

Cases for the nugget in modeling computer experiments

Robert B. Gramacy · Herbert K.H. Lee

Received: 26 July 2010 / Accepted: 22 December 2010 / Published online: 30 December 2010
© Springer Science+Business Media, LLC 2010

Abstract Most surrogate models for computer experiments are interpolators, and the most common interpolator is a Gaussian process (GP) that deliberately omits a small-scale (measurement) error term called the *nugget*. The explanation is that computer experiments are, by definition, “deterministic”, and so there is no measurement error. We think this is too narrow a focus for a computer experiment and a statistically inefficient way to model them. We show that estimating a (non-zero) nugget can lead to surrogate models with better statistical properties, such as predictive accuracy and coverage, in a variety of common situations.

Keywords Computer simulator · Surrogate model · Gaussian process · Interpolation · Smoothing

1 Introduction

To some, interpolation is *the* defining feature that distinguishes surrogate models (or emulators) for computer experiments from models for ordinary experiments. We think this is old-fashioned at best and misguided at worst. It is certainly true that a large swath of computer experiments are “deterministic”, in the sense that once $y(x)$ is known there

can be no uncertainty in the output $Y(x')$ if $x' = x$, because the simulator does not behave stochastically. Interpolation would seem natural in this case, and this is typically facilitated by a zero nugget in a Gaussian process (GP) prior for $Y(x)$. Our first observation is that many of the more recent computer experiments are indeed stochastic. A typical formulation is as an agent based model or finite element simulation where the purpose is to study cohort/community effects in independent organisms/agents whose behavior is governed by simple stochastic rules which cannot be understood analytically. Examples abound in biology (Johnson 2008; Henderson et al. 2009), chemistry (Gillespie 2001), and industrial design and engineering (Ankenman et al. 2010) to name just a few. It is in this sense that the defining feature of zero-nugget GPs for computer experiments is old-fashioned. Many computer experiments these days are not deterministic, so in those cases you would include a nugget without hesitation. The definition of *surrogate model* for a computer experiment needs to be updated.

But that is not what this paper is really about. We shall concentrate on those computer experiments that really are “deterministic”—in a sense similar to its usage above but whose decomposition of meaning in modern experiments is one of the main foci of this paper—and argue that you should use a nugget anyway. Our arguments for this are not computational, although the numerical instabilities of zero-nugget models are well-documented (Ababou et al. 1994; Neal 1997). Another established criticism of zero-nugget models, upon which we will not focus, involves theoretical aspects of smoothness and derivatives. Stein (1999, p. 96) proves that the smoother the spatial process, the smaller any error or variability needs to be in order for it to have negligible effect. Since the standard assumption in the computer modeling literature is a Gaussian correlation function, and therefore of infinite differentiability, means that the results

Part of this work was done while R.B.G. was at the Statistical Laboratory, University of Cambridge.

R.B. Gramacy (✉)
Booth School of Business, University of Chicago, Chicago, USA
e-mail: rbgramacy@chicagobooth.edu

H.K.H. Lee
Applied Math & Statistics, University of California, Santa Cruz,
Santa Cruz, USA
e-mail: herbie@ams.ucsc.edu

are highly sensitive to any possible deviations and thus Stein strongly cautions against omitting a nugget term.

As larger nugget values can impact the fitted values of other parameters (Gramacy and Lee 2008b; Pepelyshev 2010), some authors go to great lengths to reconcile numerical stability and zero-nugget-like interpolation, usually by using as small a nugget as possible (Ranjan et al. 2010). Instead, we argue that issues of numerical stability, while they are strong arguments in favor of a nugget, are a bit of a red herring in the face of more serious conceptual issues. We aim to separate the ideology of forcing interpolation from some important (and undesirable) consequences of the zero-nugget model. We shall argue that when the data are sparse or when model assumptions are violated (e.g., stationarity)—and they typically are—the nugget is crucial for maintaining good statistical properties for the emulator (e.g., coverage). Essentially, when modeling computer experiments, we must be pragmatic about how assumptions map to conclusions (surrogate model fits), and this leads us to conclude that the most sensible default is to estimate a (nonzero) nugget.

The remainder of the paper is outlined as follows. We conclude this section with a brief review of GP basics, with further reference to their application as surrogate models for computer experiments. In Sect. 2 we elaborate on several conceptual problems with the zero-nugget approach. Section 3 provides numerical examples, showing how sparseness of the sample (Sect. 3.1) or violations of standard (and uncheckable) assumptions (Sect. 3.2) can lead to inferior predictive surfaces with the zero-nugget approach. The issue of “determinism” is explored in Sect. 3.3 to similar effect. And in Sect. 4 we revisit these points on a real-world computer experiment involving CFD simulations of a rocket booster re-entering the atmosphere. Finally, we conclude with a discussion.

1.1 GP basics

The canonical choice of surrogate model for computer experiments is the stationary Gaussian process (Sacks et al. 1989; O’Hagan et al. 1999; Santner et al. 2003), which defines a random process whose evaluation at any finite collection of locations has a multivariate Gaussian distribution with a specified mean and covariance function that depend only on the relative positions of the locations. A typical specification of the covariance function is the Gaussian correlation, so that the covariance between any two points is

$$\begin{aligned} C(\mathbf{x}_j, \mathbf{x}_k) &= \sigma^2 K(\mathbf{x}_j, \mathbf{x}_k) \\ &= \sigma^2 \exp \left\{ - \sum_{\ell=1}^m \frac{|x_{ij} - x_{ik}|^2}{d_\ell} \right\}, \end{aligned} \quad (1)$$

where m is the dimension of the space and \mathbf{d} is a vector (the range parameter) which scales the correlation length in

each dimension. This model will interpolate the data, fitting a smooth curve between observed outputs of the computer simulator. The predictive distribution (or so-called *kriging* equations) at new inputs \mathbf{x}^* are conditionally Gaussian given (\mathbf{x}, y) pairs $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)$ and settings of the parameters σ^2 and \mathbf{d} . Our references contain the relevant equations; we simply remark here that the variance of this distribution has the distinctive property that it is zero when $\mathbf{x}^* = \mathbf{x}_i$ for one of $i = 1, \dots, n$, and increases away from zero as the distance from \mathbf{x}^* to the nearest \mathbf{x}_i increases. When all elements of \mathbf{d} are equal, the process is called *isotropic*.

