



RESEARCH ARTICLE

10.1029/2023MS003926

Special Section:

Machine learning application to
Earth system modeling

Key Points:

- We introduce FaIRGP, a Bayesian machine learning emulator for global and local mean surface temperatures that builds upon a physically based simple climate model
- The model improves upon both purely physically-driven and purely data-driven baseline emulators on several metrics across realistic future scenarios
- The model is fully mathematically tractable, which makes it a convenient and easy-to-use probabilistic tool for the emulation of surface temperatures, but also for downstream applications such as detection and attribution or precipitation emulation

Correspondence to:

S. Bouabid,
shahine.bouabid@stats.ox.ac.uk

Citation:

Bouabid, S., Sejdinovic, D., & Watson-Parris, D. (2024). FaIRGP: A Bayesian energy balance model for surface temperatures emulation. *Journal of Advances in Modeling Earth Systems*, 16, e2023MS003926. <https://doi.org/10.1029/2023MS003926>

Received 19 JULY 2023

Accepted 7 MAR 2024

Author Contributions:

Conceptualization: Shahine Bouabid,
Dino Sejdinovic, Duncan Watson-Parris

Data curation: Shahine Bouabid,
Duncan Watson-Parris

Formal analysis: Shahine Bouabid

Funding acquisition: Dino Sejdinovic,
Duncan Watson-Parris

Investigation: Shahine Bouabid

Methodology: Shahine Bouabid,
Dino Sejdinovic, Duncan Watson-Parris

© 2024 The Authors. Journal of Advances in Modeling Earth Systems published by Wiley Periodicals LLC on behalf of American Geophysical Union.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

FaIRGP: A Bayesian Energy Balance Model for Surface Temperatures Emulation

Shahine Bouabid¹ , Dino Sejdinovic² , and Duncan Watson-Parris³
¹Department of Statistics, University of Oxford, Oxford, UK, ²School of CMS & AIML, University of Adelaide, Adelaide, SA, Australia, ³Scripps Institution of Oceanography and Halicioğlu Data Science Institute, University of California, San Diego, San Diego, CA, USA

Abstract Emulators, or reduced complexity climate models, are surrogate Earth system models (ESMs) that produce projections of key climate quantities with minimal computational resources. Using time-series modeling or more advanced machine learning techniques, data-driven emulators have emerged as a promising avenue of research, producing spatially resolved climate responses that are visually indistinguishable from state-of-the-art ESMs. Yet, their lack of physical interpretability limits their wider adoption. In this work, we introduce FaIRGP, a data-driven emulator that satisfies the physical temperature response equations of an energy balance model. The result is an emulator that (a) enjoys the flexibility of statistical machine learning models and can learn from data, and (b) has a robust physical grounding with interpretable parameters that can be used to make inference about the climate system. Further, our Bayesian approach allows a principled and mathematically tractable uncertainty quantification. Our model demonstrates skillful emulation of global mean surface temperature and spatial surface temperatures across realistic future scenarios. Its ability to learn from data allows it to outperform EBM, while its robust physical foundation safeguards against the pitfalls of purely data-driven models. We also illustrate how FaIRGP can be used to obtain estimates of top-of-atmosphere radiative forcing and discuss the benefits of its mathematical tractability for applications such as detection and attribution or precipitation emulation. We hope that this work will contribute to widening the adoption of data-driven methods in climate emulation.

Plain Language Summary Emulators are simplified climate models that can be used to rapidly explore climate scenarios—they can run in less than a minute on an average computer. They are key tools used by the Intergovernmental Panel on Climate Change to explore the diversity of possible future climates. Data-driven emulators use advanced machine learning techniques to produce climate predictions that look very similar to the predictions of complex climate models. However, they are not easy to interpret, and therefore to trust in practice. In this work, we introduce FaIRGP, a data-driven emulator based on physics. The emulator is flexible and can learn from data to improve its predictions, but is also grounded on physical energy balance relationships, which makes it robust and interpretable. The model performs well in predicting future global and local temperatures under realistic future scenarios, outperforming purely physics-driven or purely data-driven models. Further, the probabilistic nature of our model allows for mathematically tractable uncertainty quantification. By gaining trust in such a data-driven yet physically grounded model, we hope the climate science community can benefit more widely from their potential.

1. Introduction

Earth system models (ESMs) (Flato, 2011) are key tools to understand current climate dynamics and climate change responses to greenhouse gas emissions. They constitute an extensive physical simulation of Earth's atmosphere and ocean fluid dynamics, used for example, in the Couple Model Intercomparison Project (Eyring et al., 2016; Meehl et al., 2007; Taylor et al., 2012) to study past and future climate. As such, they offer the most comprehensive view of what future climate could look like. They are also used as an idealized fully controlled environment to study climate dynamics and understand its underlying drivers. In particular, they play a central role in the estimation of key properties of the climate system such as timescales and equilibrium responses to the change in carbon dioxide concentration in the atmosphere (Allen et al., 2009; Collins et al., 2014) and the effect of aerosols (Levy et al., 2013).

Running simulations with an ESM requires an astute understanding of the climate science background, of the numerical schemes used to simulate climate dynamics, and access to an adequate computational infrastructure.

Software: Shahine Bouabid
Supervision: Dino Sejdinovic,
Duncan Watson-Parris
Validation: Shahine Bouabid
Visualization: Shahine Bouabid
Writing – original draft:
Shahine Bouabid
Writing – review & editing:
Shahine Bouabid, Dino Sejdinovic,
Duncan Watson-Parris

As an order of magnitude, running the CESM2 model (Danabasoglu et al., 2020) for a single year ahead takes about 2,000 core hours on a supercomputer. Therefore, only a limited number of research teams around the world can realistically afford to perform climate simulations. A direct consequence is that a variety of scientific applications relying on future climate projections—such as agricultural studies (Rosenzweig et al., 2014), energy infrastructure models (Turner et al., 2017) or global socio-economic human models (Akhtar et al., 2013)—must settle for publicly available precomputed climate projections that have been selected based on their policy-relevance, and may not be tailored to the application needs. The need for expensive computational resources also serves as an important barrier, making it inaccessible for independent researchers and less well-equipped research teams to run experiments with ESMs. This exacerbates the already unequal representation in high-impact climate science research, where the global north is disproportionately represented (Tandon, 2021).

Furthermore, even when the resource needs are met to run experiments with an ESM, their computational cost remains a critical obstacle. Indeed, the uncertainty over the ESM parametrization, the climate system internal variability and the emission pathway the world will choose, together span a high-dimensional uncertainty space. Therefore, obtaining a comprehensive coverage of this uncertainty requires running numerous climate simulations, which quickly meets computational cost limitations. As a result, much of the climate variability and potential socio-economic pathways remain in practice unexplored.

Together, these limitations have fostered the emergence of simpler surrogate ESMs which are inexpensive to compute and referred to as *emulators*. Unlike ESMs, emulators do not explicitly model the fluid dynamics of the atmosphere and oceans and focus on a limited number of climate features, such as surface temperatures or precipitations. They can run thousands of years of simulation in less than a minute on an average personal computer (Leach et al., 2021), hence making accessible the emulation of climate projections under configurations unexplored by ESMs.

Simple climate models (SCMs) (Leach et al., 2021; Meinshausen et al., 2011, 2020; Millar et al., 2017; Smith et al., 2018) are an important class of models that can emulate climate projections. They propose a reduced-order representation of the climate system, modeling the carbon cycle, radiative forcing, and temperature response with simplified first order dynamics. A key component of SCMs is their *energy balance model* (EBM) (Fredriksen & Rypdal, 2017; Geoffroy et al., 2013; Lovejoy et al., 2021; Millar et al., 2017; Rypdal, 2012), which describes the changes in global surface temperature through imbalances in Earth's energy budget (Held et al., 2010). It represents the atmosphere-ocean system as a set of connected boxes forced by radiative flux at the top of the atmosphere. Whilst this constitutes a drastic simplification of the climate system, Energy balance models (EBMs) are robust physically-motivated models, and therefore remain the main tool used to connect the IPCC working physical basis research (Lee et al., 2021) with the adaptation and mitigation efforts (Pörtner et al., 2022; Shukla et al., 2022).

Another important class of emulators are *data-driven emulators*. In contrast to SCMs, they do not explicitly model climate dynamics and rather rely on statistical modeling techniques to emulate key climate variables such as temperature or precipitation (Beusch et al., 2020b; Castruccio et al., 2014; Del Rio Amador & Lovejoy, 2019; Goodwin et al., 2020; Holden & Edwards, 2010; Link et al., 2019). Statistical modeling enjoys greater flexibility and has produced powerful emulators, capable of successfully approximating ESMs' spatially resolved climate projections with visually indistinguishable outputs. More recently, statistically driven emulators drawing from advances in statistical machine learning have demonstrated a remarkable capacity at regressing global emission profiles onto spatially resolved temperature and precipitation maps (Watson-Parris et al., 2022).

However, both SCMs and data-driven emulators display fundamental limitations. Whilst SCMs provide a robust physical framework to reason about the climate system, they remain a simplistic model which may display a poor fit to ESM's outputs in certain scenarios (Geoffroy et al., 2013; Jackson et al., 2022; Meinshausen et al., 2011) and can only operate at the global level. Reasoning only in terms of global mean temperatures fails to capture the difference in exposure of different world regions and limits use for scientific applications that require regional projections. On the other hand, data-driven emulators are also limited in their ability to provide a complete, reliable picture of the climate system. Indeed, their lack of robust physical grounding limits their capacity to make inferences about the climate system (Beusch et al., 2020b; Link et al., 2019). Further, the outputs from these emulators are often subject to qualitative evaluation and may not be trustworthy, in part because of limited




Figure 1. Schematic representation of 3-box energy balance model, the arrows represent the flow of heat. *Left:* Temperature response model representation. Thicker boxes have a larger heat capacity (i.e., $C_1 < C_2 < C_3$) and the arrows represent the flow of heat. $\kappa_1, \kappa_2, \kappa_3$ are the heat transfer coefficients between the boxes and $\epsilon > 0$ denotes the deep ocean heat uptake efficacy. *Right:* Thermal response model representation. Thicker boxes have larger response timescale (i.e., $d_1 < d_2 < d_3$). Adapted from Figure 1, Cummins et al. (2020).

understanding on how they might behave outside the training data regimes. Whilst statistical explainability methods (Shapley, 1953; Štrumbelj & Kononenko, 2014) may help understand the contributions to predictions, we argue that the lack of a physically grounded model would still harm trust in the predictions.

In this work, we address these limitations by formulating a hybrid physical/statistical emulator that will both enjoy the robust physical grounding of SCMs, and the flexibility of statistical machine learning methods, thereby combining the advantages of both classes of emulators. We propose a probabilistic emulator for the task of reproducing an ESM temperature response to greenhouse gas and aerosol emissions. Our model builds upon a simple EBM, but draws from the latent force model paradigm (M. Alvarez et al., 2009; M. A. Alvarez et al., 2013) by placing a Gaussian process (GP) prior over the radiative forcing, thereby inducing a stochastic temperature response model. We show that the resulting emulated temperature turns out to also be a GP, with a physically-informed covariance structure that reflects the dynamics of the EBM. In consequence, we can update the EBM with temperature data to learn a posterior distribution over global temperatures, but also over radiative forcing. Further, we demonstrate how our model can be easily extended to emulate spatially-resolved maps of annual surface temperatures.

Experiments demonstrate skillful prediction of global and spatial mean surface temperatures, with improvements over a simple EBM and a simple GP model. We obtain robust predictions even when only training with historical data, or when predicting over scenarios outside the range of emissions of greenhouse gases and aerosols observed during training. Additionally, we show that the model improves over baselines to emulate temperature changes induced by anthropogenic aerosol emissions, and can provide useful estimates of the top-of-atmosphere radiative forcing.

2. Background

2.1. Energy Balance Models

EBMs (Holden & Edwards, 2010; Sellers, 1969) are a key component of SCMs that provide a lower-order representation of the climate system, where the changes in global temperature are explained by the imbalance in the Earth's energy budget. The most common and established class of EBM are box models. They represent the atmosphere and ocean as a set of vertically stacked boxes, where the uppermost box is exposed to top-of-atmosphere radiative flux.

In a box model EBM, each box has a different heat capacity C_i , heat transfer coefficients κ_i with the adjacent boxes, and its own temperature $T^{(i)}(t)$, as depicted in Figure 1. The uppermost box represents the fast components of the climate system, generally limited to the atmosphere and the land surface, while the following boxes are used to represent slower components of the climate system such as shallow ocean and deep ocean.

Let $\mathbf{X}(t) = [T^{(1)}(t) \dots T^{(k)}(t)]^T$ be the vector concatenation of the temperatures within the boxes in a model with k boxes. The change in temperature of a k -box EBM is described by the following simple first order linear ordinary differential equation (ODE)

$$\frac{d\mathbf{X}(t)}{dt} = \mathbf{A}\mathbf{X}(t) + \mathbf{b}F(t), \quad (1)$$

where \mathbf{A} is a tridiagonal temperature feedback matrix that depends on heat capacities C_i , heat transfer coefficients κ_i , and deep ocean uptake efficacy ε (explicit form of the matrix is provided in Appendix B2), and \mathbf{b} is a radiative forcing feedback vector given by

$$\mathbf{b} = [1/C_1 \ 0 \ \dots \ 0]^T. \quad (2)$$

$F(t)$ denotes the top of atmosphere effective radiative forcing—radiative forcing for short—that is, the change in energy flux caused by natural or anthropogenic factors of climate change. It is only applied to the surface box, that is, the uppermost box.

Whilst physically motivated, box-models remain an abstract representation of the climate system. Therefore their parameters, such as boxes heat capacities, are not realistic quantities that can be calculated, but rather need to be tuned against data using calibration methods (Meinshausen et al., 2011; Millar et al., 2017) or maximum likelihood strategies (Cummins et al., 2020).

The simplest box models only use $k = 2$ boxes, thereby splitting the climate system into two groups: fast and slow components. Whilst reductive, this split has proven to be a robust approximation of the climate system (Held et al., 2010). In fact, 2-box EBMs are today the primary tool used for linking the IPCC physical basis research (Lee et al., 2021) with the adaptation and mitigation efforts (Pörtner et al., 2022; Shukla et al., 2022). It is however worth noting recent advocacy for 3-box models (Leach et al., 2021), highlighting the insufficiency of 2-box models to capture the full range of behavior observed in coupled model intercomparison project (CMIP6) models (Cummins et al., 2020; Tsutsui, 2020).

2.2. Impulse Response Formulation

The main box of interest in a box-model is the uppermost box since it describes the global mean surface temperature response $T^{(1)}(t)$ to the radiative forcing $F(t)$. In general, computing $T^{(1)}(t)$ can be simplified by diagonalizing the ODE (1). The result is an equivalent impulse response formulation of the temperature response model, depicted on the right in Figure 1, and referred to as the *thermal response ODE* (Millar et al., 2017).

Let us rename the temperature of the first box as $T(t)$ for simplicity. Then, for a k -box model, the thermal response ODE is formally given by

$$\begin{cases} \frac{dS_i(t)}{dt} = \frac{1}{d_i}(q_i F(t) - S_i(t)), & 1 \leq i \leq k \\ T(t) = \sum_{i=1}^k S_i(t), \end{cases} \quad (3)$$

where $S_i(t)$ is the i th thermal response, d_i is the i th response timescale and q_i is the i th equilibrium response. Table 1a provides a brief description of named parameters and units.

Whilst it is harder to interpret the impulse response formulation in terms of physical processes, its primary benefit is that each thermal response ODE can be solved independently, thereby avoiding the intricacies of solving a coupled system of ODE. Further, the timescale and equilibrium parameters d_i and q_i can be expressed in terms of the original boxes heat capacities C_i and heat transfer coefficients κ_i . Detailed derivations can be found in Geoffroy et al. (2013).

Throughout this work, we will use as a reference SCM FaIRv2.0.0 (Leach et al., 2021) (for Finite amplitude Impulse Response), a recent update of a well established SCM (Millar et al., 2017; Smith et al., 2018), that offers a

Table 1
Notation, Description and Unit of Named Parameters

| Notation | Unit | Description |
|------------------------------|---------------------------------|--|
| (a) Deterministic parameters | | |
| $F(t)$ | W m^{-2} | Effective top of atmosphere radiative forcing |
| $T^{(i)}(t)$ | K | Temperature of the i th box |
| C_i | J K^{-1} | Heat capacity of the i th box |
| κ_i | $\text{W m}^{-2} \text{K}^{-1}$ | Heat exchange coefficient of the i th box |
| ε | 1 | Deep ocean heat uptake efficacy |
| $T(t)$ | K | Global mean surface temperature ($T^{(1)}(t)$) |
| $S_i(t)$ | K | Response of the i th thermal box |
| d_i | Years | Response timescale of i th thermal box |
| q_i | $\text{KW}^{-1} \text{m}^2$ | Equilibrium response of i th thermal box |
| χ | – | Atmospheric agent (e.g., CO_2 , SO_2) |
| $E(t)$ | – | Emission of agents in the atmosphere |
| (b) Stochastic parameters | | |
| $F(t)$ | W m^{-2} | Effective top of atmosphere radiative forcing |
| $T(t)$ | K | Global mean surface temperature |
| $S_i(t)$ | K | Response of the i th thermal box |

minimal level of structural complexity. FaIRv2.0.0 is effectively composed of 3 submodels: a gas cycle model, which converts emissions to concentrations, a radiative forcing model, which converts concentrations to radiative forcing, and a temperature response model, which converts radiative forcing into temperatures. The temperature response model of FaIRv2.0.0 exactly corresponds to the impulse response EBM described in Equation 3. We refer to reader to the work of Leach et al. (2021) for a comprehensive presentation of FaIRv2.0.0. In the rest of the paper, we permit ourselves to drop v.2.0.0 and refer to the model as FaIR.

2.3. Gaussian Processes

GPs (Rasmussen & Williams, 2005) are a ubiquitous class of Bayesian priors over real-valued functions. They enjoy convenient closed-form expressions, principled uncertainty quantification, and cover a rich class of complex functions. As a result, they have been widely used in various nonlinear and nonparametric regression problems in geosciences (Camps-Valls et al., 2016).

We say that a real-valued stochastic process function $f(t)$ is a GP if any finite collection of its evaluations has a joint multivariate normal distribution. The parameters of this distribution are fully determined by the GP mean function $m(t) := \mathbb{E}[f(t)]$ and its covariance function $k(t, t') := \text{Cov}(f(t), f(t'))$. For example, for two inputs t_1, t_2 we have that

$$\begin{bmatrix} f(t_1) \\ f(t_2) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} m(t_1) \\ m(t_2) \end{bmatrix}, \begin{bmatrix} k(t_1, t_1) & k(t_1, t_2) \\ k(t_2, t_1) & k(t_2, t_2) \end{bmatrix}\right). \quad (4)$$

Because this extends to any finite collection of inputs t_1, \dots, t_n , we say that GPs induce a distribution over the entire function $f(t)$. Figure 2a shows examples of functions sampled from this distribution, referred to as *sample paths*. Formally, since this distribution is fully determined by $m(t)$ and $k(t, t')$, we write

$$f(t) \sim \text{GP}(m, k). \quad (5)$$

When using GPs as a Bayesian prior over a function, the mean and covariance functions are typically user-specified. The mean $m(t)$ can be set as any function specifying the average behavior of the function. The covariance function $k(t, t')$, commonly called *kernel*, is chosen to specify the correlation between $f(t)$ and $f(t')$ by

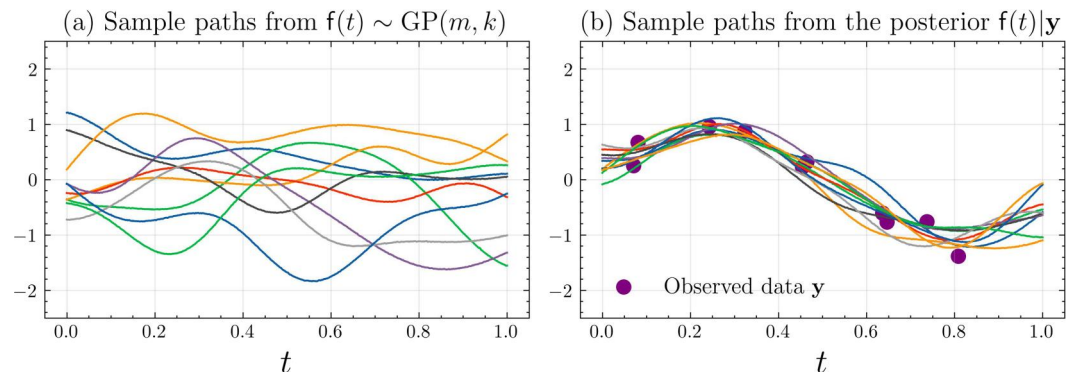


Figure 2. (a) 10 sample paths from $f(t) \sim \text{GP}(m, k)$ with $m(t) = 0$ and $k(t, t') = \sigma_f^2 \exp(-|t - t'|^2 / \ell^2)$, $\sigma_f^2 = 0.43$, $\ell = 0.17$. (b) 10 sample paths from the posterior $f(t)|y$ obtained by updating the prior. The data y is sampled from $y = \sin(2\pi t) + \epsilon$ with $\epsilon \sim \mathcal{N}(0, 0.04)$.

computing a notion of similarity between t and t' . For example, a widely used family of kernels are the Matérn kernels (Stein, 1999), where the one-dimensional Matérn-1/2 kernel is given by

$$k(t, t') = \sigma_f^2 \exp\left(-\frac{|t - t'|}{\ell}\right), \quad (6)$$

with variance parameter σ_f^2 and lengthscale parameter ℓ . When t and t' are similar, that is, $|t - t'| \ll \ell$, this kernel specifies a strong correlation between $f(t)$ and $f(t')$, that is, $\text{Cov}(f(t), f(t'))/\sigma_f^2 \approx 1$. Conversely, if t and t' are dissimilar, that is, $|t - t'| \gg \ell$, the kernel tends to zero and specifies almost no correlation between $f(t)$ and $f(t')$.

Other popular choices of kernel include the squared exponential kernel $k(t, t') = \sigma_f^2 \exp(-|t - t'|^2/\ell)$ which specifies a smooth correlation, or the cosine kernel $k(t, t') = \sigma_f^2 \cos(|t - t'|/\ell)$ which specifies a periodic correlation (Rasmussen & Williams, 2005, Chapter 4). More broadly, kernels (and hence GPs) can also be defined over multidimensional inputs, for example, by substituting $|t - t'|$ by the Euclidean distance between input vectors. When each dimension has a different lengthscale parameter, the kernel is referred to as an anisotropic, or automatic relevance determination, kernel (Rasmussen & Williams, 2005, Section 5.1).

Let $\mathbf{y} = [y_1 \dots y_n]^\top$ denote independent data samples at inputs $\mathbf{t} = [t_1 \dots t_n]^\top$ from the noisy data generating process $y = f(t) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. It is possible to inform the GP with this data, thereby updating the prior into a posterior. In particular, the posterior distribution gives rise to a *posterior GP*, with a posterior mean function $\bar{m}(t)$ and a posterior covariance function $\bar{k}(t, t')$,

$$f(t) | \mathbf{y} \sim \text{GP}(\bar{m}, \bar{k}). \quad (7)$$

In other words, the posterior GP defines a new distribution over functions, parametrized by $\bar{m}(t)$ and $\bar{k}(t, t')$, which is informed by collected data. Figure 2b shows examples of posterior sample paths, which are now concentrated around the collected data points.

A great advantage of GPs is that the posterior mean and covariance can be analytically computed following

$$\bar{m}(t) = m(t) + k(t, \mathbf{t})(\mathbf{K} + \sigma_\epsilon^2 \mathbf{I}_n)^{-1}(\mathbf{y} - m(\mathbf{t})), \quad (8)$$

$$\bar{k}(t, t') = k(t, t') - k(t, \mathbf{t})(\mathbf{K} + \sigma_\epsilon^2 \mathbf{I}_n)^{-1}k(\mathbf{t}, t), \quad (9)$$

where \mathbf{I}_n denotes the identity matrix of size n , $k(t, \mathbf{t}) = [k(t, t_1) \dots k(t, t_n)]$, $k(\mathbf{t}, t) = k(t, \mathbf{t})^\top$ and $\mathbf{K} = k(\mathbf{t}, \mathbf{t}) = [k(t_i, t_j)]_{1 \leq i, j \leq n}$ is the covariance matrix.

We provide in Appendix A1 a detailed illustrated walk-through of GP regression, and refer the reader to (Rasmussen & Williams, 2005) for a comprehensive overview of GPs in machine learning.

3. FaIRGP

In this section, we present FaIRGP, a GP emulator for global mean surface temperature that builds upon the thermal response model from FaIR. We begin by motivating a Bayesian treatment of the radiative forcing and formulate a GP prior over the forcing. We show this modification naturally results in a stochastic formulation of the thermal response model, which admits a GP with a physically-informed covariance structure for solution. In addition, we show how this framework can seamlessly account for the climate internal variability. Finally, we provide closed-form expressions for the posterior distributions over temperature and forcing, which can readily be used for emulation.

In what follows, we adopt the following notational conventions: deterministic variables are denoted with serif font (Z), stochastic variables are denoted with sans-serif font (Z) and vector/matrix versions are denoted in bold (\mathbf{Z} , \mathbf{Z}).

3.1. Epistemic Uncertainty Over the Radiative Forcing Model

The top-of-atmosphere radiative forcing is the fundamental quantity used to describe Earth's energy imbalance. The value of this forcing is primarily determined by the emissions of greenhouse gases and aerosols in the atmosphere. In FaIR, the radiative forcing term $F(t)$ drives the temperature response model (Equation 1). Having access to a reliable estimate of the forcing, and in particular, to the forcing response to greenhouse gas and aerosol emissions, is therefore critical to produce trustworthy emulated temperatures.

