

Bayesian optimization as an efficient tool for QSP model calibration: a proof-of-concept

Paulo Paneque Galuzio^a, Weirong Wang^b, Shalla Hanson^b

^a Johnson & Johnson Innovative Medicine, Beerse, Belgium

^b Johnson & Johnson Innovative Medicine, Spring House, PA, USA;

Key Takeaway: Bayesian optimization can be efficiently used to perform QSP model calibration

Conclusions: Bayesian optimization was successfully used for QSP virtual population calibration with considerably small running times, and it is able to handle non-deterministic objective functions, and to actively constrain regions of the parameter space that lead to unstable solutions, avoiding the selection of non-convergent parameters in the final Virtual Population.

Abstract

Introduction: Quantitative Systems Pharmacology (QSP) is being increasingly adopted as a strategy for Model Informed Drug Development (MIDD). QSP uses high-dimensional, multi-scale models that leverage data from multiple sources, and is not always amenable to traditional Pharmacometrics-style model fitting strategies, due to the scale and complexity of the models. QSP models also often involve generation of virtual patients (VPs) and virtual populations (Vpops), and different methods for their calibration exist across the industry, many of which are computationally intensive and time consuming. In this work we propose a workflow for the use of Bayesian optimization for QSP model calibration. The time efficiency of Bayesian optimization makes it a suitable method for accelerated model calibration, which is particularly useful in the early stages of model development.

Objectives: Develop and evaluate the feasibility and limitations of Bayesian optimization for QSP model calibration.

Methods: A proof of concept of the Bayesian optimization method was established for a two-dimensional search space for a large, multi-scale QSP model in immuno-oncology. The model includes cell and antibody PBPK at the organ level, killing and activation at the cellular level, and receptor binding and internalization at the molecular level. The model was calibrated against clinical response data from a Phase 1 dose escalation study. A log-likelihood function of the response rates for the different cohorts was used as the optimization objective function. Distributed parameters of the generated virtual population were generated from pre-specified (known) log-normal distributions. For each objective function evaluation, cohorts of N randomly sampled virtual patients were used to estimate the model predicted probability of response per cohort.

Steps for Model Calibration

- QSP model calibration requires many steps of successive calibration and fitting of parameters. Frequently leveraging information from in-vitro data or clinical biomarkers, which are fit to mini-models, and the parameters obtained are used in the large-scale QSP model.
- Clinical response data allows the final calibration of any remaining model parameters, such as in vivo drug potency, using Bayesian optimization. In this step, model predictions of the generated virtual population are compared to population level response rates from the Clinical data.

Results

- Results were generated in a workstation with an Intel(R) Xeon(R) Platinum 8160T CPU @ 2.10GHz, 2095 Mhz, 24 Core(s), 48 Logical Processors, with 192GB of installed memory, running Matlab 2022b with a 48 core parallel pool.
- For a two-dimensional space, and 22 cohorts simulated, using 100 VPs per cohort (leading to a total of 2200 model simulations per objective function evaluation) the Bayesian optimization achieves satisfactory convergence, with 5x48 iterations of the method, with run times in the order of approximately 1h. By reducing the number of iterations and the number of VPs per simulated cohort, runtimes can be reduced to as low as approximately 15min. In the interest of comparison, it is not uncommon for runtimes to reach or exceed 10 hours when calibrating large QSP models.
- The ability of Bayesian optimization to handle non-deterministic objectives ensures good convergence properties even with reduced virtual populations. Additionally, the Bayesian optimizer can also constrain regions of the parameter space that fail to integrate successfully.

The number of evaluations of the objective function necessary for convergence is dependent on the dimensionality of the parameter space of the calibration problem. Initial assessment on the simulations performed for QSP model calibration seem to suggest a guideline of approximately 30 iterations per dimension of the parameter space to be sufficient.