An extension of this model is to include a nugget term, specifying the covariance function as

$$\begin{aligned} C(\mathbf{x}_j, \mathbf{x}_k) &= \sigma^2 K(\mathbf{x}_j, \mathbf{x}_k) \\ &= \sigma^2 \left[\exp \left\{ - \sum_{\ell=1}^m \frac{|x_{ij} - x_{ik}|^2}{d_\ell} \right\} + g \delta_{j,k} \right], \end{aligned}$$

where $\delta_{.,.}$ is the Kronecker delta function and g is the nugget term. Originally introduced to model small-scale variation in geostatistical models, it is also mathematically equivalent to adding random noise term into the likelihood. Thus with $g > 0$, this model no longer interpolates the data, and returns us to a situation analogous to fitting a mean function with noisy data. The predictive distribution has many features in common with that obtained from the zero-nugget/no-nugget model (above). For example, the variance increases away from the nearest input \mathbf{x}_i , but is not zero at $\mathbf{x}^* = \mathbf{x}_i$; rather, it is $g\sigma^2$ at those locations.

There are many ways to infer the parameters σ^2, \mathbf{d}, g given data $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)$. For example, the resulting multivariate normal likelihood can be maximized numerically, although the presence of the nugget may lead to a bimodal surface (Gramacy and Lee 2008b). In this paper we happen to take a Bayesian approach, but all of our arguments hold true under the frequentist paradigm as well. Our implementation is in R using the GP code from the `tgpp` library on CRAN (Gramacy 2007). The package makes use of a default Inverse-Gamma prior for σ^2 that yields a multivariate Student- t marginal likelihood (integrating out σ^2) for y_1, \dots, y_n given $\mathbf{x}_1, \dots, \mathbf{x}_n$ and the parameters \mathbf{d} and g . Standard Gamma priors are placed independently on the d_ℓ components and g , and the Metropolis-within-Gibbs MCMC method is used to sample from the posterior via the Student- t marginal likelihood. The inputs \mathbf{x}_i are pre-scaled and the proposals are designed to deal with the bimodal posteriors that can result.

2 Examining the model assumptions

Most papers obsess on the zero-nugget model. When a nugget is needed for computational reasons, one aims to

make it as small as possible while still maintaining numerical stability. The argument is that the closer the nugget is to zero, the more accurate the surrogate model approximation is to the computer code output. This may be true if there is sufficient data, but is it even the right thing to be worried about? The measurement error captured by the nugget (which is presumed to be zero for deterministic computer simulations) is but one of many possible sources of error. Here we discuss four such sources of uncertainty which are likely to be of greater importance, so it is boggling why so much attention is paid to the nugget.

2.1 Simulator bias

No computer simulator is a perfect representation of the real world. All simulators are mathematical models and thus only approximate the real world, so they have some “bias”. How we deal with this discrepancy depends on whether or not real world data are available. We take those two cases in turn.

When real data are available, it is well-established that the simulator can be calibrated using the data, that is, the discrepancy between the simulator and reality can be modeled using an additional Gaussian process (Kennedy and O’Hagan 2001; Santner et al. 2003). While this addresses the simulator bias, it introduces a source of noise—that of the real data. A measurement error term in the likelihood can be shown to be a re-parameterization of a nugget term in the covariance function (Gramacy 2005, Appendix B). Thus we get the same model at the end whether we interpolate then add noise, or whether we just fit a nugget term. Since the two approaches end up at the same place, we might as well embrace the nugget while fitting the model.

If real data are not available, then the bias cannot be estimated, and that term is typically ignored and swept under the rug. Yet pretending that the simulator is perfect, even though we know it is not is clearly ignoring a major source of error. Rather than insist that the statistical model interpolate the simulator, why not allow the model to smooth the simulator output? We make this point philosophically, in that we do not see why forcing the model to interpolate something that is deterministically wrong would be any better than allowing smoothing, and that smoothing can offer a measure of protection and robustness.

2.2 The stationarity assumption

Nearly every analysis in the computer modeling literature makes an assumption of stationarity, second-order stationarity, or at least piece-wise stationarity. More precisely, a *residual* process, arrived at after subtracting off some mean which might be estimated using a fairly substantial number of regressors (e.g., Rougier et al. 2009; Martin and Simpson 2005), is assumed stationary. Such a two-step process

can, indeed, lead to good fits. But it can also be fairly involved as choosing the best regressors is a non-trivial task, and there is no guarantee that the residuals thus obtained are stationary. In any case, the stationarity assumption is easily challenged. While it may be reasonable as a close approximation to the truth, assuming stationarity will not be exactly correct in most cases.

Typically there is not enough data available to fit a fully nonstationary model, and if there is enough data, then the model becomes too difficult to fit efficiently. Like the bias case, when there is unknown error, a general statistical principle is that smoothing (or shrinking) can give better results. Thus a nugget can help protect us in the case of moderate deviations from stationarity, which would be hard to detect in practice. In Sect. 3.2 we show that even minor violations in the stationarity assumption lead to emulators with poor statistical properties without a nugget.

