DEEP GAUSSIAN PROCESSES AND INVERSION FOR DECISION SUPPORT IN MODEL-BASED CLIMATE CHANGE MITIGATION AND ADAPTATION PROBLEMS

Bertrand Nortier, Daniel Williamson, Mattia Mancini, Ian Bateman, Brett Day, Amy Binner
University of Exeter

ABSTRACT

To inform their decisions, policy makers often rely on models developed by researchers that are computationally intensive and complex and that frequently run on High Performance Computers (HPC). These decision-support models are not used directly by deciders and the results of these models tend to be presented by experts as a limited number of potential scenarios that would result from a limited number of potential policy choices. Machine learning models such as Deep Gaussian Processes (DGPs) can be used to radically re-define how decision makers can use models by creating a 'surrogate model' or 'emulator' of the original model. Surrogate models can then be embedded into apps that decisions makers can use to directly explore a vast array of policy options corresponding to potential target outcomes (model inversion). To illustrate the mechanism, we give an example of application that is envisaged as part of the UK government's Net Zero strategy. To achieve Net Zero CO_2 emissions by 2050, the UK government is considering multiple options that include planting trees to capture carbon. However, the amount of CO_2 captured by the trees depend on a large number of factors that include climate conditions, soil type, soil carbon, tree type, ... Depending on these factors the net balance of carbon removal after planting trees may not necessarily be positive. Hence, choosing the right place to plant the right tree is very important. A decision-helping model has been developed to tackle this problem. For a given policy input, the model outputs its impact in terms of CO_2 sequestration, biodiversity and other ecosystem services. We show how DGPs can be used to create a surrogate model of this original afforestation model and how these can be embedded into an R shiny app that can then be directly used by decision makers.

1 INTRODUCTION

We consider a model, or network of models being used to support decision making as a black box with vector-valued inputs and outputs. Surrogate modelling aims to create a statistical model that quickly and accurately predicts the original model with a well-calibrated measure of uncertainty, for example providing a mean and variance for any value of the model inputs, or a full probability distribution Gramacy (2020). A surrogate model can be used to support decision makers, enabling them to understand the predicted consequences of any policy that can be evaluated by the model in real-time. Another main advantage of the surrogate modelling approach is that the model can be "inverted", i.e. the space of inputs compatible with given targets on the outputs can be obtained and explored. Though surrogates based on Neural Networks are possible if sufficiently many model evaluations are available to densely cover the input space, environmental models are typically too computationally expensive to provide sufficiently rich training sets. We use deep Gaussian Processes (hereafter DGPs Damianou & Lawrence (2013), Ming et al. (2023)), which have been shown to perform well for environmental models Williamson et al. (2013). Gaussian processes (GPs) are distributions over functions s.t. any finite collection of function evaluations follows a multivariate Normal distribution with mean and variance matrix given by the mean and covariance functions of the process. A DGP composes GPs in layers (much like a Neural Network composes transfer functions) to enable flexible learning of non-stationary functions. Although DGPs can be fitted to some outputs of the model directly, some outputs are too high-dimensional for this to be efficient. A general approach, based on Salter et al. (2019); Chang et al. (2016) involves projecting the output onto a low-dimensional basis that captures the majority of the output variability in response to changes in the input. Denote the N-dimensional model output, $\mathbf{Y}(\mathbf{x})$, at input locations $\mathbf{x}$, where $\mathbf{Y}(\mathbf{x})$ is assumed to represent a single variable, but perhaps over space and time (note the approach below can be repeated for each variable). The approach finds decompositions such that

$$\mathbf{Y}\left(\mathbf{x}\right) \approx g\left(\boldsymbol{\mu} + \mathbf{K}\mathbf{W}\left(\mathbf{x}\right)\right),$$

