Bayesian Emulation for Computer Models with Multiple Partial Discontinuities

Ian Vernon^{*}, Jonathan Owen^{*}, and Jonathan Carter[†]

Abstract. Computer models are widely used across a range of scientific disciplines to describe various complex physical systems, however to perform full uncertainty quantification we often need to employ emulators. An emulator is a fast statistical construct that mimics the slow to evaluate computer model, and greatly aids the vastly more computationally intensive uncertainty quantification calculations that an important scientific analysis often requires. We examine the problem of emulating computer models that possess multiple, partial discontinuities occurring at known non-linear locations. We introduce the Torn Embedding Non-Stationary Emulation (TENSE) framework, based on carefully designed correlation structures that respect the discontinuities while enabling full exploitation of any smoothness/continuity elsewhere. This leads to a single emulator object that can be updated by all runs simultaneously, and also used for efficient design. This approach avoids having to split the input space into multiple subregions. We apply the TENSE framework to the TNO Challenge II, emulating the OLYMPUS reservoir model, which possesses multiple such discontinuities.

Keywords: uncertainty quantification, Gaussian process, Bayes linear.

MSC2020 subject classifications: Primary 62F15; secondary 62K20, 62G08.

1 Introduction

The use of computer models, or *simulators*, to describe the dynamics of complex physical systems is now commonplace in a wide variety of scientific disciplines. Often such simulators possess high numbers of input and/or output dimensions and, due to their complexity, take a substantial amount of time to evaluate. This presents an immediate challenge, as the responsible use of a simulator (e.g. for model calibration, prediction, decision support, etc.), usually demands Bayesian uncertainty quantification, to capture all major sources of uncertainty, which typically requires a vast number of simulator evaluations. For complex simulators possessing even a modest runtime, this is utterly infeasible. Emulators represent a solution to this problem. An emulator is a statistical construct that seeks to mimic the behaviour of the simulator over its input space, but which is several orders of magnitude faster to evaluate. As the emulator provides both a prediction and an uncertainty statement about the simulator's behaviour at unexplored input locations (an attribute that elevates it above interpolation or other proxy modelling approaches), it can naturally be incorporated in a wider Bayesian uncertainty analysis.

© 2024 International Society for Bayesian Analysis https://doi.org/10.1214/24-BA1456

^{*}Department of Mathematical Sciences, Durham University, Science Laboratories, Durham, DH1 3LE, UK, i.r.vernon@durham.ac.uk; j.owen1@leeds.ac.uk

 $^{^{\}dagger}\text{Research}$ Centre for Fluid and Complex Systems, Coventry University, Coventry, UK, ac8191@ coventry.ac.uk

Emulation with Multiple Partial Discontinuities

Early uses of Gaussian process emulators for computer models were given by Sacks et al. (1989); Currin et al. (1991). For an early example using multilevel emulation combined with structural discrepancy modelling in a Bayesian history matching context see Craig et al. (1997), and for a fully Bayesian calibration of a complex nuclear radiation model, see Kennedy and O'Hagan (2001). Emulators have now been successfully employed across several scientific disciplines, including cosmology (Vernon et al., 2010a,b; Bower et al., 2010; Schneider et al., 2008; Heitmann et al., 2009; Kaufman et al., 2011; Vernon et al., 2014; Rodrigues et al., 2017), climate modelling (Williamson et al., 2013; Johnson et al., 2015; Holden et al., 2016; Edwards et al., 2019, 2021), engineering (Du et al., 2021), epidemiology (Andrianakis et al., 2015, 2017; McKinley et al., 2018; McCreesh et al., 2017; Vernon et al., 2022), systems biology (Vernon et al., 2018; Jackson et al., 2020), oil reservoir modelling (Cumming and Goldstein, 2010, 2009), environmental science (Goldstein et al., 2013), vulcanology (Bayarri et al., 2009; Gu and Berger, 2016; Marshall et al., 2019) and even to Bayesian analysis itself (Vernon and Gosling, 2023). The development of improved emulation strategies therefore has the potential to benefit multiple scientific areas, allowing more accurate analyses with lower computational cost (Higdon et al., 2008).

Most emulator constructions exploit prior judgements about the behaviour of the simulator in terms of its smoothness/differentiability/continuity etc. In this work, however, we are confronted with a problem arising in the TNO Challenge II: a joint industrial and academic challenge posed in the oil industry (see Section 4 for details). A key part of this problem requires the emulation of simulators that are anticipated to be smooth over much of the input space, but that also possess multiple, partial discontinuities at known, non-linear locations. We use the term "partial" in the sense that the locations of the discontinuities begin within the input space, typically ending on the boundary, and hence are not closed, nor do they necessarily bisect the space. Examples of the locations of these discontinuities are shown in Figure 4a (with toy versions in Figures 1 to 3).

A possible way to incorporate discontinuities is to partition the input space into various subregions, and then fit separate, independent emulators in each subregion. For example. Treed GPs (Gramacy and Lee, 2008) which use rectangular, axis aligned subregions, or Pope et al. (2021) who use Voronoi tessellations. Although flexible, these approaches typically require substantial numbers of simulator evaluations, especially in higher dimensions, and critically will not exploit the smoothness around the discontinuity endpoints, which we wish to do here. In addition, many subregions may be required to handle curved discontinuities (especially for Treed GPs). Caiado and Goldstein (2015) use emulators to identify discontinuities caused by tipping points, and then emulate the output separately in each region. This however, is used for discontinuities that bisect the input space, unlike the case here, and the identification of the discontinuities is reported to be time-consuming (see also Ghosh et al., 2018). Deep GPs (see e.g. Dunlop et al., 2018 and references therein), whereby either the correlation lengths or GP inputs are modelled by a second layer GP, with inputs or dependent parameters in turn modelled by the next layer GP etc. have almost unlimited flexibility, but this comes at a cost. requiring substantial numbers of runs to train, whereas for our application run numbers will be extremely limited. Deep GPs also typically have non-analytic uncertainty propagation, which poses problems for full UQ (Sauer et al., 2022). More importantly, even

a deep GP based on smooth layers may fail diagnostics on closer examination, as the impact of the discontinuity will percolate down the layers and still be evident at each level e.g. mimicking rapid (i.e. discontinuous) change in the simulator on the top layer would require rapid (also discontinuous) change of inputs or correlation lengths on the second layer, and so on. Mohammadi et al. (2020) attempted to emulate across simple 1D step functions using a variety of interesting covariance structures with moderate success, although most structures used were either still essentially continuous and hence couldn't fully represent the discontinuity, or induced additional unwanted features.

We instead introduce the TENSE framework, based around carefully designed covariance structures that respect the discontinuities while fully exploiting any smoothness/continuity elsewhere, leading to a single emulator object that can be updated by all runs simultaneously. The layout of the article is as follows. In Section 2 we construct emulators that exhibit partial discontinuities using torn embeddings, before showing how to correct for various induced warpings in Section 3. In Section 4 we apply the TENSE framework to the TNO Challenge II. Example code to reproduce the plots in Sections 2 and 3 can be found at https://github.com/ivernon/TENSE.git.

2 Emulating Computer Models with Partial Discontinuities using Torn Embeddings

2.1 Emulation of Computer Models

We now summarise the standard emulation of computer models approach. We consider a complex computer model represented by a function $f(\mathbf{x})$, where $\mathbf{x} \in \mathcal{X}$ denotes a *d*dimensional vector containing the computer model's input parameters, and $\mathcal{X} \subset \mathbb{R}^d$ is a pre-specified input parameter space of interest. We imagine that due to its complexity, a single evaluation of the computer model will take a substantial amount of time to complete, and due to limited computational resources we will only be able evaluate it at a relatively small number of locations across the input space. Here we assume $f(\mathbf{x})$ is univariate, but the methods we develop should in principle generalise to the multivariate case. Following the Bayesian paradigm, we represent our beliefs about the unknown $f(\mathbf{x})$ at unevaluated input \mathbf{x} via an emulator. A typical approach is to use a pure Gaussian process (GP) for the emulator, such that

$$f(\cdot)|m(\cdot), c(\cdot, \cdot) \sim GP(m(\cdot), c(\cdot, \cdot)), \tag{2.1}$$

for some mean function $m(\cdot)$ and covariance function $c(\cdot, \cdot)$ (Kennedy and O'Hagan, 2001), chosen corresponding to any prior beliefs we hold about the properties of the function $f(\mathbf{x})$. While this form of GP emulator has been successfully employed in a large number of applications, it is sometimes argued that it is the core second-order structure of the GP that is its most important feature, a structure which aligns more closely with our actual beliefs about the behaviour of $f(\mathbf{x})$. The additional distributional assumptions that use of a GP entails, namely that any finite collection of outputs $\{f(\mathbf{x}^{(1)}), \ldots, f(\mathbf{x}^{(n)})\}$ have specifically a multivariate normal distribution is, in some cases, too strong an assumption, which can have unintended consequences.

Emulation with Multiple Partial Discontinuities

Therefore, we often prefer to focus directly on the second-order structure itself, and employ Bayes linear emulators instead of the above GP version. Bayes linear methods follow the foundational work of DeFinetti (De Finetti, 1974) by treating expectation instead of probability as primitive, and respect the subjectivist Bayesian paradigm, but require only a second-order specification (Goldstein, 1999; Goldstein and Wooff, 2007). In this framework, instead of a GP we represent $f(\mathbf{x})$ as a weakly stationary stochastic process. A simple prior specification appropriate for some computer models (see Appendix A (Vernon et al., 2024) for a more complex version) would be to set $E[f(\mathbf{x})] = m(\mathbf{x})$ for some mean function $m(\mathbf{x})$, and to specify the covariance structure as

$$\operatorname{Cov}[f(\mathbf{x}), f(\mathbf{x}')] = \sigma^2 r(\mathbf{x} - \mathbf{x}')$$
(2.2)

where σ^2 represents the prior variance of $f(\mathbf{x})$, and $r(\mathbf{x} - \mathbf{x}')$ defines a stationary correlation structure, of which there are many possible options (see Rasmussen and Williams (2006)). A popular choice for smooth (i.e. infinitely differentiable) functions being the squared exponential:

$$r(\mathbf{x} - \mathbf{x}') = \exp\{-(\mathbf{x} - \mathbf{x}')^T \Sigma^{-1} (\mathbf{x} - \mathbf{x}')\}$$
(2.3)

where Σ is a covariance matrix governing general Mahalanobis distances. Setting $\Sigma = \text{diag}\{\theta, \ldots, \theta\}$, regains the usual isotropic form, where θ is the standard correlation length. Another widely used choice is the Matérn correlation function:

$$r(\mathbf{x} - \mathbf{x}') = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu} \|\mathbf{x} - \mathbf{x}'\|}{\theta}\right)^{\nu} K_{\nu} \left(\frac{\sqrt{2\nu} \|\mathbf{x} - \mathbf{x}'\|}{\theta}\right),$$
(2.4)

where K_{ν} is a modified Bessel function of the second kind and θ and ν are parameters to be specified that govern the correlation length and the derivatives of the computer model respectively (ν rounded up to the next integer gives the number of derivatives that exist).