2.3 Correlation assumptions

There is an underlying assumption that the specified (typically Gaussian) correlation structure is correct. While this is a nice modeling assumption, it is yet another convenient approximation to reality. Parameters for the form of the correlation function can be difficult to fit in practice, and so it is often necessary to simply specify a reasonable guess. Since it is only an approximation, this is a further reason for allowing smoothing in the model.

2.4 The assumption of a deterministic simulator

The modeling assumptions addressed above may indeed be reasonable for a particular true physical process, but the computer implementation of the solution may still behave in unpredictable ways. The assumption of a deterministic simulator may itself be a problem. Here we discuss two related possible issues, nonmodelable determinism and theoretical but not numerical determinism, among other possible problems with the assumption of deterministic behavior in practice.

Some deterministic functions really are better treated as nondeterministic. As a simple example, consider a pseudo-random number generator where, for any given seed, an output is returned deterministically (if not also unpredictably unless you know a lot about numerical analysis). A version of a computer simulator $f(x)$ approximating a function $g(x)$ numerically might effectively behave as follows (coded in R):

```
f <- function(x) {
  set.seed(x)
  return(g(x) + rnorm(1))
}
```

This function $f(x)$ is theoretically deterministic, but knowing how it relates to the true function $g(x)$, it would be irrational to interpolate it. Clearly one would want to smooth out the pseudo-random noise which would give us a much better fit of the underlying $g(x)$, and this is exactly what the nugget is designed to do. You may argue about the “deterministic” nature of the f simulator, but that is to get bogged down in philosophical matters and miss the practical point. In this “cartoon”, g might represent the mathematical model/equations that describe a system (perhaps only on paper) whose solutions or realizations are not available analytically and require numerically approximate solutions. This is where computer implementation in f comes in, which may be deterministic if not otherwise ill-behaved. In our experience, such scenarios are more the rule than the exception in modern computer experiments.

Along similar lines, the f and g relationship might just as easily represent a function with chaotic behavior, which can happen in complex systems of differential equations, or the Perlin noise function (Perlin 2002), which is a deterministic method for generating random-looking smooth surfaces in computer graphics. Alternatively, the `rnorm(1)` term may stand in for the amount by which an iterative approximation algorithm steps over the convergence threshold. Although we usually we assume that this amount can be made to be arbitrarily small, this might not always be justified. One reason is the lack of uniformity in machine-representable floating-point numbers.

Now, the above example may seem pathological, but in Sects. 3.3 & 4 we give a synthetic and real example, respectively, which are essentially the following adaptation:

```
f2 <- function(x) {
  set.seed(x)
  y <- runif(1)
  if(y < 0.9) return(g(x))
  else return(h(x))
}
```

for some new $h(x)$. The pseudo-random (but deterministic) y is intended to represent the chance that the computer code was (poorly) initialized such that it may end up converging to a sub-optimal (but locally converged) solution $h(x)$ 10% of the time rather than the true globally converged approximation to $g(x)$. This is not an uncommon feature of a modern computer simulator, i.e., where the final output depends upon an initial “solution” for which there are defaults that usually work, but sometimes lead to a converged solution which is different from the one intended. It is clearly sub-optimal to use a zero-nugget model in this case, because some of the outputs are not the correct values. Despite their deterministic nature, we show in Sect. 3.3 that uncertainty about the true function is best modeled with a random process that smooths rather than interpolates.

3 Statistically better fits with the nugget

3.1 Protecting against misfits with sparse data

Many computer experiments are expensive to run and the number of datapoints is limited. As many experiments have higher dimensional input spaces, the curse of dimensionality implies that the data will be sparse in the input region. When the data are sparse, interpolation can have unpleasant results (Taddy et al. 2008, Sect. 2.2). We present here a simulated example where the data are sparse in one dimension, but this represents the concept of sparseness in higher dimensions with a simpler function or with more data.

Consider the function

$$Z = \frac{\sin(10\pi X)}{2X} + (X - 1)^4.$$

Suppose we only have 20 datapoints available (in practice, we would have more points but more dimensions). We randomly generated 10,000 such datasets (with X generated from a uniform distribution each time) and fit models both with a nugget and without one. The table in Fig. 1 gives the distribution of the mean square errors of fits under each model, and the model that includes a nugget does better on average (a paired t -test gives a p -value of less than 2.2×10^{-16}). The plots in Fig. 1 show one of the runs. The data are too sparse to get a good fit of the function for smaller input values. While the nugget model smooths and produces reasonable confidence bands, in order to interpolate smoothly the no-nugget model ends up making predictions well outside the range of the actual data in that region, and its confidence bands are all over the place.

In practice, if we only had one-dimensional inputs, we would use evenly-spaced points for small samples, and much of the issue here would go away. However, as the dimension of the space increases, it becomes impossible to use a regular grid, and random Latin hypercubes are often used. It can be impossible to understand how well a design really covers a high-dimensional space. Thus the nugget provides good protection against the strange fits that interpolation can produce.

3.2 Poor coverage

In fact, the nugget offers protection from a slew of problematic scenarios. Here we shall illustrate that the no-nugget model under-covers the true computer simulator response when the stationarity assumption is not satisfied. We use three examples.