where $\mathbf{K}$ is a matrix of basis vectors, $\mathbf{W}(\mathbf{x})$ is a vector of J coefficients with $J << N$, $\boldsymbol{\mu}$ is the ensemble mean and $\mathbf{g}$ is a function transforming $(-\infty, \infty)^N$ to the appropriate space for the model output (e.g. strictly positive $\mathbf{Y}$ might use an exponential function for g). Constructing $\mathbf{K}$ and $\mathbf{W}(\mathbf{x})$ so that the entries of $\mathbf{W}(\mathbf{x})$ are uncorrelated, we can model each component as:

$$W_k\left(\mathbf{x}\right) \sim \mathrm{DGP}\left(\mu_k\left(\mathbf{x}\right), \sigma_k^2\left(\mathbf{x}\right)\right); \qquad k = 1, ..., J$$

In the rest of this document, we introduce a model for exploring the efficacy of tree-planting incentivisation policies on Greenhouse Gas (GHG) emissions and biodiversity, and show how surrogates for this model can be deployed to provide a novel form of decision support. In section 2, we introduce the model. In section 3, we then describe our surrogate models and how they can be used in inversion for decision support. Results and conclusions are presented in Sections 4 and 5.

2 Data Generating Model: Treescaping Incentivisation Schemes for enhanced ecosystem services

The UK government aims to reduce its net CO_2 emissions to zero by 2050 Government of the United Kingdom (2021). To achieve this, significant GHG removal is required to offset carbon emitting activities such as agricultural production. Tree planting is one scalable and affordable type of GHG removal and the UK has mandated an increase of its woodland by a further 2.5% of UK land. Delivering land use change on this scale requires private land holders, mainly farmers, to be incentivised to change their existing land use in favour of planting trees. We explore the use of a connected suite of models Ritchie et al. (2020); Day et al. (2020) that simulate farmer response to a suite of plausible prices offered by government, the associated change in land use, the change in GHG emissions and the impacts of that change on ecosystem services, including food production and biodiversity. We consider varying the prices offered to farmers per hectare for different types of tree species (deciduous/conifer), including different prices for planting unmanaged woodland (‘unfelled’), and managed trees for timber (‘felled’). We also consider policies paying per ton of CO_2 captured by the scheme and the influence of an annual budget for paying farmers out of the scheme. The model outputs include: a) Categorical maps in 2030, 2040, 2050 indicating where the trees of different types are planted in the UK depending on the inputs of the model; b) Biodiversity net gain for different species; c) Time series of different important indicators for 150 years: cumulative CO_2 captured, total trees planted, total spent and food production loss. The dataset used to fit our surrogate models then consist in data generated by the original model. As we have access to it, we can generate the amount of data necessary to obtain a well-fitted surrogate models.

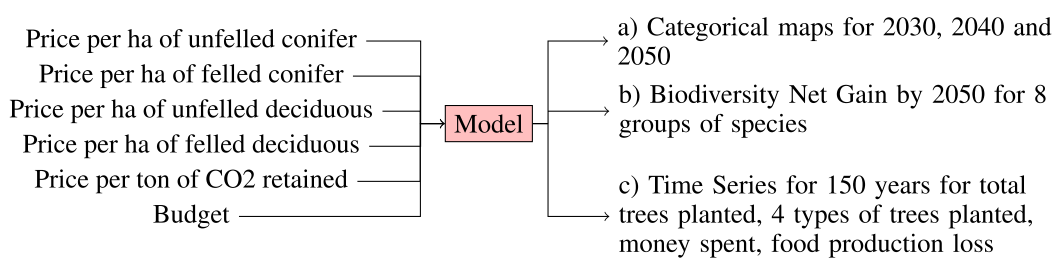

3 Methodology: Surrogate Models and model inversion

The general structure described in Section 1 was used to fit emulators to the key outputs of the model. Aggregate biodiversity metrics could be emulated with DGPs without transformation and without projection onto a basis. Figure 1 shows leave one out cross-validation for a 2-layer DGP fitted to the increase in bee populations (y-axis) as the policies are changed. Setting $g(\cdot)$ as the identity and using the singular value decomposition to derive $\mathbf{K}$ and $\mathbf{W}(\mathbf{x})$ (see Salter et al. (2019), for example) worked well for time series output. Bespoke methodology we developed for emulating categorical maps with DGPs is the topic of a forthcoming submission (as far as the authors are aware, the methodology for categorical output models is novel (GP-based surrogate models for binary output have been investigated in Chang et al. (2016))). One of the major advantages of using surrogate