Computing an accurate estimate of the radiative forcing requires modeling how greenhouse gases and aerosols interact with radiations in the atmosphere, accounting for atmospheric adjustments and uncertainty in the measured parameters which may further complicate the calculation. The intricacy of such tasks has led to the development of forcing estimation methods which rely on simplifying assumptions. For example, in FaIR, the forcing model is formulated as a combination of logarithmic, linear and square-root terms of greenhouse gas and aerosol concentrations. Namely, let χ denote a given atmospheric agent (e.g., CO₂, CH₄, SO₂), the radiative forcing induced by χ is modeled as

$$F^\chi(t) = \alpha_{\log}^\chi \log\left(\frac{C^\chi(t)}{C_0^\chi}\right) + \alpha_{\text{lin}}^\chi (C^\chi(t) - C_0^\chi) + \alpha_{\text{sqr}}^\chi \left(\sqrt{C^\chi(t)} - \sqrt{C_0^\chi}\right), \quad (10)$$

where $C^\chi(t)$ denotes the concentration in the atmosphere of agent χ , C_0^χ the agent concentration at pre-industrial period and $\alpha_{\log}^\chi, \alpha_{\text{lin}}^\chi, \alpha_{\text{sqr}}^\chi$ are scalar sensitivity coefficients. The total radiative forcing is then obtained by combining the contribution of each agent

$$F(t) = \sum_{\chi} F^\chi(t). \quad (11)$$

This choice of forcing model has a physical motivation drawing from model calculations and empirical studies of temperature and historical trajectories of these concentrations (Forster et al., 2007; Manabe & Wetherald, 1967; Myhre et al., 2014), which have provided substantial evidence that the concentration-to-forcing relationship can be reasonably approximated by the combination of terms in Equation 10. However, it is important to recall that such a functional form is primarily based on empirical physics, in contrast with the EBM which has a greater theoretical motivation. In fact, the values of the sensitivity coefficients $\alpha_{\log}^\chi, \alpha_{\text{lin}}^\chi, \alpha_{\text{sqr}}^\chi$ can vary substantially depending on the data used and the fitting procedure (Leach et al., 2021).

From a statistical machine learning perspective, the concentration-to-forcing relationship is arguably the component of FaIR which is most akin to statistical modeling. Indeed, $\Phi(t) = [\log C^\chi(t) \quad C^\chi(t) \quad \sqrt{C^\chi(t)}]^\top$ can be interpreted as a feature engineering vector, where the features relevance have been identified by physics domain knowledge. The relationship (Equation 10) can then be rearranged as postulating a linear model $F^\chi(t) = \alpha^\top \Phi(t) + \alpha_0$, where the coefficients α are fitted against climate model data and the intercept α_0 is determined by the pre-industrial concentration C_0^χ .

Because the radiative forcing is a key driving quantity of the climate system, and therefore of the emulator, we argue it is critical to account for the uncertainty introduced by the FaIR model of radiative forcing. Specifically, this choice of a simplified model of forcing reflects uncertainty caused by the lack of knowledge about the phenomenon we wish to describe, referred to as *epistemic uncertainty* (Hüllermeier & Waegeman, 2021). Drawing from the latent force modeling literature (M. Alvarez et al., 2009), we propose to acknowledge the epistemic uncertainty over the forcing model functional form by treating the radiative forcing as an under-specified latent force. Specifically, following the work of M. A. Alvarez et al. (2013), we propose to account for the epistemic uncertainty over the forcing model through a Bayesian formalism and place a GP prior over the radiative forcing.

3.2. Radiative Forcing as a Gaussian Process

To specify a GP prior over the radiative forcing, we must first specify a choice of mean function. We propose to use the deterministic forcing function $F(t)$ as the mean function of our GP. This choice ensures that our prior will




Figure 3. Example of sample paths from the radiative forcing prior $F(t, E(t)) \sim \text{GP}(F, K)$ over the 1850–2014 period. The FaIR forcing response (dashed red) corresponds to the average sample path. The covariance function is taken as a Matérn-3/2 covariance.

behave on average like the FaIR forcing model. Second, we must introduce a covariance function that specifies how different forcing levels should correlate as a function of some input variable.

The radiative forcing level is primarily determined by the emission rate of short-lived forcing agents and the cumulative emissions of long-lived forcing agents. Therefore, we anticipate a correlation of forcing levels for similar long-lived and short-lived forcing agent emissions trajectories. Conversely, for dissimilar emissions trajectories, we expect the resulting forcing levels to be fairly different, resulting in a lower correlation. For this reason, we propose to specify our covariance function as a function of long-lived forcing agents' cumulative emissions and short-lived forcing agents' emission rates.

Namely, let $E_r^X(t)$ and $E_c^X(t)$ be respectively the rate of emission and the cumulative emission of agent χ in the atmosphere at time t . Let $E_c^{\text{LL}}(t) = [E_c^{\chi_1}(t) \dots E_c^{\chi_d}(t)]^\top$ be the vector concatenation of long-lived agents $\{\chi_i\}_{i=1}^p$ cumulative emissions, and let $E_r^{\text{SL}}(t) = [E_r^{\chi_{p+1}}(t) \dots E_r^{\chi_d}(t)]^\top$ be the vector concatenation of short-lived agents $\{\chi_i\}_{i=p+1}^d$ rate of emission. The covariance function we introduce takes the form

$$K(t, t') = \rho(E(t) - E'(t')), \quad (12)$$

where $E(t) = [E_c^{\text{LL}}(t) \ E_r^{\text{SL}}(t)]^\top$ and $E'(t') = [E_c^{\text{LL}}(t') \ E_r^{\text{SL}}(t')]^\top$ are two vectors encapsulating the cumulative emissions of long-lived forcing agents and emission rates of short-lived forcing agents for two different scenarios at times t and t' , and ρ is a user-specified kernel.

With these notations, we can then formally define the prior we specify over the forcing model as

$$F(t, E(t)) \sim \text{GP}(F, K), \quad (13)$$

where we emphasize that F is now a stochastic process by using a sans-serif notation. This Bayesian formulation of the forcing effectively introduces a distribution over the functional forms the forcing response can take, and each sample path drawn from this distribution may describe a different function. Figure 3 shows examples of sample paths from $F(t, E(t))$.

Taking the mean function as the forcing model F from Equation 10 aligns sample paths with the FaIR forcing response on average. However, the probabilistic nature of this prior introduces a relaxation from exactly replicating the FaIR forcing response: sample paths can deviate from the mean, insofar as the forcing trajectories they describe maintain a correlation specified by $K(t, t')$ given their emission trajectories. This degree of freedom introduces a notion of variability over the forcing response, thereby allowing us to account for the epistemic uncertainty over the functional form of the FaIR forcing response.

To concretely understand how the covariance function affects F , recall that for GPs, $\text{Cov}(F(t, E(t)), F(t', E'(t'))) = K(t, t')$. Therefore, if $E(t)$ and $E'(t')$ describe similar emission trajectories, the covariance $K(t, t')$ will take larger values, and specify a strong correlation between $F(t, E(t))$ and $F(t', E'(t'))$. On the other hand, if $E(t)$ and $E'(t')$ describe very different emission trajectories, $K(t, t')$ will take values close to zero, and thus specify a low correlation between $F(t, E(t))$ and $F(t', E'(t'))$. Figure 4 shows an example of a covariance matrix between two scenarios with different CO_2 emissions trajectories.

The choice of kernel ρ is an important choice as it specifies how correlation in forcings should be encoded as a function of emissions. For example, choosing a dirac kernel $\rho(E(t) - E'(t')) = \delta(E(t) - E'(t'))$ makes the forcing levels corresponding to $E(t)$ and $E'(t')$ independent, except if the emission trajectories are identical. Another extreme is to choose $\rho(E(t) - E'(t')) = 1$, which causes the forcing values taken to correlate regardless of emissions. In between these extremes, a widely used class of kernels are the Matérn family (Stein, 1999). They specify an exponentially decaying covariance as a function of the distance between emission trajectories, and




Figure 4. *Left:* Cumulative CO₂ emissions for a scenario with increasing emissions $E(t)$ and a scenario with reduced emissions $E'(t')$. *Right:* Cross-covariance matrix between emission scenarios for a Matérn-3/2 covariance, diagonals give covariances for $t = t'$ and off-diagonals give covariances for $t \neq t'$. For $t, t' \leq 100$ the emission trajectories are similar, therefore $K(t, t')$ specifies a correlation between $F(t, E(t))$ and $F(t', E'(t'))$. For $t \geq 130$, $E(t)$ becomes dissimilar from the reduced emission scenario, therefore $K(t, t')$ specifies a low correlation. For $t' \geq 150$, $E'(t')$ stabilizes to levels similar to E ($t \approx 130$), therefore $K(t, t')$ specifies a correlation between $F(t', E'(t'))$ and $F(t \approx 130, E(t \approx 130))$.

most notably allow to control the smoothness of the GP sample path. Appendix A2 provides detailed expressions of the Matérn covariance family and illustrations of their sample paths.

A popular choice within the Matérn family is the squared exponential kernel, which produces infinitely differentiable GP samples paths (to be precise, the squared exponential kernel corresponds to the limiting case of a Matérn kernel). We argue that this may not be faithful to real-world physical processes, which can exhibit abrupt changes. For this reason, we propose in this work to specify ρ as an anisotropic Matérn-3/2 kernel given by

$$C_{3/2}(E(t) - E'(t')) = \sigma_F^2 \left(1 + \sqrt{3 \sum_{i=1}^d \left(\frac{\delta_i}{\ell_i} \right)^2} \right) \exp \left(-\sqrt{3 \sum_{i=1}^d \left(\frac{\delta_i}{\ell_i} \right)^2} \right), \quad (14)$$

where δ_i is the difference between the i th component of $E(t)$ and $E'(t')$, σ_F^2 is a variance parameter, ℓ_1, \dots, ℓ_d are independent lengthscale parameters associated to each atmospheric agent χ_1, \dots, χ_d . The Matérn-3/2 covariance yields GP sample paths that are differentiable, thus allowing for smoothness, and yet permitting more abrupt changes compared to a squared exponential kernel.

In the remainder of the paper, we permit ourselves to drop notations and we denote the prior $F(t)$ as a function of time only for the sake of conciseness. However, we emphasize that $F(t)$, along with everything derived from it, fundamentally depends on emissions.

3.3. Thermal Response Model With GP Forcing

We will now show that by combining the GP prior over the radiative forcing and the FaIR thermal response model, we obtain a GP prior over temperatures, which we name FaIRGP. Recall the thermal impulse response model presented in Section 2.2, described by a system of k independent linear first order ODEs

$$\frac{dS_i(t)}{dt} = \frac{1}{d_i} (q_i F(t) - S_i(t)), \quad (15)$$

where the thermal responses $S_i(t)$ are such that $T(t) = \sum_{i=1}^k S_i(t)$ is the global mean surface temperature.

We propose to substitute the deterministic forcing function $F(t)$ in Equation 15 by its Bayesian counterpart $F(t)$. In doing so, we naturally induce a stochastic version of the thermal impulse response model. Namely, the resulting i th stochastic thermal response, which we denote $S_i(t)$, will satisfy a differential equation forced by $F(t) \sim \text{GP}(m, K)$, and given by

$$dS_i(t) = \frac{1}{d_i}(q_i F(t) - S_i(t)) dt. \quad (16)$$

The resolution of this stochastic thermal response differential equation is similar to the resolution of the thermal response ODE. Namely, assuming $S_i(0) = 0$ at pre-industrial time, the solution to Equation 16 takes the canonical form

$$S_i(t) = \frac{q_i}{d_i} \int_0^t F(s) e^{-(t-s)/d_i} ds. \quad (17)$$

However, because $F(t)$ is random, the solution is now a stochastic process. In particular, because $F(t)$ is a GP to which we apply a linear convolution with an exponential function $e^{-(t-s)/d_i}$, we can show that $S_i(t)$ must also be a GP, with its mean and covariance functions shaped by the form of this linear convolution.

Namely, the i th thermal response can be characterized as a GP with the following mean and covariance functions

$$\begin{cases} S_i(t) \sim \text{GP}(m_i, k_{ii}) \\ m_i(t) = \frac{q_i}{d_i} \int_0^t F(s) e^{-(t-s)/d_i} ds \\ k_{ii}(t, t') = \left(\frac{q_i}{d_i}\right)^2 \int_0^t \int_0^{t'} K(s, s') e^{-(t-s)/d_i} e^{-(t'-s')/d_i} ds ds' \end{cases} \quad (18)$$

We observe that the mean function $m_i(t)$ exactly corresponds to the solution of the deterministic thermal response ODE (Equation 15). This is consistent with our expectation: because the forcing sample paths from $F(t)$ are specified to behave on average like the FaIR forcing response, the GP thermal response sample paths $S_i(t)$ should behave on average like the FaIR thermal response model. Further, the covariance $k_{ii}(t, t')$ is expressed as a function of the forcing prior covariance K , but also of the parameters of the EBM, d_i and q_i . As such, $k_{ii}(t, t')$ defines a physically-informed covariance structure that propagates the uncertainty over the forcing $F(t)$ —specified by our Bayesian prior—into uncertainty over the thermal response $S_i(t)$.

Under this model, the correlation $k_{ii}(t, t')$ between two stochastic thermal responses $S_i(t)$ and $S_i(t')$ is influenced by how fast the response of the i th thermal box is, and by how much the radiative forcings trajectories leading to $S_i(t)$ and $S_i(t')$ did correlate. Indeed, the covariance in Equation 18, expressed in integral form, specifies a stronger correlation if the past correlation $K(s, s')$ between stochastic forcing trajectories $(F(s))_{s \leq t}$ and $(F(s'))_{s' \leq t'}$ is greater. This aligns with the idea that similar forcing processes $F(s)$ and $F(s')$ should lead to similar thermal responses. Further, a fast thermal box (small d_i) primarily influences the integral with local correlations in forcings just before t and t' , reflecting the short-term impact of fast components of the climate system. Conversely, a slow thermal box (large d_i) accumulates forcing covariance over an extended period before t and t' , reflecting the long-term influence of slow components of the climate system. Figure 5 shows an example of cross-covariance matrices for a stochastic two-box thermal response model.

If we now define the global mean surface temperature as the sum of thermal response GPs $T(t) = \sum_{i=1}^k S_i(t)$, then we can show that $T(t)$ must also be a GP. Namely, let us define the cross-covariance function between thermal responses $S_i(t)$ and $S_j(t')$ as

$$k_{ij}(t, t') := \text{Cov}(S_i(t), S_j(t')) = \frac{q_i q_j}{d_i d_j} \int_0^t \int_0^{t'} K(s, s') e^{-(t-s)/d_i} e^{-(t'-s')/d_j} ds ds', \quad (19)$$




Figure 5. Cross-covariance matrices for the stochastic thermal responses $S_1(t)$ and $S_2(t)$ from a 2-box model under the two emission scenarios from Figure 4. *Left:* Cross-covariance matrix of a fast thermal box with $d_1 = 8$ years and $q_1 = 2 \text{ KW}^{-1} \text{ m}^{-2}$. Only the forcings correlations at times right before (t, t') influence $k_{11}(t, t')$, hence the specified correlation is moderate and the form of the thermal box covariance matrix is similar to the form of the forcing covariance matrix $K(t, t')$ in Figure 4. *Right:* Cross-covariance matrix of a slow thermal box with $d_2 = 80$ years and $q_2 = 20 \text{ KW}^{-1} \text{ m}^{-2}$. The covariance $k_{22}(t, t')$ integrates forcing correlations over an extended period, which specifies a stronger correlation, and additionally models the memory effect of the climate system: even after the two emission scenarios start to diverge for $t, t' \geq 130$, the cross-covariance remains important.

which describes how the stochastic thermal responses from different boxes should correlate within our model. Then, the global mean surface temperature is a GP specified by the following mean and covariance functions

$$\begin{cases} T(t) \sim \text{GP}(m_T, k_T), \\ m_T(t) = \sum_{i=1}^k m_i(t), \\ k_T(t, t') = \sum_{i=1}^k \sum_{j=1}^k k_{ij}(t, t'). \end{cases} \quad (20)$$

The mean function $m_T(t)$ exactly corresponds to the FaIR temperature response model from Equation 3. This is again consistent with the expectation that since $F(t)$ behaves on average like the FaIR forcing model, our GP temperature response $T(t)$ should behave on average like the FaIR temperature response. The covariance function $k_T(t, t')$, which specifies the correlation between $T(t)$ and $T(t')$, includes the covariances of individual thermal boxes $k_{ii}(t, t')$, aligning with the intuition that if thermal responses correlate, so should global mean surface temperatures. However, we note $k_T(t, t')$ also considers correlations between different thermal boxes $k_{ij}(t, t')$ for $i \neq j$. Therefore, if certain components of the climate system exhibit correlated responses across different time-scales, this contributes positively to the correlation between $T(t)$ and $T(t')$ within our model.

By specifying a GP prior over the radiative forcing, we have obtained a GP prior over the temperature that uses the FaIR thermal response model as its backbone, which we name FaIRGP. Using FaIRGP serves as a principled measure of epistemic uncertainty over the emulator design. In particular, the integral form of its covariance function allows us to account for past trajectories, thereby capturing the climate system memory effect, and producing robust uncertainty estimates. In addition, as we will see in Section 3.5, FaIRGP can go beyond a standard impulse response model by learning from data using standard GP regression techniques. We emphasize that whilst we abuse notations for conciseness, the forcing prior $F(t)$ is effectively a function of emissions (or cumulative emissions) through its covariance function $\rho(E(t) - E'(t'))$. Therefore, $T(t)$ is also a function of emissions, and its covariance function can be understood as “ $k_T(E(t), E'(t'))$ ”.

3.4. Accounting for Climate Internal Variability

An important component of the climate system is its internal variability, which integrates the effects of weather phenomena—typically operating on the scale of days—into elements of the climate system, such as the ocean, cryosphere and land vegetation—which rather operate on the scale of months, years or decades (Hasselmann, 1976).

The climate internal variability can classically be modeled in a k -box model by introducing a white noise forcing disturbance over the uppermost box, that is, the atmosphere and land surface box (Cummins et al., 2020). Formally, let $B(t)$ be the standard one-dimensional Brownian motion and let $\mathbf{X}(t)$ denote a stochastic version of the k -box model temperatures from Equation 1. The temperature response model with internal variability is given by

$$d\mathbf{X}(t) = \mathbf{A}\mathbf{X}(t) dt + \mathbf{b}F(t) dt + \sigma \mathbf{b} dB(t), \quad (21)$$

where $\sigma > 0$ is a variance term that controls the amplitude of the white noise, and we recall that \mathbf{b} has zero everywhere but its first entry, therefore the white noise disturbance is only applied to the uppermost box. When the radiative forcing $F(t)$ is deterministic, this is equivalent to adding red noise onto the global mean surface temperature, or equivalently, modeling it as an Ornstein-Uhlenbeck process. In the literature, it is common to encounter its discrete time analog, the autoregressive process of order 1 AR(1). The corresponding diagonalized impulse response system is given by

$$dS_i(t) = \frac{1}{d_i}(q_i F(t) - S_i(t)) dt + \sigma \frac{q_i}{d_i} dB(t), \quad (22)$$

where derivations are detailed in Appendix B3. If we now again substitute the deterministic forcing $F(t)$ with the GP forcing $F(t)$, it can be shown that the long time regime solution to this stochastic impulse response system can be expressed as

$$S_i(t) \sim \text{GP}(m_i, k_{ii} + \sigma^2 \gamma_i), \quad (23)$$

where $\gamma_i(t, t')$ is an exponential, or Matérn-1/2, kernel function given by

$$\gamma_i(t, t') = \frac{q_i^2}{2d_i} \exp\left(-\frac{|t - t'|}{d_i}\right). \quad (24)$$

Therefore, FaIRGP can easily model the internal variability of the climate system simply by modifying its covariance structure. We observe that accounting for internal variability essentially corresponds to adding an independent autocorrelated noise process $\xi_i(t) \sim \text{GP}(0, \sigma^2 \gamma_i)$ to the thermal response GP obtained in Equation 18.

Ultimately, we sum the thermal response together to obtain a surface temperature GP $T(t) = \sum_{i=1}^k S_i(t)$, which in the long time regime takes the form

$$\begin{cases} T(t) \sim \text{GP}(m_T, k_T + \sigma^2 \gamma_T) \\ \gamma_T(t, t') = \sum_{i=1}^k \nu_i \gamma_i(t, t'), \end{cases} \quad (25)$$

where ν_i is a dimensionless weight determined by q_i, d_i which accounts for cross-covariances between thermal responses internal variabilities. Its detailed expression can be found in Appendix B3.

Therefore, we can account for the climate internal variability simply by adding an independent noise process $\xi_T(t) = \sum_{i=1}^k \xi_i(t) \sim \text{GP}(0, \sigma^2 \gamma_T)$ to the temperature response GP obtained in Equation 20. Figure 6 shows examples of sample paths from the FaIRGP prior with and without accounting for internal variability.

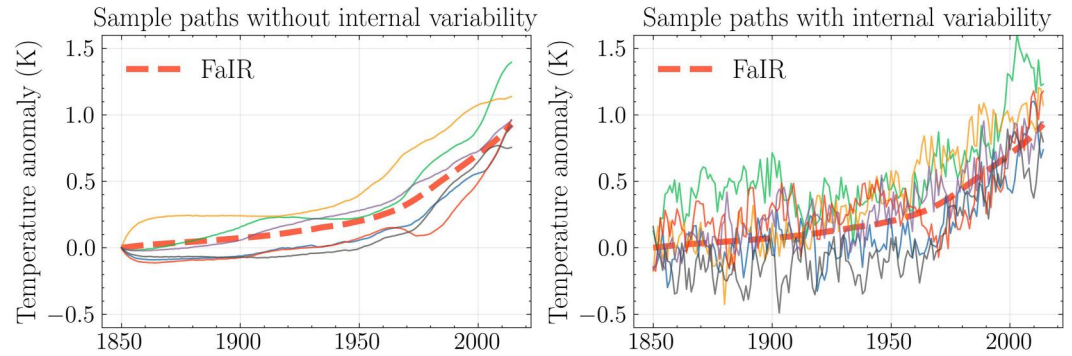


Figure 6. *Left:* Example of sample paths from the FaIRGP prior $T(t) \sim GP(m_T, k_T)$ without internal variability over the 1850–2014 period. *Right:* Example of sample paths from the FaIRGP prior with internal variability $T(t) \sim GP(m_T, k_T + \sigma^2 \gamma_T)$ over the same period. The FaIR temperature response (dashed red) corresponds in both cases to the average sample path.

3.5. Posterior Distribution Over Temperature and Radiative Forcing

In FaIRGP, the surface temperature response and the radiative forcing model can both be informed by global temperature data. Using standard GP regression techniques (Rasmussen & Williams, 2005), FaIRGP can learn from data how to deviate from its backbone impulse response model to best account for actual temperature simulated by a climate model. The resulting model can then be used to emulate future temperatures, benefiting from both the robustness of its prior and the flexibility of GP regression. As we will demonstrate in Section 5, we can inform FaIRGP with historical observations and climate model data.

Assume we are under a fixed emission scenario for example, working with historical emissions only. For timesteps $t_1 < \dots < t_n$, suppose we simulate global mean surface temperatures T_1, \dots, T_n , and also have access to greenhouse gas and aerosol emissions data $E_1, \dots, E_n \in \mathbb{R}^d$, where d corresponds to the number of atmospheric agents χ_1, \dots, χ_d . We concatenate this data into $\mathbf{t} = [t_1 \dots t_n]^T$, $\mathbf{E} = [E_1 \dots E_n]^T$ and $\mathbf{T} = [T_1 \dots T_n]^T$.

With notation abuse, the kernel k_T allows to specify how the FaIRGP prior correlates between times t_i and t_j by computing the similarity between emission data, that is, $k_T(t_i, t_j) = k_T(E_i, E_j)$, because the forcing prior kernel ρ is a function of emissions. The internal variability kernel γ_T allows us to specify how the additive internal variability noise process correlates between times t_i and t_j by computing $\gamma_T(t_i, t_j)$. By extending this to the entire data set, we can define the covariance matrices induced by k_T and γ_T over \mathbf{E} and \mathbf{t} following

$$\mathbf{K} = k_T(\mathbf{E}, \mathbf{E}) = [k_T(E_i, E_j)]_{1 \leq i, j \leq n} \quad (26)$$

$$\mathbf{\Gamma} = \gamma_T(\mathbf{t}, \mathbf{t}) = [\gamma_T(t_i, t_j)]_{1 \leq i, j \leq n} \quad (27)$$

Intuitively, \mathbf{K} captures how much the values taken by the stochastic global mean temperature response from the FaIRGP prior will correlate between different times. Similarly, $\mathbf{\Gamma}$ captures how much the values taken by the internal variability process ξ_T correlate between different time steps. Combining them into $\mathbf{K} + \sigma^2 \mathbf{\Gamma}$, we obtain a covariance matrix that captures the total correlation we have between different times for the values taken by the FaIRGP prior with internal variability.