Given such a second-order specification and a set of model evaluations at locations $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(n)}$, yielding simulator outputs $D = (f(\mathbf{x}^{(1)}), \ldots, f(\mathbf{x}^{(n)}))^T$, we can update our second-order beliefs about $f(\mathbf{x})$ at unevaluated location \mathbf{x} via the Bayes linear adjustment formulae:

$$\mathbf{E}_{D}[f(\mathbf{x})] = \mathbf{E}[f(\mathbf{x})] + \mathbf{Cov}[f(\mathbf{x}), D]\mathbf{Var}[D]^{-1}(D - \mathbf{E}[D])$$
(2.5)

$$\operatorname{Var}_{D}[f(\mathbf{x})] = \operatorname{Var}[f(\mathbf{x})] - \operatorname{Cov}[f(\mathbf{x}), D] \operatorname{Var}[D]^{-1} \operatorname{Cov}[D, f(\mathbf{x})]$$
(2.6)

where $E_D[f(\mathbf{x})]$ and $\operatorname{Var}_D[f(\mathbf{x})]$ are the expectation and variance of $f(\mathbf{x})$ adjusted by D. See Goldstein (1999); Goldstein and Wooff (2007) for details and discussion of the benefits of using a Bayes linear approach, and Vernon et al. (2010a,b, 2018) for the benefits within a computer model setting. The fully specified Bayesian GP based calculation, would of course yield similar update formulae for the analogous posterior mean and variance quantities (conditioned upon various hyperparameters in the definitions of $c(\cdot, \cdot)$ and $m(\cdot)$). While the results derived in this article apply to both the Bayes linear and the fully specified GP emulator frameworks, we will most often refer to the Bayes linear case, as the core arguments concern the second order covariance structure itself, and how we adapt it to the presence of discontinuities. Additionally, for clarity of exposition, we will mainly focus on the standard emulator specification as given by equations (2.2) and (2.3), however see Appendix A for more advanced emulator specifications.

2.2 Emulation Problems caused by Partial Discontinuities

It is worth discussing the specific difficulties that partial discontinuities pose for standard emulators of the form described in the previous section. An example toy computer model that exhibits a partial discontinuity is given by the function:

$$f(\mathbf{x}) \equiv f(x,y) = 0.4\sin(5x) + 0.4\cos(5y) + 0.8(x - 0.75)^2 \sin(y - 1)\mathbb{1}_{\{x>0.75\}}$$
(2.7)

where the two-dimensional $\mathbf{x} = (x, y)^T$ and $\mathbb{1}_A$ is the indicator function that takes value 1 when statement A is true and 0 otherwise, and "sign" just returns the sign of its argument. The form of this function is shown in Figure 1a for the region $\mathcal{X} = \{0 < x < 2, 0 < y < 2\}$. We see that it has a discontinuity across the line y = 1, for x > 0.75, shown as the black horizontal line, and that the discontinuity begins in the interior of \mathcal{X} at the point (x = 0.75, y = 1), and ends on the boundary at (x = 2, y = 1). It is also clear that the function is smooth everywhere else apart from the discontinuity, an attribute that we would wish to exploit in the emulation process.

However, if we naively attempt to apply standard GP or Bayes Linear emulation procedures to f(x, y) they will fail, as they will attempt to smooth over the discontinuity leading to two problems (i) the emulator predictions close to the discontinuity will be highly inaccurate resulting in poor emulator diagnostics, and (ii) the estimation of global emulation parameters (e.g. the correlation lengths θ) may produce strange results that are very sensitive to the design, leading to possible global issues with the emulator. We see that the main problem here is that a discontinuity of this form severely violates the assumption of stationarity and also the common assumption of some form of smoothness/differentiability/continuity implicit in the standard emulator covariance structures. As argued in the Introduction, attempts to alter these assumptions e.g. by breaking stationarity via input dependent correlation lengths or resorting to full deep GPs, do not adequately address this issue as they are still using essentially continuous structures to represent a discontinuity. Our approach in contrast, uses torn embeddings that naturally capture the essence of the discontinuity.

Another approach worth mentioning would be to tinker with the correlation structure of the emulator directly, to reduce the correlation between outputs either side of the discontinuity. For example, one suggestion is to use the geodesic distance between input points in the correlation function, defined such that viable geodesics do not cross the discontinuity (and hence have to go around it). However, this fails as it does not provide a valid covariance structure. This is easy to demonstrate e.g. by using equations (2.2) and (2.3) to construct the 4×4 covariance matrix formed from the four outputs $f(x_A)$, $f(x_B)$, $f(x_C)$, $f(x_D)$ corresponding to the four input points $x_A = (0.5, 1)$, $x_B = (0.75, 1)$, $x_C = (1, 1^+)$, $x_D = (1, 1^-)$, and noting that it is not positive semi-definite (see Appendix B for details). This shows that altering the covariance structure of an emulator to deal with a discontinuity in such ad hoc ways is fraught with danger, even more so for multiple discontinuities of possibly complex, non-linear shape.

Our proposed approach however, guarantees the validity of the emulator's covariance structure, even in the presence of multiple discontinuities of arbitrary shape, while still providing a flexible choice of emulator form, as we now describe.



Figure 1: (a) An example toy 2-dimensional function f(x) with partial discontinuity located along the black horizontal line. (b) The embedding surface v(x, y), torn along the location of the discontinuity. (c) Emulator expectation $E_D[f(x,y)]$ with induced partial discontinuity. (d) Emulator standard deviation $\sqrt{\operatorname{Var}_D[f(x,y)]}$ with induced partial discontinuity (note the horizontal compression for larger x).

2.3 Torn Embedding in a Higher Dimension

The challenge is therefore clear: to develop more sophisticated emulators that exploit regions of smoothness/differentiability/continuity while also respecting the effects of multiple partial discontinuities at known, but possibly non-linear, locations, as seen in the TNO Challenge II. In the interest of clarity, we introduce our approach in terms of a 2-dimensional computer model, but note that the generalisation to higher dimensions is straightforward. To incorporate discontinuities we employ the following procedure:

1. We embed the emulator's 2-dimensional input space $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^2$ into a higher 3-dimensional input space $\mathbf{v}(\mathbf{x}) \in \mathcal{V} \subset \mathbb{R}^3$ using the embedding surface v(x, y) such that we have

$$\mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix}$$
 and $\mathbf{v}(\mathbf{x}) = \begin{pmatrix} x \\ y \\ v(x, y) \end{pmatrix}$. (2.8)

- 2. We tear the otherwise smooth 2-dimensional embedding surface v(x, y) along the known locations of the discontinuities.
- 3. We then set up the emulator as usual using equations (2.2), (2.3), (2.5) and (2.6), but now in the full 3-dimensional space, using the 3-dimensional $\mathbf{v}(\mathbf{x})$ as its input. Specifically, we can design a space-filling collection of runs at locations $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(n)}$ that are embedded in 3-dimensional space as $\mathbf{v}(\mathbf{x}^{(1)}), \ldots, \mathbf{v}(\mathbf{x}^{(n)})$, where the design process can now respect the presence of the discontinuities.
- 4. To evaluate the emulator's expectation and variance at a new point \mathbf{x} we simply evaluate the emulator on the projection of \mathbf{x} onto the embedding surface, that is evaluate $E_D[f(\mathbf{v}(\mathbf{x}))]$ and $\operatorname{Var}_D[f(\mathbf{v}(\mathbf{x}))]$ using equations (2.5) and (2.6).

The tears in the embedding surface v(x, y) will induce a discontinuity, of as yet uncertain size, in the unknown output $f(\mathbf{x})$, and also in our uncertainty statements for $f(\mathbf{x})$, just as we require.

So for example, the covariance structure of the original non-embedded 2-dimensional emulator using the squared exponential covariance function of equation (2.3) was:

$$\operatorname{Cov}[f(\mathbf{x}), f(\mathbf{x}')] = \sigma^{2} \exp\{-(\mathbf{x} - \mathbf{x}')^{T} \Sigma_{2D}^{-1}(\mathbf{x} - \mathbf{x}')\}.$$
(2.9)

After the embedding into 3-dimensions, the covariance becomes simply:

$$\operatorname{Cov}\left[f(\mathbf{v}(\mathbf{x})), f(\mathbf{v}(\mathbf{x}'))\right] = \sigma^{2} \exp\left\{-\left(\mathbf{v}(\mathbf{x}) - \mathbf{v}(\mathbf{x}')\right)^{T} \Sigma_{3D}^{-1}\left(\mathbf{v}(\mathbf{x}) - \mathbf{v}(\mathbf{x}')\right)\right\}$$
(2.10)

i.e. it depends on distances in the new 3-dimensional space via $\mathbf{v}(\mathbf{x}) \in \mathcal{V}$, where Σ_{3D} governs the general 3D Mahalanobis distances. The freedom to choose from various allowable forms for Σ_{3D} will be an important part in the full embedded emulator development as we shall discuss in Section 3. In practical terms, when constructing the emulator using the embedding surface, we simply replace equations (2.2) and (2.3) by equation (2.10).

Toy Example: Returning to the toy model of Figure 1a and equation (2.7), we specify an embedding surface as

$$v(x,y) = -0.4(x - 0.75)^2 \operatorname{sign}(y - 1)\mathbb{1}_{\{x > 0.75\}}$$
(2.11)

which is shown in Figure 1b. The main requirement of the embedding surface at this stage is that it is locally smooth, whilst also being torn along the discontinuity such that the regions above and below the discontinuity are sufficiently different in height in the third dimension in order to decorrelate outputs either side of the discontinuity. Note that v(x, y) does not have to track the form of the actual computer model function

f(x, y) at all: in this example v(x, y) above/below the discontinuity goes low/high while the function f(x, y) does the opposite.

To demonstrate, we design a simple grid of 16 runs $\mathbf{x}^{(i)}$, $i = 1, \ldots, 16$ in the 2dimensional region \mathcal{X} , shown as the black points in Figures 1c and 1d, and raise them into 3 dimensions using $\mathbf{v}^{(i)} = \mathbf{v}(\mathbf{x}^{(i)})$. Note that we choose a grid here as its symmetries help to illustrate the emulator's behaviour. We then emulate in the 3-dimensional space as usual, using equations (2.10), (2.5) and (2.6), with $D = \{f(\mathbf{x}^{(1)}), \dots, f(\mathbf{x}^{(16)})\}$, and using isotropic $\Sigma_{3D} = \text{diag}\{\theta, \theta, \theta\}$, with $\theta = 0.5$ and $\sigma = 0.7$. The emulator expectation $E_D[f(\mathbf{x})] \equiv E_D[f(\mathbf{v}(\mathbf{x}))]$ evaluated across a dense grid of 80×80 points over \mathcal{X} , is shown in Figure 1c. We see that the emulator expectation is smooth away from the discontinuity, but displays a suitable jump across the discontinuity, as desired, hence mimicking the discontinuous behaviour of the real function f(x), given in Figure 1a, reasonably well. Note that we do not claim that this emulator is particularly accurate (especially given the simple grid design), just that it has the desired capability to represent smooth regions combined with partial discontinuities. Individual realisations of f(x) drawn from the emulator, also must have similar smooth/discontinuous behaviour, as shown in Appendix C. The emulator standard deviation $\sqrt{\operatorname{Var}_D[f(\mathbf{x})]} \equiv \sqrt{\operatorname{Var}_D[f(\mathbf{v}(\mathbf{x}))]}$ is shown in Figure 1d, and shows the desired behaviour, in that the further we go along the discontinuity (in the positive x direction) the more uncorrelated the two regions (above and below the discontinuity) become. For example, the point $(1.75, 1^{-})$ just below the discontinuity has a similarly low level of emulator standard deviation as the point (1.75, 0)i.e. a point on the lower boundary. This shows that the emulator at the point $(1.75, 1^{-})$ is just as uninformed as on the lower boundary, and is therefore hardly learning anything from the runs above the discontinuity: it is almost uncorrelated with them, as desired. There will be a more detailed discussion of this point and an examination of the underlying induced 2D correlation structure in Section 4.2.