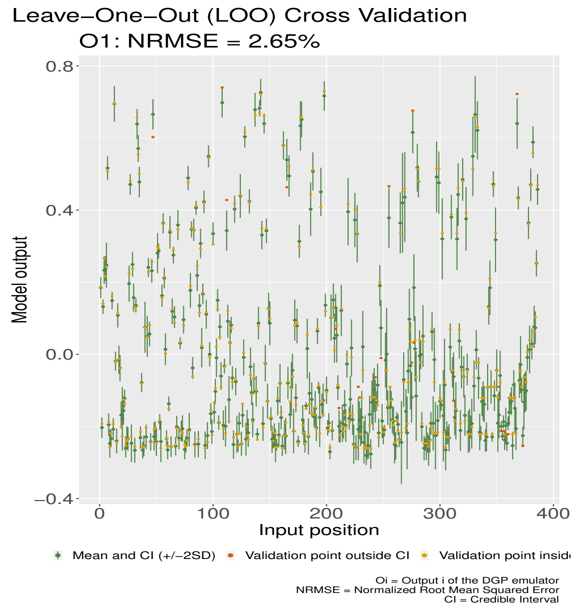

Figure 1: Leave-one-out cross validation for a 2-layer DGP fitted to the normalized increase in bee populations (y-axis) in our training set. The x-axis gives the index of the simulation in the training set, the predicted outputs (green) are compared with actual model runs (orange/red if within/outside of the 95% credible interval).

models is that we can 'invert the model', i.e. we can set targets on the outputs of the model (such as achieving sufficient GHG reduction by tree planting to contribute to net zero targets) and we can obtain the space of possible inputs (policies) that is potentially compatible with the targets. Consider a scalar output from a surrogate model $Y(\mathbf{x})$ for which we want to meet a target $\mathcal{T}$. Our inversion seeks to rule out the space of inputs where $\mathbb{P}(Y(\mathbf{x}) \geq \mathcal{T}) \leq \alpha$ where $\alpha > 0$ is a small value. Similar to Baker (2021), we define Implausibility, $I(\mathbf{x})$, via

$$I(\mathbf{x}) = \frac{\mathcal{T} - E[\mathbf{Y}(\mathbf{x})]}{\sqrt{var[\mathbf{Y}(\mathbf{x})]}}.$$

We can then show that

$$\mathbb{P}(Y(\mathbf{x}) \geq \mathcal{T}) \leq \alpha \iff \sqrt{\frac{1-\alpha}{\alpha}} \leq I(\mathbf{x})$$

To obtain the target-compatible space, we can then choose a level α and Rule Out all the values of $\mathbf{x}$ such that $I(\mathbf{x}) \geq \sqrt{(1-\alpha)/\alpha}$. Note that if we have multiple targets $\mathcal{T}_1, ..., \mathcal{T}_a$ on multiple variables $Y_1(\mathbf{x}), ..., Y_a(\mathbf{x})$, we can calculate multiple implausibilities $I_1(\mathbf{x}), ..., I_a(\mathbf{x})$ and use $I(\mathbf{x}) = \max_{i=1,...,a} I_i(\mathbf{x})$ as the overall Implausibility matching all the targets.

4 Results: Rshiny app display

Our surrogate models are embedded into an R shiny app that allows a user to select a policy on the left-hand side and visualize immediately the estimated results of this policy, with uncertainty. Figure 2 shows an example of display of the R shiny app. One example for communicating the target compatible space consists of representing multiple bivariate plots showing what the percentage of policy space is retained given that 2 of the inputs are fixed within a small area (Vernon et al. (2010)) (see Figure 3).