Because it captures both the effect of internal variability and of the global mean temperature response to changes in emissions, the covariance matrix $\mathbf{K} + \sigma^2 \mathbf{\Gamma}$ is well suited to capture the correlation structure we get from global mean surface temperature data \mathbf{T} simulated with a climate model. Using this new covariance matrix, we can then apply the update rules from Equations 8 and 9 to update our prior over $T(t)$ with climate model data \mathbf{T} , and obtain a posterior GP over global mean surface temperatures given by

$$T(t) | \mathbf{T} \sim GP(\bar{m}_T, \bar{k}_T), \quad (28)$$




Figure 7. *Left:* Historical global mean surface temperature anomaly simulated with NorESM2-LM from which we randomly extract 10 data points (purple). *Right:* Example of sample paths from the FaIRGP posterior $T(t)|T$ conditioned on the data. In contrast with the FaIRGP prior sample paths (see Figure 6), posterior sample paths are more concentrated around data points and deviate from the FaIR temperature response (dashed red).

where the posterior mean $\bar{m}_T(t)$ and the posterior covariance $\bar{k}_T(t, t')$ have the following closed-form expressions

$$\underbrace{\bar{m}_T(t)}_{\text{posterior mean}} = \underbrace{m_T(t)}_{\text{prior mean}} + \underbrace{k_T(t, \mathbf{t})(\mathbf{K} + \sigma^2 \mathbf{I})^{-1}(\mathbf{T} - m_T(\mathbf{t}))}_{\text{posterior correction}} \quad (29)$$

$$\underbrace{\bar{k}_T(t, t')}_{\text{posterior covariance}} = \underbrace{k_T(t, t')}_{\text{prior covariance}} - \underbrace{k_T(t, \mathbf{t})(\mathbf{K} + \sigma^2 \mathbf{I})^{-1}k_T(\mathbf{t}, t')}_{\text{posterior correction}}. \quad (30)$$

The posterior mean $\bar{m}_T(t)$ explicitly corrects the prior mean $m_T(t)$, which we recall corresponds to the FaIR deterministic temperature response model. The correction introduced by the posterior mean in Equation 29 is a data-driven correction term, informed by the temperature data \mathbf{T} , times \mathbf{t} , and the emissions \mathbf{E} (which are used to compute the covariance matrix \mathbf{K}). Intuitively, the correction term regresses the discrepancy between the simulated temperature data and the FaIR temperature response onto the emission data. Thus, it learns how much to deviate from the FaIR temperature response model to better fit climate model data, given emissions data. Similarly, the posterior covariance $\bar{k}_T(t, t')$ explicitly adjusts the prior covariance $k_T(t, t')$ with a data-driven correction term. Intuitively, this correction term acknowledges the new information we have updated our model with, such that when predicting over emission trajectories similar to the ones from the training data, the variability of the FaIRGP global mean surface temperature response is reduced. Figure 7 shows examples of sample paths from the FaIRGP posterior when updated with simulated data of global mean surface temperature.

This posterior can then be used to emulate temperatures by making predictions at future time steps. A major interest of this formulation is that whilst the emulator hinges on the robustness of the EBM for its prior, it can also benefit from the flexibility of GP regression by informing it with data in its posterior, thereby going beyond a simple impulse response model. In addition, the GP approach enjoys full probabilistic tractability, with the analytical expression of the posterior probability distribution provided in Appendix B5. This means we can analytically compute the mean, the 95% credible range and quantiles of the predicted surface temperature anomaly as shown in Figure 8.

Using the same data, we can also update the GP prior over the radiative forcing with temperature data to formulate a posterior distribution over the forcing. Indeed, because the radiative forcing $F(t)$ is also GP, it is jointly Gaussian with the temperature $T(t)$, and thus admits a cross-covariance function that we can derive following

$$k_{FT}(t, t') := \text{Cov}(F(t), T(t')) = \sum_{i=1}^k \frac{q_i}{d_i} \int_0^{t'} K(t, s') e^{-(t'-s')/d_i} ds'. \quad (31)$$




Figure 8. Left: FaIRGP posterior mean and 95% credible range. Right: FaIRGP posterior quantiles. The posterior distribution $T(t)|T$ is conditioned on the same 10 random data points used in Figure 7.

Intuitively, this cross-covariance function plays a mediation role between the temperature and radiative forcing domains: it translates how variations in surface temperatures correlate with variations in radiative forcing under our model. Using this cross-covariance, it is therefore possible to update our prior over $F(t)$ with temperature data T , resulting in a posterior distribution over the radiative forcing given by

$$\begin{cases} F(t)|T \sim GP(\bar{m}_F, \bar{k}_F) \\ \bar{m}_F(t) = F(t) + k_{FT}(t, \mathbf{t})(\mathbf{K} + \sigma^2 \mathbf{\Gamma})^{-1}(\mathbf{T} - m_T(\mathbf{t})) \\ \bar{k}_F(t, t') = K(t, t') - k_{FT}(t, \mathbf{t})(\mathbf{K} + \sigma^2 \mathbf{\Gamma})^{-1}k_{FT}(t', \mathbf{t})^\top. \end{cases} \quad (32)$$

As for temperatures, the posterior mean $\bar{m}_F(t)$ corrects the FaIR forcing model $F(t)$ with a data-informed correction term. The mean correction term again regresses out the discrepancy between T and the FaIR temperature response onto emissions, but converts this temperature discrepancy signal into a forcing discrepancy signal through the cross-covariance function k_{FT} . The uncertainty quantification given by the posterior covariance $\bar{k}_F(t, t')$ also adjusts the forcing prior covariance $K(t, t')$ to reduce the variability of the stochastic forcing response when predicting over emission trajectories similar to the ones from the training data E . Figure 9 shows examples of sample paths from the forcing posterior when updated with global mean surface temperature data.

Finally, we can verify that solving the thermal response stochastic differential equation (SDE) for the posterior forcing $F(t)|T$ yields the posterior temperature $T(t)|T$ as the solution. Therefore, the forcing posterior and the temperature posterior are consistent with each other. It is important to note that the posterior forcing correction proposed in Equation 32 may correct the forcing prior with sources of temperature discrepancies that are not imputable to the forcing. Indeed, beyond differences between the FaIR and the climate model forcing response, other aspects of the climate system, such as the climate feedback parameter or the ocean inertia, can also contribute to discrepancies between T and the FaIR temperature response. Because the posterior correction reports all of the temperature response discrepancies onto the forcing, there is a risk of simply compensating errors, which needs cautious examination.

3.6. Parameters Tuning

Access to a closed-form probability density for the prior allows us to tune the model parameters using a maximum likelihood strategy. To illustrate this, consider again the data \mathbf{t} , E , T from the previous section. We can derive a closed-form expression for the *marginal log-likelihood* $\log p(T|E, \mathbf{t})$ of the data under our FaIRGP prior, where the exact formula can be found in Appendix B5. The marginal log-likelihood represents (the logarithm of) how




Figure 9. Example of sample paths from the radiative forcing posterior $F(t)|T$ conditioned on the same 10 random data points used in Figure 7. In contrast with the forcing prior sample paths (see Figure 3), the posterior sample paths are informed by temperature data and deviate from the FaIR forcing response (dashed red).

likely is the simulated data \mathbf{t} , \mathbf{E} , \mathbf{T} to occur under the FaIRGP prior we have specified. Hence, a well-specified prior yields a greater marginal log-likelihood. This motivates a tuning strategy for the model parameters that seeks parameter values that maximize $\log p(\mathbf{T}|\mathbf{E}, \mathbf{t})$.

This can be achieved by using a gradient-based optimization algorithm to maximize the marginal log-likelihood $\log p(\mathbf{T}|\mathbf{E}, \mathbf{t})$ with respect to the model parameters. Namely, let $\theta = \{\sigma, \sigma_F, \ell_1, \dots, \ell_d\}$ be the internal variability standard deviation and the kernel $K(t, t')$ parameters we wish to tune. We can compute the direction in which a small shift in θ can be taken to increase $\log p(\mathbf{T}|\mathbf{E}, \mathbf{t})$. This direction corresponds to the gradient of the marginal log-likelihood with respect to the model parameters. By iteratively performing this operation, each *gradient step* brings us closer to an optimal choice of model parameters θ that maximizes $\log p(\mathbf{T}|\mathbf{E}, \mathbf{t})$. This procedure is known in machine learning as *gradient descent* and illustrated in Algorithm 1. The parameter $\eta > 0$ is typically a small positive number called the learning rate which controls the size of the steps taken in the direction of the gradient at each step. A detailed discussion of its role and of the parameter tuning procedure with gradient descent is provided in Appendix B5.

Algorithm 1. Maximum Likelihood Tuning of θ With Gradient Descent

```

1: Input:  $\mathbf{t}, \mathbf{E}, \mathbf{T}, \theta_0, \eta > 0, n\_steps$ 
2: for  $t \in \{1, \dots, n\_steps\}$  do
3:   Compute  $\log p(\mathbf{T}|\mathbf{E}, \mathbf{t})$ 
4:   Take  $\theta_t \leftarrow \theta_{t-1} + \eta \nabla_{\theta} \log p(\mathbf{T}|\mathbf{E}, \mathbf{t})$ 
5: end for
6: Return:  $\theta_{n\_steps}$ 

```

Finally, note that whilst we limit ourselves to tuning the GP and internal variability parameters in this work, we may also want to tune FaIR parameters such as d_i , q_i , the forcing model coefficients $\alpha_{\log}^{\ell}, \alpha_{\text{lin}}^{\ell}, \alpha_{\text{sqr}}^{\ell}$ or even the temperature response parameters \mathbf{A} , \mathbf{b} . This can be achieved using the same procedure by including these parameters in θ .

3.7. Spatial FaIRGP

Whilst informative, the global mean surface temperature fails to capture the difference in exposure of world regions to a changing climate. It is therefore a necessity to invest efforts in obtaining spatially-resolved climate projections. We propose an extension of FaIRGP to emulate spatially-resolved surface temperature maps that grounds itself on a *pattern scaling* prior.

Pattern scaling is a well-established technique to model changes in local surface temperature as a function of changes in global mean surface temperature (Huntingford & Cox, 2000; Mitchell et al., 1999; Santer et al., 1990; Tebaldi & Arblaster, 2014). It consists of a simple scaling of a fixed spatial pattern by global mean temperature changes. Whilst very simple, such approach is supported by findings that regional changes in temperature scale robustly with global temperature (S. I. Seneviratne et al., 2016), and has been successfully used in existing spatial temperature emulation models (Beusch et al., 2020a, 2020b, 2022; Goodwin et al., 2020; Link et al., 2019).

Formally, pattern scaling assumes that for a given spatial location x , the local temperature response $T(x, t)$ is given by

$$T(x, t) = \beta^{(x)} T(t) + \beta_0^{(x)}, \quad (33)$$

with regression coefficient $\beta^{(x)}$ and intercept $\beta_0^{(x)}$. These coefficients are typically obtained by fitting independent local linear regression models of global temperature $T(t)$ onto local temperature $T(x, t)$. If we substitute the deterministic global temperature response $T(t)$ with its GP version $\mathbf{T}(t)$, we obtain a local FaIRGP prior temperature response given by

$$\mathbf{T}(x, t) \sim \text{GP}\left(\beta^{(x)} m_{\mathbf{T}} + \beta_0^{(x)}, (\beta^{(x)})^2 k_{\mathbf{T}}\right), \quad (34)$$

Table 2
Details of the Experiments From ClimateBench v1.0 (Watson-Parris et al., 2022) Used in the Data Set

| Experiment | Period | Description |
|------------|-----------|--|
| Historical | 1850–2014 | A simulation using historical emissions of all forcing agents designed to recreate the historically observed climate |
| SSP126 | 2015–2100 | A high ambition scenario designed to produce significantly less than 2° warming by 2100 |
| SSP245 | 2015–2100 | A medium forcing future scenario aligned with current trajectories |
| SSP370 | 2015–2100 | A medium-high forcing scenario with high emissions of near-term climate forcers such as methane and aerosols |
| SSP585 | 2015–2100 | A high forcing scenario leading to a large 8.5 W m^{-2} forcing in 2100 |

which admits the local pattern scaling temperature response for its mean, and a locally scaled covariance.

As for the global response, this local prior can be updated with local temperature data to obtain a posterior temperature response at spatial location x . This gives FaIRGP the flexibility to learn from data how to deviate from a fixed spatial pattern, in order to better account for climate model simulated data.

4. Experimental Setup

In this section, we introduce the data set used in our emulation experiments, the baseline models we benchmark FaIRGP against, and the evaluation metrics. Code and data to reproduce experiments are publicly available (<https://github.com/shahineb/FaIRGP>).

4.1. Data Set Description

The data is obtained from the ClimateBench v1.0 (Watson-Parris et al., 2022) climate emulation benchmark data set. ClimateBench v1.0 proposes a curated data set of local annual mean surface temperature ($\sim 2^\circ$ horizontal resolution), paired with annual emissions for four of the main anthropogenic forcing agents: carbon dioxide (CO_2), methane (CH_4), sulfur dioxide (SO_2) and black carbon (BC). The temperature data is generated from the latest version of the Norwegian Earth System Model (NorESM2) (Seland et al., 2020) as part of the sixth CMIP6 (Eyring et al., 2016). The emission data is constructed from the Community Emissions Data System (Hoesly et al., 2018), and corresponds to the input data used to drive the original NorESM2-LM simulations.

The data set we use includes paired time series of global annual emissions and local annual mean surface temperature for five experiments. We include the CMIP6 *historical* experiment, and four experiments corresponding to different possible shared socio-economic pathways (SSPs) (Riahi et al., 2017) from the ScenarioMIP protocol (O'Neill et al., 2016): *SSP126*, *SSP245*, *SSP370* and *SSP585*. These scenarios are designed to span a range of emissions trajectories corresponding to plausible mitigation scenarios and end-of-century forcing possibilities. Table 2 gives a description of the experiments and the periods covered.

In all experiments, we use as inputs for FaIRGP the global annual cumulative emissions of CO_2 and global annual emissions of CH_4 , SO_2 and BC. In the experiments of Section 5, we focus on emulating global annual mean surface temperature. Therefore, we preprocess the temperature data by computing the global annual mean surface temperatures using an area-weighted mean of the local annual mean surface temperatures data. In Section 6 we focus on emulating local annual mean surface temperatures, thus no preprocessing is required.

Our experiments are effectively conducted on data from a single climate model, and therefore do not reflect the diversity of response of different climate models to a common scenario. However, we argue they still provide a valid assessment of an emulator's ability to reproduce climate models outputs. Indeed, whilst different climate models exhibit variations in their response to a similar scenario, it has been shown that SCMs are still capable of capturing the diversity of forcing responses spanned (Smith et al., 2018). Therefore, there are reasons to believe that insights drawn from experiments where emulators are tuned on NorESM2 data should carry over to experiments where emulators are tuned on data from different climate models.

Table 3
Details on Evaluation Metrics Used in Experiments

| | Notation | Description | Best when |
|---------------|----------|-------------------------------------|--------------|
| Deterministic | RMSE | Root mean square error | Close to 0 |
| | MAE | Mean absolute error | Close to 0 |
| | Bias | Mean of (prediction – groundtruth) | Close to 0 |
| Probabilistic | LL | Log-likelihood | Higher |
| | Calib95 | 95% calibration score | Close to 95% |
| | CRPS | Continuous ranked probability score | Close to 0 |

4.2. Baseline Emulators

To develop intuition on the benefit of combining FaIR with a GP, we propose to compare our model to the temperature anomaly projections obtained with (a) FaIR only and (b) with a purely data-driven plain GP regression model only. By comparing to the emulation with FaIR, we hope to highlight the benefit that there is to combine the flexibility of a data-driven approach to an impulse response model, and answer the question “*Does introducing a GP in FaIR add value?*”. On the other hand, by comparing FaIRGP with a plain GP regression model, we hope to demonstrate the importance of having a robust physical prior underlying a data-driven approach, and answer the question “*Does a FaIR prior in a GP add value?*”. Both baseline models take as inputs greenhouse gas and aerosol emissions data, and predict annual mean surface temperature anomalies with respect to the preindustrial period.

For FaIR, we use parameter values that have been tuned against NorESM2 simulations (Leach et al., 2021). Therefore, the FaIR model we use in experiments is fully deterministic with fixed parameter values.

The plain GP model is entirely physics-free, and is specified with a zero mean and an anisotropic Matérn-3/2 kernel. This is exactly the same kernel ρ we use in our prior over the forcing in FaIRGP, which takes as inputs the global annual cumulative emissions of CO₂ and global annual emissions of CH₄, SO₂ and BC. This kernel is to contrast with the physics-informed kernel k_T used for temperatures in FaIRGP. As it is standard practice in GP regression, we assume Gaussian noise with constant variance over the data. The noise variance and the kernel parameters are tuned using marginal likelihood maximization by gradient descent.

4.3. Evaluation Metrics

We use two kinds of metrics to evaluate the predicted probabilistic surface temperatures: deterministic metrics, which compare only the posterior mean prediction to the ClimateBench temperature data, and probabilistic metrics, which evaluate the entire posterior probability distribution against ClimateBench temperature data. Table 3 provides a brief description of the metrics used.

The log-likelihood (LL) score evaluates the log probability of the groundtruth temperature data under the predictive posterior probability distribution predicted by our model. Therefore, greater LL means that the predicted distribution is a good fit for the test data. The 95% calibration score (Calib95) computes the percentage of groundtruth temperature data that fall within the 95% credible interval of the predicted posterior distribution. Therefore, if the predicted probability distribution is well calibrated, this score should be close to 95%.

The Continuous Ranked Probability Score (CRPS) is an extension of the root mean square error (RMSE) for probability distributions, which measures the discrepancy between the cumulative distribution functions (cdfs) of the predicted posterior distribution and the empirical cdf of the test data. Namely, if Φ denotes the cdf of the predicted posterior distribution and y is a scalar test data point, the CRPS between Φ and y is given by

$$\text{CRPS}(\Phi, y) = \int_{-\infty}^{+\infty} (\Phi(x) - \mathbb{1}_{[y, +\infty)}(x))^2 dx, \quad (35)$$

where $\mathbb{1}_{[y, +\infty)}(x)$ is the indicator function which takes value 0 if $x < y$ and 1 if $x \geq y$.




Figure 10. Emulation of NorESM2-LM SSP370 global mean surface temperature anomaly. *Left:* FaIRGP prior temperature response over SSP370. The prior mean response exactly corresponds to the FaIR temperature response, but with prior uncertainty quantification. *Middle:* Posterior correction for the mean and variance that learns from data how to deviate from the prior. The pink area corresponds to deviations toward greater values and the green area deviations toward smaller values. *Right:* FaIRGP posterior temperature response when updated with the other scenarios. The posterior response provides a sounder fit to the climate model data, with better calibrated uncertainty. Shading indicates a 5%–95% range.

When evaluating predictions of local annual mean surface temperatures, we compute global mean metrics using an area-weighted mean that accounts for the decreasing grid-cell area toward the poles. Namely, we take

$$\langle \text{Score} \rangle = \frac{1}{N_{\text{lon}} \sum_{i=1}^{N_{\text{lat}}} w_i} \sum_{i=1}^{N_{\text{lat}}} \sum_{j=1}^{N_{\text{lon}}} w_i \text{Score}_{i,j}, \quad (36)$$

where $w_i = \cos(\text{lat}(i))$.

5. Application: Global Surface Temperatures Emulation

In this section, we benchmark FaIRGP against baseline models for the task of emulating mean global surface temperatures over SSP scenarios. We first briefly illustrate how the model concretely applies to an example. Then, we evaluate the emulated global temperatures trajectories when the model is trained on historical and SSP data, and when trained on historical data only. Finally, we probe the potential of FaIRGP to infer the top-of-atmosphere radiative forcing.

5.1. Example of Global Temperature Emulation With FaIRGP

We begin with an example, and propose in Figure 10 a concrete illustration of global mean surface temperature emulation for SSP370 with FaIRGP, that parallels the posterior mean and covariance expressions from Equations 29 and 30.

The prior temperature response over SSP370 admits as its mean the FaIR response, but with an additional layer of uncertainty quantification that arises from the GP. We then use temperature data from scenarios $D_{\text{train}} = \{\text{historical}, \text{SSP126}, \text{SSP245}, \text{SSP585}\}$ as training data to learn a posterior correction. The posterior correction allows deviation from the prior—both in mean and variance—to provide a better fit to the training data. Finally, by linearly adding the posterior correction to the prior, we obtain a posterior temperature response over SSP370.

5.2. Shared Socio-Economic Pathways Emulation

In this experiment, we consider the experiment data set given by $D = \{\text{historical}, \text{SSP126}, \text{SSP245}, \text{SSP370}, \text{SSP585}\}$. We iteratively remove one SSP experiment from the data set to construct a training set (e.g., retain SSP245 to obtain $D_{\text{train}} = \{\text{historical}, \text{SSP126}, \text{SSP370}, \text{SSP585}\}$) and use it to train FaIRGP and the baseline GP model. We run predictions over the retained SSP experiment and evaluate the emulated global temperatures against NorESM2-LM data. We find that on average, FaIRGP outperforms the baseline models in every evaluation metric, as reported in Table 4. These results suggest that FaIRGP provides a better emulation of the global temperature response to anthropogenic forcing compared to both FaIR and the plain GP model. We include in Appendix C2 a comparison with the GP emulator from ClimateBench (Watson-Parris et al., 2022).

Table 4

Scores of Baseline Emulators and FaIRGP for the Task of Emulating Shared Socio-Economic Pathways (SSPs) Global Temperatures Over 2015–2100 Period

| Test experiment | Emulator | RMSE ↓ | MAE ↓ | Bias | LL ↑ | Calib95 | CRPS ↓ |
|-----------------|---------------------|--------------------|--------------------|---------------------|--------------------|--------------------|--------------------|
| SSP126 | FaIR | 0.17 | 0.13 | 0.06 | – | – | – |
| | Plain GP | 0.18 | 0.14 | 0.08 | 0.23 | 1.0 | 0.10 |
| | FaIRGP [†] | 0.15 | 0.12 | 0.02 | 0.48 | 0.97 | 0.08 |
| SSP245 | FaIR | 0.14 | 0.12 | 0.06 | – | – | – |
| | Plain GP | 0.09 | 0.07 | 0.03 | 0.67 | 1.0 | 0.06 |
| | FaIRGP [†] | 0.09 | 0.08 | −0.01 | 0.71 | 1.0 | 0.06 |
| SSP370 | FaIR | 0.22 | 0.19 | 0.10 | – | – | – |
| | Plain GP | 0.22 | 0.17 | −0.16 | 0.12 | 1.0 | 0.12 |
| | FaIRGP [†] | 0.17 | 0.14 | 0.06 | 0.39 | 0.94 | 0.10 |
| SSP585 | FaIR | 0.28 | 0.22 | −0.07 | – | – | – |
| | Plain GP | 0.30 | 0.21 | −0.12 | 0.16 | 1.0 | 0.14 |
| | FaIRGP [†] | 0.21 | 0.17 | −0.08 | 0.18 | 0.88 | 0.12 |
| Mean | FaIR | 0.20 ± 0.06 | 0.16 ± 0.05 | 0.04 ± 0.08 | – | – | – |
| | Plain GP | 0.20 ± 0.09 | 0.15 ± 0.06 | −0.04 ± 0.11 | 0.30 ± 0.25 | 1.0 ± 0.0 | 0.11 ± 0.03 |
| | FaIRGP [†] | 0.16 ± 0.05 | 0.12 ± 0.04 | −0.01 ± 0.06 | 0.44 ± 0.22 | 0.95 ± 0.05 | 0.09 ± 0.02 |

Note. The test experiment is retained from the training data and only used for evaluation; the best emulator scores for each test experiment are highlighted in bold; ↑/↓ indicates higher/lower is better; we report 1 standard deviation over the mean scores; [†] indicates our proposed method.

Figure 11 shows that FaIRGP provides a better temperature projection than FaIR in the near future—over the 2015–2050 period—for all scenarios. This is because being informed by data grants FaIRGP the flexibility to deviate from the impulse response prior on which it hinges, and provide predictions that are better aligned with the historical and near future data from the training set. However, when further away from the training data, over the 2080–2100 period, FaIRGP reverts back to the prior behavior of FaIR. This suggests that FaIRGP is better suited

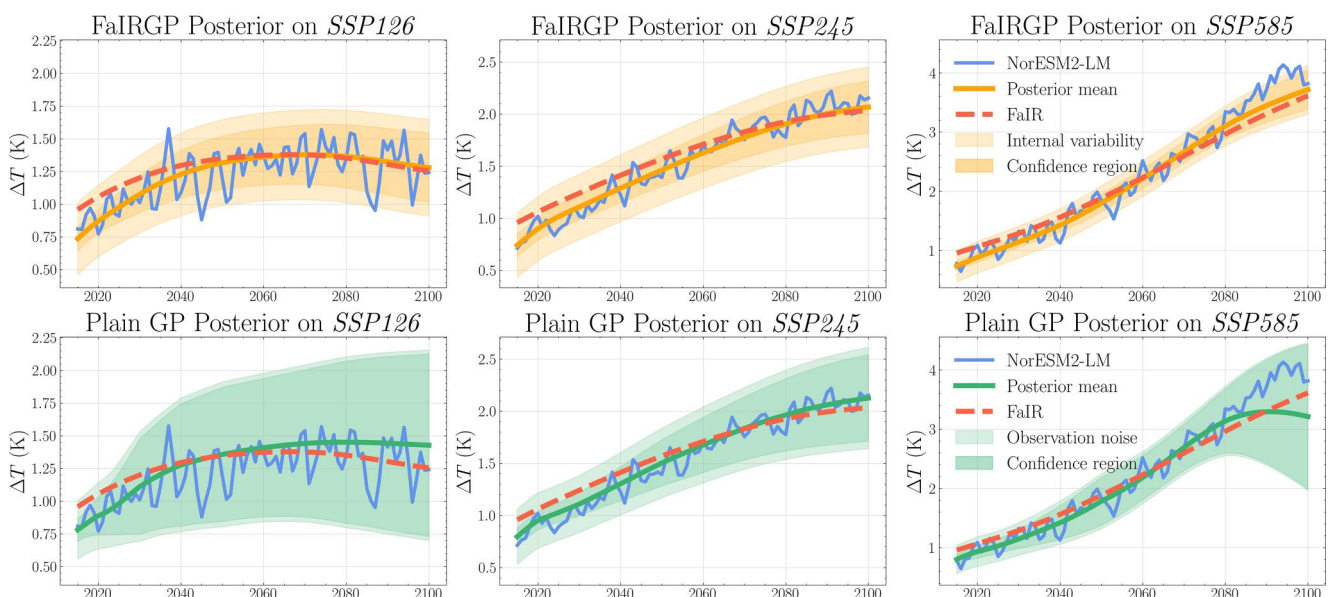


Figure 11. Emulated global mean surface temperature anomaly over 2015–2100 period with our FaIRGP emulator (top) and with a baseline Gaussian process (GP) emulator (bottom) for a low-forcing future scenario (SSP126), a medium forcing future scenario (SSP245) and a high forcing future scenario (SSP585).