However, there is a problem: the emulator standard deviation (and expectation) seem compressed slightly, in the x direction, for larger values of x. This issue is more clearly seen in Figure 2 which shows a similar toy model example but now with two discontinuities of different length. Here we have (see Figure 2a):

$$f(x,y) = 0.4\sin(5x) + 0.4\cos(5y) + 1.2\mathbb{1}_{\{x>1\}}(x-1)^2\mathbb{1}_{\{y>1.25\}} - 0.6(x-0.6)^2\mathbb{1}_{\{x>0.6\}}\mathbb{1}_{\{y<0.75\}}.$$

Now we have to use a more complex embedding surface to accommodate the discontinuities of differing length (we postpone discussion of embedding strategies to Section 3.2):

$$v(x,y) = 0.6(x - b(y))^2 \mathbb{1}_{\{x > b(y)\}} \mathbb{1}_{\{y < 1.25\}} \mathbb{1}_{\{y > 0.75\}} - 0.6(x - 0.6)^2 \mathbb{1}_{\{x > 0.6\}} \mathbb{1}_{\{y < 0.75\}}$$

with $b(y) = 0.6 + (1 - 0.6)(y - 0.75)/(1.25 - 0.75)$ (2.12)

where b(y) represents the x coordinate of the line that interpolates the two interior end points (0.6, 0.75) and (1, 1.25) of the discontinuities (see Figure 2b).

Now the emulator standard deviation $\sqrt{\operatorname{Var}_D[f(\mathbf{x})]}$, shown in Figure 2d displays clear compression/warping effects in the middle and lower regions for larger x, which



Figure 2: (a) A 2-dimensional function f(x, y) with two partial discontinuities of differing length. (b) The embedding surface v(x, y). (c) The naive emulator expectation $E_D[f(x, y)]$ (note warping). (d) Naive emulator standard deviation $\sqrt{\operatorname{Var}_D[f(x, y)]}$ (note warping due to the embedded surface v(x, y)). (e) TENSE emulator expectation $E_D[f(x, y)]$ with the warping induced by the use of the embedding surface v(x, y) shown in (b), corrected using Non-Stationary Covariance Structures (compare with the uncorrected version given in (c)). (f) TENSE emulator standard deviation $\sqrt{\operatorname{Var}_D[f(x, y)]}$, again with the warping corrected using Non-Stationary-Covariance Structures (compare with the uncorrected version in (d)).

can be seen to be a direct consequence of the chosen form of v(x, y), as shown in Figure 2b. This compression is a natural consequence of using a stretched embedding surface (that for example does not conserve 2D distances) whilst using a stationary (isotropic) 3-dimensional covariance structure: paths on steep regions of the embedded surface move "too fast" into the 3rd dimension, and lead to an induced compression in 2-dimensions. Equivalently, pairs of points in 2-dimensions end up further apart in 3-dimensions for regions of the embedding v(x, y) that possess large partial derivatives.

However, we really wish to keep the flexibility of stretched embeddings to ensure that we can always create large enough jumps across discontinuities, and to handle more complex cases, for example, discontinuities that begin and end within the space \mathcal{X} , or multiple sets of discontinuities of non-linear form that could be closed, or may even intersect, neither of which could be addressed using say a distance conserving embedding (which notably would mitigate such compression effects, but not entirely remove them). Therefore the compression resulting from use of stretched embeddings represents a serious problem that we will address in the next section.

3 Controlling the Warping Effect of the Embedding

3.1 Reversing the Local Impact of the Embedding

Problem: The use of the stretched embedding surface v(x, y) warps the emulator, compressing the variances and expectations in the examples we have seen, inducing unwanted **x** dependent correlation lengths (and more). This may lead to inefficient emulators and multiple unintended consequences, and will not reflect our actual prior beliefs about the 2-dimensional computer model. Additionally, we want the freedom to choose a wide variety of embedding surfaces v(x, y) without this possibly damaging warping effect occurring.

<u>Solution</u>: We can control this issue using carefully chosen non-stationary covariance structures (NS-CS) defined over the 3-dimensional space.

We now detail a proposed form of the 3-dimensional correlation matrix Σ_{3D} , used in equation (2.10), that is guaranteed to reverse the *local* effect of the embedding, that is for input points close together compared to the curvature of the embedding surface. We discuss how to incorporate this choice across the whole input space using necessarily NS-CS in the next section. We first focus on a reference input point \mathbf{x}_0 , and wish to specify a form for Σ_{3D} that induces the desired squared exponential 2D covariance structure locally around this point, that is such that $\text{Cov}[f(\mathbf{x}), f(\mathbf{x}_0)]$ approximately has the form given by equation (2.9):

$$\operatorname{Cov}[f(\mathbf{x}), f(\mathbf{x}_0)] \simeq \sigma^2 \exp\{-(\mathbf{x} - \mathbf{x}_0)^T \Sigma_{2D}^{-1}(\mathbf{x} - \mathbf{x}_0)\}$$
(3.1)

for inputs \mathbf{x} close to \mathbf{x}_0 . For definiteness we choose the standard isotropic form of

$$\Sigma_{2D} = \begin{pmatrix} \theta^2 & 0\\ 0 & \theta^2 \end{pmatrix}$$
(3.2)

although everything that follows can be applied to general Σ_{2D} by using a simple pretransformation. As the actual covariance structure will be calculated via the embedding $\mathbf{v}(\mathbf{x})$, using equations (2.10) and (3.1) we see that we simply require:

$$\operatorname{Cov}[f(\mathbf{x}), f(\mathbf{x}_0)] \equiv \operatorname{Cov}[f(\mathbf{v}(\mathbf{x})), f(\mathbf{v}(\mathbf{x}_0))] \simeq \sigma^2 \exp\{-(\mathbf{x} - \mathbf{x}_0)^T \Sigma_{2D}^{-1}(\mathbf{x} - \mathbf{x}_0)\} \\ \Leftrightarrow \quad \sigma^2 \exp\{-(\mathbf{v}(\mathbf{x}) - \mathbf{v}(\mathbf{x}_0))^T \Sigma_{3D}^{-1}(\mathbf{v}(\mathbf{x}) - \mathbf{v}(\mathbf{x}_0))\} \simeq \sigma^2 \exp\{-(\mathbf{x} - \mathbf{x}_0)^T \Sigma_{2D}^{-1}(\mathbf{x} - \mathbf{x}_0)\} \\ \Leftrightarrow \qquad \left(\mathbf{v}(\mathbf{x}) - \mathbf{v}(\mathbf{x}_0)\right)^T \Sigma_{3D}^{-1}(\mathbf{v}(\mathbf{x}) - \mathbf{v}(\mathbf{x}_0)) \simeq (\mathbf{x} - \mathbf{x}_0)^T \Sigma_{2D}^{-1}(\mathbf{x} - \mathbf{x}_0). \tag{3.3}$$

We now approximate $\mathbf{v}(\mathbf{x})$ by its linear Taylor expansion around the point \mathbf{x}_0 . This is equivalent to approximating the embedding surface v(x, y) by the tangent plane to v(x, y) at the point \mathbf{x}_0 (we will require the tangent plane below for the construction of \sum_{3D}^{-1}). Hence we approximate:

$$v(x,y) - v(x_0,y_0) = v_x(x-x_0) + v_y(y-y_0) + \mathcal{O}(\mathbf{x}^2)$$
(3.4)

where $v_x = \partial v(x, y)/\partial x$ and $v_y = \partial v(x, y)/\partial y$ are the partial derivatives of v(x, y) evaluated at \mathbf{x}_0 , and $\mathcal{O}(\mathbf{x}^2)$ represents second-order terms and above. Similarly for the vector quantity $\mathbf{v}(\mathbf{x})$, we have that, using equations (2.8) and (3.4):

$$\mathbf{v}(\mathbf{x}) - \mathbf{v}(\mathbf{x}_0) = \begin{pmatrix} x - x_0 \\ y - y_0 \\ v(x, y) - v(x_0, y_0) \end{pmatrix} = \begin{pmatrix} x - x_0 \\ y - y_0 \\ v_x(x - x_0) + v_y(y - y_0) + \mathcal{O}(\mathbf{x}^2) \end{pmatrix}$$
$$= \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ v_x & v_y \end{pmatrix} \begin{pmatrix} x - x_0 \\ y - y_0 \end{pmatrix} + \mathcal{O}(\mathbf{x}^2)$$
$$= A(\mathbf{x} - \mathbf{x}_0) + \mathcal{O}(\mathbf{x}^2), \quad \text{where } A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ v_x & v_y \end{pmatrix}.$$
(3.5)

Replacing this into equation (3.3) and dropping second-order terms and above, we get

$$\Leftrightarrow \quad (\mathbf{x} - \mathbf{x}_0)^T A^T \Sigma_{3D}^{-1} A(\mathbf{x} - \mathbf{x}_0) \simeq (\mathbf{x} - \mathbf{x}_0)^T \Sigma_{2D}^{-1} (\mathbf{x} - \mathbf{x}_0)$$
(3.6)
$$\Leftrightarrow \qquad A^T \Sigma_{3D}^{-1} A \simeq \Sigma_{2D}^{-1}.$$
(3.7)

Hence we see the intuitive result that in order to counter the linear effect of the embedding surface in the vicinity of \mathbf{x}_0 , we just need to choose a form for Σ_{3D} that satisfies equation (3.7), where A represents the linear embedding operator that raises the 2-dimensional position vector \mathbf{x} onto its corresponding location on the 3-dimensional tangent plane given by $A(\mathbf{x} - \mathbf{x}_0)$.

Constructing Σ_{3D}

There are several forms one could choose for Σ_{3D} in order to satisfy equation (3.7), however, many of these will not facilitate sufficient decorrelation of the emulator across discontinuities in the embedding surface v(x, y). We hence choose a form for Σ_{3D} that is aligned with the tangent plane to v(x, y) at the point \mathbf{x}_0 , a form which is specifically selected to provide substantial (possibly maximal) and controllable decorrelation across the discontinuities.