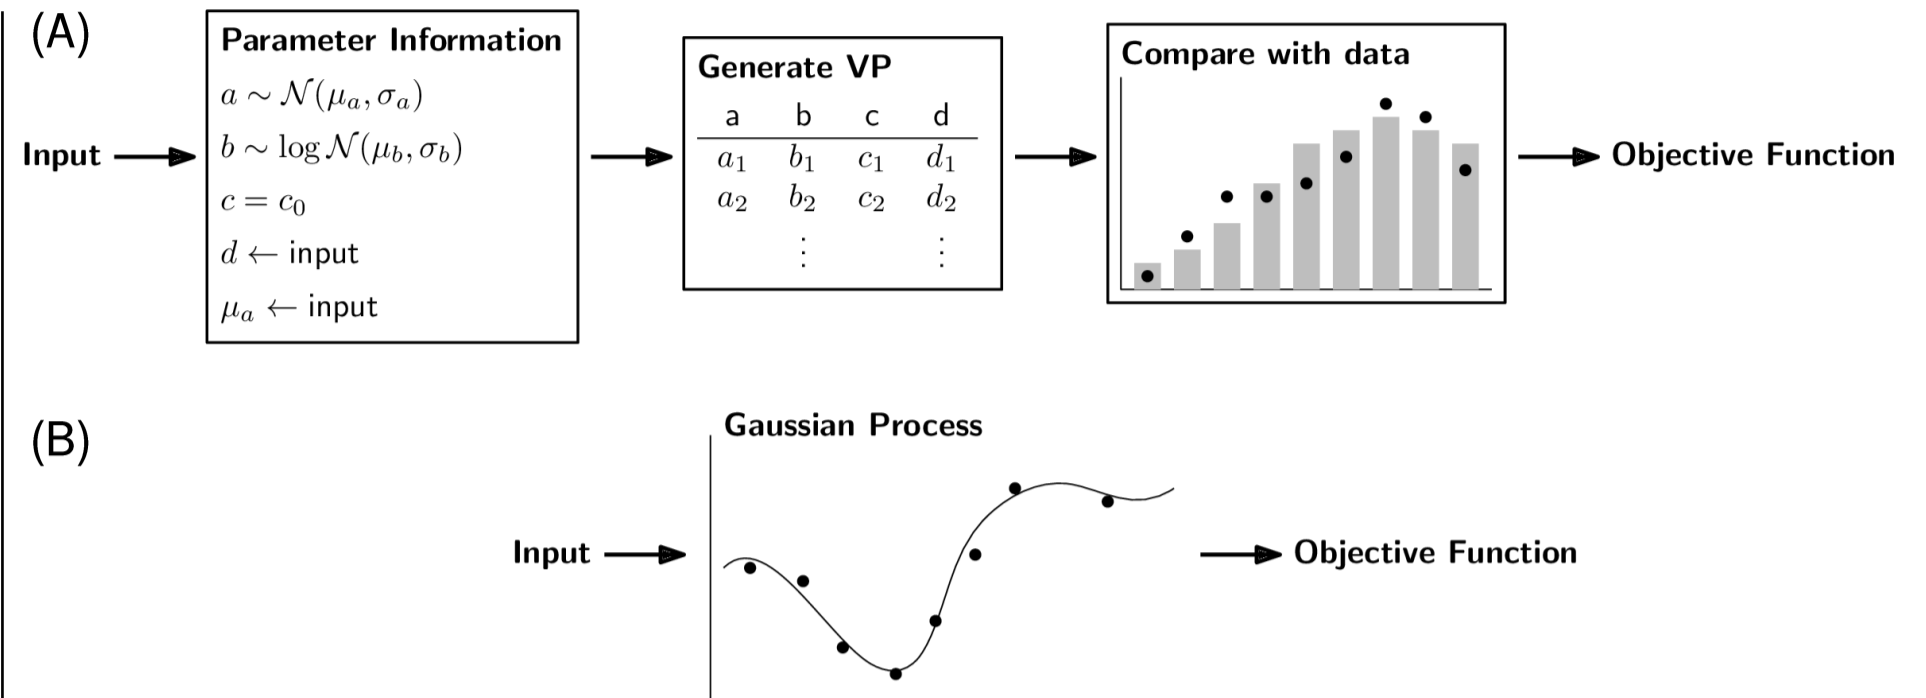


Figure 1: Simplified schematic representation of the method. Figure A represents the calculation of the objective function. Input parameters are used to sample from pre-specified parameter distributions, from which a plausible virtual population can be sampled from. Model predictions from this Vpop are compared to the Clinical and an appropriate error function is calculated. Bayesian optimization replaces this process with a non-linear regression of the objective function that is calculated using Gaussian processes (Figure B), substantially reducing the number of times it is necessary to simulate a plausible Vpop in the QSP model.

