, albeit anecdotic, discussion of the bias introduced in any
144 sequential simulation by the fact that conditioning data are always visited first in
145 the simulation path, before the chapter ends with an excursion into the computer

146 graphics literature to review some texture generating algorithms that can also be
147 applied in MPS: patchwork and image quilting.

148

149 Chapter I.6 closes the first part of the book by returning to the case study posed
150 in the first chapter. Conditional realizations of the study area are generated using
151 three different approaches: multiGaussian simulation, direct sampling and image
152 quilting. The authors conclude that, in practice, no model can be validated since
153 there is never a “true” reference with which to compare the results, even though,
154 in the case at hand, the multiGaussian model is poorer than the other two
155 models. They also argue that the underlying multivariate random function model
156 is not primordial, but rather the ensemble of realizations, and therefore those
157 realizations should have the characteristics we wish in terms of continuity,
158 texture, etc. And finally, they remind that when the realizations have to be fed to
159 a non-linear transfer function, the results will most likely depend on the higher-
160 order moments of those realizations, and, consequently, the modeler should use
161 those methods that best capture those moments.

162

163 Part II, on methods, provides the tools.

164

165 Chapter II.1 starts with a clear statement: it is not about theoretical
166 developments, but about how to implement ways to borrow information from
167 training images; because, at the end, the user needs a computer code, which is
168 made of algorithms, and these algorithms are built based on decisions which may
169 give rise to different interpretations of the same concept.

170

171 In this introduction the concept of TI is reinforced, and some of its potential
172 pitfalls already pointed out, such as the possible incoherence between image and
173 conditioning data, or the need that TIs must constrain the images only up to a
174 certain point (otherwise the realizations will be verbatim replicas of the TI or of
175 portions of it).

176

177 Chapter II.2 is the longest chapter in the book; it lies out all the elementary
178 blocks necessary to build an MPS algorithm. The chapter starts with a discussion
179 on how data, for one or several variables, can be provided to the modeler,
180 differentiating between scattered data and gridded data. Then, it continues
181 discussing the concept of neighborhood, a key concept in MPS since it does not
182 only fixes the spatial extent within which data will be searched, but also
183 determines the order of the spatial statistical to be considered. (Contrary to
184 traditional geostatistics, there is no equivalence to a unique or global
185 neighborhood in MPS.) Next concept discussed is that of a “data event”, which is
186 the combination of the locations of a neighborhood and the values at those
187 locations (a concept that is valid for one or several variables). There is an
188 important remark at this point in the book: concepts in MPS are easy to
189 understand and intuitive, but they could be very difficult to implement,
190 especially for the multivariate cases.

191

192 The application of the extended normal equation requires finding and counting
193 specific data events from a TI. The book now discusses the two possible
194 approaches to extract this information from the TI: the raw approach, in which

195 the TI is kept in memory and each time a data event needs to be evaluated, the TI
196 is searched; and the tree-storage approach, in which the TI is thoroughly
197 scanned before the beginning of the simulation and all potential data events are
198 identified and counted, storing, at the end, a search tree with all this information.
199 Raw storage relies on CPU power, whereas tree search relies on RAM availability.
200

201 The notion of convolution is introduced to formalize the computation of the local
202 conditional probability distribution given a data event in the case of the raw
203 storage approach, with a natural extension to continuous variables. Next, the tree
204 storage is described in detail; tree storage is especially suited for categorical
205 variables, although the authors also point out to some extension, based on the
206 use of prototypes, for its application to continuous variables. List storage is
207 presented as a less RAM demanding alternative.
208

209 This part of the chapter ends with a discussion on how to retrieve/store
210 information from a TI when simulating by patches. Very detailed information is
211 given on how to proceed in this case: the need for clustering similar data events,
212 how to compute similarities using filters, how to assign a score to each pattern,
213 how to classify the patterns using multidimensional scaling or how to extract the
214 cluster from this scaling. At the end, a few pattern clusters are identified, each of
215 which is represented by a prototype, and which can be used to build cluster-
216 histogram of patterns.
217

218 An alternative approach to store patterns would be using a parametric approach,
219 and the next part of this chapter goes in detail in how high-order statistics can be
220 computed from cumulants, derived from a TI, and how these cumulants can be
221 used to computed any local conditional probability needed.
222

223 The next section of this chapter enters into the discussion of distances both for
224 categorical and continuous variables, a key element to decide how similar two
225 data events are. The Manhattan, Euclidean, and Hausdorff distances are
226 presented. The possibility of comparing data events after translating or
227 normalizing their values, or up to some transformation of their node coordinates
228 is also discussed, together with the discussion of other possible transformations
229 for specific data types. Finally, several distances for the comparison of two
230 distributions are mentioned.
231