We first set up a relevant orthonormal basis $\{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3\}$. Setting g(x, y, z) = z - v(x, y) and noting that g(x, y, z) = 0 defines the embedding surface z = v(x, y), we see, according to standard vector calculus results, that $\nabla g(x, y, z)$ evaluated at \mathbf{x}_0 gives the vector normal to the embedding surface (and normal to the tangent plane), which we set as the unit vector \mathbf{w}_3 :

$$\mathbf{w}_3 \propto \nabla g(x, y, z) = -v_x \mathbf{e}_x - v_y \mathbf{e}_y + \mathbf{e}_z. \tag{3.8}$$

We choose the unit basis vector \mathbf{w}_1 to lie in the tangent plane, but pointing in the direction of maximally increasing v(x, y). Hence \mathbf{w}_1 has 2-dimensional components parallel to $\nabla v(x, y) = v_x \mathbf{e}_x + v_y \mathbf{e}_y$, and hence has the form

$$\mathbf{w}_1 \propto v_x \mathbf{e}_x + v_y \mathbf{e}_y + \gamma \mathbf{e}_z \tag{3.9}$$

where as \mathbf{w}_1 lies on the tangent plane we have that $\mathbf{w}_1.\mathbf{w}_3 = 0$ which implies that $\gamma = v_x^2 + v_y^2$. The vector \mathbf{w}_2 will be orthogonal to both \mathbf{w}_1 and \mathbf{w}_3 , but as \mathbf{w}_1 was chosen to be in the direction of maximally increasing v(x, y), \mathbf{w}_2 must have zero component in the 3rd dimension and so takes the form:

$$\mathbf{w}_2 \propto \beta \mathbf{e}_x + \delta \mathbf{e}_y. \tag{3.10}$$

Applying the orthogonality relation $\mathbf{w}_2 \cdot \mathbf{w}_1 = 0$ implies $\beta v_x + \delta v_y = 0$ which in turn implies that $\beta = -v_y, \delta = v_x$ are suitable choices, up to an overall normalising constant. To summarise, we have constructed the orthonormal basis $\{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3\}$ given by

$$\mathbf{w}_{1} = \frac{1}{c_{1}} \left[v_{x} \mathbf{e}_{x} + v_{y} \mathbf{e}_{y} + \left(v_{x}^{2} + v_{y}^{2} \right) \mathbf{e}_{z} \right], \quad \text{where } c_{1}^{2} = v_{x}^{2} + v_{y}^{2} + \left(v_{x}^{2} + v_{y}^{2} \right)^{2} \quad (3.11)$$

$$\mathbf{w}_{2} = \frac{1}{c_{2}} [-v_{y}\mathbf{e}_{x} + v_{x}\mathbf{e}_{y}], \qquad \text{where } c_{2}^{2} = v_{x}^{2} + v_{y}^{2} \qquad (3.12)$$

$$\mathbf{w}_3 = \frac{1}{c_3} [-v_x \mathbf{e}_x - v_y \mathbf{e}_y + \mathbf{e}_z], \qquad \text{where } c_3^2 = v_x^2 + v_y^2 + 1 \qquad (3.13)$$

where \mathbf{w}_1 and \mathbf{w}_2 lie on the tangent plane at \mathbf{x}_0 , while \mathbf{w}_3 is orthogonal to the tangent plane.

We postulate that if we specify Σ_{3D} to be diagonal in the above $\{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3\}$ basis, then it will satisfy the desired projection constraint given by equation (3.7). We now show this to be true, subject to some additional conditions. For definiteness, say that Σ_{3D} is indeed diagonal with respect to the $\{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3\}$ basis with corresponding eigenvalues $\{\alpha_1^2, \alpha_2^2, \alpha_3^2\}$, hence we can represent Σ_{3D} as:

$$\Sigma_{3D} = \alpha_1^2 \mathbf{w}_1 \mathbf{w}_1^T + \alpha_2^2 \mathbf{w}_2 \mathbf{w}_2^T + \alpha_3^2 \mathbf{w}_3 \mathbf{w}_3^T$$
(3.14)

and similarly the inverse Σ_{3D}^{-1} as

$$\Sigma_{3D}^{-1} = \frac{1}{\alpha_1^2} \mathbf{w}_1 \mathbf{w}_1^T + \frac{1}{\alpha_2^2} \mathbf{w}_2 \mathbf{w}_2^T + \frac{1}{\alpha_3^2} \mathbf{w}_3 \mathbf{w}_3^T.$$
(3.15)

To evaluate $A^T \Sigma_{3D}^{-1} A$ as required by equation (3.7), we first note that

$$\mathbf{w}_{3}^{T}A = \frac{1}{c_{3}} \begin{pmatrix} -v_{x} & -v_{y} & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ v_{x} & v_{y} \end{pmatrix} = \begin{pmatrix} 0 & 0 \end{pmatrix}$$
(3.16)

and hence we have that

$$A^{T}\Sigma_{3D}^{-1}A = A^{T} \left(\frac{1}{\alpha_{1}^{2}} \mathbf{w}_{1} \mathbf{w}_{1}^{T} + \frac{1}{\alpha_{2}^{2}} \mathbf{w}_{2} \mathbf{w}_{2}^{T} + \frac{1}{\alpha_{3}^{2}} \mathbf{w}_{3} \mathbf{w}_{3}^{T} \right) A = \frac{1}{\alpha_{1}^{2}} A^{T} \mathbf{w}_{1} \mathbf{w}_{1}^{T} A + \frac{1}{\alpha_{2}^{2}} A^{T} \mathbf{w}_{2} \mathbf{w}_{2}^{T} A.$$
(3.17)

We see that $A^T \Sigma_{3D}^{-1} A$ does not depend on α_3 . As will be discussed further below, α_3 is a free parameter, which we can choose to control the extent of the decorrelation of the emulator across the discontinuities, and is one of the motivations for choosing the proposed form of Σ_{3D} given by equation (3.14).

To evaluate the remaining terms in equation (3.17), we have that:

$$\mathbf{w}_{2}^{T}A = \frac{1}{c_{2}} \begin{pmatrix} -v_{y} & v_{x} & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ v_{x} & v_{y} \end{pmatrix} = \frac{1}{c_{2}} \begin{pmatrix} -v_{y} & v_{x} \end{pmatrix}$$
(3.18)

$$\Rightarrow \quad \frac{1}{\alpha_2^2} A^T \mathbf{w}_2 \mathbf{w}_2^T A = \frac{1}{\alpha_2^2 c_2^2} \begin{pmatrix} -v_y \\ v_x \end{pmatrix} \begin{pmatrix} -v_y & v_x \end{pmatrix} = \frac{1}{\alpha_2^2 r^2} \begin{pmatrix} v_y^2 & -v_x v_y \\ -v_x v_y & v_x^2 \end{pmatrix}$$
(3.19)

where we have employed the simplifying notation $r^2 \equiv v_x^2 + v_y^2 = c_2^2$. Similarly, and using $c_1^2 = r^2 + r^4 = r^2(1+r^2)$, we have

$$\mathbf{w}_{1}^{T}A = \frac{1}{c_{1}} \begin{pmatrix} v_{x} & v_{y} & v_{x}^{2} + v_{y}^{2} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ v_{x} & v_{y} \end{pmatrix} = \frac{1}{c_{1}} \begin{pmatrix} v_{x}(1+r^{2}) & v_{y}(1+r^{2}) \end{pmatrix}$$

$$\Rightarrow \quad \frac{1}{\alpha_1^2} A^T \mathbf{w}_1 \mathbf{w}_1^T A = \frac{(1+r^2)^2}{\alpha_1^2 c_1^2} \begin{pmatrix} v_x \\ v_y \end{pmatrix} \begin{pmatrix} v_x & v_y \end{pmatrix} = \frac{(1+r^2)}{\alpha_1^2 r^2} \begin{pmatrix} v_x^2 & v_x v_y \\ v_x v_y & v_y^2 \end{pmatrix}.$$
(3.20)

Combining equations (3.2), (3.17), (3.19) and (3.20), we see that the projection constraint given by equation (3.7) can now be rewritten as

$$\Sigma_{2D} = A^T \Sigma_{3D}^{-1} A \tag{3.21}$$

$$\Leftrightarrow \qquad \begin{pmatrix} \frac{1}{\theta^2} & 0\\ 0 & \frac{1}{\theta^2} \end{pmatrix} = \frac{(1+r^2)}{\alpha_1^2 r^2} \begin{pmatrix} v_x^2 & v_x v_y\\ v_x v_y & v_y^2 \end{pmatrix} + \frac{1}{\alpha_2^2 r^2} \begin{pmatrix} v_y^2 & -v_x v_y\\ -v_x v_y & v_x^2 \end{pmatrix}$$
(3.22)

$$\Leftrightarrow \qquad \begin{pmatrix} \frac{1}{\theta^2} & 0\\ 0 & \frac{1}{\theta^2} \end{pmatrix} = \frac{1}{r^2} \begin{pmatrix} \frac{v_y^r}{\alpha_2^2} + \frac{(1+r^2)v_x^r}{\alpha_1^2} & (\frac{(1+r^2)}{\alpha_1^2} - \frac{1}{\alpha_2^2})v_x v_y\\ (\frac{(1+r^2)}{\alpha_1^2} - \frac{1}{\alpha_2^2})v_x v_y & \frac{v_x^2}{\alpha_2^2} + \frac{(1+r^2)v_y^2}{\alpha_1^2} \end{pmatrix}.$$
(3.23)

Equating the off-diagonal terms gives:

$$\left(\frac{(1+r^2)}{\alpha_1^2} - \frac{1}{\alpha_2^2}\right)v_x v_y = 0 \tag{3.24}$$

Emulation with Multiple Partial Discontinuities

 $\Leftrightarrow \qquad \mathbf{Case 1:} \ \alpha_1^2 = \alpha_2^2 (1+r^2) \quad \text{or} \quad \mathbf{Case 2:} \ v_x = 0 \quad \text{or} \quad \mathbf{Case 3:} \ v_y = 0.$

For **Case 1** we replace $\alpha_1^2 = \alpha_2^2(1+r^2)$ into equation (3.23) giving

$$\begin{pmatrix} \frac{1}{\theta^2} & 0\\ 0 & \frac{1}{\theta^2} \end{pmatrix} = \frac{1}{r^2} \begin{pmatrix} \frac{v_y^2}{\alpha_2^2} + \frac{(1+r^2)v_x^2}{\alpha_2^2(1+r^2)} & 0\\ 0 & \frac{v_x^2}{\alpha_2^2} + \frac{(1+r^2)v_y^2}{\alpha_2^2(1+r^2)} \end{pmatrix} = \begin{pmatrix} \frac{1}{\alpha_2^2} & 0\\ 0 & \frac{1}{\alpha_2^2} \end{pmatrix} \quad \Leftrightarrow \quad \alpha_2^2 = \theta^2.$$

$$(3.25)$$

For **Case 2** we replace $v_x = 0$ (which implies $r^2 = v_y^2$) into equation (3.23) giving

$$\begin{pmatrix} \frac{1}{\theta^2} & 0\\ 0 & \frac{1}{\theta^2} \end{pmatrix} = \begin{pmatrix} \frac{1}{\alpha_2^2} & 0\\ 0 & \frac{(1+r^2)}{\alpha_1^2} \end{pmatrix} \quad \Leftrightarrow \quad \alpha_2^2 = \theta^2 \text{ and } \alpha_1^2 = \alpha_2^2 (1+r^2) \tag{3.26}$$

which is exactly the same result as **Case 1**. **Case 3** gives the same answer also, due to the symmetry between x and y. Therefore, we finally see that the projection requirement given by equation (3.7) is satisfied by specifying the first two eigenvalues α_1^2 and α_2^2 of Σ_{3D} to be

$$\alpha_1^2 = \theta^2 (1+r^2) \text{ and } \alpha_2^2 = \theta^2$$
 (3.27)

with $r^2 = v_x^2 + v_y^2$. This demonstrates that the choice of form of Σ_{3D} as proposed in equation (3.14) is indeed valid. The constraint on the eigenvalues is intuitive from a geometric perspective especially when considering the choice of the basis $\{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3\}$: as \mathbf{w}_2 points along a direction in which the embedding surface v(x, y) is not (locally) increasing, there will be no warping/compression of the emulator along this direction, in which case α_2^2 must equal the desired 2D correlation length of θ^2 . Conversely, \mathbf{w}_1 was defined to point in the direction of maximally increasing v(x, y), and the gradient of v(x, y) in this direction is $|\nabla v(x, y)| = r$ hence α_1^2 must be increased to counteract the warping/compression along this direction that would otherwise be induced by the use of such a stretched embedding surface. Finally, as \mathbf{w}_3 by construction is orthogonal to v(x, y) at \mathbf{x}_0 , and as we are only interested in points that lie on v(x, y), there must be no constraint imposed at this stage on α_3^2 , and hence it will be a free parameter that we can choose or indeed infer.