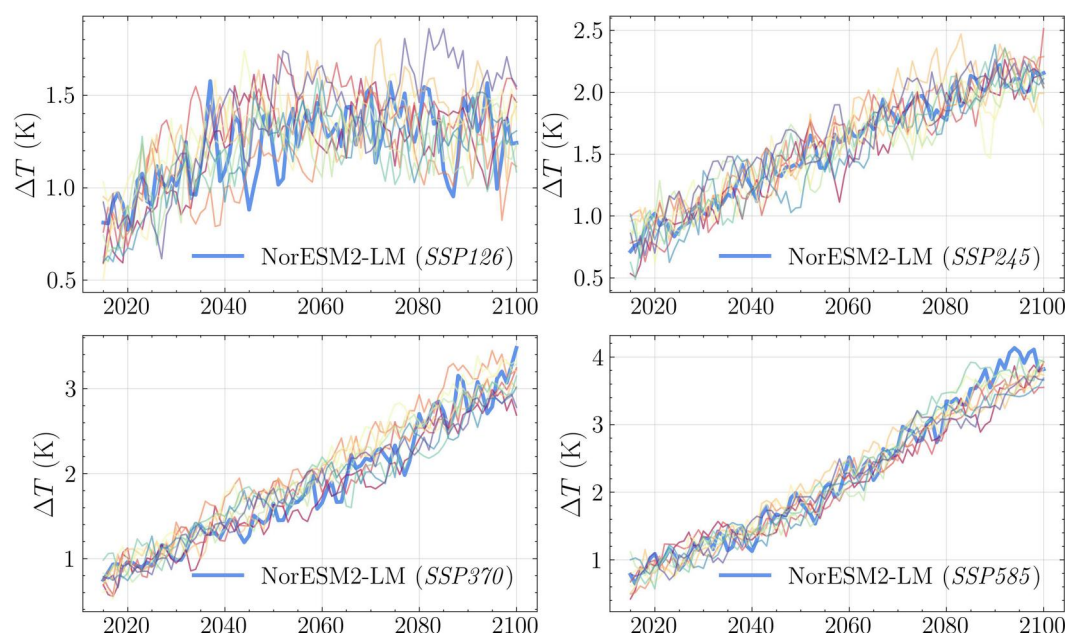


Figure 12. Example of sample paths from the FaIRGP posterior for the emulation of *SSP126* (top-left), *SSP245* (top-right), *SSP370* (bottom-left) and *SSP585* (bottom-right) global mean surface temperature anomaly over the 2015–2100 period.

for emulating global temperature response in the near future, and is at least as good as FaIR for longer term emulation.

The plain GP model provides excellent predictions on *SSP245*, and even outperforms FaIRGP by a slight margin on deterministic metrics for this scenario. This is because GPs are notorious for excelling at interpolation tasks, and the *SSP245* scenario is a medium forcing scenario that perfectly lies within the range of the other low and high forcing scenarios used in the training data. However, when the evaluation scenario lies outside of the range of the training data, the plain GP model predictions will simply revert to the prior, and are therefore much less reliable. This is an important drawback of purely data-driven emulation, which by essence are interpolation models, and struggle to extrapolate outside the range of the training set. In Figure 11, this is particularly evident in the plain GP prediction on *SSP585*, where the model struggles at the end of the century due to unprecedented forcing levels that have not been seen in the training data. FaIRGP addresses this shortcoming by using FaIR as its prior, and therefore naturally reverts to an impulse response model when the prediction scenario is too distant from the training data.

Finally, FaIRGP provides a tighter and more robust uncertainty quantification over emulated temperatures compared to the plain GP model, which tends to overestimate the variance. This is reflected in Table 4 by a substantial improvement in LL of FaIRGP against the Plain GP, and suggests that the physics-informed covariance structure of FaIRGP is a sound choice to quantify emulator uncertainty. Further, the model formulated is able to quantify the uncertainty introduced by the internal variability separately from the uncertainty introduced by the emulator design. As shown in Figure 11, the uncertainty due to internal variability dominates in the near future, but the emulator uncertainty becomes predominant toward the end of the century, in accordance with the well-established assessment of Hawkins and Sutton (2009).

In Figure 11, the posterior mean lines and the shaded area depict the mean function and the 95% credible region of the predicted posterior distribution over temperatures using FaIRGP. Since this represents a probability distribution, we can generate sample paths from the posterior GP (which will fall within the shaded area with probability 0.95). Figure 12 displays examples of sample paths from the FaIRGP posterior for each scenario. It shows that the smooth shaded areas hide the specific form of the functions described by the FaIRGP posterior. The sample paths are in fact capable of reproducing the variability of abrupt annual changes in global mean surface temperature anomalies. Further, the sample paths are generally well-concentrated around the NorESM2 simulated temperature anomaly, with the exception of the end of century temperature anomaly under *SSP585* which tends to

Table 5
Scores of Baseline Emulators and FaIRGP When Models Are Only Trained Using the Historical Experiment

| Emulator | RMSE ↓ | MAE ↓ | Bias | LL ↑ | Calib95 | CRPS ↓ |
|---------------------|--------------|--------------|---------------|--------------|--------------|--------------|
| FaIR | 0.208 | 0.163 | 0.038 | — | — | — |
| Plain GP | 1.875 | 1.690 | −1.690 | −30.979 | 0.009 | 1.558 |
| FaIRGP [†] | 0.191 | 0.140 | −0.037 | 0.314 | 0.948 | 0.102 |

Note. Scores are evaluated and averaged for the task of emulating *SSP126*, *SSP245*, *SSP370*, *SSP585* Global surface temperatures over 2015–2100 period; the best emulator for each metric is highlighted in bold; ↑/↓ indicates higher/lower is better; [†] indicates our proposed method.

be underestimated—consistently with the shaded area from Figure 11. Appendix C1 provides additional plots of individual sample paths from the FaIRGP posterior over SSPs.

5.3. Emulation From Historical Data

In this experiment, we propose to investigate how well can the emulators generalize when they are solely trained on historical data, and no simulated future data (with the exception of the FaIR parameters which have been calibrated using NorESM2 outputs for future scenarios.). We train the plain GP and FaIRGP using the *historical* experiment only, and evaluate predictions on all SSP scenarios. The results are reported in Table 5. We find that FaIRGP outperforms other models in almost every metric, and displays performance similar to FaIR only in mean bias.

In this context, a purely data-driven method struggles to extrapolate because it cannot marry the new previously unseen emission values with the underlying physical model. Figure 13 shows that the plain GP posterior immediately reverts to its prior after 2015, therefore providing uninformative temperature projections. On the other hand, FaIRGP manages to learn from historical data to provide a better fit to the SSP over the 2015–2050 period, and slowly reverts back to the prior behavior of FaIR toward the end of the century.

This demonstrates the importance of having a robust underlying physical model for an emulator, and further suggests that FaIRGP is a well suited candidate for this task, and can provide meaningful temperature projections based only on historical data.

5.4. Inferring Radiative Forcing From Temperatures

Whilst FaIRGP is intended to emulate the surface temperature response to changes in emission, its GP nature also allows to make inference of radiative forcing given temperature data as demonstrated in Section 3.5. In this experiment, we propose to assess how global surface temperature anomaly data can be used to inform an estimate of the effective radiative forcing using FaIRGP.

We use the complete data set of global temperature time series from the *historical* and *SSP126*, *SSP245*, *SSP370*, *SSP585* experiments, which are depicted in Figure 14. We update the GP prior placed over the radiative forcing $F(t)$ with temperature data to obtain a posterior radiative forcing response which incorporates information from temperature data. We emphasize that the posterior over $F(t)$ does not use any forcing data, only simulated temperature data. As for temperature emulation, the posterior radiative forcing response is given in closed-form by Equation 32. We evaluate the posterior historical radiative forcing response against NorESM2-LM historical top-of-atmosphere radiative forcing.

Figure 14 illustrates that the posterior forcing response learns from temperature time series how to deviate from the FaIR forcing trend to better account for the NorESM2 simulated data. This is particularly evident between 1960 and 1980 where the posterior reproduces a decrease in forcing to better account for the global cooling trend




Figure 13. Emulated global mean surface temperature anomaly over 2015–2100 for the *SSP245* scenario with emulators trained over historical temperatures only.




Figure 14. *Left:* Global temperature anomaly responses used to inform the posterior forcing response; *Right:* FaIRGP posterior radiative forcing predicted over experiments; the temperature response and radiative forcing obtained with FaIR are plotted in dashed red lines.

in that time period, which FaIR struggles to capture. The results reported in Table 6 show that FaIRGP helps improve over the FaIR forcing in RMSE and Bias. Whilst FaIR performs better at inferring the stable forcing before 1950, informing the inferred forcing with temperature data with FaIRGP helps better account for the forcing variations after 1950.

FaIRGP therefore proves to be useful to infer global radiative forcing trends informed by surface temperatures. This is possible because our prior is specified within an EBM which explicitly connects the changes in surface temperatures to the changes in radiative forcing. On the contrary, this would not be possible with the plain GP model because it ignores forcing dynamics.

5.5. FaIR Multi-Model Uncertainty in FaIRGP

In this section, we explore how the uncertainty bounds introduced by the FaIR model compare with those of the FaIRGP uncertainty. FaIR has been widely employed to assess the uncertainty surrounding global temperature projections. The methodology employed by FaIR for generating probabilistic projections is rooted in a constrained ensembling process. This involves specifying prior distributions over model parameters (such as carbon cycle sensitivity, forcing model sensitivity and equilibrium climate sensitivity). These prior distributions are determined from the inter-model variability. Then, a large ensemble of climate projections is drawn using these distributions, and is subsequently constrained based on the likelihood of projections aligning with the simulated present-day level and rate of anthropogenic warming (Leach et al., 2021). The resulting *constrained ensemble* of projections allows to quantify the effect of *FaIR multi-model uncertainty* on temperature projections.

Following the protocol of Leach et al. (2021), we generate a constrained ensemble of 1,000 FaIR surface temperature anomaly projections. For each ensemble member, we then generate a FaIRGP projection. The total

Table 6
Scores of FaIR and FaIRGP at Predicting the Global Radiative Forcing Over Historical Periods

| Period | Emulator | RMSE ↓ | MAE ↓ | Bias | LL ↑ | Calib95 | CRPS ↓ |
|-----------|---------------------|--------------|--------------|---------------|--------|---------|--------|
| 1850–1950 | FaIR | 0.396 | 0.261 | −0.008 | — | — | — |
| | FaIRGP [†] | 0.421 | 0.285 | 0.041 | −0.797 | 0.99 | 0.25 |
| 1950–2014 | FaIR | 0.622 | 0.461 | 0.350 | — | — | — |
| | FaIRGP [†] | 0.568 | 0.466 | −0.315 | −0.925 | 0.984 | 0.328 |
| 1850–2014 | FaIR | 0.496 | 0.339 | 0.129 | — | — | — |
| | FaIRGP [†] | 0.484 | 0.357 | −0.101 | −0.847 | 0.988 | 0.281 |

Note. Scores are evaluated against NorESM2-LM top-of-atmosphere radiative forcing; the best emulator scores for each period are highlighted in bold; ↑/↓ indicates higher/lower is better; [†] indicates our proposed method.

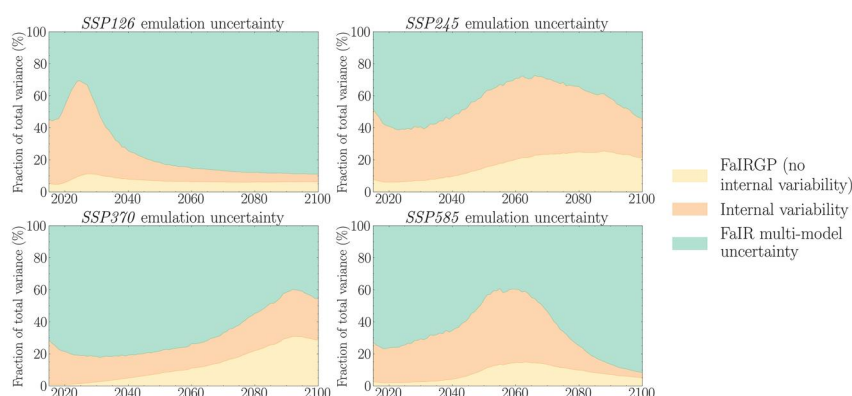


Figure 15. Proportional stacked area charts of the contribution to total variance of the FaIRGP uncertainty without internal variability (yellow), the internal variability modeled in FaIRGP as in Section 3.4 (orange) and the FaIR multi-model uncertainty (green), for the emulation of the global annual mean surface temperature anomaly under *SSP126* (top-left), *SSP245* (top-right), *SSP370* (bottom-left) and *SSP585* (bottom-right).

variability of this projection can be decomposed into two components: (a) the average variability introduced by FaIRGP across ensemble members and (b) the variability in the mean FaIRGP response across ensemble members. The first component corresponds to the variability that FaIRGP's stochasticity introduces on the deterministic FaIR temperature response, averaged across ensemble member, and can be decomposed into the variability introduced by FaIRGP without internal variability, and the contribution of internal variability alone. For this reason component (a) is viewed as the uncertainty imputable to FaIRGP. The second component corresponds to the variability that sampling across the constrained ensemble introduces in the FaIRGP mean response. For this reason, component (b) is viewed as the uncertainty imputable to the FaIR multi-model uncertainty. Figure 15 shows how the different sources of uncertainty contribute to the total variance when emulating surface temperature anomalies for SSPs.

A higher FaIR multi-model uncertainty contribution to the total variance (green) indicates significant sensitivity of the FaIRGP mean response to parameter choices within the constrained ensemble. Conversely, a greater FaIRGP uncertainty contribution (orange and yellow) suggests limited sensitivity to parameter choices, with most variability arising from our stochastic treatment of the temperature response. We find that in general, the FaIR multi-model uncertainty contributes the largest fraction of the total variance in projections of future warming. This is particularly evident in 2100 for low/high end forcing scenarios *SSP126* and *SSP585*, since they lie outside of the training data regime used in FaIRGP and the model's need for extrapolation makes it more sensitive to FaIR's structural uncertainty.

We find that for intermediate forcing scenarios *SSP245* and *SSP370*, the FaIRGP uncertainty contribution to the total variance remains significant. Because these scenarios lie in between other scenarios from the training data, FaIRGP's reliance on interpolation between its training data results in lower sensitivity to the multi-model uncertainty, making the idea of an uncertain future warming better constrained. This is supported by Figure 16, depicting how the average contribution of each source of uncertainty to the standard deviation stacks up in predictions. It shows a consistently stable and comparable magnitude of uncertainty introduced by FaIRGP with internal variability across scenarios.

It is important to caveat that the FaIR constrained ensemble samples from inter-model variability, whereas FaIRGP is here only calibrated on NorESM2-LM data. Therefore, one could reasonably expect that a multi-model version of FaIRGP, allowing calibration across multiple climate models, would likely contribute more uncertainty to projections of end-of-century warming.

6. Application: Spatial Surface Temperatures Emulation

In this section, we pursue the comparison of FaIRGP to baseline models, but for the task of emulating spatially-resolved temperatures. As for global emulation, we start by briefly illustrating how the model concretely applies. Then, we benchmark it against baseline models for SSP emulation and evaluate the emulation of surface

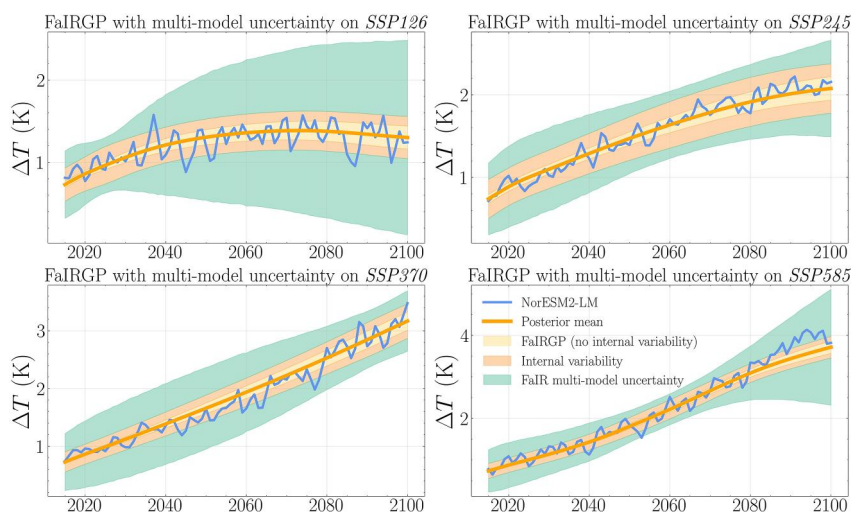


Figure 16. Global annual mean surface temperature anomaly projections over shared socio-economic pathways (SSPs) with contributions to uncertainty from FaIRGP without internal variability (yellow), the internal variability modeled in FaIRGP as in Section 3.4 (orange) and the FaIR multi-model uncertainty (green). Shading indicates 2 standard deviations and is computed by taking the average contribution to the standard deviation of each source of uncertainty.

temperatures forced by anthropogenic aerosols only. We conclude by investigating how FaIRGP can help infer spatial top-of-atmosphere radiative forcing maps.

6.1. FaIRGP for Spatial Temperatures Emulation

Figure 17 illustrates the spatial surface temperature anomaly emulation with FaIRGP for test scenario *SSP245*. As in the global case, a prior response based on FaIR is first specified, and then shifted by a data-informed posterior correction map. The posterior correction is learned from training scenarios $\mathcal{D}_{\text{train}} = \{\text{historical}, \text{SSP126}, \text{SSP370}, \text{SSP585}\}$.

The prior spatial response is constructed using a pattern scaling model trained on $\mathcal{D}_{\text{train}}$, and therefore introduces a fixed spatial pattern that is only rescaled by changes in global mean temperature. The posterior correction at location x is obtained by updating the prior over $T(x, t)$ from Equation 34 with local surface temperature data from $\mathcal{D}_{\text{train}}$. The posterior correction maps effectively provide a data-driven way to deviate from this fixed spatial pattern, and better account for the possibly varying spatial temperature patterns of the emulated scenario. Finally, by linearly adding the posterior correction map to the prior, we obtain a posterior spatial temperature response over *SSP245*.

6.2. Shared Socio-Economic Pathways Emulation

In this experiment, we follow the same procedure as in the global emulation experiment and iteratively train models on the data set deprived of one SSP scenario, then use the retained scenario as a test scenario for evaluation. We benchmark FaIRGP against a FaIR pattern scaling model and a purely data-driven plain GP emulator. The pattern scaling model is obtained by fitting a linear regression model using the same training data as the other models. The plain GP model is analogous to the baseline GP emulator in Watson-Parris et al. (2022), but differs in two aspects: (a) we adopt a simpler construction for the covariance with an anisotropic Matérn-3/2 kernel, and (b) our model takes as input global aerosols emissions whereas Watson-Parris et al. (2022) use spatially-resolved aerosols emission maps. Scores are computed over the 2080–2100 period since the start of all SSPs is quite similar. Mean scores are reported in Table 7.

We find that FaIRGP has on average lower error than baseline models. Figure 18 shows that the spatial bias patterns of FaIRGP are similar to the bias patterns obtained with the FaIR pattern scaling model. This indicates that the prior has a strong influence on the predicted posterior. Nonetheless, we observe that the posterior correction in FaIRGP helps mitigate the spatial inaccuracies of its prior. This is particularly evident for *SSP126* and *SSP245* where the magnitude of the spatial bias in FaIRGP is overall smaller than the one for the FaIR pattern

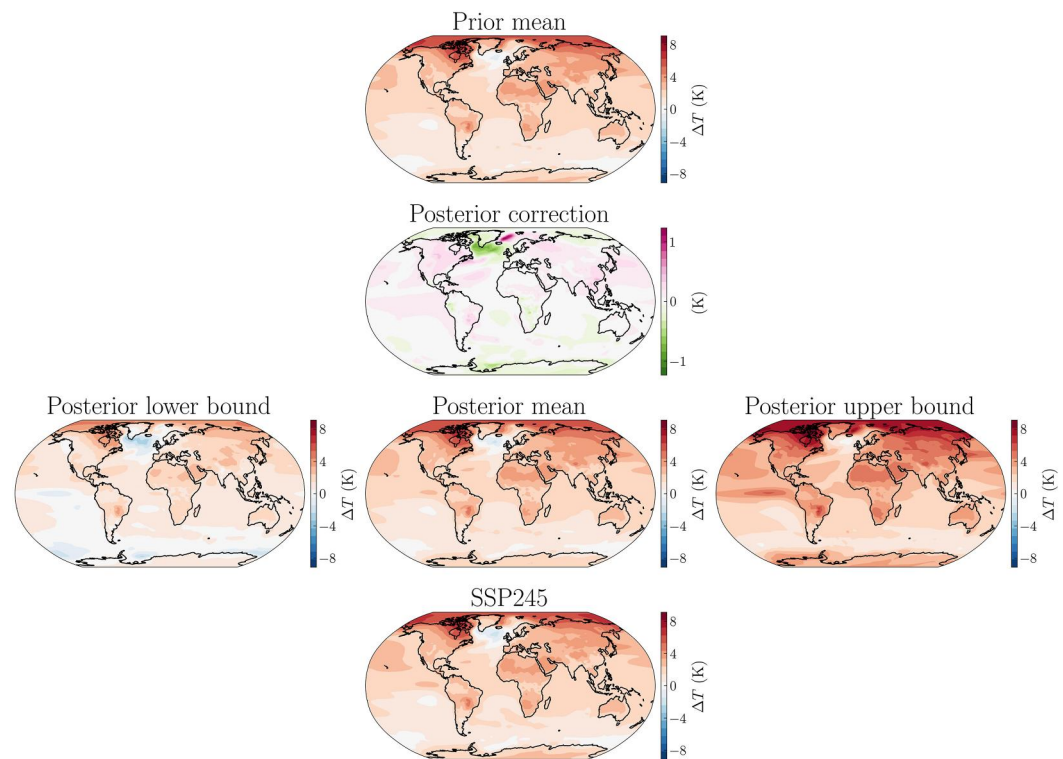


Figure 17. Spatial emulation of NorESM2-LM SSP245 temperatures. The prior mean (top) corresponds to a pattern scaling model used on the FaIR global temperature response. The posterior correction (middle-top) learns from data how to deviate from the pattern scaling prior to produce a posterior response (middle-bottom) that better emulates the groundtruth surface temperatures (bottom). The posterior lower/upper bounds capture the 95% credible interval. Maps are averaged over the 2080–2100 period.

scaling model. Regarding the high forcing scenario SSP585, the spatial correction is more subtle. This is due to it becoming an extrapolation task, and as a result, FaIRGP exhibits behavior closer to its pattern scaling prior.

Figure 18 also shows that, as for global temperature emulation, the plain GP model predicts sound surface temperature maps for low and medium forcing scenarios, but struggles at extrapolating over high forcing scenarios.

6.3. Emulating Anthropogenic Aerosols Forcing

In this experiment, we want to evaluate the emulation of temperature changes induced by anthropogenic aerosol emissions. We use the *hist-aer* experiment from the ClimateBench v1.0 data set. The *hist-aer* experiment is generated using NorESM2-LM, using only historical anthropogenic aerosol emissions, and setting long-lived

Table 7

Scores of Baseline Emulators and FaIRGP for the Task of Emulating Shared Socio-Economic Pathways Spatial Temperatures Over 2080–2100 Period

| Emulator | RMSE ↓ | MAE ↓ | Bias | LL ↑ | Calib95 | CRPS ↓ |
|---------------------|----------------------|----------------------|-----------------------|-----------------------|----------------------|----------------------|
| Pattern scaling | 0.637 ± 0.189 | 0.467 ± 0.139 | −0.162 ± 0.194 | — | — | — |
| Plain GP | 1.069 ± 0.616 | 0.761 ± 0.441 | −0.362 ± 0.350 | −1.625 ± 1.192 | 0.850 ± 0.141 | 0.560 ± 0.341 |
| FaIRGP [†] | 0.619 ± 0.188 | 0.451 ± 0.138 | −0.113 ± 0.192 | −0.853 ± 0.398 | 0.873 ± 0.093 | 0.341 ± 0.099 |

Note. The best emulator for each metric is highlighted in bold; ↑/↓ indicates higher/lower is better; we report 1 standard deviation; † indicates our proposed method.