The first example is a 1-d function which is clearly non-stationary, but otherwise mimics typical features of a computer code. The response is given by $y(x) = \sin(x) - 0.02 \cdot t_1(x, 1.57, 0.05)$ where $t_1(\cdot, \mu, \sigma)$ is a Cauchy density with




Fig. 1 (Color online) Plots of one example of a fit (dark solid line), confidence bands (red), and true function (light grey). The left plot shows the fit using a nugget, the right plot without a nugget. The table

on the right is the summary of the mean square errors under both models for 10,000 repeated uniform designs

mean μ and spread σ . The two rows of Fig. 2 show fits for two typical random uniform designs of size ten. The difference between smoothing (estimated nugget; left panels) and interpolation (no nugget; right panels) is clear. We see that the no-nugget model under-covers the truth (in gray) and can have wildly different (i.e., narrow or wide) 90% predictive credible intervals. This experiment was repeated 100 times and the percentage of the area of the input space where the true $y(x)$ was covered (pointwise) by the 90% interval was recorded. A table summarizing the results numerically is on the right in the figure. We see from this table that the under-coverage of the no-nugget model can be drastic. For one of the random designs it only covered 6.5% of $y(x)$ and 3/4 of the trials under-covered by more than 10%. By contrast, the model which estimates a nugget has good coverage properties. Its median and mean coverages are close to 90% and the central 50% region tightly brackets the truth. Clearly, connecting the dots comes at the expense of other, arguably more important, statistical measures of goodness of fit.

Although it is intuitive, pointwise coverage via might not be an ideal metric for comparing model fit because it ignores correlation aspects of the posterior predictive distribution—mis-coverage at one input is indicative of miscoverage at a continuum of nearby inputs. A metric based on the Mahalanobis distance, proposed by Bastos and O’Hagan (2009) for GP models, allows such correlation in the predicted outputs to be taken into account. For completeness, we report a summary of the square root of these distances in Fig. 2 as well. The conclusions based on this metric are largely the same: estimating a nugget leads to a superior fit (smaller distances) compared to the zero-nugget interpolating approach.

Table 1 Left are coverages and square-root Mahalanobis distances for the 2-d exponential data; right for the 5-d Friedman data

coverage	exp data		fried data	
	nug	nonug	nug	nonug
Min.	0.5479	0.3965	0.5480	0.4580
1st Qu.	0.8623	0.8242	0.8930	0.8350
Median	0.9185	0.8936	0.9320	0.8890
Mean	0.8962	0.8691	0.9205	0.8762
3rd Qu.	0.9492	0.9395	0.9580	0.9310
Max.	1.0000	1.0000	0.9990	1.0000

$\sqrt{\text{mah}}$	exp data		fried data	
	nug	nonug	nug	nonug
Min.	3.181	3.857	14.312	15.816
1st Qu.	20.870	25.850	24.203	30.760
Median	34.840	48.270	28.676	36.100
Mean	78.020	96.960	29.271	36.794
3rd Qu.	67.260	100.500	33.621	41.968
Max.	1971.000	21820.000	55.435	69.111

We performed similar experiments on two higher-dimensional data sets. The first is a 2-d exponential function $y(x) = x_1 \exp\{-x_1^2 - x_2^2\}$, which less clearly violates the stationarity assumption. From the left side of Table 1 we see a similar under-coverage of the no-nugget model with repeated uniform designs of size 20. Mahalanobis distances are included for completeness. Our second experiment involved the first Friedman data function (Friedman 1991) with five inputs where the response is $y(x) =$