We need an explicit representation for Σ_{3D} (in the standard Cartesian basis) for use in the non-stationary emulators employed in the next section, and we now have all the pieces required to build this representation, using equations (3.14), (3.27) and the definition of the basis vectors (equations (3.11), (3.12) and (3.13)), as follows. We have

$$\begin{aligned} \alpha_1^2 \mathbf{w}_1 \mathbf{w}_1^T &= \frac{\alpha_1^2}{c_1^2} \begin{pmatrix} v_x \\ v_y \\ v_x^2 + v_y^2 \end{pmatrix} \begin{pmatrix} v_x & v_y & v_x^2 + v_y^2 \end{pmatrix} = \frac{\theta^2}{r^2} \begin{pmatrix} v_x^2 & v_x v_y & v_x r^2 \\ v_x v_y & v_y^2 & v_y r^2 \\ v_x r^2 & v_y r^2 & r^4 \end{pmatrix}, \\ \alpha_2^2 \mathbf{w}_2 \mathbf{w}_2^T &= \frac{\alpha_2^2}{c_2^2} \begin{pmatrix} -v_y \\ v_x \\ 0 \end{pmatrix} \begin{pmatrix} -v_y & v_x & 0 \end{pmatrix} = \frac{\theta^2}{r^2} \begin{pmatrix} v_y^2 & -v_x v_y & 0 \\ -v_x v_y & v_x^2 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \\ \alpha_3^2 \mathbf{w}_3 \mathbf{w}_3^T &= \frac{\alpha_3^2}{c_3^2} \begin{pmatrix} -v_x \\ -v_y \\ 1 \end{pmatrix} \begin{pmatrix} -v_x & v_y & 1 \end{pmatrix} = \frac{\alpha_3^2}{r^2 + 1} \begin{pmatrix} v_x^2 & v_x v_y & -v_x \\ v_x v_y & v_y^2 & -v_y \\ -v_x & -v_y & 1 \end{pmatrix}. \end{aligned}$$

We can hence explicitly construct $\Sigma_{3D} = \alpha_1^2 \mathbf{w}_1 \mathbf{w}_1^T + \alpha_2^2 \mathbf{w}_2 \mathbf{w}_2^T + \alpha_3^2 \mathbf{w}_3 \mathbf{w}_3^T$ giving

$$\Sigma_{3D}(\mathbf{x}_0) = \begin{pmatrix} \theta^2 + \frac{\alpha_3^2 v_x^2}{r^2 + 1} & \frac{\alpha_3^2 v_x v_y}{r^2 + 1} & v_x (\theta^2 - \frac{\alpha_3^2}{r^2 + 1}) \\ \frac{\alpha_3^2 v_x v_y}{r^2 + 1} & \theta^2 + \frac{\alpha_3^2 v_y^2}{r^2 + 1} & v_y (\theta^2 - \frac{\alpha_3^2}{r^2 + 1}) \\ v_x (\theta^2 - \frac{\alpha_3^2}{r^2 + 1}) & v_y (\theta^2 - \frac{\alpha_3^2}{r^2 + 1}) & \theta^2 r^2 + \frac{\alpha_3^2}{r^2 + 1} \end{pmatrix}$$
(3.28)

where we make the dependence on \mathbf{x}_0 explicit. Using this expression for $\Sigma_{3D}(\mathbf{x}_0)$ in the covariance structure of the embedded emulator as given in equation (2.10), will yield for points close to \mathbf{x}_0 , the desired induced covariance structure as represented by Σ_{2D} in equation (3.2). For embeddings with zero curvature, this correction based on linear approximations to the embedding surface v(x, y), is exact.

3.2 Controlling the Global Impact of the Embedding Using Non-Stationary Emulation

The above form of $\Sigma_{3D}(\mathbf{x}_0)$ as given by equation (3.28), will correct for the impact of the embedding surface on the emulator's covariance structure, but only locally around the point \mathbf{x}_0 , as v_x , v_y and r^2 are all evaluated at \mathbf{x}_0 . This is not enough for our needs, as we wish to correct the whole emulator globally over all of \mathcal{X} . Therefore we employ a non-stationary covariance structure as follows. We define an \mathbf{x} dependent covariance matrix $\Sigma_{3D}(\mathbf{x})$ exactly of the form given by equation (3.28), but now evaluated at general point \mathbf{x} . As this covariance matrix $\Sigma_{3D}(\mathbf{x})$ will vary over the input space for general embeddings v(x, y) (except in the trivial case of a linear embedding), we employ the non-stationary apparatus recently used by Dunlop et al. (2018), first derived by Paciorek (2003), in order to define a valid covariance structure.

In the standard non-stationary scenario (i.e. without any embedding surface) Dunlop et al. (2018) use the generalised non-stationary squared exponential covariance function which essentially averages an \mathbf{x} dependent covariance matrix $\Sigma(\mathbf{x})$ as follows, while guaranteeing a valid covariance structure over the whole input space. They define the quadratic form $Q(\mathbf{x}, \mathbf{x}')$ for an \mathbf{x} dependent covariance matrix $\Sigma(\mathbf{x})$ as

$$Q(\mathbf{x}, \mathbf{x}') = (\mathbf{x} - \mathbf{x}')^T \left(\frac{\Sigma(\mathbf{x}) + \Sigma(\mathbf{x}')}{2}\right)^{-1} (\mathbf{x} - \mathbf{x}')$$
(3.29)

and then the corresponding non-stationary squared exponential covariance function for use in the emulator is given, for d-dimensional \mathbf{x} , as

$$\operatorname{Cov}\left[f(\mathbf{x}), f\left(\mathbf{x}'\right)\right] = \sigma^2 \frac{2^{\frac{d}{2}} |\Sigma(\mathbf{x})|^{\frac{1}{4}} |\Sigma(\mathbf{x}')|^{\frac{1}{4}}}{|\Sigma(\mathbf{x}) + \Sigma(\mathbf{x}')|^{\frac{1}{2}}} \exp\left\{-Q\left(\mathbf{x}, \mathbf{x}'\right)\right\}.$$
(3.30)

For our use we simply elevate this non-stationary structure to lie on the embedding surface v(x, y) in the 3D space, hence we instead define the quadratic form via the position vector on the embedding surface $\mathbf{v}(\mathbf{x})$ as

$$Q(\mathbf{v}(\mathbf{x}), \mathbf{v}(\mathbf{x}')) = (\mathbf{v}(\mathbf{x}) - \mathbf{v}(\mathbf{x}'))^T \left(\frac{\Sigma_{3D}(\mathbf{v}(\mathbf{x})) + \Sigma_{3D}(\mathbf{v}(\mathbf{x}'))}{2}\right)^{-1} (\mathbf{v}(\mathbf{x}) - \mathbf{v}(\mathbf{x}')) \quad (3.31)$$

and similarly the corresponding non-stationary squared exponential covariance function in the embedded 3D input space is given as

$$\operatorname{Cov}\left[f(\mathbf{v}(\mathbf{x})), f(\mathbf{v}(\mathbf{x}'))\right] = \sigma^2 \frac{2^{\frac{3}{2}} |\Sigma_{3D}(\mathbf{v}(\mathbf{x}))|^{\frac{1}{4}} |\Sigma_{3D}(\mathbf{v}(\mathbf{x}'))|^{\frac{1}{4}}}{|\Sigma_{3D}(\mathbf{v}(\mathbf{x})) + \Sigma_{3D}(\mathbf{v}(\mathbf{x}'))|^{\frac{1}{2}}} \exp\left\{-Q\left(\mathbf{v}(\mathbf{x}), \mathbf{v}(\mathbf{x}')\right)\right\}$$
(3.32)

where $\Sigma_{3D}(\mathbf{v}(\mathbf{x})) \equiv \Sigma_{3D}(\mathbf{x})$ is given by equation (3.28) with \mathbf{x}_0 replaced by \mathbf{x} . This again guarantees a valid covariance structure throughout both the 3D space and the induced 2D space. Note that this construction generalises to a wide class of covariance structures (Dunlop et al., 2018).

We see that now for *any* pair of input points \mathbf{x} and \mathbf{x}' that are close together relative to the curvature of the embedding surface, the non-stationary covariance structure as given by equation (3.32), which essentially averages the covariance matrices $\Sigma_{3D}(\mathbf{x})$ and $\Sigma_{3D}(\mathbf{x}')$ defined at each of the points, will counteract the local warping effect of the embedding surface, to first order. For pairs of input points that are further apart, non-linear effects may become noticeable, however, for modest choices of correlation length θ these effects will typically be suppressed as the covariance rapidly drops to zero for points that are further apart than the correlation length. Therefore, an emulator constructed using the non-stationary covariance structure given by equation (3.32) will a) allow us the freedom to choose from a wide class of torn embedding surfaces v(x, y)to handle unlimited numbers of discontinuities of complex configuration and to ensure that the emulator is decorrelated across them, as discussed in Section 2.3, and b) will approximately induce the desired stationary 2D covariance structure across local regions that do not contain discontinuities. We refer to this general framework as the Torn Embedding Non-Stationary Emulation (TENSE) approach.

In Figures 2e and 2f we apply the TENSE approach to the toy model discussed in Section 2.3. Comparing with the uncorrected version, as seen in Figures 2c and 2d, we see that the emulator standard deviation $\sqrt{\operatorname{Var}_D[f(x)]}$ now displays no noticeable warping effects and maintains the symmetry we would expect around each of the run locations (the black points) especially in the top, middle and lower regions for large x, while also displaying suitable uncorrelated behaviour across the discontinuities. The emulator expectation also looks far more reasonable, displaying no noticeable warping, as desired. See Appendix J for more detail where we present emulator diagnostics defined over various test sets and a further examination of the induced TENSE covariance structure.

We see that the embedding surface v(x, y), here given by equation (2.12), is clearly important for the TENSE emulator and must be constructed to satisfy the following principles: a) there must be a large enough tear along each discontinuity to ensure suitable decorrelation of the emulator, but b) the curvature of v(x, y) must not be too large in order to avoid unwanted non-linear artefacts in the induced covariance structure. We have utilised quadratic surfaces to achieve these goals noting that they are continuous in both v(x, y) and its first derivatives across the transition from the flat region of v(x, y) = 0 to the quadratic region, and smooth (i.e. infinitely differentiable) elsewhere, however many other embedding surfaces would be suitable. See Appendix K for more details, where we show the effects on the emulator of different embedding choices and discuss the impact of the decorrelation parameter α_3 . We leave the interesting problem of inferring the embedding surface to future work.

Although we demonstrate this framework in 2D/3D and for squared exponential covariances, it is simple to extend in various ways. For example, the above calculations extend to any covariance structure of the form r(a) where a is the general Mahalanobis distance between \mathbf{x} and \mathbf{x}' and r(.) is a valid covariance function, e.g. the Matérn (Rasmussen and Williams, 2006), using the general form for equation (3.32) (Dunlop et al., 2018). If one desires a non-stationary induced 2D covariance structure we can achieve this using a similar strategy by inserting a point-wise 2D pre-transformation. Similarly this torn embedding strategy can be extended to higher dimensional input spaces with more complex discontinuities, e.g. a d-dimensional input space containing discontinuities residing on d-1 dimensional hypersurfaces would be embedded in a d+1 dimensional space. Note that more complex networks of m discontinuities may require embedding in a higher dimensional space, e.g. of dimension d+m, which would provide more freedom to decorrelate the emulator and therefore to avoid unwanted effects due to neighbouring discontinuities. However, as we would still be operating on a d-dimensional surface, emulator accuracy may not be penalised too severely by the use of m extra dimensions.

