Preface

Since the term "random field" has a variety of different connotations, ranging from agriculture to statistical mechanics, let us start by clarifying that, in this book, a random field is a stochastic process, usually taking values in a Euclidean space, and defined over a parameter space of dimensionality at least 1.

Consequently, random processes defined on countable parameter spaces will not appear here. Indeed, even processes on $\mathbb{R}^1$ will make only rare appearances and, from the point of view of this book, are almost trivial. The parameter spaces we like best are manifolds, although for much of the time we shall require no more than that they be pseudometric spaces.

With this clarification in hand, the next thing that you should know is that this book will have a sequel dealing primarily with applications.

In fact, as we complete this book, we have already started, together with KW (Keith Worsley), on a companion volume [8] tentatively entitled *RFG-A*, or *Random Fields and Geometry: Applications*. The current volume—*RFG*—concentrates on the theory and mathematical background of random fields, while *RFG-A* is intended to do precisely what its title promises. Once the companion volume is published, you will find there not only applications of the theory of this book, but of (smooth) random fields in general.

Making a clear split between theory and practice has both advantages and disadvantages. It certainly eased the pressure on us to attempt the almost impossible goal of writing in a style that would be accessible to all. It also, to a large extent, eases the load on you, the reader, since you can now choose the volume closer to your interests and so avoid either "irrelevant" mathematical detail or the "real world," depending on your outlook and tastes. However, these are small gains when compared to the major loss of creating an apparent dichotomy between two things that should, in principle, go hand-in-hand: theory and application. What is true in principle is particularly true of the topic at hand, and, to explain why, we shall indulge ourselves in a paragraph or two of history.

The precusor to both of the current volumes was the 1981 monograph *The Geometry of Random Fields* (*GRF*) which grew out of RJA's (i.e., Robert Adler's) Ph.D. thesis under Michael Hasofer. The problem that gave birth to the thesis was an applied one, having to do with ground fissures due to water permeating through the earth under a building site. However, both the thesis and *GRF* ended up being directed more to theoreticians than to subject-matter researchers. Nevertheless, the topics there found many applications over the past two decades, in disciplines as widespread as astrophysics and medical imaging.

These applications led to a wide variety of extensions of the material of *GRF*, which, while different in extent to what was there, were not really different in kind. However, in the late 1990s KW found himself facing a brain mapping problem on the cerebral cortex (i.e., the brain surface) that involved looking at random fields on manifolds. Jonathan Taylor (JET) looked at this problem and, in somewhat of a repetition of history, took it to an abstract level and wrote a Ph.D. thesis that completely revolutionized[1] the way one should think about problems involving the geometry generated by smooth random fields. This, and subsequent material, makes up Part III of the current, three-part, book.

In fact, this book is really about Part III, and it is there that most of the new material will be found. Part I is mainly an adaptation of RJA's 1990 IMS lecture notes, *An Introduction to Continuity, Extrema, and Related Topics for General Gaussian Processes*, considerably corrected and somewhat reworked with the intention of providing all that one needs to know about Gaussian random fields in order to read Part III. In addition, Part I includes a chapter on stationarity. En passant, we also included many things that were not really needed for Part III, so that Part I can be (and often has been) used as the basis of a one-quarter course in Gaussian processes. Such a course (and, indeed, this book as a whole) would be aimed at students who have already taken a basic course in measure-theoretic probability and also have some basic familiarity with stochastic processes.

Part II covers material from both integral and differential geometry. However, the material here is considerably less standard than that of Part I, and we expect that few readers other than professional geometers will be familiar with all of it. In addition, some of the proofs are different from what is found in the standard geometry literature in that they use properties of Gaussian distributions.[2]

There are two main aims to Part II. One is to set up an analogue of the critical point theory of Marston Morse in the framework of Whitney stratified manifolds. What makes this nonstandard (at least in terms of what most students of mathematics see as part of their graduate education) is that Morse theory is usually done for smooth manifolds, preferably without boundaries. Whitney stratified manifolds are only piecewise smooth, and are permitted any number of edges, corners, etc. This brings them closer to the objects of integral geometry, to which we devote a chapter. While the results of this specific chapter are actually subsumed by what we shall have to say about Whitney stratified manifolds, they have the advantage that they are easy to state and prove without heavy machinery.