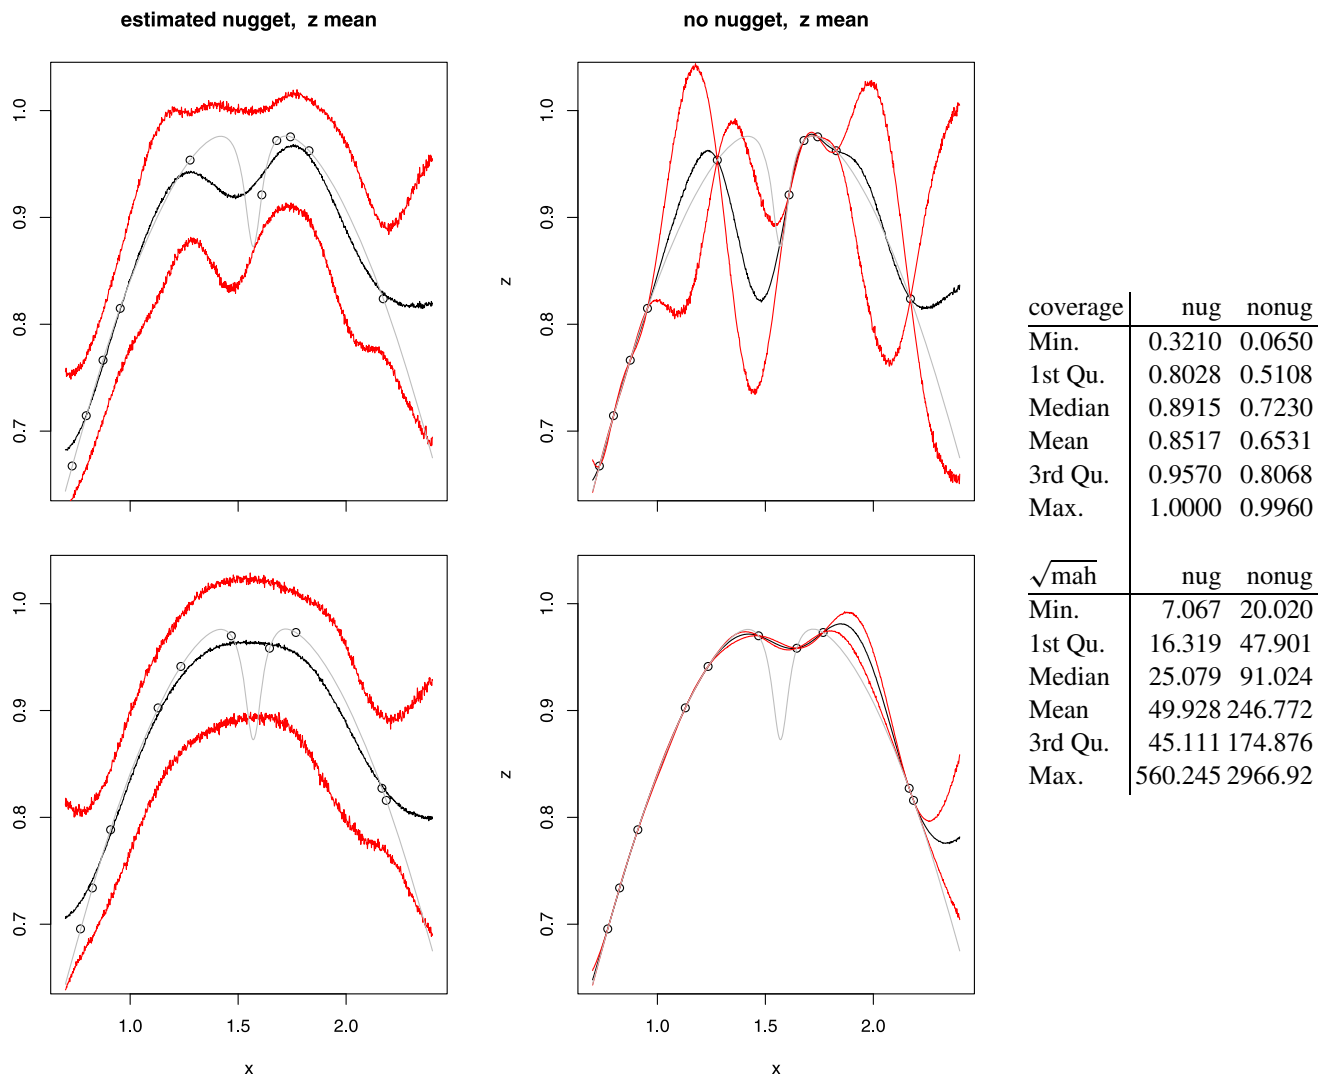


Fig. 2 The plots on the left are examples of fits under two uniform designs in nugget (left column) and no-nugget (right column) models. The table on the right summarizes of the coverages and (square root) Mahalanobis distances under both models for 100 repeated uniform designs

$10 \sin(\pi x_1 x_2) + 20(x_3 - 0.5)^2 + 10x_4 + 5x_5$. This function is better behaved (i.e., stationarity may be a reasonable assumption). However, it is apparent that correlation in the response would decay at different rates along the five coordinates—clear anisotropy. To illustrate how the effect of an inappropriate choice of correlation function is felt more strongly in the no-nugget model we used an isotropic Gaussian correlation function and uniform designs of size 25. The results are shown on the right in the table.

It is worth pointing out that as the size of the designs are increased, and/or as the data less obviously violate assumptions, both models (nugget and no-nugget) will tend to over cover in practice. This is because we are fitting a model (GP) which always yields positive posterior predictive error away from the (discrete set of) design points, i.e., over an uncountably large region. Since the function we are modeling is deterministic, we know that as the size of the

design tends to (countable) infinity we should be able to obtain a “perfect” fit with a high degree polynomial. So a GP is the wrong model in this case. Since over-coverage is inevitable, under-coverage should be our primary concern, and to avoid under-covering we can see that a nugget is needed.

3.3 Challenging determinism in computer simulation

Some computer experiments are deterministic in a technical sense, but not necessarily in a way that translates into sensible assumptions for the building of a surrogate model. We may reasonably presume that codes implementing the algorithms and calculations behind the experiment are nontrivial. They are expensive to program and expensive to execute, requiring long iterations to convergence and the (sometimes arbitrary) specification of tuning parameters, tolerances, and

grid/mesh sizes. As a rule more than an exception, the resulting apparatus works better for some choices of inputs than for others. The most important issue is in detecting global convergence of the code, whose properties usually depend crucially on other implementation choices. It is essentially impossible to guarantee good global convergence proper-

ties, and so this the main target of our attack on the modeling of such “deterministic” computer simulations without a nugget.

Consider the following computer simulator coded in R below.

```
## 2-d function
f2d <- function(x1, x2) {
  w <- function(y) {
    return(exp(-(y-1)^2) + exp(-0.8*(y+1)^2) - 0.05*sin(8*(y+0.1)))
  }
  return(-w(x2)*w(x1))
}

## find the minimum of a projection of the 2-d function
f <- function(x) {
  return(optim(par=x, fn=f2d, x2=x)$value)
}
```

The true underlying function $f(x)$, evaluated by $\mathbb{f}(x)$ in R, is $\arg \min_{x_1} f(x_1, x)$ where

$$f(x_1, x_2) = -w(x_1)w(x_2), \quad \text{and}$$

$$w(y) = \exp\left(-(y-1)^2\right) + \exp\left(-0.8(y+1)^2\right) - 0.05 \sin(8(y+0.1)).$$

The optimization method used by the code above is the `optim` function in R initialized at $x_1 = x$.

Figure 3 shows the true $f(x)$ (dashed-green) and the output of the simulator $\mathbb{f}(x)$ (gray) for $x \in [-1.5, 1.5]$. We can see that the result of numerically finding the optimal value of

the objective function (initialized somewhat arbitrarily, but not pathologically) is that the simulations $\mathbb{f}(x)$ are biased, and behave badly/unpredictably in some parts of the input space. It is worth noting that both behaviors persist with a different static initialization scheme; the $f(x_1, x_2)$ surface has about a dozen local minima. Also, the implementation $\mathbb{f}(x)$ is completely deterministic in a technical sense. However, $\mathbb{f}(x)$ is exhibiting “random” behavior of the sort alluded to in Sect. 2 as the initialization scheme causes the algorithm to converge to different local minima in a way that is not (easily) predictable. There are three places where the initialization causes it to have discontinuities (even though the true $f(x)$ is smooth everywhere), and it is particularly unstable near $x = 0$ since $(0, 0)$ is more or less equidistant from the many local minima of $f(x_1, x_2)$ in the 2-d space.