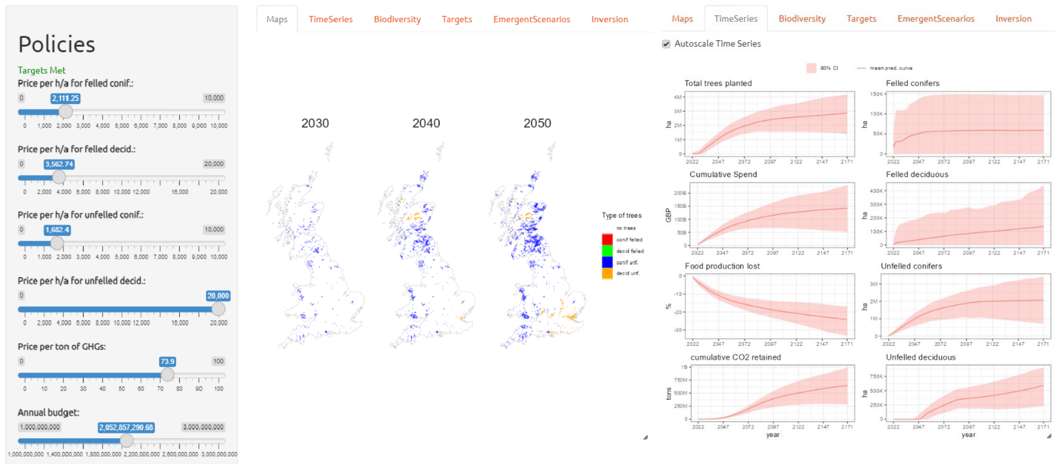

Figure 2: Screenshots of our R shiny app. The left hand side panel shows the policy inputs as described in the main text. The central visualization are expected categorical maps of planted type of trees. The right hand side panel shows examples of time series predicted. A 95% credible interval is shown (red shading) together with the expected value (solid red line) for the chosen policy.

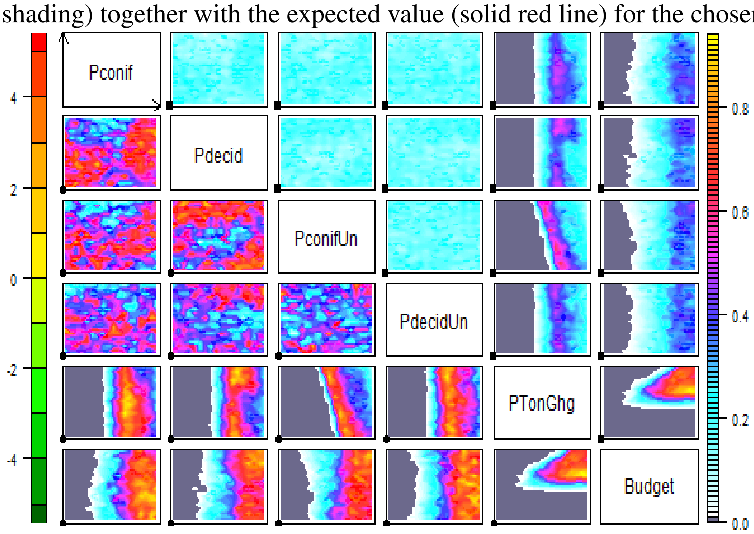

Figure 3: Panels depicting the density of policy space that is target-compatible behind each pixel. For each panel, 50,000 inputs/outputs are sampled using the surrogate models. For each pixel in each panel of the upper triangle, the proportion of samples corresponding to that pixel that are not ruled out is given the colour indicated by the legend. Panels on the lower triangle are given a panel specific colour scale to highlight any patterns masked by the global scale.