3.3 Emulating with Discontinuities on Non-Linear Locations

An attractive feature of the Torn Embedding Non-Stationary Emulation (TENSE) approach is that it can be applied to a broad class of discontinuities, for example, when the discontinuities are situated on non-linear locations. An example of this is provided by the function f(x, y), shown in Figure 3a (see Appendix D for the full definition). A suitable embedding surface v(x, y) is shown in Figure 3b. Note again the difference in form between v(x, y) and f(x, y): e.g. in the top/bottom regions v(x, y) is flat while f(x, y) tends downwards, and in the right/left regions v(x, y) tends downwards/upwards respectively whilst f(x, y) tends upwards in each case.

TENSE emulator expectation $E_D[f(x, y)]$ and standard deviation The $\sqrt{\operatorname{Var}_D[f(x,y)]}$ with v(x,y) induced warping corrected, are shown in Figures 3c and 3d respectively, based on a 16 point grid design given as the black points. Comparing Figure 3c with 3a we see that the emulator expectation captures the form of f(x, y) well, and handles the curved discontinuities with ease. We note that one could apply the Treed GP method (Gramacy and Lee, 2008) here, that divides the input space up by partitioning on individual inputs, effectively creating rectangular subregions in which independent GPs are trained. However, although this method may learn the locations of the discontinuities, it may perform poorly here, as it is very inefficient to represent curved discontinuities using rectangular regions, and many more runs may be required to train the independent GPs, instead of the single emulator used in the TENSE approach. Similarly, we note that a brief investigation showed that even using Treed GPs for the simple examples in Figures 1 and 2, which possess linear discontinuity locations. required far more runs in order to detect the existence of the discontinuities (36 and 64 runs respectively), than we used here (16 and 12 runs respectively). This is not unexpected and demonstrates that learning about even simple discontinuities without any prior information is expensive.



Figure 3: (a) A 2-dimensional function f(x, y) with curved discontinuity locations given by the curved black lines. (b) The embedding surface v(x, y). (c) The TENSE emulator expectation $E_D[f(x, y)]$ with v(x, y) warping corrected using NS-CS. (d) The TENSE emulator variance $\operatorname{Var}_D[f(x, y)]$ with v(x, y) warping corrected using NS-CS.

4 Application: TNO 2 Well Placement Challenge

4.1 Problem Setup: Multiple Partial Discontinuities

The motivation for developing the TENSE framework is in direct response to the following problem posed within the oil industry. The TNO OLYMPUS Field Development Optimisation Challenge was devised by the Netherlands Organisation for Applied Scientific Research (TNO) in collaboration with Delft University of Technology (TU Delft), and industrial partners Eni S.p.A, Equinor ASA and Petrobras. The TNO challenge is based around the fictitious oil reservoir model named OLYMPUS (TNO, 2017), and was designed to mimic realistic simulation, optimisation and decision problems faced by the oil industry. It has attracted much attention from industry and academia with results from the active competition period presented and compared at the EAGE/TNO Workshop on OLYMPUS Field Development Optimization (TNO, 2018).



(a) Olympus reservoir in physical coordinates. (b) Olympus re-

(b) Olympus reservoir in transformed coordinates.

Figure 4: The TNO II Challenge Olympus oil reservoir model. (a) An image in physical coordinates, with the blue areas representing higher oil volume per unit area. The black lines show the locations of curved geological faults that will cause discontinuities in the Net Present Value (NPV) surface defined over the 2D map. (b) In transformed grid aligned coordinates, showing the oil volume per unit area of one of the 50 geological realisations. Note that the geological faults (black horizontal lines) are now straight: this is not required for the TENSE methodology, but is useful and worth exploiting. The non-oil containing region is coloured blue.

The TNO Challenge I concerns well control, however the TNO Challenge II, which we exclusively focus on here, concerns well placement. The challenge is to choose a configuration of oil well placement to optimise the Net Present Value (NPV) over a 20 year period for the OLYMPUS reservoir model. NPV essentially represents the discounted profits over the 20 year period. As the reservoir model, used to calculate the NPV, has complex features including geological uncertainty and is expensive to evaluate, and as multiple wells may be used, this represents a demanding task. Figure 4a shows an image from above of the Olympus reservoir in physical coordinates, coloured by oil volume per unit area. We can choose to locate production wells or injection wells at any location over this 2D map, with each configuration yielding a certain NPV value. Note however in Figure 4a the black lines extending into the map from the northern edge: these are geological faults in the model, with know location, that will inhibit the flow of oil and water across them. This will induce a sharp discontinuity in the NPV response as the possible well is moved either side of the fault. Away from such faults, we anticipate the NPV surface to be far smoother.

In the Olympus model the location of the faults is fixed and known, however many other geological aspects (e.g. the permeability/porosity fields) are treated as uncertain and represented via 50 geological realisations provided by the TNO consortium, derived from an underlying geology model which was not made freely available. An example of one of the geological realisations is given in Figure 4b, coloured by the oil volume per unit area, and more realisations are given in Figure 8, Appendix F, along with additional

plots of the mean and SD of the oil volume per unit area of the 50 realisations. In these plots the physical 2D coordinates have been transformed into grid aligned 2D coordinates, which has the added effect of transforming the faults so that they lie along constant horizontal (black) lines. As demonstrated in Section 3.3, the TENSE approach does not require linear discontinuity locations, but this transformation, available due to the way the OLYMPUS model was constructed, simplifies subsequent specifications e.g. for the embedding surface, hence it would be remiss of us not to exploit it here.

The precise remit of the TNO Olympus Challenge II is to choose well locations to optimise the mean NPV over the 50 geological realisations (each geological realisation will generate its own NPV). For example, for a single vertical producer well located at position $\mathbf{x} = (x, y)$, we could evaluate the NPV for any \mathbf{x} and for any of the $i = 1, \ldots, 50$ geological realisations, giving output $NPV^{(i)}(\mathbf{x})$. We hence define our primary deterministic computer model of interest $f(\mathbf{x})$ to be the mean over 50 realisations for a single producer well located at \mathbf{x} , in accordance with the challenge:

$$f(\mathbf{x}) \equiv \overline{NPV}(\mathbf{x}) = \frac{1}{50} \sum_{i=1}^{50} NPV^{(i)}(\mathbf{x})$$
(4.1)

(see Appendix E for details). Obviously there are several uncertainties and features that one might want to include in a more detailed analysis, that are missing from the original TNO Challenge. These include the effects of the finite sample size of geological realisations, uncertainties due to oil price and water cost, model discrepancy due to the imperfection of the reservoir (and geology) model itself, the benefits of sequential decision making, and indeed whether the NPV should even be identified with the utility of the decision makers. See Owen et al. (2020) for discussion of several of these issues, and also House et al. (2009) for a relevant treatment of exchangeable computer models. However, here we are primarily interested in the following emulation problem.

Concern has been expressed in the oil industry over the transparency of various black-box optimisers that can produce counterintuitive well configurations of unfamiliar form (and of unknown optimality), that made some engineers nervous. We were hence approached and asked if we could help visualise the NPV surface, to aid interpretation and insight in various situations that may occur within a more human informed optimisation process. Specifically a key request was to visualise the mean NPV surface for a single producer well, as represented by $f(\mathbf{x})$, over the full reservoir map \mathcal{X} in the presence of multiple discontinuities, using only a limited set of evaluations of the expensive OLYMPUS model. This was the original motivation for developing TENSE.

4.2 Constructing the Embedding Surface v(x, y)

We proceed to apply the TENSE framework to the deterministic function $f(\mathbf{x})$ representing the mean NPV of a single producer well as follows. We specify an embedding surface $v(\mathbf{x}) = v(x, y)$ by tearing along the five discontinuities shown in Figure 4b and bending alternate regions higher and lower into the 3D space using quadratic forms, exploiting a similar strategy to that employed in Section 2.3. The embedding surface is shown in Figure 9 (Appendix G) with the full definition also given in Appendix G.

To check that this choice of embedding will produce the desired behaviour of allowing the emulator to exhibit discontinuous jumps over the discontinuities with minimal warping, we examine the induced covariance structure of vertical lines of points that cross all the discontinuities, as shown in Figure 5. For example, Figure 5a shows a zoomed in section of the embedding surface v(x, y) with the discontinuities as horizontal black lines (as in Figure 4b), but also highlights a green vertical line of points at x = 42, while Figure 5b shows the induced 2D emulator correlation matrix corresponding to this green line of points. The correlation matrix is formed from $Cov[f(\mathbf{v}(\mathbf{x})), f(\mathbf{v}(\mathbf{x}'))]$ using equations (3.32), (3.31) and (3.28).

We see that the two regions y > 123.5 and y < 123.5 either side of the highest fault are uncorrelated as desired, and that the correlation structure resorts to the usual squared exponential form within each region. Figures 5c and 5d are defined similarly, but for the line x = 78. Now we see that the regions either side of the fault at y = 85.5are almost entirely uncorrelated, while either side of the fault at y = 99.5 the regions have suppressed correlation, as the start of the fault is relatively close to the green line. In Figures 5e and 5f the more extreme case of x = 116 is examined, where we see six uncorrelated regions separated by the five faults, precisely as desired. Due to the TENSE approach of embedding in a higher dimension, all these correlation matrices are guaranteed to be valid. Note that we choose to directly specify the form of the embedding surface v(x,y) here, as it is feasible to do this in a controlled way as to ensure each region either side of a discontinuity is well separated in the third dimension. One could of course treat v(x,y) as uncertain, possibly of parameterised form but still torn along the locations of the discontinuities, and then use the run data to learn about v(x, y). However, this may lead to several identifiability issues, and there may not be a strong signal as to the particular form for v(x, y), so we leave such considerations to future work. See also Appendices J and K for further discussions and analysis regarding the covariance structures typically induced by TENSE.

4.3 Emulating the Net Present Value Surface

Having defined the embedding surface v(x, y), we are now able to construct an emulator for the NPV output as represented by $f(\mathbf{x})$, corresponding to a single producer well at location $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^2$, in the presence of the discontinuities caused by the geological faults. However, there is additional prior information about the Olympus model that we can include. We know, without performing any model evaluations, that if a well is placed outside of the oil containing region of the reservoir, there will be no oil production and the NPV will be zero (or a small negative value). For linear boundaries, one can in fact incorporate known model behaviour on the boundary, within the emulator analytically in any dimension (see for example Vernon et al., 2019 and Jackson and Vernon, 2023). However, here the boundary around the edge of the oil containing region is complex, and so we simply add a set of 36 "ghost points" just outside the oil containing region, with the NPV value of each set to $f(\mathbf{x}) = 0$. The effect of this prior information is shown in Figure 6a which gives the prior emulator expectation $\mathbb{E}[f(\mathbf{x})]$ over \mathcal{X} , and shows the ghost points as red points located within the grey non-oil region.



Figure 5: Left panels (a), (c) and (e): the torn surface v(x, y) embedded in 3D used to induce the discontinuities along the five geological faults, shown as the horizontal black lines, in the Olympus model. Right panels (b), (d) and (f): the induced emulator correlation matrix of the set of points along the green vertical line highlighted in the corresponding left panel, at locations x = 42,78 and 116 respectively. See Figure 9, Appendix G, for a 3D image of the full torn embedded surface v(x, y).