The figure also shows a fit to the computer simulator ($\mathbb{f}(x)$) output obtained from a gridded design of 100 input–output pairs using a GP with an estimated nugget. The fit is sensible given the discontinuities and otherwise “noisy” behavior of the simulator. It is not possible to fit this data without a nugget, or even with a small one, due to numerical instabilities. However, it is possible to do so with a reduced design of about 20 points or so. To connect with the coverage results in the last section (where stationarity was the issue) we calculated the coverage of $\mathbb{f}(x)$ with 100 repeated uniform random designs of size 20 under the estimated nugget and no-nugget models, and the story is much the same as before. The results are shown in Table 2. Square-root Mahalanobis distances are also shown, as are the distances to the true $f(x)$. They show that when “determinism” is challenged as an assumption on the nature of the data-generating mechanism, a nugget for smoothing is clearly preferred to interpolation in the surrogate model.




Fig. 3 GP fit (with an estimated nugget) to a deterministic function which is the result of an iterative procedure

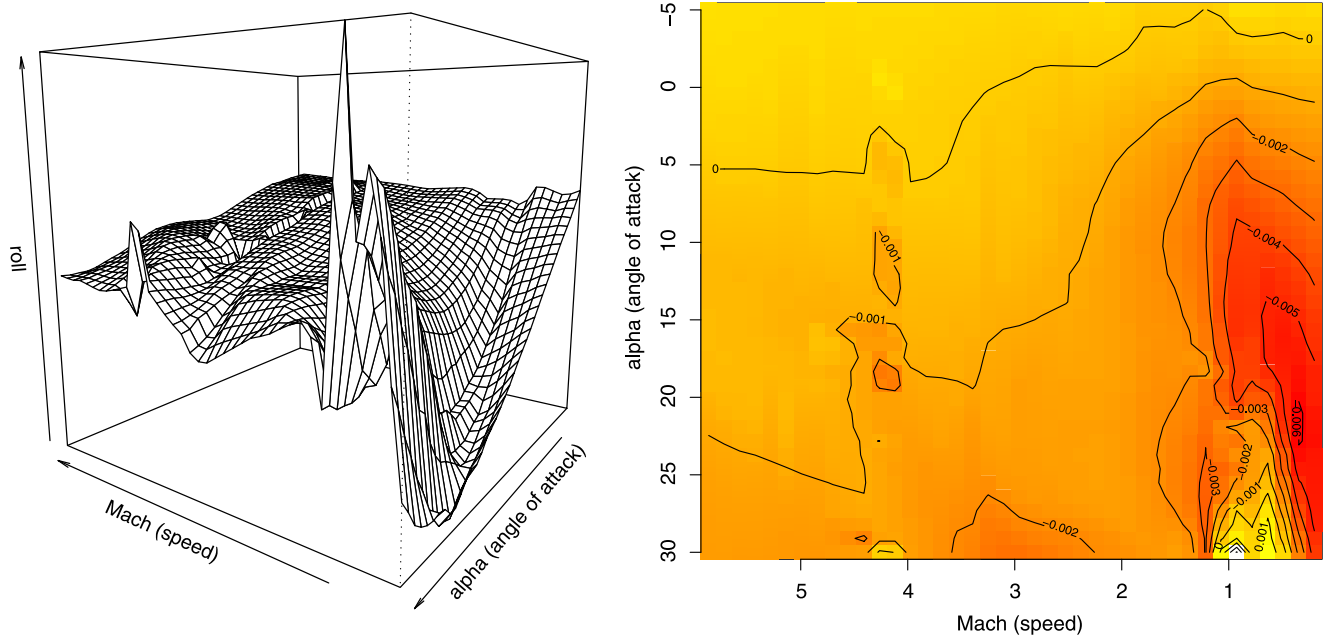


Fig. 4 (Color online) Linearly interpolated slice of the roll response plotted in perspective (*left*) and image/contour (*right*) as a function of speed (Mach) and angle of attack (α), with the slide slip angle (β) fixed to 2. In the image plot, *dark/red* values are lower and

light/yellow values higher in the image plot; the perspective plot is rotated for visualization purposes so that the closest corner corresponds to low speed and high angle of attack

Table 2 Summaries of coverage of the “deterministic” computer simulated $\hat{f}(x)$ data, and the square-root Mahalanobis distances to $\hat{f}(x)$ and the analytic solution $f(x)$

coverage	computer \hat{f}		$\sqrt{\text{mah}}$	computer \hat{f}		true f	
	nug	nonug		nug	nonug	nug	nonug
Min.	0.433	0.2280	Min.	14.07	36.88	41.90	279.17
1st Qu.	0.787	0.6665	1st Qu.	33.24	139.39	104.44	8544.02
Median	0.875	0.7345	Median	56.87	292.17	213.78	12932.55
Mean	0.846	0.7276	Mean	183.45	1163.17	855.44	15402.05
3rd Qu.	0.938	0.8362	3rd Qu.	136.59	811.23	617.03	19392.16
Max.	0.993	0.9760	Max.	3212.00	14561.06	12271.41	49362.90

4 A modern computer experiment

The Langley Glide-Back Booster (LGBB) is a rocket booster that underwent design phases at NASA primarily through the use of computational fluid dynamics simulators that numerically solve the relevant inviscid Euler equations over a mesh of 1.4 million cells (Rogers et al. 2003). The simulator models the forces felt by the rocket at the moment it is re-entering the atmosphere as a function of three inputs describing its state: speed (measured by Mach number), angle of attack (the α angle), and sideslip angle (the β angle). As a free body in space, there are six degrees of freedom, so the six relevant forces/outputs are lift, drag, pitch, side-force, yaw, and roll. While theoretically deterministic, the simulator can fail to converge. Some nonconvergent runs are caught by an automated checker, and re-run with a new schedule of initial conditions, but some are erroneously ac-

cepted even after converging to a clearly inferior solution. Input configurations arbitrarily close to one another can fail to achieve the same estimated convergence, even after satisfying the same stopping criterion.