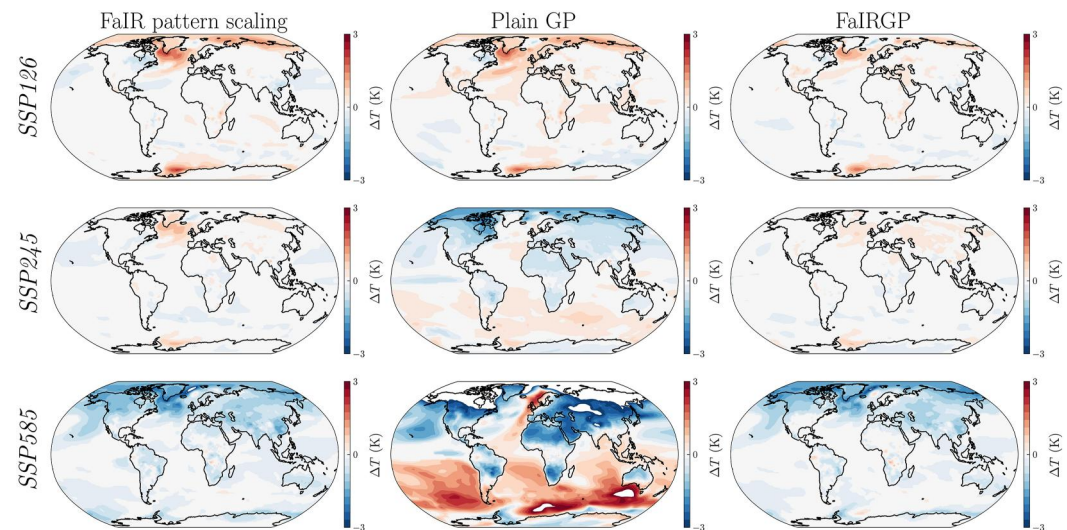


Figure 18. Maps of the mean difference between emulated spatial temperatures and target Norwegian Earth System Model (NorESM2) values for each emulator under a low forcing future scenario (*SSP126*), a medium forcing future scenario (*SSP245*) and a high forcing future scenario (*SSP585*). Red/Blue means the emulator overshoots/undershoots NorESM2 temperatures. Maps are averaged over the 2080–2100 period. Differences insignificant at the $p < 0.05$ level are masked.

greenhouse gas emissions to zero. We emulate surface temperature anomalies over this scenario using emulators trained on all available *historical* and SSPs experiments. The FaIR forcing model is modified to include the aerosol-cloud interaction parametrization from Leach et al. (2021, Section 2.2.1).

Figure 19 shows the emulated global mean surface temperature anomaly with FaIRGP. Whilst FaIR faces challenges in reproducing the magnitude of the global temperature response to anthropogenic emissions, the posterior correction introduced by FaIRGP allows for a deviation from FaIR, resulting in predicted temperature anomalies that more accurately reflect the NorESM2-LM response. This suggests that FaIRGP can leverage information from the *historical* and SSPs experiments to enhance its representation of the temperature response to anthropogenic aerosol emissions.

We now turn to the emulation of local surface temperature anomaly maps for the *hist-aer* experiment. Figure 20 presents a comparison of predictions with the FaIR pattern scaling baseline and FaIRGP over three time periods. Consistently with the global trend from Figure 19, the predicted temperature anomaly magnitudes for 1900–1950 are similar between the pattern scaling model and FaIRGP, and align with the NorESM2-LM response whereas,

over periods 1950–1980 and 1980–2014, the magnitude of the FaIRGP spatial response more accurately reflects the magnitude of the NorESM2-LM temperature response.




Figure 19. Emulated global mean surface temperature anomaly over 1850–2014 for the *hist-aer* experiment.

Regarding predicted spatial patterns, we observe that FaIRGP is able to deviate from its fixed pattern prior to partially reproduce the North Atlantic warming observed in NorESM2-LM after 1950. However, the predicted spatial temperature anomaly pattern with FaIRGP remains strongly influenced by its pattern scaling prior. This is particularly evident in the South Atlantic where FaIRGP reproduces the warming predicted by the pattern scaling model. This task is notoriously difficult for pattern scaling models, which excel at emulating responses to greenhouse gas emissions but struggle under strong aerosol forcing scenarios (Levy et al., 2013; May, 2012). Considering the significant impact of the prior choice on FaIRGP, this fosters advocacy for a prior local response model that goes beyond pattern scaling models. Additionally, using spatially-resolved emissions maps for aerosols, instead of global aerosols emissions, should enhance emulated temperatures given the strong influence spatial patterns of aerosols have over radiative forcing (Williams et al., 2022).

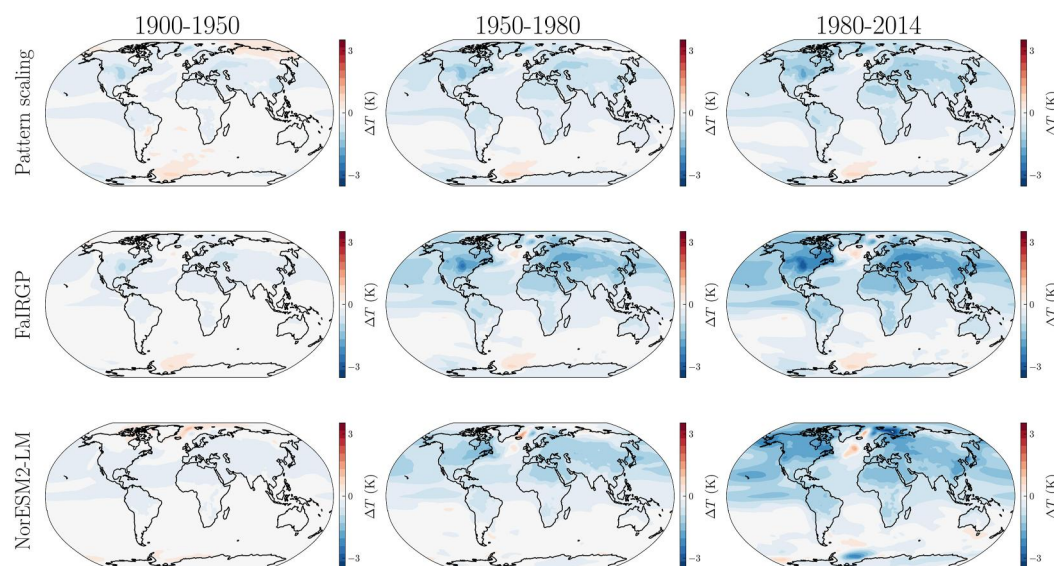


Figure 20. Maps of local surface temperature anomalies for the *hist-aer* experiment averaged over the 1900–1950 period (left), 1950–1980 (left-right center) and 1980–2014 (right). *Top*: Prediction with the FaIR pattern scaling baseline. *Top-Bottom Center*: Prediction with the FaIRGP posterior mean. *Bottom*: Groundtruth NorESM2-LM surface temperature anomaly map.

We conclude this experiment with a quantitative evaluation of the predicted local temperature anomaly maps over the 1950–2014 period. This is the period where the temperature response to aerosol emissions is most discernible. Because the anthropogenic aerosols forcing pattern is spatially heterogeneous and mostly localized around industrialized regions (Carslaw & Pringle, 2022), we assess the emulators' performance over four major industrialized regions: North America, East Asia, South Asia and Europe. The regions' extents are depicted in Figure 21 and are constructed from world region delineation proposed by (S. Seneviratne et al., 2012). Table 8 shows that, consistently with the prediction maps, the FaIRGP posterior mean spatial response outperforms the FaIR pattern scaling baseline for every region considered. The plain GP model, as anticipated, faces difficulties in emulating local surface temperature anomalies solely based on aerosol emissions inputs. We believe this is primarily because every scenario from its training data includes greenhouse gas emissions. Appendix C4 provides maps emulated with the plain GP model.

6.4. Inferring Radiative Forcing From Temperatures

As in Section 5.4, we propose to probe whether FaIRGP can be used to infer spatial forcing maps. Since we do not have a spatial forcing model, we simply use a spatially constant forcing prior for $F(x,t)$. We update it with spatially-resolved temperature data from *historical* and SSPs scenarios. Figure 22 compares the obtained posterior mean forcing with historical forcing maps simulated with NorESM2-LM.

The magnitude of the inferred posterior mean forcing with FaIRGP is very conservative, and does not align with the magnitude of the forcing levels observed in the groundtruth NorESM2-LM maps. The sign of the inferred forcing with FaIRGP also disagrees with the groundtruth maps in several regions. This is most evident around 90° latitudes, where FaIRGP infers a positive forcing whilst NorESM2-LM simulations suggest a negative forcing.

We attribute this discrepancy to FaIRGP relying only on local temperature anomalies to infer local radiative forcing. Whilst high latitudes tend to warm faster than around the equator (due among things to surface albedo feedback




Figure 21. The polygons show the industrialized regions used for the evaluation of the predicted local annual mean surface temperature anomaly under the *hist-aer* experiment: North America (red), Europe (blue), South Asia (purple), East Asia (green).

Table 8

Scores of Baseline Emulators and FaIRGP for the Task of Emulating Hist-Aer Local Annual Mean Surface Temperature Anomaly for Four Industrialized Regions Over the 1950–2014 Period

| Region | Emulator | RMSE (10^{-1}) | MAE (10^{-2}) | Bias (10^{-2}) |
|---------------|---------------------|--------------------|-------------------|--------------------|
| Europe | Pattern scaling | 1.29 | 1.71 | 1.29 |
| | Plain GP | 2.26 | 3.37 | 3.07 |
| | FaIRGP [†] | 1.05 | 1.35 | −0.313 |
| North America | Pattern scaling | 1.10 | 1.63 | 1.21 |
| | Plain GP | 2.40 | 3.70 | 3.65 |
| | FaIRGP [†] | 0.998 | 1.35 | −0.795 |
| South Asia | Pattern scaling | 0.523 | 0.594 | 0.389 |
| | Plain GP | 1.26 | 1.53 | 1.52 |
| | FaIRGP [†] | 0.382 | 0.404 | −0.0100 |
| East Asia | Pattern scaling | 1.01 | 1.32 | 1.20 |
| | Plain GP | 2.22 | 3.08 | 3.08 |
| | FaIRGP [†] | 0.631 | 0.772 | −0.173 |

Note. The best emulator for each metric is highlighted in bold; † indicates our proposed method.

and atmospheric circulation for equator to poles), the top-of-atmospheric effective radiative forcing at high latitudes can remain negative (due to higher albedo, larger zenith-angle). This underlines an important point: FaIRGP performs probabilistic inference of local radiative forcing given local temperature anomalies, which alone is not sufficient to account for the complex spatiotemporal dynamics of the climate system. However, there are reason to believe that informing this probabilistic inference with additional covariates, such as surface albedo or net top-of-atmosphere data, could enhance the accuracy of the inferred spatial forcing maps.

Finally, we note that whilst the spatial forcing maps inferred with FaIRGP inadequately capture the groundtruth forcing magnitude and pattern, they still manage to partially reproduce two important features of the climate system: the overall temporal increase in forcing, as well as the mid-1900s hemispheric contrast. In particular, Figure 23 shows that whilst predictions largely underestimate the forcing magnitude, they also display some latitudinal correlation with the NorESM2-LM simulated forcing over the 1940–1980 period. This is encouraging considering that the predicted spatial patterns are solely inferred from temperature patterns.

7. Discussion

7.1. About the GP Approach

7.1.1. Comparison to Placing Bayesian Priors Over the SCM Parameters

A simple way to introduce model variability in SCMs is to place Bayesian priors over model parameters, such as the carbon cycle feedback terms or the forcing model coefficients (Leach et al., 2021). Such priors can then be updated with global temperature data to formulate posterior distributions. This specifies a probabilistic climate model emulator calibrated against climate model data. Whilst aligned in spirit with the work proposed in this paper, it is important to highlight the distinctions in objectives and technical scope with our GP-based approach.

Placing Bayesian priors over parameters allows to sample from diverse sources of uncertainty such as uncertainty on the temperature response parameters, forcing sensitivities or carbon cycle sensitivities. Importantly, this allows to explore climate sensitivity and the structural uncertainty of models. In contrast, the framework we propose only addresses the forcing uncertainty and considers fixed values for the SCM parameters. In that sense, FaIRGP, as it is formulated in our work, provides a more restrictive exploration of the uncertainty than methods placing priors over an SCM parameters.

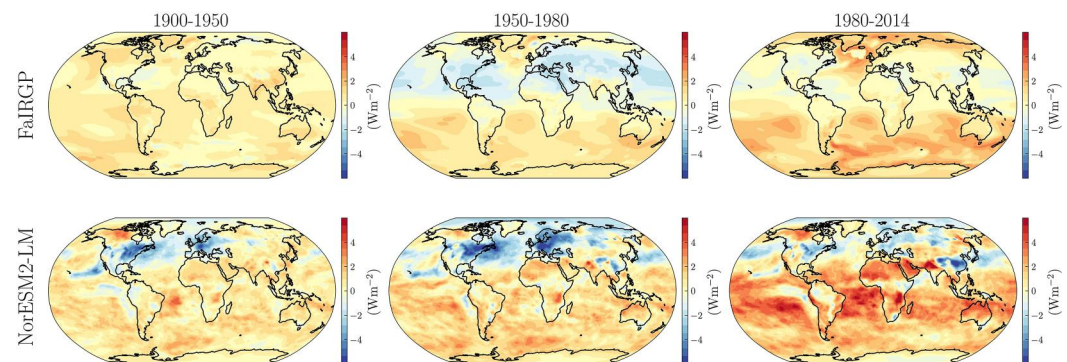


Figure 22. Maps of top-of-atmosphere radiative forcing averaged over 1900–1950, 1950–1980 and 1980–2014 periods of the historical scenario. *Top:* Mean posterior radiative forcing inferred with FaIRGP. *Bottom:* NorESM2-LM radiative forcing maps.




Figure 23. Latitudinal average of FaIRGP inferred radiative forcing and NorESM2-LM radiative forcing over 1940–1980 period. The shaded area reports two standard deviations of longitudinal variability. $R^2 = 0.404$ is computed between latitudinal averages.

The main advantage of our GP-based approach rather lies in its technical convenience. Chiefly, FaIRGP formulates analytical expressions for the posterior distribution over forcing and temperatures, which have an intuitive interpretation as the sum of a physics-driven prior and a data-driven correction. In contrast, when placing a Bayesian prior over parameters, we do not in general have access to a probability distribution for temperatures. Therefore, probabilistic emulation may need to be sampling-based, and one must resort to more complex Markov-Chain Monte-Carlo techniques to sample from the posterior distribution.

In addition, beyond particular modeling objectives, sampling-based approaches have two main shortcoming from a purely technical perspective: (a) a thorough uncertainty quantification requires storing a tremendous amount of scenarios, which can be limited by memory capacity—Beusch et al. (2022) require an ensemble of 9 million emulations; (b) it reduces the statistical representation of uncertainty to summary statistics such as the mean, standard deviation or quantiles. In contrast, having a closed-form expression for the posterior distribution with GPs allows to analytically conduct probabilistic studies using the full probability density, and additionally draw samples, if needed. Finally, having access to the probability density expression allows to evaluate the likelihood of observations. This can critically be used as a maximization objective to tune the model parameters against observations, but also in the context of Bayesian optimization routines to find optimal emission trajectories to meet climate goals.

Because of the technical convenience that GP-based approaches display, and the richer structural uncertainty of sampling-based approaches, we argue that rather than opposing them, the two paradigms should be viewed as complementary. Namely, endeavors to formulate a hierarchical Bayesian model that both places priors over the SCM parameters and uses a FaIRGP backbone are promising directions to simultaneously address many sources of uncertainty while retaining some technical advantages of working with GPs.

7.1.2. Connection With Stochastic Energy Balance Models

Our work is related to the work of Cummins et al. (2020), which formulates a stochastic EBM by introducing a white noise variability term in the temperature response, but also in the forcing model. This allows them to account for climate internal variability, and formulate a Kalman filtering strategy to obtain maximum likelihood estimators of the EBM parameters. Our work similarly introduces a white noise in the temperature response model to account for climate internal variability, which results in an additional temporal Matérn-1/2 covariance term in the prior over temperatures. Whilst an extended discussion goes beyond the scope of this work, it can be shown that in the long term regime, these two approaches are equivalent, and that more broadly, Kalman filtering models are in fact equivalent to temporal GPs with Matérn covariance functions (Hartikainen & Särkkä, 2010; Sarkka & Hartikainen, 2012; Sarkka et al., 2013).

Our work differs from Cummins et al. (2020) in that beyond the stochasticity arising from internal variability, we also introduce a GP prior over the radiative forcing. This GP prior introduces stochasticity over the SCM design, which is not only a function of time, but also of emission levels. Because our modeling is not purely temporal (the GP is also a function of emissions), we cannot employ Kalman filtering strategies and instead choose to use GP regression techniques.

7.1.3. Choice of Kernel ρ

The choice of covariance function ρ in Equation 12 is an important choice that allows the user to incorporate their domain knowledge into the prior over the radiative forcing. Let u and u' be generic notations for input data (in our work, greenhouse gas and aerosol emissions), the kernel $\rho(u, u')$ specifies how will the prior covary between these two inputs. For example, choosing $\rho(u, u') = \delta(u - u')$ makes the GP independent at any two inputs. On the other hand, choosing $\rho(u, u') = 1$ causes the GP to covary equally between any two inputs.

The Matérn family are a common family of kernel parameterized by a degree ν . They allow to control for the functional regularity of the GP. For $\nu = 1/2$, draws from the GP are continuous functions, for $\nu = 3/2$ they are once differentiable, and in the limit $\nu = \infty$ they become infinitely differentiable and corresponds to the squared exponential kernel. A detailed presentation of the Matérn kernels is provided in Appendix A2.

Additions or multiplications of kernels can be used to construct more elaborate covariance functions that reflect an additive or multiplicative structure in the forcing. Periodic kernels can also be introduced to model seasonality. Going further, more complex choices of kernels include the spectral mixture kernel (A. Wilson & Adams, 2013) which attempts to learn the spectral density of the data, or even kernels parametrized as neural networks (A. G. Wilson et al., 2016; Law et al., 2019).

Whilst kernel selection plays a key role in GP regression, our focus in this work is on the development of the FaIRGP framework rather than refining the kernel itself. The kernel is treated throughout as a modular component, with potential for refinement. As a result, we choose to work with a simple anisotropic Matérn-3/2 kernel throughout. Preliminary findings indicate that the choice of kernel does not significantly degrade the results. Hence, dedicating efforts to constructing more elaborate kernels is likely to yield comparable or better results than our current approach.

7.1.4. Computational Efficiency and Scalability

We report that to emulate 100 years of surface temperature anomaly with FaIRGP, it takes less than a second for global and spatial emulation on an “average” personal laptop (16Go memory), without requiring any parallelization methods.

Scalability issues are commonly associated with GPs when the training set grows in size. This is because computing their posterior distribution involves a matrix inversion, which has a cubic computational cost in the number of training samples. Fortunately, unlike neural networks which require large amounts of data (Watson-Parris et al., 2022), GPs excel in scenarios with limited data (Rasmussen & Williams, 2005). Consequently, it is possible to develop skillful GP emulators with limited training data.

In cases where using a larger training data set becomes a necessity, one can still employ linear conjugate gradient methods and parallelization schemes (Wang et al., 2019) to scale exact GPs to millions of data points. Alternatively, sparse approximation techniques can be used to obtain a scalable estimate of the posterior distribution (Matthews et al., 2016; Rahimi & Recht, 2007; Titsias, 2009).

7.2. Climate Modeling Considerations

7.2.1. Beyond FaIR: Broader Applicability of the Method

While we have chosen to use FaIR as the backbone climate model for our work, the rationale behind the development of FaIRGP is easily transferable to other commonly used SCMs such as MAGICC (Meinshausen et al., 2011) or OSCAR (Gasser et al., 2017). These models share linear time invariant dynamics, allowing us to incorporate a GP prior into the forcing term of these dynamics. By doing so, we can obtain a GP-based solution that is informed by the model parameters, exactly like in FaIRGP.

The dynamical systems of interest can naturally describe the temperature response to radiative forcing, as is the case in our work. However, we could also imagine extending this to carbon cycle models, where emission levels prescribe the forcing function, and the output of the dynamical system are atmospheric concentrations. This GP framework over dynamical systems is in fact highly general, and is known as latent force modeling. This paradigm was initially introduced by M. Alvarez et al. (2009) and M. A. Alvarez et al. (2013) in the context of linear dynamical systems where the forcing function is unknown.

For more complex dynamical systems involving non-linearities, more sophisticated latent force modeling techniques may be applied, such as methodologies based on Volterra series approximation of the dynamical system (M. A. Álvarez et al., 2019; Ross et al., 2021), variational approximations of the posterior response (Moss et al., 2021; Ward et al., 2020), with the option of using deep probabilistic modeling of the dynamical system (Baldwin-McDonald & Álvarez, 2023; McDonald & Álvarez, 2021).

7.2.2. Pixel Independence Assumption

The pattern scaling model used in our prior effectively uses independent linear regressions at each location to map changes in global mean temperature onto changes in local temperature. However, this modeling approach challenges our intuition as it overlooks the spatial dependence of temperature fields, despite our expectation that temperatures at nearby locations should covary.

To address this modeling concern, a common solution is to incorporate spatially correlated innovations into the pattern scaling response, which represents the spatial expression of climate internal variability (Beusch et al., 2020a, 2020b, 2022; Castruccio & Stein, 2013; Goodwin et al., 2020). Alternatively, Link et al. (2019) design a procedure based on the Wiener-Khinchin theorem (Champeney & Fuoss, 1973) to emulate a climate variability field with the same variance and spatiotemporal correlation structure as the one in ESMs outputs.

From a statistical modeling standpoint, these approaches introduce such spatially correlated innovations as a multivariate Gaussian variable, their main differences lying in the procedure used to determine its spatial covariance structure. Because of the Gaussian nature of the modeling we propose in this work, a Gaussian spatial correlated innovation can easily be incorporated in the FaIRGP spatial emulation model we propose. However, for the sake of clarity, we chose not to delve into these additional considerations and instead focus on the exposition of the Bayesian EBM.

Further, whilst it has been pointed out that, in general, regional changes in temperature scale robustly with global temperature (S. I. Seneviratne et al., 2016), this may not be true under strong mitigation scenarios or under strong aerosol forcing (Levy et al., 2013; May, 2012; Tebaldi & Knutti, 2018; Tebaldi et al., 2020). This fosters advocacy for the development of “spatialized” SCMs going beyond pattern scaling.

7.2.3. Opportunities for Precipitation Emulation

Going beyond surface temperature emulation, we can explore how FaIRGP can be used to emulate precipitations. One approach is to combine the GP GP emulator for precipitation described in Watson-Parris et al. (2022) with FaIRGP. By leveraging Gaussian conjugacy relationships, a natural cross-covariance between the two emulators is induced. This enables the exchange of information between precipitation and temperature fields. This is in line with recent advocacy for joint emulation of temperatures and precipitations (Schöngart et al., 2022, 2023; Snyder et al., 2019).

Another option is to use the emulated temperatures from FaIRGP as input for a statistical emulator of precipitation, such as a Gamma regression model (Martinez-Villalobos & Neelin, 2019). Indeed, many climate impacts are routinely assumed to be a function of temperatures. Additionally, the full probability distribution predicted by FaIRGP can be incorporated into the model, allowing for the propagation of epistemic uncertainty associated with FaIRGP.

7.2.4. Application to Detection and Attribution

With FaIRGP, we have access to the analytical expression of the probability density distribution of emulated temperatures. Therefore, we can emulate surface temperatures under historical scenarios, both with and without anthropogenic forcing, and analytically compute the probability of temperature occurrences in each scenario. By comparing these probabilities, we can assess the extent to which human activity has made a certain temperature range more likely, enabling us to conduct attribution studies.

Conducting detection and attribution studies with emulators is not exclusive to FaIRGP, and could in principle be conducted with any emulator as discussed in Watson-Parris et al. (2022). However, the strength of FaIRGP lies in its ability to input temperature ranges directly into a known probability density function, providing a precise probability between 0 and 1 of such temperatures to occur under a given emission scenario.

8. Conclusion and Outlooks

SCMs are robust physically-motivated emulators of changes in global mean surface temperatures. GPs are powerful Bayesian machine learning models capable of learning complex relationships, and emulate from data how changes in emissions affect changes in surface temperatures. By combining them together, we reconcile these two paradigms of emulator design, which mutually address their respective limitations. We introduce

FaIRGP, a Bayesian EBM that (a) maintains the robustness and interpretability of a SCM, (b) gains the flexibility of modern statistical machine learning models with the ability to learn from data, and (c) provides principled uncertainty quantification over the emulator design.

We demonstrate skillful emulation of global mean surface temperatures over realistic emission scenarios. Unlike GPs, FaIRGP has a robust physical grounding which allows it to provide reliable predictions even on out-of-sample scenarios. On the other hand, unlike SCMs, FaIRGP can learn complex non-linear relationships to deviate from an SCM and improve predictions. In particular, FaIRGP better accounts for the temperature response to anthropogenic aerosol emissions. We further show that these findings carry over to the task of emulating spatially-resolved surface temperature maps. In addition, we find that FaIRGP can also be used to produce estimates of top-of-atmosphere radiative forcing given temperature data.

The full mathematical tractability, with analytical expressions for probability distributions, provides great control over the modeling, and a rich framework to reason about probability distributions over temperatures. This is of great relevance to detection and attribution studies. Further, whilst our work focuses on temperature emulation—which has already been thoroughly studied—we envision FaIRGP as a foundation for the development of robust data-driven emulators for more complex climate variables, such as precipitations. Harnessing the mathematical properties of GPs, we believe that emulating climate impacts using FaIRGP will provide additional control over pure machine learning methods, whilst being able to capture complex non-linear relationships to forcing.

We hope this work will contribute to building trust in data-driven models, and thereby allow the climate science community to benefit more widely from their potential.