The second aim of Part II is to develop Lipschitz–Killing curvatures in the setting of Whitney stratified manifolds and to describe their role in what are known as "tube

[1] This verb was chosen by RJA and not JET.

[2] After all, since we shall by then have the Gaussian Part I behind us, it seems wasteful not to use it when it can help simplify proofs.

formulas." We shall spend quite some time on this. Some of the material here is "well known" (albeit only to experts) and some, particularly that relating to tube and Crofton formulas in Gauss space, is new. Furthermore, we derive the tube formulas for locally convex Whitney stratified manifolds, which is both somewhat more general than the usual approach for smooth manifolds, and somewhat more practical, since most of the parameter spaces we are interested in have boundaries. In addition, the approach we adopt is often unconventional.

These two aims make for a somewhat unusual combination of material and there is no easily accessible and succinct[3] alternative to our Part II for learning about them. In the same vein, in order to help novice differential geometers, we have included a one-chapter primer on differential geometry that runs quickly, and often unaesthetically, through the basic concepts and notation of this most beautiful part of mathematics.

However, although Parts I and II of this book contain much material of intrinsic interest we would not have written them were it not for Part III, for which they provide necessary background material. What is it in Part III that justifies close to 300 pages of preparation? Part III revolves around the *excursion sets* of smooth, $\mathbb{R}^k$-valued random fields f over piecewise smooth manifolds M. Excursion sets are subsets of M given by

$$A_D \equiv A_D(f, M) \stackrel{\Delta}{=} \{t \in M : f(t) \in D\}$$

for $D \subset \mathbb{R}^k$.

A great deal of the sample function behavior of such fields can be deduced from their excursion sets and a surprising amount from the Euler, or Euler–Poincaré, characteristics of these excursion sets, defined in Part II. In particular, if we denote the Euler characteristic of a set A by $\varphi(A)$, then much of Part III is devoted to finding the following expression for their expectation, when f is Gaussian with zero mean and unit constant variance:

$$\mathbb{E}\{\varphi(A_D)\} = \sum_{j=0}^{\dim M} (2\pi)^{-j/2} \mathcal{L}_j(M) \mathcal{M}_j^k(D). \tag{0.0.1}$$

Here the $\mathcal{L}_j(M)$ are the Lipschitz–Killing curvatures of M with respect to a Riemannian metric induced by the random field f, and the $\mathcal{M}_j^k(D)$ are certain Minkowski-like functionals (closely akin to Lipschitz–Killing curvatures) on $\mathbb{R}^k$ under Gauss measure.

If all of this sounds terribly abstract, the truth is that it both is, and is not. It is abstract, because while (0.0.1) has had many precursors over the last 60 years or so, it has never before been established in the generality described above. It is also abstract in that the tools involved in the derivation of (0.0.1) in this setting require some rather heavy machinery from differential geometry. However, this level of abstraction has

[3] The stress here is on "succinct." With the exception of the material on Gauss space, almost everything that we have to say can be found somewhere in the literatures of integral and differential geometry, for which there are many excellent texts, some of which we shall list later. However, all presume a background knowledge that is beyond what we shall require, and each contains only a subset of the results we shall need.

turned out to pay significant dividends, for not only does it yield insight into earlier results that we did not have before, but it also has practical implications. For example, the approach that we shall employ works just as well for nonstationary processes as it does for stationary ones.[4] However, nonstationarity, even on manifolds as simple as $[0, 1]^2$, was previously considered essentially intractable. Simply put, this is one of those rare but constantly pursued examples in mathematics in which abstraction leads not only to a complete and elegant theory, but also to practical consequences.

An extremely simple and very down-to-earth application of (0.0.1) arises when the manifold is the unit interval $[0, 1]$, f is real-valued, and $D = [u, \infty)$. In that case, $\mathbb{E}\{\varphi(A_D)\}$ is no more than the mean number of upcrossings of the level u by the process f, along with a boundary correction term. Consequently, modulo the boundary term, (0.0.1) collapses to no more than the famous Rice formula, undoubtedly one of the most important results in the applications of smooth stochastic processes. If you are unfamiliar with Rice's formula, then you might want to start reading this book at Section 11.1, where it appears in some detail, together with heuristic, but instructional, proofs and applications.