The initial space filling set of (wave 1) runs was designed respecting the following considerations. The Olympus model is computationally intensive and our collaborator was uncertain as to how much (cloud) computational resources would be available, implying early termination of the design was possible. We therefore constructed the design one point at a time, with each point $x^{(i)}$ chosen to minimise the current emulator variance $\operatorname{Var}_{D^{(i)}}[f(x)]$ integrated over the oil reservoir region \mathcal{X} , given $x^{(i)}$ and the previous design points, where $D^{(i)} = \{f(x^{(1)}), \dots, f(x^{(i-1)}), f(x^{(i)})\}$. This criterion was calculated approximately using a large test set of regularly spaced input points spanning \mathcal{X} . from which $x^{(i)}$ was also selected. The test set was composed of a fine grid of points with points removed that did not lie within the oil containing region \mathcal{X} , shown in Figure 4b. As this calculation uses the emulator's correlation structure, it respects the discontinuities and specifically the low correlation between certain regions as highlighted in Figure 5. In addition, due to the sequential nature of the design construction, even early termination would result in a well spaced and informative set of runs. Some pragmatic choices were used in the design calculation e.g. within an isotropic Σ_{2D} we specified a fixed 2D correlation length of $\theta = 12$, a judgement informed by the local correlation seen in the oil volume per unit area of the geological realisations (Figure 4b). We also employed a nearest neighbour approximation in the emulator variance calculation, to greatly improve efficiency: this exploited the smaller correlation lengths, which imply that candidate points would mainly affect and be affected by the runs in their local vicinity. Finally, we added three pairs of points to the design, either side of three of the major faults to give more direct information regarding the discontinuities in those regions. The resulting 47 point wave 1 design is shown in Figure 6b as the green points. It displays good space filling properties, while adequately exploring each of the uncorrelated regions in between the faults. At each of the 47 points $\mathbf{x}^{(i)}$, all 50 of the geological realisations were evaluated giving $NPV^{(j)}(\mathbf{x}^{(i)}), j = 1, ..., 50$, and the mean calculated, giving $f(\mathbf{x}^{(i)})$ and hence the first batch of runs, denoted $D_1 = \{f(\mathbf{x}^{(1)}), \dots, f(\mathbf{x}^{(47)})\},\$ for use in the emulator equations.

The TENSE framework was then applied to D_1 , employing the embedding surface v(x, y) given in Appendix G, using equations (3.32), (3.28), (2.5) and (2.6), with details given in Appendix H. The resulting emulator expectation $E_{D_1}[f(\mathbf{x})]$ adjusted by the model evaluations D_1 is shown in Figure 6b as the coloured contours. We see that the emulator incorporates jumps in $f(\mathbf{x})$ due to the discontinuities caused by the faults, while remaining smooth in all other parts of the space \mathcal{X} , as desired. In addition, a clear visualisation of the (expected) NPV surface across the oil reservoir is obtained, and the regions of suspected high NPV identified for further investigation.

Our primary goal is to visualise this surface, and to identify and examine in more detail regions of higher NPV for consideration by the relevant expert/decision maker. We hence use an upper credible interval (UCI) approach to define a region \mathcal{X}_1 of possibly high NPV, worthy of further investigation, as

$$\mathcal{X}_{1} \equiv \left\{ x \in \mathcal{X} \mid \mathrm{E}_{D}[f(\mathbf{x})] + c\sqrt{\mathrm{Var}_{D}[f(\mathbf{x})]} > f^{+} - \delta \right\}$$
(4.2)

where f^+ is the highest NPV seen so far, δ is a tolerance based on uncertainties in the decision process itself (Owen et al., 2020) and on our desire to explore the region



Figure 6: The output of the TENSE emulator as applied to the TNO Challenge II Olympus reservoir model. (a) The prior emulator expectation $E[f(\mathbf{x})]$ trained only on the ghost runs (red points) located in the non-oil producing (grey) regions. (b) Wave 1 TENSE emulator expectation $E_{D_1}[f(\mathbf{x})]$ trained on 47 wave 1 runs denoted D_1 (green points) in addition to the ghost points. Plots (a) and (b) share the same key. (c) The wave 2 TENSE emulator expectation $E_{D_1 \cup D_2}[f(\mathbf{x})]$ trained on an additional set of 48 wave 2 runs denoted D_2 . (d) The wave 2 TENSE emulator expectation of panel (c) now with the high oil production regions highlighted. Plots (c) and (d) share the same key. In all panels the horizontal black lines show the location of the geological faults which induce the discontinuities.

of high NPV and not just to identify a single possibly non-robust maxima, and c is typically chosen to be 3 based on Pukelsheim's 95% 3-sigma rule for arbitrary unimodal distributions (Pukelsheim, 1994). Figures 11a and 11b (in Appendix I) show respec-

tively the prior UCI: $E[f(\mathbf{x})] + c\sqrt{\operatorname{Var}[f(\mathbf{x})]}$, and also the wave 1 UCI adjusted by D_1 : $E_{D_1}[f(\mathbf{x})] + c\sqrt{\operatorname{Var}_{D_1}[f(\mathbf{x})]}$.

Following a general history matching strategy (Vernon et al., 2010a), we proceeded by designing a second wave of space filling runs over the \mathcal{X}_1 region, now chosen to minimise the mean emulator variance over \mathcal{X}_1 only. This design is shown as the light blue points in Figure 6c. Evaluation of this design using the Olympus model creates a second vector of model outputs denoted D_2 . We can then adjust the TENSE emulators by $D_1 \cup D_2$, giving the emulator expectation $E_{D_1 \cup D_2}[f(\mathbf{x})]$ for the mean NPV, were a producer well to be placed at location \mathbf{x} , which is shown in Figure 6c as the coloured contours. We now have a detailed representation of the high NPV areas of the \mathcal{X} map, naturally incorporating the fault discontinuities. Examination of the UCI after the wave 2 runs (Appendix I, Figure 11c) shows that there is little to be learned about this region by performing further runs. Figure 6d highlights the high NPV region in question, with the solid contours corresponding to thresholds of $f^+ - \delta = 2.2 \times 10^7$ and 2.3×10^7 respectively. This achieves our objective of locating and visualising the high NPV areas for the position of a single producer well. The TENSE framework can be employed for several further types of analysis e.g. for quantile emulation to examine the uncertainties in the NPV induced by the unknown geology, which we demonstrate in Appendix I.

4.4 Extensions and Generalisations

These initial investigations of the TNO challenge using TENSE can be extended in multiple ways. The full problem of optimising the joint location of multiple producer and injector wells is of course the long-term goal. This is a very challenging problem especially when combined with an appropriate level of uncertainty quantification (Owen et al., 2020). However, due to the localised structure of oil reservoirs, for early/medium times, often small groups of wells (e.g. one injector combined with two producers) are optimised on particular sub-regions of the map, to break the full problem into smaller, tractable pieces. The TENSE framework can be directly extended to such cases. For the example of three wells, a 6-dimensional problem, we would need to employ a torn embedding in a 9-dimensional space to account for the discontinuity effects on each of the three wells. This 9-dimensional space would look like the direct product of three versions of the 3-dimensional space used here in the single well example.

In principle the TENSE approach can be generalised to far more wells that just three, however, constructing an accurate emulator over the full input space for larger numbers may require infeasible numbers of runs (and we would waste a lot of runs exploring the low NPV parts of the space). So a more targeted approach, optimising sets of three wells, combining them and then employing a final wave or two of optimisation on the full set of wells, may be a sensible strategy. We leave this, and the various associated design strategies, to future work.

5 Conclusion and Future Plans

We have introduced the Torn Embedding Non-Stationary Emulation (TENSE) approach for emulating expensive functions that possess partial discontinuities of known location and general non-linear form, which possibly begin and/or end within the input space of interest. This method utilises a torn embedding surface to induce the required discontinuities, combined with a carefully chosen non-stationary covariance structure over the embedding space, to correct for the local impact of the use of the non-linear embedding. While we have introduced this in the context of a squared exponential covariance structure in 2D/3D, it can be applied to a wide class of covariance structures and emulator forms, and in principle, extended into higher dimensions. We demonstrated this approach on various example functions, and then applied it to the realistic OLYMPUS reservoir model, showing how it facilitated the design of model evaluations and the construction of appropriate emulators to visualise the NPV surface, both of which respected the presence of the discontinuities. It was also applied to quantile emulation, and the extension to multiple wells and higher dimensions was discussed.

There are many possible extensions of this methodology. While we have employed fixed embedding surfaces v(x, y) chosen to suitably decorrelate outputs either side of the discontinuities, one could instead use the TENSE framework to learn about such surfaces to find more accurate embeddings. This could be combined with methods to learn about the location of the discontinuities themselves to provide a more complete analysis, in a fully Bayesian framework. The extension to more complex networks of discontinuities is also very interesting, as it may require embedding in higher dimensional spaces to provide the necessary freedom to ensure sufficient decorrelation across all discontinuities, especially those that intersect, resulting in a challenging embedding problem.

Funding

I. V. gratefully acknowledges UKRI (EP/W011956/1) and Wellcome (218261/Z/19/Z) funding, and an AWE William Penney Fellowship. J. O. gratefully acknowledges EPSRC iCase Studentship (Smith Institute) funding. We thank Rock Flow Dynamics for use of the tNavigator simulator.

Supplementary Material

Appendices (DOI: 10.1214/24-BA1456SUPP; .pdf). The Appendices A–K referred to in the main article.

References

Andrianakis, I., McCreesh, N., Vernon, I., McKinley, T., Oakley, J., Nsubuga, R., Goldstein, M., and White, R. (2017). "Efficient history matching of a high dimensional individual based HIV transmission model." *SIAM/ASA Journal of Uncertainty Quantification*, 5(1): 694–719. MR3681376. doi: https://doi.org/10.1137/16M1093008.

- Andrianakis, I., Vernon, I., McCreesh, N., McKinley, T., Oakley, J., Nsubuga, R., Goldstein, M., and White, R. (2015). "Bayesian history matching of complex infectious disease models using emulation: a tutorial and a case study on HIV in Uganda." *PLoS Computational Biology*, 11(1): e1003968. 2
- Bayarri, M. J., Berger, J. O., Calder, E. S., Dalbey, K., Lunagomez, S., Patra, A. K., Pitman, E. B., Spiller, E. T., and Wolpert, R. L. (2009). "Using statistical and computer models to quantify volcanic hazards." *Technometrics*, 51(4): 402–413. MR2756476. doi: https://doi.org/10.1198/TECH.2009.08018.
- Bower, R. G., Vernon, I., Goldstein, M., Benson, A. J., Lacey, C. G., Baugh, C. M., Cole, S., and Frenk, C. S. (2010). "The parameter space of galaxy formation." *Monthly Notices of the Royal Astronomical Society*, 96(454): 717–729.
- Caiado, C. and Goldstein, M. (2015). "Bayesian uncertainty analysis for complex physical systems modelled by computer simulators with applications to tipping points." *Communications in Nonlinear Science and Numerical Simulation*, 26(1–3): 123–136. 2
- Craig, P. S., Goldstein, M., Scheult, A. H., and Smith, J. A. (1997). "Pressure matching for hydrocarbon reservoirs: a case study in the use of Bayes linear strategies for large computer experiments (with discussion)." In Gatsonis, C., Hodges, J. S., Kass, R. E., McCulloch, R., Rossi, P., and Singpurwalla, N. D. (eds.), *Case Studies in Bayesian Statistics*, volume 3, 36–93. New York: Springer-Verlag. MR1425400. 2
- Cumming, J. A. and Goldstein, M. (2009). "Small sample Bayesian designs for complex high-dimensional models based on information gained using fast approximations." *Technometrics*, 51(4): 377–388. MR2756474. doi: https://doi.org/10.1198/TECH. 2009.08015.
- Cumming, J. A. and Goldstein, M. (2010). "Bayes linear uncertainty analysis for oil reservoirs based on multiscale computer experiments." In O'Hagan, A. and West, M. (eds.), *The Oxford Handbook of Applied Bayesian Analysis*, 241–270. Oxford, UK: Oxford University Press. MR2790346. 2
- Currin, C., Mitchell, T., Morris, M., and Ylvisaker, D. (1991). "Bayesian prediction of deterministic functions with applications to the design and analysis of computer experiments." *Journal of the American Statistical Association*, 86(416): 953–963. MR1146343. 2
- De Finetti, B. (1974). Theory of Probability, volume 1. London: Wiley. MR0440640. 4
- Du, H., Sun, W., Goldstein, M., and Harrison, G. (2021). "Optimization via statistical emulation and uncertainty quantification: hosting capacity analysis of distribution networks." *IEEE Access*, 9: 118472–118483. 2
- Dunlop, M. M., Girolami, M. A., Stewart, A. M., and Teckentrup, A. L. (2018). "How deep are deep Gaussian processes?" *Journal of Machine Learning Research*, 19(54): 1–46. MR3874162. 2, 15, 16, 17
- Edwards, T. L., Brandon, M. A., Durand, G., Edwards, N. R., Golledge, N. R., Holden,