232 Next, the issue of how the path visiting the nodes of a realization in the context of
233 sequential simulation should be constructed is analyzed, with two main
234 alternatives, to use a random path or a unilateral path. Each one has its
235 advantages and disadvantages, which are discussed, together with variants. An
236 important point brought up next is the need to use multiple nested grids in order
237 to capture the heterogeneity patterns at all scales.
238

239 Conditioning to data is next. There is a distinction between hard data and other
240 types. Hard data can be assigned to the grid nodes, although there may be an
241 issue of support. Global proportions can be enforced using a servo-system
242 approach. When indirect data can be translated onto a conditional probability,

243 this probability can be aggregated to the conditional probability given the hard
244 data using different models, such as odd ratios, or the tau model.
245
246 Chapter II.3 is probably the most interesting chapter of the book since it
247 describes, in a very clear way, the different algorithms currently available for
248 MPS simulation. There are many algorithms but they have a lot of similarities. All
249 algorithms are presented in a colored box with three components: the inputs
250 needed, a pseudo-code of the algorithm, and the output produced. The chapter
251 starts with the presentation of what the authors call the archetypal MPS
252 algorithm, and a warning on the need to perform some sensitivity analysis of the
253 algorithmic parameters to evaluate how they influence the quantification of the
254 final uncertainty. (Tools for this latter task will be presented later in the book.)
255
256 The algorithms are divided into pixel based and pattern based. The pixel-based
257 ones are, in chronological order, ENESIM, SNESIM, Direct sampling, and
258 Simulated Annealing. The pattern-based ones, also in chronological order, are
259 SIMPAT, FILTERSIM, Patchwork simulation, CCSIM, and IQ. For all algorithms,
260 some comments on their weaknesses and strengths are given, and then, all
261 algorithms are compiled in a table where, at a glance, the reader can know the
262 ability of each algorithm to deal with categorical and continuous variables, to
263 handle multivariate simulations, to use hard data, soft probabilities or non-
264 stationarity, plus an estimate of its CPU performance.
265
266 This chapter ends with some advice on how to post process the realizations, a
267 task needed when there are inconsistencies between the values in a data event
268 and the TI, inconsistencies that may lead to the simulation of spurious values at
269 some locations.
270
271 Chapter II.4 falls a little bit outside the main theme of the book. It presents the
272 theory behind Markov chain models, and Markov mesh models, and discusses
273 how they could be applied in MPS.
274
275 Chapter II.5 discusses at depth how to use non-stationary models in MPS. It
276 starts with a discussion on what non-stationarity means to follow with a
277 description of different approaches to inject non-stationarity into the
278 realizations, and how to control it.
279
280 First, the authors focus on the use of stationary TIs to generate non-stationary
281 realizations, one way is by using a zonation of the realization, with each zone to
282 be simulated with a different TI; another way is by using data event
283 transformations such as rotations and affinities, specifying locally in the
284 realization the affinity factor or the rotation angle to apply; another alternative is
285 by using soft probabilities (derived, for instance, from a seismic map) which are
286 later aggregated with the probabilities obtained after the solution of the
287 extended normal equations.
288
289 Next, the authors continue describing how to use non-stationary images to
290 generate non-stationary realizations. Although it is possible to use zonation here,
291 too, the best alternative is to incorporate control maps, which are continuous

