Bayesian model calibration with a non-additive phase discrepancy

Paul Castéras[†,1,2,3], Josselin Garnier[‡,2], Julien Bect[§,3], Gwenael Salin[§,1]

[†] Étudiant en thèse, [‡] Directeur de thèse, [§] Co-encadrant
Durée de la thèse : 01/12/2024 - 30/11/2027

[1] CEA / DAM / DIF, France
[2] CMAP, École polytechnique, Institut Polytechnique de Paris, France
[3] Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des signaux et systèmes, France

Calibrating a simulation code consists in determining the parameters that allow the code to reproduce experimental outputs as accurately as possible. In the context of uncertainty quantification, a popular method for solving this problem is Bayesian calibration [3]. The idea is to link experiments and simulations by taking into account the various sources of error (measurement, model, and surrogate errors) in order to then determine the distribution of the calibration parameters β conditionally on both simulations and experimental data collected under different experimental settings x.

We focus here on calibration problems where the outputs are time series. In this setting, experimental and simulation outputs are usually related through additive errors [2]. Recent studies have highlighted the importance of accounting for temporal misalignment between signals as well [1, 4]. To address this issue, we introduce a novel Bayesian calibration framework that incorporates, in addition to an additive model error δ, a phase discrepancy γ, allowing the simulation output to be flexibly aligned with experimental observations:

$$y_{\text{exp}}(t) = y_{\text{sim}}(x, \beta, \gamma(x,t)) + \delta(x,t) + \varepsilon.$$

This formulation enables a richer representation of discrepancies between simulations and experiments and provides a natural mechanism for handling temporal shifts. Introducing this phase discrepancy raises several challenges. First, an appropriate prior must be assigned to γ. We compare different priors proposed in the literature for temporal misalignment. Second, the error terms can no longer be combined into a single multivariate Gaussian distribution, which leads to a more challenging sampling problem. We show that, under appropriate approximations, our framework can be reformulated as a tractable variant of the method proposed in [1]. Finally, we demonstrate the practical relevance of the approach through the calibration of the Johnson damage model, which describes material degradation induced by shock-wave propagation.

References

[1] D. Francom, J. D. Tucker, G. Huerta, K. Shuler, and D. Ries. Elastic Bayesian model calibration. *SIAM/ASA J. Uncertain. Quantif.*, 13(1):195–227, 2025.

[2] D. Higdon, J. Gattiker, B. Williams, and M. Rightley. Computer model calibration using high-dimensional output. *J. Am. Stat. Assoc.*, 103(482):570–583, 2008.

[3] M. C. Kennedy and A. O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.

[4] Y. Lu, R. Herbei, and S. Kurtek. Bayesian registration of functions with a Gaussian process prior. *Journal of Computational and Graphical Statistics*, 26(4):894–904, 2017.