Appendix A: Complementary Material on Gaussian Processes

A1. Illustrated Walk-Through Gaussian Process Regression

We provide in this section an illustrated walk-through of how a regression task can be achieved using GPs. Begin by considering a GP specified by

$$\begin{cases} f(x) \sim \text{GP}(m, k) \\ m(x) = 0 \\ k(x, x') = \sigma_f^2 \exp\left(-\frac{|x - x'|^2}{\ell}\right). \end{cases} \quad (\text{A1})$$

Further consider a set of $p = 500$ regularly spaced points between 0 and 1, $\text{grid} = [z_1 \dots z_p]$ where

$$z_i = \frac{i}{p}, \quad 1 \leq i \leq p. \quad (\text{A2})$$

Then, by definition of GPs, the evaluation of f on grid admits a joint multivariate normal distribution. Namely, if we define the covariance matrix

$$\mathbf{K} = \begin{bmatrix} k(z_1, z_1) & \dots & k(z_1, z_p) \\ \vdots & \ddots & \vdots \\ k(z_p, z_1) & \dots & k(z_p, z_p) \end{bmatrix}, \quad (\text{A3})$$

then we have

$$\begin{bmatrix} f(z_1) \\ \vdots \\ f(z_p) \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \mathbf{K}). \quad (\text{A4})$$




Figure A1. Plot of the $f(x) \sim \text{GP}(0, k)$ prior mean function and 95% confidence region for $\sigma_f^2 = 1$ and $\ell = 0.1$.

Therefore, the zero vector $\mathbf{0}$ provides the mean of $f(x)$ at every grid point, and the diagonal elements of the covariance matrix provide the standard deviation of $f(x)$ at every grid point following $\sqrt{k(z_i, z_i)} = \sigma_f$. Figure A1 plots the mean function and the 95% credible interval of $f(x)$ over the grid.

Furthermore, we are also able to draw samples from the multivariate normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{K})$ by taking

$$\mathbf{f} = \mathbf{K}^{1/2} \mathbf{u}, \quad \mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p). \quad (\text{A5})$$

Such a sample $\mathbf{f} \in \mathbb{R}^p$ corresponds to a sample path from the GP $f(x)$ evaluated on the grid points. Figure A2 shows the plots of 10 sample paths, from this multivariate normal distribution.

Consider now the following data generating process

$$y = \sin(2\pi x) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, 0.04), \quad (\text{A6})$$

from which we simulate $n = 10$ observations $\mathbf{y} = [y_1 \ \dots \ y_n]$ at inputs $\mathbf{x} = [x_1 \ \dots \ x_n]$. Figure A3 shows a plot of the observed data.

Suppose that the form of the data generating process is unknown to a user who wishes to use these observations to regress y onto x . One possibility is to use the previously defined GP $f(x)$ as a Bayesian prior over the space of regression models $x \rightarrow y$ and assume the relationship

$$y_i = f(x_i) + \epsilon_i, \quad \epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2), \quad (\text{A7})$$

for an unknown noise variance σ_ϵ^2 .




Figure A2. Same as Figure A1, but plotting on top 10 sample paths from $\mathbf{f} \sim \mathcal{N}(\mathbf{0}, \mathbf{K})$. $\sigma_f^2 = 1$ and $\ell = 0.1$.




Figure A3. Scatter plot of (\mathbf{x}, \mathbf{y}) .

By definition of GPs, we know that the evaluation of f on \mathbf{x} will again admit a joint multivariate normal distribution, that is, by defining the covariance matrix

$$\mathbf{K}_{\mathbf{x}} = \begin{bmatrix} k(x_1, x_1) & \dots & k(x_1, x_n) \\ \vdots & \ddots & \vdots \\ k(x_n, x_1) & \dots & k(x_n, x_n) \end{bmatrix} \quad (\text{A8})$$

we get that

$$\begin{bmatrix} f(x_1) \\ \vdots \\ f(x_n) \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \mathbf{K}_{\mathbf{x}}), \quad (\text{A9})$$

and therefore by incorporating the additive noise, it follows that

$$\mathbf{y}|\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{K}_{\mathbf{x}} + \sigma_{\varepsilon}^2 \mathbf{I}_n). \quad (\text{A10})$$

First, using this distribution, we can find values of the hyperparameters $\theta = (\sigma_f, \ell, \sigma_{\varepsilon})$ that maximize $\log p(\mathbf{y}|\mathbf{x})$. This corresponds to finding their maximum likelihood estimates under $p(\mathbf{y}|\mathbf{x})$, and can be achieved for example, following the gradient descent procedure outlined in the following Algorithm 2. Such a procedure offers a principled and robust way to calibrate the model hyperparameters against observations. Figure A4 shows a plot of




Figure A4. Plot of the $f(x) \sim \text{GP}(0, k)$ prior mean function and 95% confidence region on the grid for calibrated hyperparameters $\sigma_f^2 = 0.43$ and $\ell = 0.17$, with 10 sample paths and observed data (\mathbf{x}, \mathbf{y}) .

the GP mean and confidence region with sample paths that have been calibrated against observations. Calibrating the prior against (\mathbf{x}, \mathbf{y}) yields a smaller variance $\sigma_f^2 = 0.43$, which translates in a tighter confidence region, and a larger lengthscale $\ell = 0.17$, which translates into slower variations of the sample paths.

Algorithm 2. Maximum Likelihood Calibration of $\theta = (\sigma_f, \ell, \sigma_\epsilon)$ With Gradient Descent

```

1: Input:  $\mathbf{x}, \mathbf{y}, \theta_0, \eta > 0, N \in \mathbb{N}$ 
2: for  $t \in \{1, \dots, N\}$  do
3:   Compute  $\log p(\mathbf{y}|\mathbf{x}) = \log \mathcal{N}(\mathbf{y}|\mathbf{0}, \mathbf{K}_\mathbf{x} + \sigma_\epsilon^2 \mathbf{I}_n)$ 
4:   Take  $\theta_t \leftarrow \theta_{t-1} + \eta \nabla_{\theta} \log p(\mathbf{y}|\mathbf{x})$ 
5: end for
6: Return:  $\theta_t$ 

```

Second, we can update the prior $f(x)$ placed over the predictive model into a posterior informed with the observed data $f(x)|\mathbf{y}$. Formally, the mean and covariance of the posterior can be computed analytically following

$$\begin{cases} f(x)|\mathbf{y} \sim \text{GP}(\tilde{\mathbf{m}}, \tilde{\mathbf{k}}) \\ \tilde{\mathbf{m}}(x) = k(x, \mathbf{x})(\mathbf{K}_\mathbf{x} + \sigma_\epsilon^2 \mathbf{I}_n)^{-1} \mathbf{y} \\ \tilde{\mathbf{k}}(x, x') = k(x, x') - k(x, \mathbf{x})(\mathbf{K}_\mathbf{x} + \sigma_\epsilon^2 \mathbf{I}_n)^{-1} k(\mathbf{x}, x'), \end{cases} \quad (\text{A11})$$

where $k(x, \mathbf{x}) = [k(x, x_1) \dots k(x, x_n)]$ and $k(\mathbf{x}, x') = [k(x_1, x') \dots k(x_n, x')]^\top$.

Since this is still a GP, we can plot it on the `grid` using the exact same procedure as for the prior. Namely, define the posterior mean vector and posterior covariance matrix on `grid`

$$\tilde{\mathbf{m}} = \begin{bmatrix} \tilde{\mathbf{m}}(z_1) \\ \vdots \\ \tilde{\mathbf{m}}(z_p) \end{bmatrix}, \quad \tilde{\mathbf{K}} = \begin{bmatrix} \tilde{\mathbf{k}}(z_1, z_1) & \dots & \tilde{\mathbf{k}}(z_1, z_p) \\ \vdots & \ddots & \vdots \\ \tilde{\mathbf{k}}(z_p, z_1) & \dots & \tilde{\mathbf{k}}(z_p, z_p) \end{bmatrix}. \quad (\text{A12})$$

Then

$$\begin{bmatrix} f(z_1)|\mathbf{y} \\ \vdots \\ f(z_p)|\mathbf{y} \end{bmatrix} \sim \mathcal{N}(\tilde{\mathbf{m}}, \tilde{\mathbf{K}}). \quad (\text{A13})$$

Therefore, the posterior mean vector $\tilde{\mathbf{m}}$ provides the mean of $f(x)|\mathbf{y}$ at every grid point and the diagonal elements of the covariance matrix $\tilde{\mathbf{K}}$ provide the standard deviation of $f(x)|\mathbf{y}$ at every grid point (which does not simplify to σ_f this time). Figure A5 illustrates how updating the prior with 1, 2, and 10 observations affects the posterior mean and the posterior credible interval.

Furthermore, we are also able to draw samples from the multivariate normal distribution $\mathcal{N}(\tilde{\mathbf{m}}, \tilde{\mathbf{K}})$ by taking

$$\tilde{\mathbf{f}} = \tilde{\mathbf{m}} + \tilde{\mathbf{K}}^{1/2} \mathbf{u}, \quad \mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p). \quad (\text{A14})$$

Such a sample $\tilde{\mathbf{f}} \in \mathbb{R}^p$ corresponds to a sample path from the posterior GP $f(x)|\mathbf{y}$ evaluated on the grid points. Figure A6 shows the plots of 10 sample paths from this multivariate normal distribution. Unlike for the prior, the sample paths are now concentrated around the observations and therefore provide reasonable candidates for a

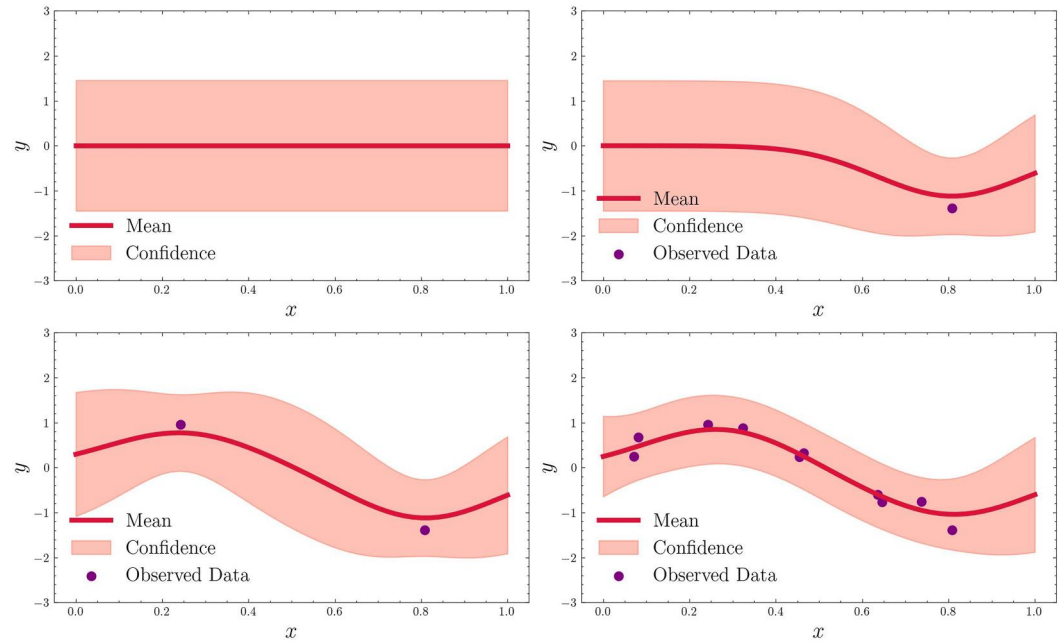


Figure A5. Plots of the posterior mean function and posterior 95% credible interval when updating the prior with 0, (top-left), 1 (top-right), 2 (bottom-left) and 10 (bottom-right) observations.




Figure A6. Plots of the posterior mean function and posterior 95% credible interval + plots of 10 sample paths from $\tilde{\mathbf{f}} \sim \mathcal{N}(\tilde{\mathbf{m}}, \tilde{\mathbf{K}})$.

predictive model $x \rightarrow y$. In other words, the posterior GP induces a probability distribution over functions that provide a sound fit to observations, and are reasonable predictive models.

A2. Matérn Covariance

The Matérn covariances are a class of stationary covariance functions widely used in spatial statistics. The Matérn- ν covariance between two points $x, x' \in \mathbb{R}$ is given by

$$C_\nu(x, x') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{|x - x'|}{\ell} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{|x - x'|}{\ell} \right), \quad (\text{A15})$$

where Γ is the gamma function, K_ν is the modified Bessel function, σ^2 is a variance parameter and $\ell > 0$ is a lengthscale hyperparameter.

The covariance function expression considerably simplifies for $\nu = p + 1/2$ where $p \in \mathbb{N}$. For example, for $\nu = 1/2$ ($p = 0$), $\nu = 3/2$ ($p = 1$) and $\nu = 5/2$ ($p = 2$) we have

$$C_{1/2}(x, x') = \sigma^2 \exp\left(-\frac{|x - x'|}{\ell}\right) \quad (\text{A16})$$

$$C_{3/2}(x, x') = \sigma^2 \left(1 + \sqrt{3} \frac{|x - x'|}{\ell}\right) \exp\left(-\sqrt{3} \frac{|x - x'|}{\ell}\right) \quad (\text{A17})$$

$$C_{5/2}(x, x') = \sigma^2 \left(1 + \sqrt{5} \frac{|x - x'|}{\ell} + \frac{5|x - x'|^2}{3\ell^2}\right) \exp\left(-\sqrt{5} \frac{|x - x'|}{\ell}\right) \quad (\text{A18})$$

When using the Matérn covariance as the covariance function of a GP, larger values of ν bestow greater smoothness on the GP sample paths. Namely, when $\nu = 1/2$ sample paths are continuous, when $\nu = 3/2$ sample paths are differentiable and when $\nu = 5/2$ sample paths are twice differentiable. More generally, for a Matérn- $p + 1/2$ covariance function, sample paths from a GP are p times continuously differentiable (with the convention that 0 times means simply continuous). Figure A7 shows sample paths drawn from GPs with Matérn kernels for different values of ν .

In the limit where $\nu \rightarrow \infty$, the Matérn covariance converges to the squared exponential covariance given by

$$C_\infty(x, x') = \sigma^2 \exp\left(-\frac{|x - x'|^2}{2\ell^2}\right). \quad (\text{A19})$$

The sample paths of a GP with a squared exponential covariance function are infinitely differentiable.

When $x, x' \in \mathbb{R}^d$, the distance $|x - x'|$ can be substituted by the norm $\|x - x'\| = \sqrt{\sum_{i=1}^d (x_i - x'_i)^2}$. The covariance is called an anisotropic—or automatic relevance determination—kernel when each dimension has its own independent lengthscale parameter $\ell_i > 0$. For example, the Matérn-1/2 and Matérn-3/2 anisotropic kernels write




Figure A7. Example of sample paths drawn from a Gaussian process with kernel set as a Matérn-1/2, Matérn-3/2 and Matérn-5/2.

$$C_{1/2}(x, x') = \sigma^2 \exp\left(-\sqrt{\sum_{i=1}^d \frac{(x_i - x'_i)^2}{\ell_i^2}}\right) \quad (\text{A20})$$

$$C_{3/2}(x, x') = \sigma^2 \left(1 + \sqrt{3 \sum_{i=1}^d \frac{(x_i - x'_i)^2}{\ell_i^2}}\right) \exp\left(-\sqrt{3 \sum_{i=1}^d \frac{(x_i - x'_i)^2}{\ell_i^2}}\right). \quad (\text{A21})$$

Appendix B: Supporting Materials for FaIRGP Derivation

B1. Useful Results

Lemma B1 Let $a, b > 0$ and $u, v \geq 0$. We have

$$(a + b)\min(u, v) - (au + bv) = \begin{cases} -b|u - v| & \text{if } u \leq v \\ -a|u - v| & \text{if } u \geq v. \end{cases} \quad (\text{B1})$$

Proof If $u \leq v$,

$$\begin{aligned} (a + b)\min(u, v) - (au + bv) &= (a + b)u - au - bv \\ &= bu - bv \\ &= -b|u - v|. \end{aligned}$$

If $u \geq v$,

$$\begin{aligned} (a + b)\min(u, v) - (au + bv) &= (a + b)v - au - bv \\ &= av - au \\ &= -a|u - v|. \end{aligned}$$

B2. Derivation of FaIRGP

Consider a k -box EBM specified by the stochastic temperature response model

$$d\mathbf{X}(t) = \mathbf{A}\mathbf{X}(t) dt + \mathbf{b}F(t) dt, \quad (\text{B2})$$

where $F(t) \sim \text{GP}(F, K)$, a forcing feedback vector $\mathbf{b} = [1/C_1 \ 0 \ \dots \ 0]^T$, and a forcing feedback matrix given by

$$\mathbf{A} = \begin{bmatrix} -(\kappa_1 + \kappa_2)/C_1 & \kappa_2/C_1 & 0 & \dots & 0 & 0 & 0 \\ \kappa_2/C_2 & -(\kappa_2 + \kappa_3)/C_2 & \kappa_3/C_2 & \dots & 0 & 0 & 0 \\ 0 & \kappa_3/C_3 & -(\kappa_3 + \kappa_4)/C_3 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & -(\kappa_{k-2} + \kappa_{k-1})/C_{k-2} & \kappa_{k-1}/C_{k-2} & 0 \\ 0 & 0 & 0 & \dots & \kappa_{k-1}/C_{k-1} & -(\kappa_{k-1} + \epsilon\kappa_k)/C_{k-1} & \epsilon\kappa_k/C_{k-1} \\ 0 & 0 & 0 & \dots & 0 & \kappa_k/C_k & -\kappa_k/C_k \end{bmatrix}$$

Denote $\mathbf{A} = \mathbf{\Phi}\mathbf{D}\mathbf{\Phi}^{-1}$ the diagonalization of the feedback matrix, where \mathbf{D} is a diagonal matrix with diagonal elements

$$\mathbf{D} = \text{Diag}(D_1, \dots, D_k). \quad (\text{B3})$$

Then we have

$$d\mathbf{X}(t) = \mathbf{A}\mathbf{X}(t) dt + \mathbf{b}F(t) dt \quad (\text{B4})$$

$$\Rightarrow \Phi^{-1} d\mathbf{X}(t) = \mathbf{D}\Phi^{-1}\mathbf{X}(t) dt + \Phi^{-1}\mathbf{b}F(t) dt \quad (\text{B5})$$

$$\Rightarrow d[\Phi^{-1}\mathbf{X}(t)] = \mathbf{D}[\Phi^{-1}\mathbf{X}(t)] dt + [\Phi^{-1}\mathbf{b}] F(t) dt. \quad (\text{B6})$$

Therefore, by using the notations from Millar et al. (2017)

$$D_i = -1/d_i \quad (\text{B7})$$

$$\Phi^{-1}\mathbf{X}(t) = [S_1(t) \quad \dots \quad S_k(t)]^T \quad (\text{B8})$$

$$\Phi^{-1}\mathbf{b} = [q_1/d_1 \quad \dots \quad q_k/d_k]^T, \quad (\text{B9})$$

the diagonalized system can be written as an impulse response system, where for $i \in \{1, \dots, k\}$ we have

$$dS_i(t) = -\frac{1}{d_i}S_i(t) dt + \frac{q_i}{d_i}F(t) dt = \frac{1}{d_i}(q_i F(t) - S_i(t)) dt. \quad (\text{B10})$$

Consider now the thermal response $S_i(t)$ of the i th SDE. This thermal response is given by

$$S_i(t) = \frac{q_i}{d_i} \int_0^t F(s) e^{-(t-s)/d_i} ds, \quad (\text{B11})$$

where we recall that $F(t) \sim \text{GP}(F, K)$. Because GPs are closed under linear transformations, $S_i(t)$ must also be GPs for all $i \in \{1, \dots, k\}$. Their mean functions can be computed following

$$m_i(t) := \mathbb{E}[S_i(t)] = \mathbb{E}\left[\frac{q_i}{d_i} \int_0^t F(s) e^{-(t-s)/d_i} ds\right] \quad (\text{B12})$$

$$= \frac{q_i}{d_i} \int_0^t \mathbb{E}[F(s)] e^{-(t-s)/d_i} ds \quad (\text{B13})$$

$$= \frac{q_i}{d_i} \int_0^t F(s) e^{-(t-s)/d_i} ds. \quad (\text{B14})$$

Their cross-covariance functions can be computed following

$$k_{ij}(t, t') := \text{Cov}(S_i(t), S_j(t')) \quad (\text{B15})$$

$$= \text{Cov}\left(\frac{q_i}{d_i} \int_0^t F(s) e^{-(t-s)/d_i} ds, \frac{q_j}{d_j} \int_0^{t'} F(s') e^{-(t'-s')/d_j} ds'\right) \quad (\text{B16})$$

$$= \frac{q_i q_j}{d_i d_j} \int_0^t \int_0^{t'} \text{Cov}(F(s), F(s')) e^{-(t-s)/d_i} e^{-(t'-s')/d_j} ds ds' \quad (\text{B17})$$

$$= \frac{q_i q_j}{d_i d_j} \int_0^t \int_0^{t'} K(s, s') e^{-(t-s)/d_i} e^{-(t'-s')/d_j} ds ds'. \quad (\text{B18})$$

So we have $S_i(t) \sim \text{GP}(m_i, k_{ii})$ for any $i \in \{1, \dots, k\}$. Finally, if we define $T(t) = \sum_{i=1}^k S_i(t)$, since the sum of Gaussians is still a Gaussian, we know that $T(t)$ must also be a GP. And we can compute its mean function following

$$m_T(t) := \mathbb{E}[T(t)] = \mathbb{E}\left[\sum_{i=1}^k S_i(t)\right] = \sum_{i=1}^k \mathbb{E}[S_i(t)] = \sum_{i=1}^k m_i(t), \quad (\text{B19})$$

and its covariance function following

$$k_T(t, t') := \text{Cov}(T(t), T(t')) = \text{Cov}\left(\sum_{i=1}^k S_i(t), \sum_{j=1}^k S_j(t')\right) \quad (\text{B20})$$

$$= \sum_{i=1}^k \sum_{j=1}^k \text{Cov}(S_i(t), S_j(t')) \quad (\text{B21})$$

$$= \sum_{i=1}^k \sum_{j=1}^k k_{ij}(t, t'). \quad (\text{B22})$$

We conclude that

$$\begin{cases} T(t) \sim \text{GP}(m_T, k_T) \\ m_T(t) = \sum_{i=1}^k m_i(t) \\ k_T(t, t') = \sum_{i=1}^k \sum_{j=1}^k k_{ij}(t, t'). \end{cases} \quad (\text{B23})$$

B3. Accounting for Climate Internal Variability in FaIRGP

Consider again the same problem, but introducing an additional white noise term in the temperature response SDE

$$d\mathbf{X}(t) = \mathbf{A}\mathbf{X}(t) dt + \mathbf{b}F(t) dt + \sigma\mathbf{b} dB(t), \quad (\text{B24})$$

where $B(t)$ denotes the standard Brownian motion.