One of the reasons that Rice's formula is so important is that it has long been used as an approximation, for large u, to the *excursion probability*

$$\mathbb{P}\Big\{ \sup_{t\in[0,1]} f(t) \geq u \Big\},$$

itself an object of major practical importance. The heuristic argument behind this is simple: If f crosses a high level, it is unlikely do so more than once. Thus, in essence, the probability that f crosses the level u is close to the probability that there is an upcrossing of u, along with a boundary correction term. (The correction comes from the fact that one way for $\sup_{t\in[0,1]} f(t)$ to be larger than u is for there to be no upcrossings but $f(0) \geq u$.) Since the number of upcrossings of a high level will always be small, the probability of an upcrossing is well approximated by the mean number of upcrossings. Hence Rice's formula gives an approximation for excursion probabilities.

If (0.0.1) is the main result of Part III, then the second-most-important result is that, at the same level of generality and for a wide choice of D, we can find a bound for the difference

$$|\mathbb{P}\{\exists\, t \in M : f(t) \in D\} - \mathbb{E}\{\varphi(A_D)\}|.$$

A specific case of this occurs when f takes values in $\mathbb{R}^1$, in which case not only can we show that, for large u,

$$\Big|\mathbb{P}\Big\{ \sup_{t\in M} f(t) \geq u \Big\} - \mathbb{E}\{\varphi(A_{[u,\infty)})\}\Big|$$

is small, but we can provide an upper bound to it that is both sharp and explicit.

Given that the second term here is known from (0.0.1), what this inequility gives is an excellent approximation to Gaussian excursion probabilities in a very wide setting, something that has long been a holy grail of Gaussian process theory.

[4] Still assuming marginal stationarity, i.e., zero mean and constant variance.

In the generality in which we shall be working, the bound is determined by geometric properties of the manifold M with the induced Riemannian metric mentioned above. Furthermore, unlike the handwaving argument described above for the simple one-dimensional case, the new tools provide a fully rigorous proof.

At this point we should probably say something about why we chose to take piecewise smooth manifolds as our generic parameter space. This is perhaps best explained via an example. Suppose that we were to take parameter spaces that were smooth, even C^∞, manifolds, with C^∞ boundaries.

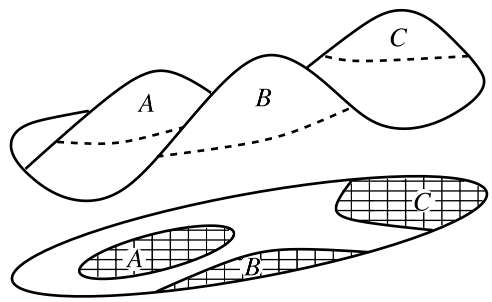

Fig. 0.0.1. A C^∞ function defined over a manifold with a C^∞ boundary gives excursion sets that have sharp, nondifferentiable, corners.

Such an example is shown in Figure 0.0.1, where the parameter space is a disk. The three excursions of a (C^∞) function f above some nominal level are marked on the function surface, and these lie above the three corresponding components of the excursion set $A_{[u,\infty)}$. Note that, despite the smoothness of each component of this example, the excursion set has sharp corners where it intersects with the boundary of the parameter space. In other words, $A_{[u,\infty)}$ is only a piecewise smooth manifold.

It turns out that since we end up with piecewise smooth manifolds for our excursion sets, there is not a lot saved by not starting with them as parameter spaces as well.[5]

So now you know what awaits you at the end of the path through this book. However, traversing the path has value in itself. Wandering, as it does, through the fields of both probability and geometry, it is a path that we imagine not too many of you will have traversed before. We hope that you will enjoy the scenery along the way as much as we have enjoyed describing it. (We also hope, for your sake, that it will be easier and faster in the reading than it was in the writing.)