292 variables defined both in the TI and in the realization, which are, essentially a
293 continuous version of zonation.
294
295 Chapter II.6 deals with multivariate modeling with training images. Conceptually,
296 the problem to solve is the same as for a single variable, but now we need as
297 many training images as there are variables. The authors mention that only raw
298 storage approaches have been used for relatively large problems and continue
299 with a demonstration in a synthetic case. The authors also describe an
300 alternative approach by interpreting multivariate modeling as a filtering
301 problem, whereby each variable has a TI, plus a filtered version of it (for
302 instance, obtained by geophysics), and the filtered version of the realization is
303 also available; the simulation problem is now cast as finding the realizations that
304 are consistent with the filtered information.
305
306 Chapter II.7 deals with the important problem of how to construct a TI. Choosing
307 a TI is the most critical modeling decision in MPS, this is why the authors devote
308 a whole chapter to study alternatives on how to construct them. The methods
309 discussed include object-based methods such as Boolean models; process-based
310 methods, in which the TI is built by modeling the physical processes underlying
311 the phenomena to simulate; process-imitating methods, in which a numerical
312 algorithm is built to imitate the structures observed in nature, much in the same
313 way as the process-based method but with no physical equations solved. Many of
314 the TIs generated in this way are difficult to use because they are clearly non-
315 stationary.
316
317 A critical problem discussed next is how to build 3D TIs from 2D ones. The book
318 discusses how to use probability aggregation to combine several 2D orthogonal
319 TIs, or how to assemble 2D orthogonal data events to obtain 3D events. In some
320 fields, the data density is so high that the data set can be used directly as a TI.
321
322 The most challenging problem is the construction of multivariate TIs, although,
323 the authors point out that, in most cases there is a primary variable, and the TIs
324 for the other variables can be derived from the primary-variable TI applying
325 some physical model.
326
327 The chapter ends with a reference to the attempts of creating some TI databases,
328 such as FAKTS and CARBDB.
329
330 Chapter II.8 describes how to validate the data versus the TI and the type of
331 quality controls that should be performed. As indicated in the book, when a
332 model is derived from data, it is impossible that there is any inconsistency
333 between data and model; however, in MPS, the TI is not built from the data, and,
334 therefore, there could be inconsistencies between TI and data that are not
335 obvious, either because of the complexity of the data or because of the density of
336 the hard data.
337
338 The most interesting part in relation to TI validation is the technique described
339 to assign probabilities to training images conditioned to the hard data. These
340 probabilities can be later used to rule out improbable TIs and to distribute the

341 number of realizations generated from each postulated TI proportionally to their
342 probabilities.

343

344 The chapter goes on with quality control; the first check is to compare
345 histograms, variograms, connectivity functions or other statistics between
346 realizations and TI. The issue of verbatim copy of big patches of the TI onto the
347 realizations is analyzed through the use of coherence maps, and examples are
348 given.

349

350 The authors then argue that the realizations should be checked by measuring
351 their spatial uncertainty. The aim is having small within-realization variability
352 (i.e., realizations stay close to the TI) but large between-realization variability
353 (i.e., realizations span the space of uncertainty). These checks can be done by
354 comparing summary statistics calculated on the realizations and on the TI.

355

356 A final check suggested by the authors is that of consistency for conditioning, in
357 the sense that conditioning data should exert some effect beyond the point
358 where they are. An expensive method to do it is proposed.

359

360 Chapter II.9 introduces the problem of inverse modeling with training images. It
361 is the second longest chapter in the book; the subject is of sufficient entity to
362 deserve a book by itself. The chapter starts acknowledging that, up to here, all
363 conditioning was direct and without iterations, but there are certain types of
364 data which require formulating and solving an inverse model. Although most of
365 the material is general, the authors warn that they will focus in the problem of
366 inverse modeling in the subsurface and the use of geophysical and dynamic data.

367

368 The chapter starts with an introduction about inverse modeling to conclude that
369 it can be formulated, in the standard Bayesian framework, as finding the
370 posterior distribution of model parameters given the data, as proportional to the
371 product of a prior distribution of the model parameters times the likelihood of
372 the model.

373

374 The authors then present a prior distribution, which is not parametric but rather
375 algorithmic and is given implicitly by the collection of realizations generated
376 from a given TI by a given algorithm. The main problem is that there is a big
377 chance that the stated prior is inconsistent with the data. The authors insist on
378 the importance of building a prior that is as informative as possible and not
379 based on data; this will likely yield to the use of several TIs, that will have to be
380 later compounded to build a prior distribution as wide as possible, while being
381 consistent with data.

382

383 The chapter continues with a description, at times too terse, of several
384 algorithms to sample the posterior distribution, and therefore, generate
385 realizations that are conditioned to this type of data that are related to the model
386 parameters through a complex non-linear forward model.

387

388 The methods described are rejection sampling, spatial resampling, Metropolis
389 sampling, sequential Gibbs sampling, and pattern frequency matching. All these

390 methods, although theoretically sound, are very CPU intensive, and, in practice,
391 they are replaced for stochastic search methods, whose aim is to generate
392 realizations that are consistent with the prior model and that are also conditional
393 to the data, in the sense that the solution of the forward model matches them.
394 The authors issue a word of caution: the stochastic search methods may generate
395 realizations, each of which is acceptable on its own, but the whole ensemble of
396 realizations may not sample exhaustively enough the space of uncertainty of the
397 posterior.

398
399 The stochastic search methods described are the probability perturbation
400 method, the gradual deformation and the neighborhood algorithm. Sometimes it
401 is difficult to follow all the details of the algorithms from the book itself, but
402 enough references are given for the interested reader.