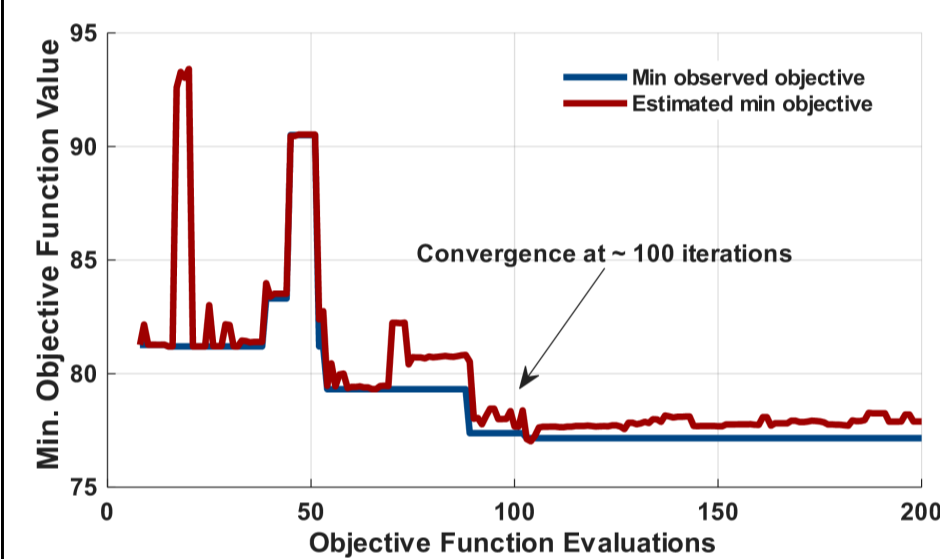


Figure 2: Convergence plot for the method on a 5-Dimensional parameter space, showing the observed and estimated (via Gaussian process) values for the objective function. Convergence is achieved after only approximately 100 iterations of the method. Each iteration of the method corresponds to an evaluation of the objective function. The value 100 is considerably smaller than what is necessary for other optimization techniques to reach convergence, such as particle swarm or genetic algorithms, which may require number of evaluations of the objective function in the order of $n \geq O(10^4)$.