Here we focus on the roll force output on a data set comprised of simulator runs at 3041 locations. See Fig. 4 for a 2-d slice of this response. Previous work has focused on the lift force (Gramacy and Lee 2008a) which exhibited many similar features, and on a sequential design task taking account of all outputs simultaneously (Gramacy and Lee 2009). The experimental design is a combination of an initial grid followed by two hand-designed finer grids focused around Mach one, as the initial run showed that the most interesting part of the input space was generally around the sound barrier, where the physics behind the simulator changes abruptly from a subsonic regime to a supersonic one. What happens close to and along the boundary is the

Table 3 Coverage of the roll response for the LGBB computer experiment data using a Gaussian process (left) and a treed Gaussian process (right)

coverage	GP		TGP	
	nug	nonug	nug	nonug
Min.	0.7547	0.5726	0.7627	0.5180
1st Qu.	0.9022	0.8427	0.8757	0.7195
Median	0.9239	0.8793	0.8978	0.7670
Mean	0.9187	0.8703	0.8954	0.7606
3rd Qu.	0.9396	0.9059	0.9186	0.8051
Max.	0.9777	0.9741	0.9771	0.9305

most difficult part of the simulation. The regime changes across this boundary cause the stationarity assumption to be violated. Also note the string of anomalies around Mach four, which appear to converge to local, rather than global, solutions. So this experiment comprises two challenging aspects—impractical determinism due to convergence issues and failed assumptions of stationarity due to physical regime changes—and we aim to show that the nugget is important in mitigating their effects when building a surrogate model.

Towards this end we calculated the coverages of predictive surfaces obtained with and without the nugget on a 20-fold partition of the 3041 input/output pairs. We iterated over the folds, training on $1/20^{\text{th}}$ of the data, about 159 pairs, and predicting at the remaining 3009-odd locations in an (inverse) cross-validation fashion. The results of this experiment, repeated 100 times for 2000 total coverages for each predictor, are shown on the *left* in Table 3. Note that this is not a uniform coverage rate (over the input area), since the design is more heavily concentrated around Mach one. However, the results here are as expected. The no-nugget model can severely under-cover in certain examples (with coverage as low as 57%), and gives the target coverage of 90% less than $1/4$ of the time. The model using an estimated nugget is much better behaved. However, it does seem to slightly over-cover. Although less of a concern, we think that the main cause of this is the nonuniformity of the design and our choice of priors for *both* the range and nugget parameters in the face of nonstationarity and nonconvergence issues.

The treed Gaussian process (TGP, Gramacy and Lee 2008a) model was designed to handle the axis-aligned nonstationarity that arises due to regime changes—exactly the sort exhibited by this data. In essence, the TGP model learns an axis-aligned partition of the data wherein the process is well-fit by separate stationary GP models. We performed an identical experiment using TGP and the results are summarized on the *right* in Table 3. We can see that the coverage of the version of TGP which estimates the nugget is improved (with better centering around 90%), but the no-nugget version is not (showing a more consistent tendency to under-

cover). We are left with the impression that the nugget is even more important when a nonstationary model is used, especially in the case of nonconvergent computer experiments where the assumption of “determinism”, while technically valid, may be challenged from a practical standpoint.

5 Discussion

Several authors have previously argued in favor of a nugget term for reasons of numerical stability even when fitting a deterministic model. We go well beyond numerical convenience, raising fundamental issues of a variety of modeling assumptions and argue that the use of a nugget helps protect against poor fits when they are violated.

In many ways, the themes harped on in this paper are just reminders of the usual statistical insights, like exploiting bias variance trade-offs and shrinkage, which are well-known to lead to improved estimators and predictors. It is true that some applications call for specific features in estimators/predictors, like interpolation for deterministic computer experiments, which might cause us to eschew those good practices. At first glance it would seem like forcing a zero nugget, versus estimating one, is just one of these situations. But our experience with fitting GPs to real computer simulation output has shown us time and time again that the ideal of an interpolating (zero-nugget) emulator is not ideal at all. In addition to showing, in this paper, a subset of examples illustrating why we prefer the nugget for statistical purposes, we have touched on high level arguments as to why the zero-nugget model will be inferior in practice.

In reflecting further on these arguments, thanks to helpful comments from our referees, we have come to the following perspective on the matter. The idealized model for the experiment is the following.

real data from physical process (a) computer model

Computer model realizations are expensive so we deploy an emulator [at (a)] based on GP interpolation of the computer model output to save on cost/time. The reality is probably more like the following.

real data (b) physical process (c) mathematical model

(d) computer code (e) emulator

In other words, the real data is a measurement of the physical process, so there is noise or measurement error [at (b)]. This is not controversial. Now, the computer model is actually two things in one. It is a mathematical model which approximates [(c)] the physical model, and although it does not perfectly describe the system—it is *biased*—it is still an idealization. The computer implementation adds a further layer

of “approximation” [(d)], often with erratic if still technically deterministic behavior, which can be an unavoidable nuisance. We try to find the knobs/settings in the code that give the best results, but it is never perfect. If we interpolate [(e)] the output of the computer code, then we are interpolating its idiosyncrasies. If we use a nugget, we might not be precisely where we want to be (closer to the mathematical model and thus the physical process [at (c)]), but we have a fighting chance since we will smooth out the “rough edges” in the computer code. While this schematic representation is a simplification, it shows the multiple stages where approximations are made and our empirical work on real and synthetic data suggest that there is something to it.