Following the same derivation steps we get

$$d\mathbf{X}(t) = \mathbf{A}\mathbf{X}(t) dt + \mathbf{b}F(t) dt + \sigma\mathbf{b} dB(t) \quad (\text{B25})$$

$$\Rightarrow \Phi^{-1} d\mathbf{X}(t) = \mathbf{D}\Phi^{-1}\mathbf{X}(t) dt + \Phi^{-1}\mathbf{b}F(t) dt + \Phi^{-1}\sigma\mathbf{b} dB(t) \quad (\text{B26})$$

$$\Rightarrow d[\Phi^{-1}\mathbf{X}(t)] = \mathbf{D}[\Phi^{-1}\mathbf{X}(t)] dt + [\Phi^{-1}\mathbf{b}] F(t) dt + \sigma[\Phi^{-1}\mathbf{b}] dB(t). \quad (\text{B27})$$

which gives an impulse response form

$$dS_i(t) = \frac{1}{d_i}(q_i F(t) - S_i(t)) dt + \sigma \frac{q_i}{d_i} dB(t), \quad \forall i \in \{1, \dots, k\}. \quad (\text{B28})$$

The solution to the i th SDE is now given by

$$S_i(t) = \underbrace{\frac{q_i}{d_i} \int_0^t F(s) e^{-(t-s)/d_i} ds}_{S_i^\circ(t)} + \underbrace{\sigma \frac{q_i}{d_i} \int_0^t e^{-(t-s)/d_i} dB(s)}_{\xi_i(t)} \quad (\text{B29})$$

$$= S_i^\circ(t) + \sigma \xi_i(t). \quad (\text{B30})$$

$S_i^\circ(t)$ corresponds to the GP we have obtained in the previous derivation without white noise, that is, $S_i^\circ(t) \sim \text{GP}(m_i, k_{ii})$. $\xi_i(t)$ is also a GP, with mean zero, and cross-covariance function given by

$$\begin{aligned} \text{Cov}(\xi_i(t), \xi_j(t')) &= \mathbb{E}[(\xi_i(t) - \mathbb{E}[\xi_i(t)])(\xi_j(t') - \mathbb{E}[\xi_j(t')])] \\ &= \mathbb{E}[\xi_i(t) \xi_j(t')] \\ &= \mathbb{E}\left[\left(\frac{q_i}{d_i} \int_0^t e^{-(t-s)/d_i} dB(s)\right) \left(\frac{q_j}{d_j} \int_0^{t'} e^{-(t'-s')/d_j} dB(s')\right)\right] \\ &= \frac{q_i q_j}{d_i d_j} e^{-t/d_i - t'/d_j} \mathbb{E}\left[\int_0^t e^{s/d_i} dB(s) \int_0^{t'} e^{s'/d_j} dB(s')\right] \\ &= \frac{q_i q_j}{d_i d_j} e^{-t/d_i - t'/d_j} \int_0^{\min(t, t')} e^{(1/d_i + 1/d_j)s} ds \\ &= \frac{q_i q_j}{d_i d_j} e^{-t/d_i - t'/d_j} \frac{d_i d_j}{d_i + d_j} \left(e^{\frac{d_i + d_j}{d_i d_j} \min(t, t')} - 1\right) \\ &= \frac{q_i q_j}{d_i + d_j} e^{-(d_i t + d_j t')/d_i d_j} \left(e^{\frac{d_i + d_j}{d_i d_j} \min(t, t')} - 1\right) \\ &= \begin{cases} \frac{q_i q_j}{d_i + d_j} (e^{-|t-t'|/d_i} - e^{-(d_j t + d_i t')/d_i d_j}) & \text{if } t \leq t' \\ \frac{q_i q_j}{d_i + d_j} (e^{-|t-t'|/d_j} - e^{-(d_j t + d_i t')/d_i d_j}) & \text{if } t \geq t' \end{cases} \quad (\text{Lemma Appendix B.1}) \\ &\sim \begin{cases} \frac{q_i q_j}{d_i + d_j} e^{-|t-t'|/d_i} & \text{if } t \leq t' \\ \frac{q_i q_j}{d_i + d_j} e^{-|t-t'|/d_j} & \text{if } t \geq t' \end{cases} \quad \text{when } t \gg d_i \text{ or } t' \gg d_j. \end{aligned}$$

Therefore, if we define

$$\gamma_{ij}(t, t') = \frac{q_i q_j}{d_i + d_j} \exp\left(\frac{-|t - t'|}{d_i \mathbb{1}_{\{t \leq t'\}} + d_j \mathbb{1}_{\{t > t'\}}}\right), \quad (\text{B31})$$

which simplifies when $i = j$ to

$$\gamma_i(t, t') = \gamma_{ii}(t, t') = \frac{q_i^2}{2d_i} \exp\left(-\frac{|t - t'|}{d_i}\right), \quad (\text{B32})$$

we obtain that in the long time regime, we can approximate $\xi_i(t) \sim \text{GP}(0, \gamma_i)$. And because $S_i^\circ(t)$ and $\xi_i(t)$ are independent processes, we obtain that

$$S_i(t) \sim \text{GP}(m_i, k_{ii} + \sigma^2 \gamma_i). \quad (\text{B33})$$

Finally, if we take again $T(t) = \sum_{i=1}^k S_i(t)$, then $T(t)$ must be a GP. Its mean function is given by

$$m_T(t) := \mathbb{E}[T(t)] = \mathbb{E}\left[\sum_{i=1}^k S_i(t)\right] = \sum_{i=1}^k \mathbb{E}[S_i(t)] = \sum_{i=1}^k m_i(t), \quad (\text{B34})$$

and its covariance function is given by

$$k_T(t, t') := \text{Cov}(T(t), T(t')) \quad (\text{B35})$$

$$= \text{Cov}\left(\sum_{i=1}^k S_i(t), \sum_{j=1}^k S_j(t')\right) \quad (\text{B36})$$

$$= \sum_{i,j=1}^k \text{Cov}(S_i(t), S_j(t')) \quad (\text{B37})$$

$$= \sum_{i,j=1}^k \text{Cov}(S_i^\circ(t) + \sigma\eta_i(t), S_i^\circ(t') + \sigma\eta_j(t')) \quad (\text{B38})$$

$$= \sum_{i,j=1}^k \text{Cov}(S_i^\circ(t), S_i^\circ(t')) + \sigma^2 \text{Cov}(\eta_i(t), \eta_j(t')) (S_i^\circ, S_i^\circ \mathbb{I}_{\xi_i}, \xi_j) \quad (\text{B39})$$

$$= \sum_{i,j=1}^k k_{ij}(t, t') + \sigma^2 \sum_{i,j=1}^k \gamma_{ij}(t, t'). \quad (\text{B40})$$

However, if $t \leq t'$ we have

$$\sum_{i,j=1}^k \gamma_{ij}(t, t') = \sum_{i,j=1}^k \frac{q_i q_j}{d_i + d_j} \exp\left(-\frac{|t - t'|}{d_i}\right) \quad (\text{B41})$$

$$= \sum_{i=1}^k \left(\sum_{j=1}^k \frac{q_i q_j}{d_i + d_j} \right) \exp\left(-\frac{|t - t'|}{d_i}\right) \quad (\text{B42})$$

$$= \sum_{i=1}^k \underbrace{\frac{2d_i}{q_i^2} \left(\sum_{j=1}^k \frac{q_i q_j}{d_i + d_j} \right)}_{\nu_i} \frac{q_i^2}{2d_i} \exp\left(-\frac{|t - t'|}{d_i}\right) \quad (\text{B43})$$

$$= \sum_{i=1}^k \nu_i \gamma_i(t, t'). \quad (\text{B44})$$

When $t \geq t'$ we can refactor terms similarly into a sum over j . By symmetry of the indices we conclude that

$$\left\{ \begin{array}{l} T(t) \sim \text{GP}(m_T, k_T) \\ m_T(t) = \sum_{i=1}^k m_i(t) \\ k_T(t, t') = \sum_{i,j=1}^k k_{ij}(t, t') + \sigma^2 \sum_{i=1}^k \nu_i \gamma_i(t, t') \\ \nu_i = \sum_{j=1}^k \frac{2d_i q_j}{q_i (d_i + d_j)} \end{array} \right. \quad (\text{B45})$$

B4. Accounting for Unforced Fluctuations in Top-Of-Atmosphere Radiative Flux

In addition to the variability arising from unforced transfers of energy modeled in the last section, Cummins et al. (2020) propose to also account for unforced fluctuations in the top-of-atmosphere radiative flux. Let us denote by $\tilde{F}(t)$ the radiative forcing model proposed by Cummins et al. (2020), to distinguish it from the FaIRGP prior $F(t)$. They propose to model $\tilde{F}(t)$ as an Ornstein-Uhlenbeck process,

$$d\tilde{F}(t) = \theta(F(t) - \tilde{F}(t)) dt + \varsigma d\tilde{B}(t) \quad (\text{B46})$$

where $\theta > 0$ is the autocorrelation parameter, $\varsigma > 0$ is the white noise variance, $\tilde{B}(t)$ is another standard Brownian motion (independent from $B(t)$ in the temperature response SDE), and $F(t)$ is a deterministic forcing component.

Assuming that $F(0) = 0$, we know that the solution of this SDE is given by

$$\tilde{F}(t) = \underbrace{\int_0^t \theta F(s) e^{-\theta(t-s)} ds}_{\tilde{F}(t)} + \underbrace{\varsigma \int_0^t e^{-\theta(t-s)} d\tilde{B}(s)}_{\eta(t)} \quad (\text{B47})$$

$$= \tilde{F}(t) + \varsigma \eta(t). \quad (\text{B48})$$

$\tilde{F}(t)$ is a deterministic quantity, whereas $\eta(t)$ is as before a GP with mean zero, and a covariance function γ_θ which can be approximated in the long time regime by

$$\gamma_\theta(t, t') = \frac{1}{2\theta} \exp(-\theta|t - t'|). \quad (\text{B49})$$

Therefore, we get that in the long term regime

$$\tilde{F}(t) \sim \text{GP}(\tilde{F}, \varsigma \gamma_\theta). \quad (\text{B50})$$

This can be combined with the FaIRGP prior by substituting the deterministic forcing $F(t)$ by the FaIRGP forcing prior $F(t) \sim \text{GP}(F, K)$. Namely, if we define

$$\tilde{K}(t, t') = \int_0^t \int_0^{t'} \theta^2 K(s, s') e^{-\theta(t-s)} e^{-\theta(t'-s')} ds ds', \quad (\text{B51})$$

then we obtain a new FaIRGP prior over the forcing

$$\tilde{F}(t) \sim \text{GP}(\tilde{F}, \tilde{K} + \varsigma \gamma_\theta), \quad (\text{B52})$$

which accounts for the unforced fluctuations in the top-of-atmosphere radiative flux like in Cummins et al. (2020).

B5. Analytical Expression for FaIRGP Probability Distribution

A key benefit of FaIRGP is its complete mathematical tractability. This tractability in particular includes access to the analytical expression of: (a) the marginal probability distribution over the training data, which can be used as an objective to tune the model hyperparameters and (b) of the probability distribution predicted over emulated temperatures, which can be useful for downstream applications of emulation.

In what follows, we assume access to a training set $\mathcal{D} = \{\mathbf{t}, \mathbf{E}, \mathbf{T}\}$ of size n , where we follow the notational conventions from Section 3.5.

B5.1. Analytical Expression of the Marginal Log-Likelihood

Using the prior mean and covariance of the FaIRGP prior over $T(t)$, we define a prior mean vector and a prior covariance matrix over \mathbf{t} , \mathbf{E} given by

$$\mathbf{m} = m_T(\mathbf{t}) = \begin{bmatrix} m_T(t_1) \\ \vdots \\ m_T(t_n) \end{bmatrix} \quad (\text{B53})$$

$$\mathbf{K} = k_T(\mathbf{E}, \mathbf{E}) = [k_T(E_i, E_j)]_{1 \leq i, j \leq n}. \quad (\text{B54})$$

Further, consider the temperatures internal variability covariance matrix defined by

$$\mathbf{\Gamma} = \gamma_T(\mathbf{t}, \mathbf{t}) = [\gamma_T(t_i, t_j)]_{1 \leq i, j \leq n}. \quad (\text{B55})$$

Then the prior distribution over temperatures is exactly given by the multivariate normal distribution $\mathcal{N}(\mathbf{m}, \mathbf{K} + \sigma^2 \mathbf{\Gamma})$. Let us denote $\mathbf{K}_{\sigma^2} = \mathbf{K} + \sigma^2 \mathbf{\Gamma}$ for conciseness. Then, we can exactly evaluate the marginal probability of $\mathbf{T} \in \mathbb{R}^n$ under the prior distribution following

$$p(\mathbf{T}|\mathbf{E}, \mathbf{t}) = \frac{1}{\sqrt{(2\pi)^n \det(\mathbf{K}_{\sigma^2})}} \exp\left(-\frac{1}{2}(\mathbf{T} - \mathbf{m})^\top \mathbf{K}_{\sigma^2}^{-1} (\mathbf{T} - \mathbf{m})\right), \quad (\text{B56})$$

which can be used as a maximization objective to tune model hyperparameters such that the observed temperatures \mathbf{T} have the greatest possible likelihood under the prior. In practice, we prefer working with the marginal log-likelihood for computational stability. It is given in closed-form by

$$\log p(\mathbf{T}|\mathbf{E}, \mathbf{t}) = -\frac{1}{2}\{n \log(2\pi) + \log \det(\mathbf{K}_{\sigma^2}) + (\mathbf{T} - \mathbf{m})^\top \mathbf{K}_{\sigma^2}^{-1} (\mathbf{T} - \mathbf{m})\}. \quad (\text{B57})$$

Let θ denote a set of model parameters we wish to tune against simulated data \mathbf{t} , \mathbf{E} , \mathbf{T} . For example, we can choose the internal variability variance and GP kernel parameters $\theta = \{\sigma, \sigma_F, \ell_1, \dots, \ell_d\}$, FaIR temperature response parameters $\theta = \{d_1, \dots, d_k, q_1, \dots, q_k\}$, the radiative forcing model parameters $\theta = \{\alpha_{\log}^{\chi}, \alpha_{\text{lin}}^{\chi}, \alpha_{\text{sqr}}^{\chi}\}$, or any combination of these options. Let θ_0 be the initial values we arbitrarily choose for these parameters, and $\eta > 0$ is a typically small positive number called the *learning rate*. Then the gradient descent algorithm outlined below proposes a principled procedure to tune θ such that it maximizes the marginal log-likelihood $\log p(\mathbf{T}|\mathbf{E}, \mathbf{t})$ (Algorithm 3).

Algorithm 3. Maximum Likelihood Calibration of θ With Gradient Descent

```

1: Input:  $\mathbf{t}, \mathbf{E}, \mathbf{T}, \theta_0, \eta > 0, N \in \mathbb{N}$ 
2: for  $t \in \{1, \dots, N\}$  do
3:   Compute  $\log p(\mathbf{T}|\mathbf{E}, \mathbf{t}) = \log \mathcal{N}(\mathbf{T}|\mathbf{m}, \mathbf{K} + \sigma^2 \mathbf{\Gamma})$ 
4:   Take  $\theta_t \leftarrow \theta_{t-1} + \eta \nabla_{\theta} \log p(\mathbf{T}|\mathbf{E}, \mathbf{t})$ 
5: end for
6: Return:  $\theta_N$ 

```

The learning rate $\eta > 0$ controls the size of the steps taken in the direction of the gradient. A higher learning rate results in larger steps, while a lower learning rate leads to smaller steps. If the learning rate is too high, the algorithm may overshoot the maximum and fail to converge. On the other hand, if the learning rate is too low, the algorithm may take a long time to converge.

In the presented gradient descent algorithm, the algorithm stops after the maximum number of iterations N has been reached. More sophisticated stopping criteria can be devised, such as stopping if the absolute change in the marginal log-likelihood falls below a predefined threshold.

Finally, it may be that the size n of the data set $\mathcal{D} = \{\mathbf{t}, \mathbf{E}, \mathbf{T}\}$ is so large that the evaluation of $\log p(\mathbf{T}|\mathbf{E}, \mathbf{t})$ is in practice computationally prohibitive. In this case, a popular variant of gradient descent called *mini-batch stochastic gradient descent* can be employed. At each iteration, instead of computing the marginal log-likelihood over the entire data set, we compute it over a randomly selected subset of the training data, called a mini-batch. This mini-batch typically has a size $m \ll n$. Because only a subset of the data is used for each update, the updates are noisier, introducing more variability in the optimization process. It is however more computationally efficient, especially for large data sets, as it processes only a small subset of the data in each iteration.

B5.2. Analytical Expression of the Posterior Distribution

Suppose that we want to emulate the global temperature response for a different emission scenario where we have access to greenhouse gas and aerosol emission data E_1^*, \dots, E_m^* at times $t_1^* < \dots < t_m^*$. We concatenate them into

$$\mathbf{t}^* = \begin{bmatrix} t_1^* \\ \vdots \\ t_m^* \end{bmatrix} \in \mathbb{R}^m, \quad \mathbf{E}^* = \begin{bmatrix} E_1^* \\ \vdots \\ E_m^* \end{bmatrix} \in \mathbb{R}^{m \times d}, \quad (\text{B58})$$

where $d \geq 0$ is the number of emission agents χ_1, \dots, χ_d we observe emissions from.

Using the posterior mean and covariance functions when $\mathbf{T}(t)$ is updated with the training set \mathcal{D} , we define a posterior mean vector and posterior covariance matrix over $\mathbf{t}^*, \mathbf{E}^*$ given by

$$\bar{\mathbf{m}}^* = \bar{m}_{\mathbf{T}}(\mathbf{t}^*) \quad (\text{B59})$$

$$\bar{\mathbf{K}}^* = \bar{k}_{\mathbf{T}}(\mathbf{E}^*, \mathbf{E}^*). \quad (\text{B60})$$

Further, consider the emulated temperatures internal variability covariance matrix defined by

$$\mathbf{\Gamma}^* = \gamma_{\mathbf{T}}(\mathbf{t}^*, \mathbf{t}^*). \quad (\text{B61})$$

Then, the posterior distribution over emulated temperatures is exactly given by the multivariate normal distribution $\mathcal{N}(\bar{\mathbf{m}}^*, \bar{\mathbf{K}}^* + \sigma^2 \mathbf{\Gamma}^*)$. Let us denote $\bar{\mathbf{K}}_{\sigma^2}^* = \bar{\mathbf{K}}^* + \sigma^2 \mathbf{\Gamma}^*$ for conciseness. This means that for a given temperature vector $\mathbf{T}^* \in \mathbb{R}^m$, we can exactly evaluate the probability of \mathbf{T}^* under the predicted distribution following

$$p(\mathbf{T}^*|\mathcal{D}, \mathbf{E}^*, \mathbf{t}^*) = \frac{1}{\sqrt{(2\pi)^m \det(\bar{\mathbf{K}}_{\sigma^2}^*)}} \exp\left(-\frac{1}{2}(\mathbf{T}^* - \bar{\mathbf{m}}^*)^{\top} (\bar{\mathbf{K}}_{\sigma^2}^*)^{-1} (\mathbf{T}^* - \bar{\mathbf{m}}^*)\right). \quad (\text{B62})$$

The vector \mathbf{T}^* may be a retained test scenario, in which case computing $p(\mathbf{T}^*|\mathcal{D}, \mathbf{E}^*, \mathbf{t}^*)$ provides an evaluation of the emulated posterior $\mathcal{N}(\bar{\mathbf{m}}^*, \bar{\mathbf{K}}_{\sigma^2}^*)$. \mathbf{T}^* may also simply correspond to temperature for which we would like to assess the probability under emission scenario $\{\mathbf{t}^*, \mathbf{E}^*\}$. A natural application is in detection attribution studies. Indeed, we can take \mathbf{T}^* to be observed historical temperatures, and $\{\mathbf{t}^*, \mathbf{E}^*\}$ to be a counterfactual historical scenario without anthropogenic forcing. By evaluating $p(\mathbf{T}^*|\mathcal{D}, \mathbf{E}^*, \mathbf{t}^*)$ we would be able to assess the probability of historical temperature observations to occur in a scenario without anthropogenic forcing.

The posterior distribution can also be used to draw a sample \mathbf{V} if needed following

$$\mathbf{V} = (\bar{\mathbf{K}}_{\sigma^2}^*)^{1/2} \mathbf{U} + \bar{\mathbf{m}}^*, \quad \mathbf{U} \sim \mathcal{N}(0, \mathbf{I}_m). \quad (\text{B63})$$

Finally, whilst the above is formulated for global mean surface temperatures, we can also derive analytical forms for the posterior distribution over the radiative forcing and for spatially-resolved emulation.

Appendix C: Complementary Experimental Results

C1. FaIRGP Posterior Sample Paths Over SSPs

Positions of NorESM2-LM simulated data: (e) in Figure C1, (f) in Figure C2, (f) in Figure C3 and (a) in Figure C4.




Figure C1. Example of five sample paths from the predicted posterior distribution over global annual mean surface temperature anomaly under *SSP126* with FaIRGP + NorESM2-LM simulated global annual mean surface temperature anomaly. The posterior is conditioned on experiments $D_{\text{train}} = \{\text{historical}, \text{SSP245}, \text{SSP370}, \text{SSP585}\}$. The position of the NorESM2-LM simulated data is given at the end of the section.




Figure C2. Example of five sample paths from the predicted posterior distribution over global annual mean surface temperature anomaly under *SSP245* with FaIRGP + NorESM2-LM simulated global annual mean surface temperature anomaly. The posterior is conditioned on experiments $D_{\text{train}} = \{\text{historical}, \text{SSP126}, \text{SSP370}, \text{SSP585}\}$. The position of the NorESM2-LM simulated data is given at the end of the section.

C2. Comparison of Plain GP Baseline With ClimateBench GP Emulator From Watson-Parris et al. (2022)

C2.1. Modeling Differences

The plain GP baseline we use is analogous to the GP emulator from ClimateBench (Watson-Parris et al., 2022), but differs in two aspects.




Figure C3. Example of five sample paths from the predicted posterior distribution over global annual mean surface temperature anomaly under *SSP370* with FaIRGP + NorESM2-LM simulated global annual mean surface temperature anomaly. The posterior is conditioned on experiments $D_{\text{train}} = \{\text{historical}, \text{SSP126}, \text{SSP245}, \text{SSP585}\}$. The position of the NorESM2-LM simulated data is given at the end of the section.




Figure C4. Example of five sample paths from the predicted posterior distribution over global annual mean surface temperature anomaly under *SSP585* with FaIRGP + NorESM2-LM simulated global annual mean surface temperature anomaly. The posterior is conditioned on experiments $D_{\text{train}} = \{\text{historical}, \text{SSP126}, \text{SSP245}, \text{SSP370}\}$. The position of the NorESM2-LM simulated data is given at the end of the section.

First, we use in our plain GP baseline an anisotropic covariance structure, denoted as $\rho(E, E') = C_{3/2}(E, E')$, where $C_{3/2}$ denotes the Matérn-3/2 covariance (see Appendix A2).

This kernel introduces a different lengthscale parameter ℓ_χ for each atmospheric agent χ , which is tuned through maximizing the marginal loglikelihood. It is explicitly given by

$$\rho(E, E') = C_{3/2}(E, E') = \left(1 + \sqrt{3 \sum_{\chi} \frac{(E^\chi - E'^\chi)^2}{\ell_\chi^2}}\right) \exp\left(-\sqrt{3 \sum_{\chi} \frac{(E^\chi - E'^\chi)^2}{\ell_\chi^2}}\right), \quad (\text{C1})$$

where E^χ denotes global emission level for agent $\chi \in \{\text{CH}_4, \text{SO}_2, \text{BC}\}$, while for $\chi = \text{CO}_2$, it denotes global cumulative emission levels.

In contrast, Watson-Parris et al. (2022) adopt an additive kernel structure given by

$$\rho(E, E') = \sum_{\chi} \sigma_{\chi} C_{3/2}(E^{\chi}, E'^{\chi}) = \sum_{\chi} \sigma_{\chi} \left(1 + \sqrt{3} \frac{|E^{\chi} - E'^{\chi}|}{\ell_{\chi}} \right) \exp \left(-\sqrt{3} \frac{|E^{\chi} - E'^{\chi}|}{\ell_{\chi}} \right). \quad (\text{C2})$$

This additive structure introduces additional complexity through the variance terms σ_{χ} , which modulate the of each atmospheric agent to the covariance, and need to be tuned alongside the lengthscales ℓ_{χ} . From a functional perspective, choosing an additive kernel is equivalent to representing the temperature response as the sum of independent GPs, where each GP models the response to a single atmospheric agent. Further details on the covariance function be found in (Watson-Parris et al., 2022, Appendix A2).

Second, in our plain GP baseline E^{SO_2} and E^{BC} correspond to global emission levels, whereas Watson-Parris et al. (2022) used the 5 principal components of spatial emission maps for SO_2 and BC. We anticipate that spatially-resolved inputs for aerosol emissions should improve the model's ability to predict spatially-resolved surface temperature features.

C2.2. Predictive Performance Comparison

To compare the predictive performance against our plain GP and FaIRGP, we evaluate the ClimateBench GP for the emulation of global and spatial mean surface temperature anomaly and report scores in Table C1. All emulators are trained on the same training data: *historical*, *SSP126*, *SSP370*, *SSP585*. Predictive performance is evaluated for the emulation of *SSP245* since this is the test data used in (Watson-Parris et al., 2022).

Table C1

Scores of ClimateBench Gaussian Process (GP), Our Plain GP Baseline and FaIRGP for the Task of Emulating Global and Spatial Mean Surface Temperatures on SSP245

| | Emulator | RMSE ↓ | MAE ↓ | Bias | LL ↑ | Calib95 | CRPS ↓ |
|---------|-----------------|-------------|-------------|--------------|--------------|-------------|-------------|
| Global | ClimateBench GP | 0.12 | 0.10 | 0.01 | −0.45 | 1.0 | 0.15 |
| | Plain GP | 0.09 | 0.07 | 0.03 | 0.67 | 1.0 | 0.06 |
| | FaIRGP | 0.09 | 0.08 | −0.01 | 0.71 | 1.0 | 0.06 |
| Spatial | ClimateBench GP | 0.43 | 0.32 | −0.13 | −0.69 | 0.98 | 0.25 |
| | Plain GP | 0.56 | 0.40 | −0.19 | −0.70 | 0.97 | 0.28 |
| | FaIRGP | 0.36 | 0.26 | −0.05 | −0.34 | 0.99 | 0.20 |

Note. Scores are computed over 2015–2100 period for global emulation and over 2080–2100 period for spatial emulation; the best emulator scores are highlighted in bold; ↑/↓ indicates higher/lower is better.

Overall, FaIRGP demonstrates improved scores compared to the ClimateBench baseline GP across various metrics, both in terms of global and spatial emulation. In terms of global emulation, the predictive performance of the Plain GP appears to be better than that of the ClimateBench GP. However, when it comes to spatial surface temperature emulation, the ClimateBench GP outperforms the Plain GP. This improvement can likely be attributed to the utilization of spatially-resolved aerosol information as input in the ClimateBench GP.

C3. Tuned Model Parameters Values

See Table C2.

Table C2

Values Obtained for Internal Variability Variance σ^2 , Forcing Kernel Variance σ_F^2 and Forcing Kernel Lengthscales $\ell_{CO_2}, \ell_{CH_4}, \ell_{SO_2}, \ell_{BC}$ After Tuning Using the Gradient Descent Maximum Likelihood Procedure Described in Section 3.6

| Training data | σ^2 | σ_F^2 | ℓ_{CO_2} | ℓ_{CH_4} | ℓ_{SO_2} | ℓ_{BC} |
|-------------------------|------------|--------------|---------------|---------------|---------------|-------------|
| All SSPs and historical | 0.754 | 0.244 | 1.60 | 3.79 | 2.98 | 3.797865 |
| All but SSP126 | 0.679 | 0.271 | 1.57 | 3.90 | 2.91 | 3.86 |
| All but SSP245 | 0.845 | 0.276 | 1.52 | 3.61 | 2.86 | 3.89 |
| All but SSP370 | 0.709 | 0.241 | 1.87 | 3.87 | 2.91 | 3.40 |
| All but SSP585 | 0.647 | 0.192 | 1.58 | 3.22 | 2.41 | 3.18 |
| Historical only | 0.385 | 0.263 | 3.33 | 2.57 | 1.02 | 3.48 |

C4. Emulating Anthropogenic Aerosols Forcing

See Figure C5.