5 Impact of this work and Conclusion

The work presented here has the potential to transform the way policy makers tackling climate change and ecosystem restoration are advised. Surrogate models can democratize state of the art, process-based models, giving decision makers the opportunity to interrogate models that must normally be run by specialists, potentially with HPC, in real-time. The inversion method presented delivers the space of promising policies (those that are predicted to meet user targets), enabling policy makers to explore the implications of their target-compatible options on other outcomes that are output by the model. In our example, though we have targets for biodiversity and GHG emissions in law to help set targets, the spatial distribution of the resulting planting may need to be 'fair' with a 'reasonable' sharing of the land use change across the regions of the UK, in order to be politically palatable. It is much harder and perhaps inappropriate to put specific targets on these notions of 'fairness', meaning that optimisation-based inversion methods may be inappropriate within a policy support context. By allowing the user to explore the target compatible planting distributions themselves, we present an application of machine learning that, rather than making the decision for the policy maker, empowers them to make it themselves. This work will be extended further first by developing tools to enhance the exploration of the target-compatible decision space, second, by developing apps that allow a more granular decision making, replacing the UK-wide decision by land parcel level decisions, last by increasing the number of outputs, including fire hazard risk, flooding risk, recreation value amongst others.

ACKNOWLEDGMENTS

The first author would like to thank Paul Harder, Akshata Kishore Moharir, Konstantin Klemmer and Frankie Cho. This work was supported by UK Research and Innovation and the Biotechnology and Biological Research Council through the Strategic Priorities Fund for Greenhouse Gas Removal. NetZeroPlus: Sustainable Treescapes Demonstrator and Decision Tools (Netzeroplus.ac.uk), Grant number BB/V011588/1.

REFERENCES

Evan Baker. *Emulation of Stochastic Computer Models with an Application to Building Design*. PhD thesis, University of Exeter, 2021.

Won Chang, Murali Haran, Patrick Applegate, and David Pollard. Calibrating an ice sheet model using high-dimensional binary spatial data. *Journal of the American Statistical Association*, 111 (513):57–72, 2016.

Andreas Damianou and Neil D Lawrence. Deep gaussian processes. In *Artificial intelligence and statistics*, pp. 207–215. PMLR, 2013.

Brett Day, Nathan Owen, Amy Binner, Ian Bateman, Frankie Cho, Anthony De-Gol, Henry Ferguson-Gow, Silvia Ferrini, Carlo Fezzi, Amy Harwood, Christopher Lee, Lorena Luizzio, Mattia Mancini, and Richard Pearson. The natural environmental valuation (nev) modelling suite: a summary technical report. Technical report, LEEP Working Paper: 01, 2020.

Government of the United Kingdom. Net zero strategy: build back greener. *Controller Her Majesty's Stationery Office*, 2021.

Robert B Gramacy. *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. CRC press, 2020.

Deyu Ming, Daniel Williamson, and Serge Guillas. Deep gaussian process emulation using stochastic imputation. *Technometrics*, 65(2):150–161, 2023.

Paul DL Ritchie, Greg S Smith, Katrina J Davis, Carlo Fezzi, Solmaria Halleck-Vega, Anna B Harper, Chris A Boulton, Amy R Binner, Brett H Day, Angela V Gallego-Sala, et al. Shifts in national land use and food production in great britain after a climate tipping point. *Nature Food*, 1(1):76–83, 2020.

James M Salter, Daniel B Williamson, John Scinocca, and Viatcheslav Kharin. Uncertainty quantification for computer models with spatial output using calibration-optimal bases. *Journal of the American Statistical Association*, 2019.

Ian Vernon, Michael Goldstein, and Richard G Bower. Galaxy formation: a bayesian uncertainty analysis. *Bayesian Analysis*, 5(4):619–670, 2010.

Daniel Williamson, Michael Goldstein, Lesley Allison, Adam Blaker, Peter Challenor, Laura Jackson, and Kuniko Yamazaki. History matching for exploring and reducing climate model parameter space using observations and a large perturbed physics ensemble. *Climate dynamics*, 41: 1703–1729, 2013.