403
404 An alternative way to perform a stochastic search is by finding a way to
405 parameterize the MPS realizations and then generate realizations of the
406 parameters. The authors first discuss principal component analysis (PCA)
407 decomposition, and, then kernel PCA.

408
409 Chapter II.10 ends part II with a discussion on how to accelerate the generation
410 of realizations by parallelization of the algorithms. After a general discussion of
411 the types of parallel architectures and the types of issues to be taken into
412 account, they present three levels of parallelization: at the realization level, at the
413 patch level, and at the node level; each level with an increased degree of
414 sophistication. The chapter ends with some references to the potential of using
415 the graphic processing unit (GPU) for parallelization purposes.

416
417 Part III, on applications, is illustrative.

418
419 The book ends with three chapters presenting practical applications of MPS to
420 the fields of petroleum engineering, mining and climate modeling. As important
421 as the examples themselves are, I found much more interesting the workflows
422 presented, especially for the petroleum engineering case.

423
424 Chapter III.1 presents the problem of building multiple MPS realizations of a
425 reservoir in which the construction of a new offshore platform is being
426 considered. The objective of the study is well established, and the amounts of
427 data available are thoroughly described. Three geological scenarios are proposed
428 by the geologists, which will yield three TIs generated by Boolean modeling. The
429 issues of coordinate transformation, and the missing scale between data points
430 and cell volumes are commented. The authors go in great detail to describe how
431 they calibrate the target proportions to the seismic information, how the
432 conditional simulations are performed in two steps, and the quality control
433 checks done. The authors also discuss the type of sensitivity analysis carried out
434 to reach some conclusions such as that the choice of TI has an important impact
435 in the final realizations, but the ratio of vertical to horizontal permeability does
436 not. Then, they explain how the probabilities associated to each TI are obtained,
437 to finalize with an application of the probability perturbation method by regions
438 to generate history-matched realizations. It is apparent that those history-

439 matched realizations should be the ones used to appraise whether to construct
440 or not the new offshore platform, but the reader is left out without that
441 information.

442

443 Chapter III.2 presents an example application of MPS in mining. This chapter is
444 co-authored by Pérez, Ortiz and Boucher. The authors explain that MPS could
445 help in rapid modeling and updating of complex geological features, and they
446 focus on MPS methods that address the problem of improving an existing
447 deterministic model. The challenge in mining is to deal with plenty of data and to
448 use MPS at the mine scale with models of many blocks. After presenting several
449 models with different degree of conditioning data and validation data, the
450 authors conclude that MPS can be applied in mining, but in my opinion, the real
451 purpose of the entire exercise was obscure and difficult to understand.

452

453 Chapter III.3 presents an application in climate modeling. It is an application of
454 downscaling of a global circulation model (GCM) to extract climatic variables at a
455 smaller scale. The exercise has several interesting components: it is multivariate,
456 it requires non-stationary modeling, and it adds the time component. The
457 authors explain how they combine exhaustive data from the past at both the 100
458 km by 100 km scale, and the 50 km by 50 km scale of three climatic variables, to
459 predict those same variables at 50 km by 50 km scale into the future using the
460 past as training image, plus the predictions, at 100 km by 100 km, from the GCM.

461

462 Each chapter ends with its own bibliography. And the book ends with an index of
463 terms.

464

465 In addition to the book, the website <http://www.trainingimages.org> contains
466 additional resources including files containing many of the training images,
467 which can be freely used for research purposes. Although in most cases it is quite
468 obvious which is the format of the files, the authors may wish to consider adding
469 a small explanation about it, as well as how the split zip files should be
470 concatenated to get a working zip file from which the training image file could be
471 extracted. The web page also contains links to research codes and an updated list
472 of bibliographic references.

473

474 In summary, this was a long review for a great book tightly packed with plenty of
475 information on multiple-point geostatistics. This book is a must for any
476 researcher or practitioner interested in putting into practice MPS. The deceptive
477 simplicity of MPS is well covered in the book, with numerous indications of
478 potential mistakes in which a user looking for a black-box computer code may
479 incur. All corners of MPS are covered, with extensive descriptions of methods,
480 tools and algorithms. It is without any doubt the reference book on multiple-
481 point geostatistics today, and I presume that it will remain as such for a long time
482 in the future.

483

484 Acknowledgements. This review was completed during a sabbatical stay of the
485 author with the Kansas Geological Survey, Kansas University, Lawrence, KS, USA,
486 which was funded by the Spanish Ministry of Education, Culture and Sports
487 through grant PRX14/00501.