We are left now with two tasks: Advising how best to read this book, and offering our acknowledgments.

[5] Of course, we could simplify things considerably by working only with parameter spaces that have no boundary, something that would be natural, for example, for a differential geometer. However, this would leave us with a theory that could not handle parameter spaces as simple as the square and the cube, a situation that would be intolerable from the point of view of applications.

The best way to read the book is, of course, to start at the beginning and work through to the end. That was how we wrote it. However, here some other possibilities, depending on what you want to get out of it.

(i) *A course on Gaussian processes*: Chapters 1 through 4 along with Sections 5.1 through 5.4 if you want to learn about stationarity as well. These chapters can be read in more or less any order; see the comments in the introduction to Part I.
(ii) *Random fields on Euclidean spaces, with an accent on geometry*: Sections 1.1–1.4.2 and Chapter 3 for basic Gaussian processes, Sections 4.1, 4.5, and 4.6 for some classical material on extremal distributions, and Chapter 5 on stationarity. Chapter 6 and Section 9.4 give the basic geometry and Chapter 11 the random geometry of Gaussian fields. Section 14.4 gives examples of how the results of Chapter 11 relate to excursion probabilities, and Section 15.10 gives examples of the non-Gaussian theory. (Note that because you have chosen to remain in the Euclidean scenario, and so avoid most of the real challenges of differential geometry, you have been relegated to reading examples instead of the general case!)
(iii) *Probabilisitic problems in, and using, differential geometry*: Sections 1.1, 1.2 and the results (but not proofs) of Chapter 3 to get a bare-bones introduction to Gaussian processes, along with Sections 5.5 and 5.6 for some important notation. As much of Chapter 7 as you need to revise differential-geometric concepts, followed by Chapters 8, 9, and 10. The punch line is then in Chapters 12 and 13 for Gaussian processes and Chapter 15 in general. It is only in this last chapter that you will get to see all the geometric preliminaries of Part II in play at once.
(iv) *Applications without the theory*: Wait for *RFG-A*. We are working on it!

Now for the acknowledgments. Both RJA and JET owe debts of gratitude to KW, and we had better acknowledge them now, since we can hardly do it in the preface of *RFG-A*.

Beyond our personal debts to KW, not least for getting the two us of together, the subject matter of this book also owes him an enormous debt of gratitude. It was during his various extensions and applications of the material of *GRF* that the passage between the old Euclidean theory and its newer manifold version began to take shape. Without his tenacious refusal to leave (applied) problems because the theory (geometry) seemed too hard, the foundations on which our Part III is based would never have been laid.

Back to the personal level, we also owe debts of gratitude to numerous students at the Technion, UC Santa Barbara, Stanford, and the ICE-EM in Brisbane who sat through courses as we put this volume together, as well as the group at McGill that went through the book as a reading course with KW. Their enthusiasm, patience, and refusal to take "it is easy to see that" for an answer when it was not all that easy to see things, not to mention all the typos and errors that they found, has helped iron a lot of wrinkles out of the final product.

In particular, we would like to thank Nicholas Chamandy, Sourav Chatterjee, Steve Huntsman, Farzan Rohani, Alessio Sancetta, Armin Schwartzman, and Sreekar Vadlamani for their questions, comments, and, embarrassingly often, corrections.

The ubiquitous anonymous reviewer also made a number of useful suggestions and we are suitably grateful to him/her.

The generous support of the U.S.–Israel Binational Science Foundation, the Israel Science Foundation, the U.S. National Science Foundation, the Louis and Samuel Seiden Academic Chair, and the Canadian Natural Sciences and Engineering Research Council over the (too long a) period that we worked on this book are all gratefully acknowledged.

Finally, don't forget, after you finish reading this book, to run to your library for a copy of *RFG-A*, to see what all of this theory is *really* good for.

Until such time as *RFG-A* appears in print, preliminary versions will be available on our home pages, which is where we shall also keep a list of typos and/or corrections for this book.

Robert J. Adler
Haifa, Israel
ie.technion.ac.il/Adler.phtml

Jonathan E. Taylor
Stanford, CA, USA
www-stat.stanford.edu/~jtaylo