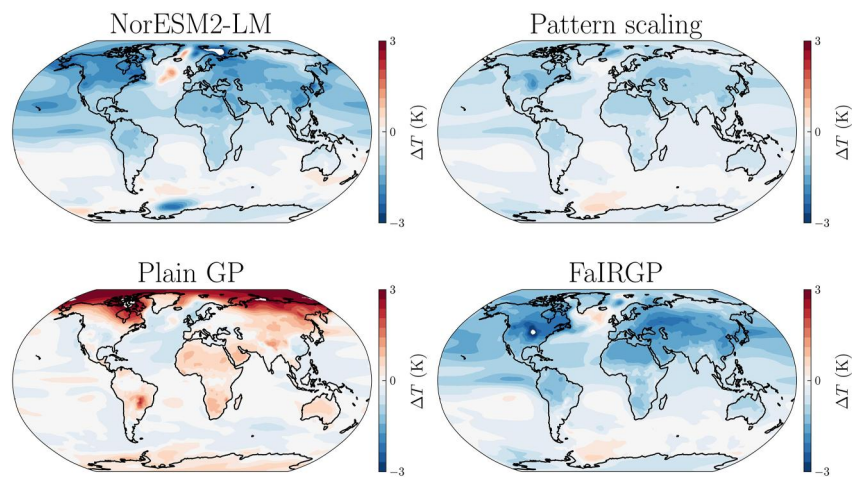


Figure C5. Maps of surface temperatures averaged for the *hist-aer* experiment averaged over 1980–2014 period. *Top-Left:* Groundtruth NorESM2-LM surface temperatures map. *Top-Right:* Prediction with the FaIR pattern scaling baseline. *Bottom-Left:* Prediction with Plain Gaussian process (GP) posterior mean. *Bottom-Right:* Prediction with the FaIRGP posterior mean.

Data Availability Statement

- Data—The data used to run experiments in this paper is obtained from ClimateBenchv1.0 (Watson-Parris et al., 2022) and available here <https://doi.org/10.5281/zenodo.5196512>.
- Software—The model, evaluation metrics and all code used to run experiments and generate the plots is available here <https://doi.org/10.5281/zenodo.8164335>. All figures were made with Matplotlib version 3.6.2 (Caswell et al., 2020; Hunter, 2007) and Cartopy (Met Office, 2010). Models are implemented with PyTorch version 1.12.1 (Paszke et al., 2019) and GPyTorch version 1.9.0 (Gardner et al., 2018).

Acknowledgments

We would like to thank Maybritt Schillinger for insightful discussions and helpful feedback which has greatly benefited this work. We would also like to thank three anonymous reviewers for their thorough reviews and comments, which contributed significantly to the quality of this work. Shahine Bouabid receives funding from the European Union's Horizon 2020 research and innovation programme under Marie Skłodowska-Curie grant agreement No 860100.

References

- Akhtar, M. K., Wibe, J., Simonovic, S. P., & MacGee, J. (2013). Integrated assessment model of society-biosphere-climate-economy-energy system. *Environmental Modelling & Software*, 49, 1–21. <https://doi.org/10.1016/j.envsoft.2013.07.006>
- Allen, M. R., Frame, D. J., Huntingford, C., Jones, C. D., Lowe, J. A., Meinshausen, M., & Meinshausen, N. (2009). Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature*, 458(7242), 1163–1166. <https://doi.org/10.1038/nature08019>
- Alvarez, M. A., Luengo, D., & Lawrence, N. D. (2009). Latent force models. In *Artificial intelligence and statistics*.
- Alvarez, M. A., Luengo, D., & Lawrence, N. D. (2013). Linear latent force models using Gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(11), 2693–2705. <https://doi.org/10.1109/tpami.2013.86>
- Álvarez, M. A., Ward, W., & Guarnizo, C. (2019). Non-linear process convolutions for multi-output Gaussian processes. In *The 22nd international conference on artificial intelligence and statistics* (pp. 1969–1977).

- Baldwin-McDonald, T., & Álvarez, M. A. (2023). Deep latent force models: ODE-based process convolutions for Bayesian deep learning. arXiv preprint arXiv:2311.14828.
- Beusch, L., Gudmundsson, L., & Seneviratne, S. I. (2020a). Crossbreeding CMIP6 Earth system models with an emulator for regionally optimized land temperature projections. *Geophysical Research Letters*, 47(15), e2019GL086812. <https://doi.org/10.1029/2019gl086812>
- Beusch, L., Gudmundsson, L., & Seneviratne, S. I. (2020b). Emulating Earth system model temperatures with MESMER: From global mean temperature trajectories to grid-point-level realizations on land. *Earth System Dynamics*, 11(1), 139–159. <https://doi.org/10.5194/esd-11-139-2020>
- Beusch, L., Nicholls, Z., Gudmundsson, L., Hauser, M., Meinshausen, M., & Seneviratne, S. I. (2022). From emission scenarios to spatially resolved projections with a chain of computationally efficient emulators: Coupling of MAGICC (v7. 5.1) and MESMER (v0. 8.3). *Geoscientific Model Development*, 15(5), 2085–2103. <https://doi.org/10.5194/gmd-15-2085-2022>
- Camps-Valls, G., Verrelst, J., Muñoz-Mari, J., Laparra, V., Mateo-Jiménez, F., & Gomez-Dans, J. (2016). A survey on Gaussian processes for earth-observation data analysis: A comprehensive investigation. *IEEE Geoscience and Remote Sensing Magazine*, 4(2), 58–78. <https://doi.org/10.1109/mgrs.2015.2510084>
- Carlsaw, K. S., & Pringle, K. (2022). Chapter 4 - Global aerosol properties. In K. S. Carlsaw (Ed.), *Aerosols and climate* (pp. 101–133). Elsevier. <https://doi.org/10.1016/B978-0-12-819766-0.00011-0>
- Castruccio, S., McInerney, D. J., Stein, M. L., Liu Crouch, F., Jacob, R. L., & Moyer, E. J. (2014). Statistical emulation of climate model projections based on precomputed GCM runs. *Journal of Climate*, 27(5), 1829–1844. <https://doi.org/10.1175/jcli-d-13-00099.1>
- Castruccio, S., & Stein, M. L. (2013). Global space–time models for climate ensembles. *Annals of Applied Statistics*, 7(3), 1593–1611. <https://doi.org/10.1214/13-aos656>
- Caswell, T. A., Droettboom, M., Lee, A., Hunter, J., Firing, E., Sales De Andrade, E., et al. (2020). matplotlib/matplotlib: REL: v3. 3.1 [Software]. Zenodo. Retrieved from <https://github.com/matplotlib/matplotlib>
- Champney, D. C., & Fuoss, R. M. (1973). Fourier transforms and their physical applications. *Physics Today*, 26(10), 57–59. <https://doi.org/10.1063/1.3128283>
- Collins, M., Knutti, R., Arblaster, J., Dufresne, J.-L., Fichet, T., Friedlingstein, P., et al. (2014). Long-term climate change: Projections, commitments and irreversibility pages 1029 to 1076. In *Climate change 2013 – The physical science basis: Working Group I contribution to the fifth assessment report of the Intergovernmental Panel on Climate Change*.
- Cummins, D. P., Stephenson, D. B., & Stott, P. A. (2020). Optimal estimation of stochastic energy balance model parameters. *Journal of Climate*, 33(18), 7909–7926. <https://doi.org/10.1175/jcli-d-19-0589.1>
- Danabasoglu, G., Lamarque, J.-F., Bacmeister, J., Bailey, D., DuVivier, A., Edwards, J., et al. (2020). The community earth system model version 2 (CESM2). *Journal of Advances in Modeling Earth Systems*, 12(2), e2019MS001916. <https://doi.org/10.1029/2019ms001916>
- Del Rio Amador, L., & Lovejoy, S. (2019). Predicting the global temperature with the stochastic seasonal to interannual prediction system (StocSIPS). *Climate Dynamics*, 53(7–8), 4373–4411. <https://doi.org/10.1007/s00382-019-04791-4>
- Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., & Taylor, K. E. (2016). Overview of the Coupled Model Inter-comparison Project Phase 6 (CMIP6) experimental design and organization. *Geoscientific Model Development*, 9(5), 1937–1958. <https://doi.org/10.5194/gmd-9-1937-2016>
- Flato, G. M. (2011). Earth system models: An overview. *WIREs Climate Change*, 2(6), 783–800. <https://doi.org/10.1002/wcc.148>
- Forster, P., Ramaswamy, V., Artaxo, P., Bernsten, T., Betts, R., Fahey, D. W., et al. (2007). Changes in atmospheric constituents and in radiative forcing. In *Climate change 2007. The physical science basis*. Chapter 2.
- Fredriksen, H.-B., & Rypdal, M. (2017). Long-range persistence in global surface temperatures explained by linear multibox energy balance models. *Journal of Climate*, 30(18), 7157–7168. <https://doi.org/10.1175/jcli-d-16-0877.1>
- Gardner, J., Pleiss, G., Weinberger, K. Q., Bindel, D., & Wilson, A. G. (2018). Gpytorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration [Software]. In *Advances in neural information processing systems* (Vol. 31). Retrieved from <https://github.com/cornellius-gp/gpytorch>
- Gasser, T., Ciais, P., Boucher, O., Quilcaille, Y., Tortora, M., Bopp, L., & Hauglustaine, D. (2017). The compact Earth system model OSCAR v2.2: Description and first results.
- Geoffroy, O., Saint-Martin, D., Olivé, D. J. L., Voltaire, A., Bellon, G., & Tytéca, S. (2013). Transient climate response in a two-layer energy-balance model. Part I: Analytical solution and parameter calibration using CMIP5 AOGCM experiments. *Journal of Climate*, 26(6), 1841–1857. <https://doi.org/10.1175/jcli-d-12-00195.1>
- Goodwin, P., Leduc, M., Partanen, A.-I., Matthews, H. D., & Rogers, A. (2020). A computationally efficient method for probabilistic local warming projections constrained by history matching and pattern scaling, demonstrated by WASP-LGRTC-1.0. *Geoscientific Model Development*, 13(11), 5389–5399. <https://doi.org/10.5194/gmd-13-5389-2020>
- Hartikainen, J., & Särkkä, S. (2010). Kalman filtering and smoothing solutions to temporal Gaussian process regression models. In *2010 IEEE international workshop on machine learning for signal processing*.
- Hasselmann, K. (1976). Stochastic climate models Part I. Theory. *Tellus*, 28(6), 473. <https://doi.org/10.3402/tellusa.v28i6.11316>
- Hawkins, E., & Sutton, R. (2009). The potential to narrow uncertainty in regional climate predictions. *Bulletin of the American Meteorological Society*, 90(8), 1095–1108. <https://doi.org/10.1175/2009bams2607.1>
- Held, I. M., Winton, M., Takahashi, K., Delworth, T., Zeng, F., & Vallis, G. K. (2010). Probing the fast and slow components of global warming by returning abruptly to preindustrial forcing. *Journal of Climate*, 23(9), 2418–2427. <https://doi.org/10.1175/2009jcli3466.1>
- Hoesly, R. M., Smith, S. J., Feng, L., Klimont, Z., Janssens-Maenhout, G., Pitkanen, T., et al. (2018). Historical (1750–2014) anthropogenic emissions of reactive gases and aerosols from the community emissions data system (CEDS). *Geoscientific Model Development*, 11(1), 369–408. <https://doi.org/10.5194/gmd-11-369-2018>
- Holden, P., & Edwards, N. (2010). Dimensionally reduced emulation of an AOGCM for application to integrated assessment modelling. *Geophysical Research Letters*, 37(21), L21707. <https://doi.org/10.1029/2010gl045137>
- Hüllermeier, E., & Waegeman, W. (2021). Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110(3), 457–506. <https://doi.org/10.1007/s10994-021-05946-3>
- Hunter, J. D. (2007). Matplotlib: A 2D graphics environment [Software]. *Computing in Science & Engineering*, 9(03), 90–95. <https://doi.org/10.1109/mcse.2007.55>
- Huntingford, C., & Cox, P. (2000). An analogue model to derive additional climate change scenarios from existing GCM simulations. *Climate Dynamics*, 16(8), 575–586. <https://doi.org/10.1007/s003820000067>
- Jackson, L. S., Maycock, A. C., Andrews, T., Fredriksen, H.-B., Smith, C. J., & Forster, P. M. (2022). Errors in simple climate model emulations of past and future global temperature change. *Geophysical Research Letters*, 49(15), e2022GL098808. <https://doi.org/10.1029/2022gl098808>

- Law, H. C., Zhao, P., Chan, L. S., Huang, J., & Sejdinovic, D. (2019). Hyperparameter learning via distributional transfer. In *Advances in neural information processing systems* (Vol. 32).
- Leach, N. J., Jenkins, S., Nicholls, Z., Smith, C. J., Lynch, J., Cain, M., et al. (2021). Fairv2.0.0: A generalized impulse response model for climate uncertainty and future scenario exploration. *Geoscientific Model Development*, 14(5), 3007–3036. <https://doi.org/10.5194/gmd-14-3007-2021>
- Lee, J.-Y., Marotzke, J., Bala, G., Cao, L., Corti, S., Dunne, J. P., et al. (2021). Future global climate: Scenario-based projections and near-term information. In *Climate change 2021: The physical science basis. Contribution of Working Group I to the sixth assessment report of the Intergovernmental Panel on Climate Change* (pp. 553–672). Cambridge University Press.
- Levy, H., Horowitz, L. W., Schwarzkopf, M. D., Ming, Y., Golaz, J.-C., Naik, V., & Ramaswamy, V. (2013). The roles of aerosol direct and indirect effects in past and future climate change. *Journal of Geophysical Research: Atmospheres*, 118(10), 4521–4532. <https://doi.org/10.1002/jgrd.50192>
- Link, R., Snyder, A., Lynch, C., Hartin, C., Kravitz, B., & Bond-Lamberty, B. (2019). Fldgen v1. 0: An emulator with internal variability and space–time correlation for earth system models. *Geoscientific Model Development*, 12(4), 1477–1489. <https://doi.org/10.5194/gmd-12-1477-2019>
- Lovejoy, S., Procyk, R., Hébert, R., & Del Rio Amador, L. (2021). The fractional energy balance equation. *Quarterly Journal of the Royal Meteorological Society*, 147(736), 1964–1988. <https://doi.org/10.1002/qj.4005>
- Manabe, S., & Wetherald, R. T. (1967). Thermal equilibrium of the atmosphere with a given distribution of relative humidity.
- Martinez-Villalobos, C., & Neelin, J. D. (2019). Why do precipitation intensities tend to follow gamma distributions? *Journal of the Atmospheric Sciences*, 76(11), 3611–3631. <https://doi.org/10.1175/jas-d-18-0343.1>
- Matthews, A. G. D. G., Hensman, J., Turner, R., & Ghahramani, Z. (2016). On sparse variational methods and the Kullback-Leibler divergence between stochastic processes. In *Proceedings of the 19th international conference on artificial intelligence and statistics*.
- May, W. (2012). Assessing the strength of regional changes in near-surface climate associated with a global warming of 2 c. *Climatic Change*, 110(3–4), 619–644. <https://doi.org/10.1007/s10584-011-0076-y>
- McDonald, T., & Álvarez, M. (2021). Compositional modeling of nonlinear dynamical systems with ode-based random features. In *Advances in neural information processing systems* (Vol. 34, pp. 13809–13819).
- Meehl, G. A., Covey, C., Delworth, T., Latif, M., McAvaney, B., Mitchell, J. F., et al. (2007). The WCRP CMIP3 multimodel dataset: A new era in climate change research. *Bulletin of the American Meteorological Society*, 88(9), 1383–1394. <https://doi.org/10.1175/bams-88-9-1383>
- Meinshausen, M., Nicholls, Z. R., Lewis, J., Gidden, M. J., Vogel, E., Freund, M., et al. (2020). The shared socio-economic pathway (SSP) greenhouse gas concentrations and their extensions to 2500. *Geoscientific Model Development*, 13(8), 3571–3605. <https://doi.org/10.5194/gmd-13-3571-2020>
- Meinshausen, M., Raper, S. C. B., & Wigley, T. M. L. (2011). Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, MAGICC6 – Part 1: Model description and calibration. *Atmospheric Chemistry and Physics*, 11(4), 1417–1456. <https://doi.org/10.5194/acp-11-1417-2011>
- Met Office. (2010). Cartopy: A cartographic python library with a matplotlib interface [Software]. Retrieved from <https://github.com/SciTools/cartopy>
- Millar, R. J., Nicholls, Z. R., Friedlingstein, P., & Allen, M. R. (2017). A modified impulse–response representation of the global near-surface air temperature and atmospheric concentration response to carbon dioxide emissions. *Atmospheric Chemistry and Physics*, 17(11), 7213–7228. <https://doi.org/10.5194/acp-17-7213-2017>
- Mitchell, J. F., Johns, T., Eagles, M., Ingram, W., & Davis, R. (1999). Towards the construction of climate change scenarios. *Climatic Change*, 41(3/4), 547–581. <https://doi.org/10.1023/a:1005466909820>
- Moss, J. D., Opolka, F. L., Dumitrascu, B., & Lió, P. (2021). Approximate latent force model inference. arXiv preprint arXiv:2109.11851.
- Myhre, G., Shindell, D., Bréon, F.-M., Collins, W., Fuglestad, J., Huang, J., et al. (2014). Anthropogenic and natural radiative forcing. In *Climate change 2013 – The physical science basis: Working Group I contribution to the fifth assessment report of the Intergovernmental Panel on Climate Change*.
- O'Neill, B. C., Tebaldi, C., Van Vuuren, D. P., Eyring, V., Friedlingstein, P., Hurtt, G., et al. (2016). The scenario model intercomparison project (ScenarioMIP) for CMIP6. *Geoscientific Model Development*, 9(9), 3461–3482. <https://doi.org/10.5194/gmd-9-3461-2016>
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). Pytorch: An imperative style, high-performance deep learning library [Software]. In *Advances in neural information processing systems* (Vol. 32). Retrieved from <https://github.com/pytorch/pytorch>
- Pörtner, H.-O., Roberts, D. C., Adams, H., Adler, C., Aldunce, P., Ali, E., et al. (2022). *Climate change 2022: Impacts, adaptation and vulnerability*. IPCC Geneva.
- Rahimi, A., & Recht, B. (2007). Random features for large-scale kernel machines. In *Advances in neural information processing systems* (Vol. 20).
- Rasmussen, C., & Williams, C. (2005). *Gaussian process for machine learning*. MIT Press.
- Riahi, K., Van Vuuren, D. P., Kriegler, E., Edmonds, J., O'Neill, B. C., Fujimori, S., et al. (2017). The shared socioeconomic pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Global Environmental Change*, 42, 153–168. <https://doi.org/10.1016/j.gloenvcha.2016.05.009>
- Rosenzweig, C., Elliott, J., Deryng, D., Ruane, A. C., Müller, C., Arneth, A., et al. (2014). Assessing agricultural risks of climate change in the 21st century in a global gridded crop model intercomparison. *Proceedings of the National Academy of Sciences*, 111(9), 3268–3273. <https://doi.org/10.1073/pnas.1222463110>
- Ross, M., Smith, M. T., & Álvarez, M. (2021). Learning nonparametric Volterra kernels with Gaussian processes. In *Advances in neural information processing systems* (Vol. 34, pp. 24099–24110).
- Rypdal, K. (2012). Global temperature response to radiative forcing: Solar cycle versus volcanic eruptions. *Journal of Geophysical Research*, 117, D06115. <https://doi.org/10.1029/2011jd017283>
- Santer, B. D., Wigley, T. M., Schlesinger, M. E., & Mitchell, J. F. (1990). Developing climate scenarios from equilibrium GCM results.
- Sarkka, S., & Hartikainen, J. (2012). Infinite-dimensional Kalman filtering approach to spatio-temporal Gaussian process regression. In *Artificial intelligence and statistics*.
- Sarkka, S., Solin, A., & Hartikainen, J. (2013). Spatiotemporal learning via infinite-dimensional Bayesian filtering and smoothing: A look at Gaussian process regression through Kalman filtering. *IEEE Signal Processing Magazine*, 30(4), 51–61. <https://doi.org/10.1109/msp.2013.2246292>
- Schöngart, S., Lejeune, Q., Gudmundsson, L., Nath, S., Seneviratne, S., & Schleußner, C.-F. (2023). Extending MESMER-M to jointly emulate earth system model temperature and precipitation realizations (Tech. Rep.). Copernicus Meetings.
- Schöngart, S., Lejeune, Q., Schleußner, C.-F., Seneviratne, S., Gudmundsson, L., & Beusch, L. (2022). A spatially explicit approach for joint temperature-precipitation emulation of earth system model simulation. In *EGU general assembly conference abstracts* (pp. EGU22–11049).

- Seland, Ø., Bentsen, M., Olivie, D., Toniazzo, T., Gjermundsen, A., Graff, L. S., et al. (2020). Overview of the Norwegian earth system model (NorESM2) and key climate response of CMIP6 deck, historical, and scenario simulations. *Geoscientific Model Development*, 13(12), 6165–6200. <https://doi.org/10.5194/gmd-13-6165-2020>
- Sellers, W. D. (1969). A global climatic model based on the energy balance of the earth-atmosphere system. *Journal of Applied Meteorology and Climatology*, 8(3), 392–400. [https://doi.org/10.1175/1520-0450\(1969\)008<0392:agcmbo>2.0.co;2](https://doi.org/10.1175/1520-0450(1969)008<0392:agcmbo>2.0.co;2)
- Seneviratne, S. I., Nicholls, N., Easterling, D., Goodess, C., Kanae, S., Kossin, J., et al. (2012). Changes in climate extremes and their impacts on the natural physical environment. In *Managing the risks of extreme events and disasters to advance climate change adaptation*. A Special Report of Working Groups I and II of the Intergovernmental Panel on Climate Change (IPCC) (pp. 109–230).
- Seneviratne, S. I., Donat, M. G., Pitman, A. J., Knutti, R., & Wilby, R. L. (2016). Allowable CO₂ emissions based on regional and impact-related climate targets. *Nature*, 529(7587), 477–483. <https://doi.org/10.1038/nature16542>
- Shapley, L. S. (1953). A value for n-person games. In *Contributions to the theory of games*.
- Shukla, P. R., Skea, J., Slade, R., Al Khourdajie, A., Van Diemen, R., McCollum, D., et al. (2022). Climate change 2022: Mitigation of climate change. In *Contribution of Working Group III to the sixth assessment report of the Intergovernmental Panel on Climate Change* (Vol. 10). 9781009157926.
- Smith, C. J., Forster, P. M., Allen, M., Leach, N., Millar, R. J., Passerello, G. A., & Regayre, L. A. (2018). Fair v1.3: A simple emissions-based impulse response and carbon cycle model. *Geoscientific Model Development*, 11(6), 2273–2297. <https://doi.org/10.5194/gmd-11-2273-2018>
- Snyder, A., Link, R., Dorheim, K., Kravitz, B., Bond-Lamberty, B., & Hartin, C. (2019). Joint emulation of earth system model temperature-precipitation realizations with internal variability and space-time and cross-variable correlation: Fldgen v2. 0 software description. *PLoS One*, 14(10), e0223542. <https://doi.org/10.1371/journal.pone.0223542>
- Stein, M. L. (1999). *Interpolation of spatial data: Some theory for kriging*. Springer Science & Business Media.
- Štrumbelj, E., & Kononenko, I. (2014). Explaining prediction models and individual predictions with feature contributions. *Knowledge and Information Systems*, 41(3), 647–665. <https://doi.org/10.1007/s10115-013-0679-x>
- Tandon, A. (2021). Analysis: The lack of diversity in climate-science research. *Carbon Brief*.
- Taylor, K. E., Stouffer, R. J., & Meehl, G. A. (2012). An overview of CMIP5 and the experiment design. *Bulletin of the American Meteorological Society*, 93(4), 485–498. <https://doi.org/10.1175/bams-d-11-00094.1>
- Tebaldi, C., & Arblaster, J. M. (2014). Pattern scaling: Its strengths and limitations, and an update on the latest model simulations. *Climatic Change*, 122(3), 459–471. <https://doi.org/10.1007/s10584-013-1032-9>
- Tebaldi, C., Armbruster, A., Engler, H., & Link, R. (2020). Emulating climate extreme indices. *Environmental Research Letters*, 15(7), 074006. <https://doi.org/10.1088/1748-9326/ab8332>
- Tebaldi, C., & Knutti, R. (2018). Evaluating the accuracy of climate change pattern emulation for low warming targets. *Environmental Research Letters*, 13(5), 055006. <https://doi.org/10.1088/1748-9326/aabef2>
- Titsias, M. K. (2009). Variational learning of inducing variables in sparse Gaussian processes. In *AISTATS*.
- Tsutsui, J. (2020). Diagnosing transient response to CO₂ forcing in coupled atmosphere-ocean model experiments using a climate model emulator. *Geophysical Research Letters*, 47(7), e2019GL085844. <https://doi.org/10.1029/2019gl085844>
- Turner, S. W., Hejazi, M., Kim, S. H., Clarke, L., & Edmonds, J. (2017). Climate impacts on hydropower and consequences for global electricity supply investment needs. *Energy*, 141, 2081–2090. <https://doi.org/10.1016/j.energy.2017.11.089>
- Wang, K., Pleiss, G., Gardner, J., Tyree, S., Weinberger, K. Q., & Wilson, A. G. (2019). Exact Gaussian processes on a million data points. In *Advances in neural information processing systems* (Vol. 32).
- Ward, W., Ryder, T., Prangle, D., & Alvarez, M. (2020). Black-box inference for non-linear latent force models. In *International conference on artificial intelligence and statistics* (pp. 3088–3098).
- Watson-Parris, D., Rao, Y., Olivie, D., Seland, Ø., Nowack, P., Camps-Valls, G., et al. (2022). ClimateBench v1. 0: A benchmark for data-driven climate projections [Data]. *Journal of Advances in Modeling Earth Systems*, 14(10), e2021MS002954. <https://doi.org/10.1029/2021ms002954>
- Williams, A. I., Stier, P., Dagan, G., & Watson-Parris, D. (2022). Strong control of effective radiative forcing by the spatial pattern of absorbing aerosol. *Nature Climate Change*, 12(8), 735–742. <https://doi.org/10.1038/s41558-022-01415-4>
- Wilson, A. G., & Adams, R. (2013). Gaussian process kernels for pattern discovery and extrapolation. In *Proceedings of the 30th international conference on machine learning*.
- Wilson, A. G., Hu, Z., Salakhutdinov, R., & Xing, E. P. (2016). Deep kernel learning. In *Artificial intelligence and statistics* (pp. 370–378).