P. B., Nias, I. J., Payne, A. J., Ritz, C., and Wernecke, A. (2019). "Revisiting Antarctic ice loss due to marine ice-cliff instability." *Nature*, 566(7742): 58–64. 2

- Edwards, T. L., Nowicki, S., and et.al. (2021). "Projected land ice contributions to twenty-first-century sea level rise." *Nature*, 593(7857): 74–82. 2
- Ghosh, S., Gavaghan, D. J., and Mirams, G. R. (2018). "Gaussian process emulation for discontinuous response surfaces with applications for cardiac electrophysiology models." arXiv: Computation. 2
- Goldstein, M. (1999). "Bayes linear analysis." In Kotz, S. et al. (eds.), Encyclopaedia of Statistical Sciences, 29–34. Wiley. 4
- Goldstein, M., Scheult, A., and Vernon, I. (2013). "Assessing Model Adequacy." In Environmental Modelling: Finding Simplicity in Complexity, Chichester, UK: John Wiley & Sons, Ltd, second edition. 2
- Goldstein, M. and Wooff, D. A. (2007). Bayes Linear Statistics: Theory and Methods. Chichester: Wiley. MR2335584. doi: https://doi.org/10.1002/9780470065662. 4
- Gramacy, R. B. and Lee, H. K. H. (2008). "Bayesian treed Gaussian process models with an application to computer modeling." *Journal of the American Statistical Association*, 103(483): 1119–1130. MR2528830. doi: https://doi.org/10.1198/016214508000000689. 2, 17
- Gu, M. and Berger, J. O. (2016). "Parallel partial Gaussian process emulation for computer models with massive output." Annals of Applied Statistics, 10(3): 1317–1347. MR3553226. doi: https://doi.org/10.1214/16-AOAS934.
- Heitmann, K., Higdon, D., et al. (2009). "The coyote universe II: Cosmological models and precision emulation of the nonlinear matter power spectrum." *The Astrophysical Journal*, 705(1): 156–174.
- Higdon, D., Gattiker, J., Williams, B., and Rightley, M. (2008). "Computer model calibration using high-dimensional output." *Journal of the American Statisti*cal Association, 103(482): 570–583. MR2523994. doi: https://doi.org/10.1198/ 016214507000000888. 2
- Holden, P. B., Edwards, N. R., Hensman, J., and Wilkinson, R. D. (2016). ABC for climate: dealing with expensive simulators. arXiv:1511.03475: Handbook of Approximate Bayesian Computation (ABC). MR3889296. 2
- House, L., Goldstein, M., and Vernon, I. (2009). "Second order exchangeable computer models." Technical Report 10/02, MUCM. 20
- Jackson, S. E. and Vernon, I. (2023). "Efficient emulation of computer models utilising multiple known boundaries of differing dimension." *Bayesian Analysis*, 18(1): 165–191. MR4515729. doi: https://doi.org/10.1214/22-ba1304. 21
- Jackson, S. E., Vernon, I., Liu, J., and Lindsey, K. (2020). "Understanding hormonal crosstalk in Arabidopsis root development via emulation and history matching." *Statistical Applications in Genetics and Molecular Biology*, 19(2): 20180053. MR4126698. doi: https://doi.org/10.1515/sagmb-2018-0053.

- Johnson, J. S., Cui, Z., Lee, L. A., Gosling, J. P., Blyth, A. M., and Carslaw, K. S. (2015). "Evaluating uncertainty in convective cloud microphysics using statistical emulation." *Journal of Advances in Modeling Earth Systems*, 7(1): 162–187.
- Kaufman, C. G., Bingham, D., Habib, S., Heitmann, K., and Frieman, J. A. (2011).
 "Efficient emulators of computer experiments using compactly supported correlation functions, with an application to cosmology." *The Annals of Applied Statistics*, 5(4): 2470–2492. MR2907123. doi: https://doi.org/10.1214/11-AOAS489.
- Kennedy, M. C. and O'Hagan, A. (2001). "Bayesian calibration of computer models." Journal of the Royal Statistical Society, Series B, 63(3): 425–464. MR1858398. doi: https://doi.org/10.1111/1467-9868.00294. 2, 3
- Marshall, L., Johnson, J. S., Mann, G. W., Lee, L., Dhomse, S. S., Regayre, L., Yoshioka, M., Carslaw, K. S., and Schmidt, A. (2019). "Exploring how eruption source parameters affect volcanic radiative forcing using statistical emulation." *Journal of Geophysical Research: Atmospheres*, 124: 964–985. 2
- McCreesh, N., Andrianakis, I., Nsubuga, R. N., Strong, M., Vernon, I., McKinley, T. J., Oakley, J. E., Goldstein, M., Hayes, R., and White, R. G. (2017). "Universal test, treat, and keep: improving ART retention is key in cost-effective HIV control in Uganda." *BMC Infectious Diseases*, 17(1): 322.
- McKinley, T., Vernon, I., Andrianakis, I., McCreesh, N., Oakley, J., Nsubuga, R., Goldstein, M., and White, R. (2018). "Approximate Bayesian computation and simulation-based inference for complex stochastic epidemic models." *Statistical Science*, 33(1): 4–18. MR3757500. doi: https://doi.org/10.1214/17-STS618.
- Mohammadi, H., Challenor, P., Goodfellow, M., and Williamson, D. (2020). "Emulating computer models with step-discontinuous outputs using Gaussian processes." URL https://arxiv.org/abs/1903.02071 3
- Owen, J., Vernon, I., and Hammersley, R. (2020). "A Bayesian statistical approach to decision support for TNO OLYMPUS well control optimisation under uncertainty." *Conference Proceedings, ECMOR XVII*, 2020(1): 1–27. 20, 23, 25
- Paciorek, C. J. (2003). "Nonstationary Gaussian processes for regression and spatial modelling." Ph.D. thesis, Carnegie Mellon University. MR2704401. 15
- Pope, C. A., Gosling, J. P., Barber, S., Johnson, J. S., Yamaguchi, T., Feingold, G., and Blackwell, P. G. (2021). "Gaussian process modeling of heterogeneity and discontinuities using Voronoi tessellations." *Technometrics*, 63(1): 53–63. MR4205691. doi: https://doi.org/10.1080/00401706.2019.1692696. 2
- Pukelsheim, F. (1994). "The three σ rule." The American Statistician, 48: 88–91. MR1292524. doi: https://doi.org/10.2307/2684253. 24
- Rasmussen, C. E. and Williams, C. K. I. (2006). Gaussian Processes for Machine Learning. MIT Press. MR2514435. 4, 17
- Rodrigues, L. F. S., Vernon, I., and Bower, R. G. (2017). "Constraints to galaxy for-

mation models using the galaxy SMF." Monthly Notices of the Royal Astronomical Society, 466(2): 2418–2435. 2

- Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. (1989). "Design and analysis of computer experiments." *Statistical Science*, 4(4): 409–435. MR1041765. 2
- Sauer, A., Gramacy, R. B., and Higdon, D. (2022). "Active learning for deep Gaussian process surrogates." *Technometrics*, 0(0): 1–15. MR4543056. doi: https://doi.org/ 10.1080/00401706.2021.2008505. 2
- Schneider, M. D., Knox, L., Habib, S., Heitmann, K., Higdon, D., and Nakhleh, C. (2008). "Simulations and cosmological inference: A statistical model for power spectra means and covariances." *Physical Review D*, 78: 063529.
- TNO (2017). OLYMPUS Oil Reservoir Model Input Decks. Netherlands Organisation for Applied Scientific Research (TNO). URL https://www.isapp2.com/optimization-challenge.html 18
- TNO (ed.) (2018). EAGE/TNO Workshop on OLYMPUS Field Development Optimization. EAGE and TNO, EAGE Publications. 18
- Vernon, I., Goldstein, M., and Bower, R. G. (2010a). "Galaxy formation: a Bayesian uncertainty analysis." *Bayesian Analysis*, 5(4): 619–670. MR2740148. doi: https:// doi.org/10.1214/10-BA524. 2, 4, 25
- Vernon, I., Goldstein, M., and Bower, R. G. (2010b). "Rejoinder for Galaxy Formation: a Bayesian Uncertainty Analysis." *Bayesian Analysis*, 5(4): 697–708. MR2740153. doi: https://doi.org/10.1214/10-BA524REJ. 2, 4
- Vernon, I., Goldstein, M., and Bower, R. G. (2014). "Galaxy formation: Bayesian history matching for the observable universe." *Statistical Science*, 29(1): 81–90. MR3201849. doi: https://doi.org/10.1214/12-STS412. 2
- Vernon, I. and Gosling, J. P. (2023). "A Bayesian computer model analysis of robust Bayesian analyses." *Bayesian Analysis*, 18(4): 1367–1399. MR4675042. doi: https:// doi.org/10.1214/22-ba1340. 2
- Vernon, I., Liu, J., Goldstein, M., Rowe, J., Topping, J., and Lindsey, K. (2018). "Bayesian uncertainty analysis for complex systems biology models: emulation, global parameter searches and evaluation of gene functions." *BMC Systems Biology*, 12(1): arXiv:1607.06358 [q-bio.MN]. 2, 4
- Vernon, I., Owen, J., Aylett-Bullock, J., Cuesta-Lazaro, C., Frawley, J., A., A. Q.-B., Sedgewick, Shi, D., Truong, H., Turner, M., Walker, J., Caulfield, T., Fong, K., and Krauss, F. (2022). "Bayesian emulation and history matching of JUNE." *Philosophical Transactions of the Royal Society of London. A*, 380(2233). 2
- Vernon, I., Jackson, S. E., and Cumming, J. A. (2019). "Known boundary emulation of complex computer models." SIAM/ASA Journal on Uncertainty Quantification, 7(3): 838–876. MR3980268. doi: https://doi.org/10.1137/18M1164457. 21
- Vernon, I., Owen, J., and Carter, J. (2024). "Supplementary Material for "Bayesian emu-

lation for computer models with multiple partial discontinuities"." *Bayesian Analysis*. doi: https://doi.org/10.1214/24-BA1456SUPP. 4

Williamson, D., Goldstein, M., Allison, L., Blaker, A., Challenor, P., Jackson, L., and Yamazaki, K. (2013). "History matching for exploring and reducing climate model parameter space using observations and a large perturbed physics ensemble." *Climate Dynamics*, 41(7-8): 1703–1729. MR3283898. doi: https://doi.org/10.1137/ 120900915. 2