Instead of using a nugget, Rougier et al. (2009) advocate using a “rougier”—but still interpolating—correlation function like the Matèrn (e.g., Stein 1999). This may lead to an improved fit, and better numerical properties/decompositions of matrices, compared to smoother correlation functions like the Gaussian (1). There are a few complications with this approach, like the burden specifying an extra smoothing parameter which is hard to infer statistically. But more importantly, one wonders whether better interpolations of idiosyncratic computer code is really what you want? We think it would be better to use a nugget and smoother correlation function because, in our experience, the true solutions to the mathematical equations underlying the model are well-behaved, i.e., smooth. It is also just an easier default option.

Acknowledgements We thank Tim Gustafson for introducing us to the Perlin function. This work was partially supported by National Science Foundation grant DMS-0906720 and EPSRC grant EP/D065704/1. We would also like to thank two referees for thoughtful comments that lead to improvements in the manuscript.

References

- Ababou, R., Bagtzoglou, A.C., Wood, E.F.: On the condition number of covariance matrices in kriging, estimation, and simulation of random fields. *Math. Geol.* **26**(1), 99–133 (1994)
- Ankenman, B., Nelson, B., Staum, J.: Stochastic kriging for simulation metamodeling. *Oper. Res.* **58**(2), 371–382 (2010)
- Bastos, L., O’Hagan, A.: Diagnostics for Gaussian process emulators. *Technometrics* **51**(4), 425–438 (2009)
- Friedman, J.H.: Multivariate adaptive regression splines. *Ann. Stat.* **19**(1), 1–67 (1991)
- Gillespie, D.: Approximate accelerated stochastic simulation of chemically reacting systems. *J. Chem. Phys.* **115**(4), 1716–1733 (2001)
- Gramacy, R.B.: Bayesian treed Gaussian process models. Ph.D. thesis, University of California, Santa Cruz (2005)
- Gramacy, R.B.: *tgp*: An R package for Bayesian nonstationary, semi-parametric nonlinear regression and design by treed Gaussian process models. *J. Stat. Softw.* **19**, 9 (2007)
- Gramacy, R.B., Lee, H.K.H.: Bayesian treed Gaussian process models with an application to computer modeling. *J. Am. Stat. Assoc.* **103**, 1119–1130 (2008a)
- Gramacy, R.B., Lee, H.K.H.: Gaussian processes and limiting linear models. *Comput. Stat. Data Anal.* **53**, 123–136 (2008b)
- Gramacy, R.B., Lee, H.K.H.: Adaptive design and analysis of super-computer experiment. *Technometrics* **51**(2), 130–145 (2009)
- Henderson, D.A., Boys, R.J., Krishnan, K.J., Lawless, C., Wilkinson, D.J.: Bayesian emulation and calibration of a stochastic computer model of mitochondrial DNA deletions in substantia nigra neurons. *J. Am. Stat. Assoc.* **104**(485), 76–87 (2009)
- Johnson, L.: Microcolony and biofilm formation as a survival strategy for bacteria. *J. Theor. Biol.* **251**, 24–34 (2008)
- Kennedy, M., O’Hagan, A.: Bayesian calibration of computer models (with discussion). *J. R. Stat. Soc. B* **63**, 425–464 (2001)
- Martin, J., Simpson, T.: Use of kriging models to approximate deterministic computer models. *AIAA J.* **43**(4), 853–863 (2005)
- Neal, R.M.: Monte Carlo implementation of Gaussian process models for Bayesian regression and classification. Tech. Rep. 9702, Department of Statistics, University of Toronto (1997)
- O’Hagan, A., Kennedy, M.C., Oakley, J.E.: Uncertainty analysis and other inference tools for complex computer codes. In: Bernardo, J.M., Berger, J.O., Dawid, A., Smith, A. (eds.) *Bayesian Statistics 6*, pp. 503–524. Oxford University Press, Oxford (1999)
- Pepelyshev, A.: The role of the nugget term in the Gaussian process method. In: *MODA 9—Advances in Model-Oriented Design and Analysis*, pp. 149–156. Springer, Berlin (2010)
- Perlin, K.: Improving noise. *ACM Trans. Graph.* **21**, 681–682 (2002)
- Ranjan, P., Haynes, R., Karsten, R.: Gaussian process models and interpolators for deterministic computer simulators. Department of Mathematics and Statistics, Acadia University (2010)
- Rogers, S.E., Aftosmis, M.J., Pandya, S.A., Chaderjian, N.M., Tejnil, E., Ahmad, J.U.: Automated CFD parameter studies on distributed parallel computers. In: 16th AIAA Computational Fluid Dynamics Conference (2003). AIAA Paper 2003-4229
- Rougier, J., Guillas, S., Maute, A., Richmond, A.: Expert knowledge and multivariate emulation: The thermosphere–ionosphere electrodynamic general circulation model (TIE-GCM). *Technometrics* **51**(4), 414–424 (2009)
- Sacks, J., Welch, W.J., Mitchell, T.J., Wynn, H.P.: Design and analysis of computer experiments. *Stat. Sci.* **4**, 409–435 (1989)
- Santner, T.J., Williams, B.J., Notz, W.I.: *The Design and Analysis of Computer Experiments*. Springer, New York (2003)
- Stein, M.L.: *Interpolation of Spatial Data*. Springer, New York (1999)
- Taddy, M., Lee, H.K.H., Gray, G.A., Griffin, J.D.: Bayesian guided pattern search for robust local optimization. Tech. Rep. ams2008-02, University of California, Santa Cruz, Department of Applied Mathematics and Statistics (2008)