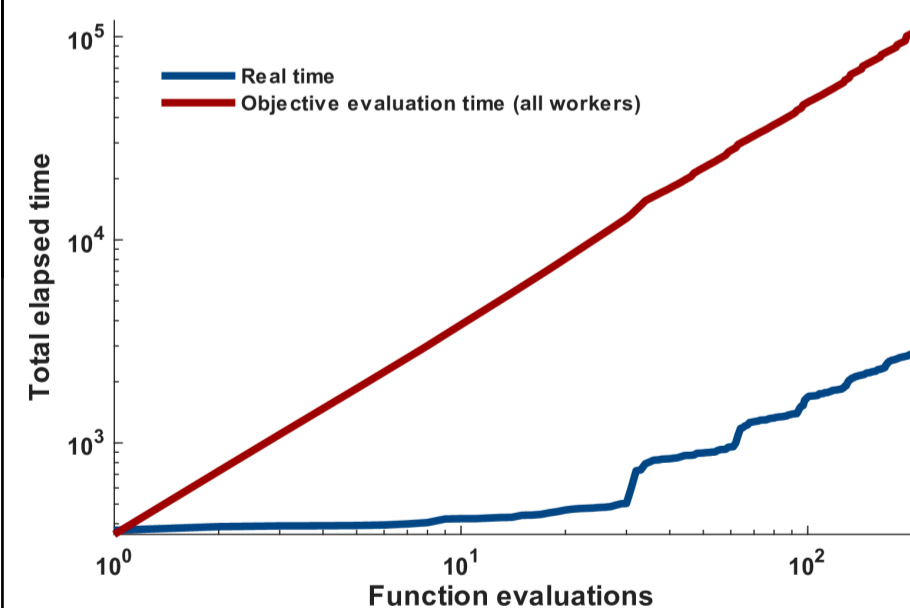


Figure 3: Total elapsed time as a function of the objective function evaluations. The blue line represents the real time, and the red line represents the total CPU evaluation time (disregarding parallelization). This results shows the linear dependency of the running time with the number of iterations in the method and also the method potential for parallelization. This results suggests good scalability of Bayesian optimization since the number of iterations scales linearly with search space dimensionality, and running times are linear in the number of iterations.

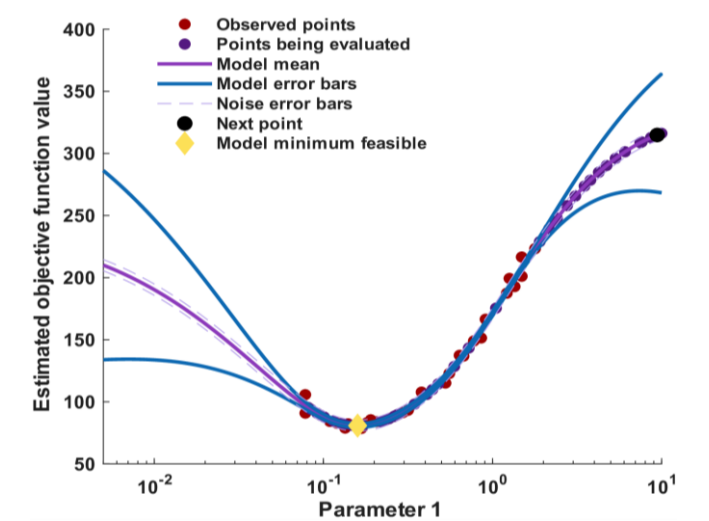


Figure 4: Gaussian Process (GP) regression on the QSP model objective function for a 1-dimensional parameter space. The Bayesian Optimization method uses GP as a surrogate of the objective function, to reduce the number needed of QSP model simulations.




Figure 5: Simulated probability of biomarker response for different doses explored in the trial. The bar plot corresponds to model simulations of 1000 virtual trials with 20 virtual patients each, and error bars correspond to 95% CIs. The points correspond to observed response rates per cohort in the clinical trial. Model predictions agree with the observed data within the confidence interval. This plot exemplifies the use of this calibration framework based on Bayesian optimization. The overall agreement of model predictions to the clinical data suggest that Bayesian Optimization can be used for efficient QSP model calibration.

Conclusion

Bayesian optimization was successfully used for QSP virtual population calibration with considerably small running times, which can help accelerate model development, especially in early stages of model development when multiple trials of model structure and comparison with data are necessary. Bayesian optimization is able to handle non-deterministic objectives, and to actively constrain regions of the parameter space that lead to unstable solutions, avoiding the selection of non-convergent parameters in the final Vpop. These characteristics make Bayesian optimization suitable for QSP calibration. Its efficient execution also makes it a good candidate to be integrated into other more complex calibration procedures (replacing burn-in phases in Monte Carlo simulations, for example).

References

- Mockus, J. "On Bayesian methods of optimization." Towards global optimization (1975): 166-181.
- Jones, Donald R., Matthias Schonlau, and William J. Welch. "Efficient global optimization of expensive black-box functions." Journal of Global optimization 13 (1998): 455-492.

Acknowledgments

Daniel Amchin, Jin Niu, Tao Peng, Halley Oyer, and other team members from QSCP in Johnson & Johnson Innovative Medicine

Disclosures

All authors are currently employed by Johnson & Johnson Innovative Medicine R&D, LLC, and the research was funded by Johnson & Johnson Innovative Medicine Research & Development, LLC.