Active Learning With Bayesian Optimization For Efficient Virtual Population Generation Using Quantitative Systems Pharmacology Models

Authors: Dennis Reddyhoff[1], Andrew Matteson[2], Andrzej Kierzek[1]
Correspondence: dennis.reddyhoff@certara.com

[1]Certara, Sheffield, UK
[2]Certara, Concord, US

Background and Objective

- Virtual population generation requires identifying feasible parameter combinations that produce physiologically realistic model outputs.
- Random sampling scales poorly with increased model complexity (number of ODEs, parameters, nonlinearity).
- Traditional batch simulation with machine learning classifiers [1] requires extensive training data and may inefficiently sample infeasible regions.
- We present a proof-of-concept active learning (AL) pipeline using Gaussian Process (GP) surrogates to efficiently learn feasible regions with reduced computational burden.

Methods

- We define a virtual patient as a set of input parameters, satisfying some biologically informed criteria, which can be simulated using an ODE model to predict clinical outputs. If clinical data is available, a feasible patient is one whose model outputs fall within clinically expected ranges. We aim to efficiently generate feasible patients for a QSP model of the Cancer Immunity Cycle (CIC) [2].

- We define physiologically-based priors for 5 inputs and 11 output constraints based on clinical data. A pool of 25,000 quasi-randomly sampled patients (14.8% feasible) was pre-simulated and split into a training pool (N=20000) and validation set (N=5000).

We compared three GP-based active learning approaches:

1. Regression GP – Signed Margin (GPR-M): learns a soft continuous label encoding the signed distance to the nearest constraint bound, normalised and transformed via sigmoid to [0,1]. The GP maximises predicted margin to target boundary-adjacent feasible points.
2. Binary Classification GP (GPC): learns P(feasible|x) from binary labels only, serving as a classifier baseline.
3. Regression GP - Raw Violation (GPR-V): learns total violation magnitude [3], serving as a regression baseline.

All methods use Matérn-5/2 kernels with automatic relevance determination (ARD) and Upper Confidence Bound acquisition (UCB = μ + βσ, where β =3.0 for first 200 iterations for exploration, β=1.0 thereafter for exploitation). Models were initialized with 100 random samples and samples were queried sequentially from the training pool for 1000 iterations. For GPR-M, margins are normalized by maximum observed violation per output and transformed via sigmoid with temperature τ selected to map the maximum feasible margin to p=0.95, ensuring balanced gradient on both sides of the boundary.

Results

- GPR-M found 1,003 feasible patients in 1,100 simulations (100 initial + 1000 AL), achieving an **83.6% efficiency gain over expected Sobol sampling** (expected ~6,710 queries), with a final hit rate of 91.2%.
- The learned surrogate achieved **AUROC=0.9782** on the validation set, with sensitivity=0.844, specificity=0.981, precision=0.882, and F1=0.863.
- The binary classification baseline (GPC) achieved a marginally higher AUROC (0.982 vs 0.978) and comparable discriminative statistics but generated fewer feasible patients overall (862/1,100; 81.0% efficiency gain; final hit rate 78.4%).
- The raw violation regression baseline (GPR-V) performed substantially worse on both feasibility yield (356/1,100) and surrogate quality (AUROC=0.784), confirming that label encoding critically determines regression surrogate performance.
- ARD kernel analysis revealed a single input parameter (mu) as the dominant feasibility determinant (88.6% relative importance), providing interpretable biological insight into constraint structure.

Conclusions

- GP-based active learning with a soft signed-margin encoding provides a principled and computationally efficient approach to virtual population generation, outperforming binary classification in feasible patient yield while achieving comparable surrogate quality.
- The signed-margin formulation preserves continuous information about constraint proximity that binary labels discard, enabling more targeted sampling near the feasibility boundary.
- The learned surrogate enables rapid feasibility screening of new parameter combinations without additional ODE simulations.

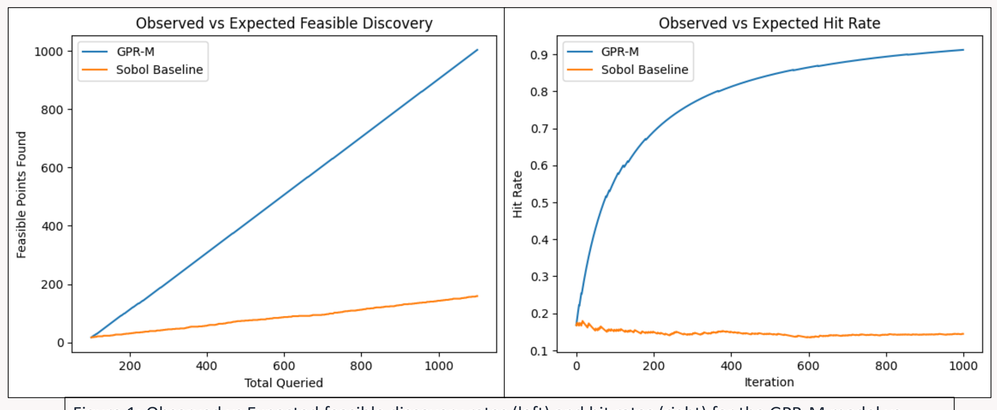

Figure 1: Observed vs Expected feasible discovery rates (left) and hit rates (right) for the GPR-M model vs Sobol sampling baseline. The GPR-M approach discovered 1003 feasible patients, while Sobol sampling discovered 159 feasible patients in 1100 queries

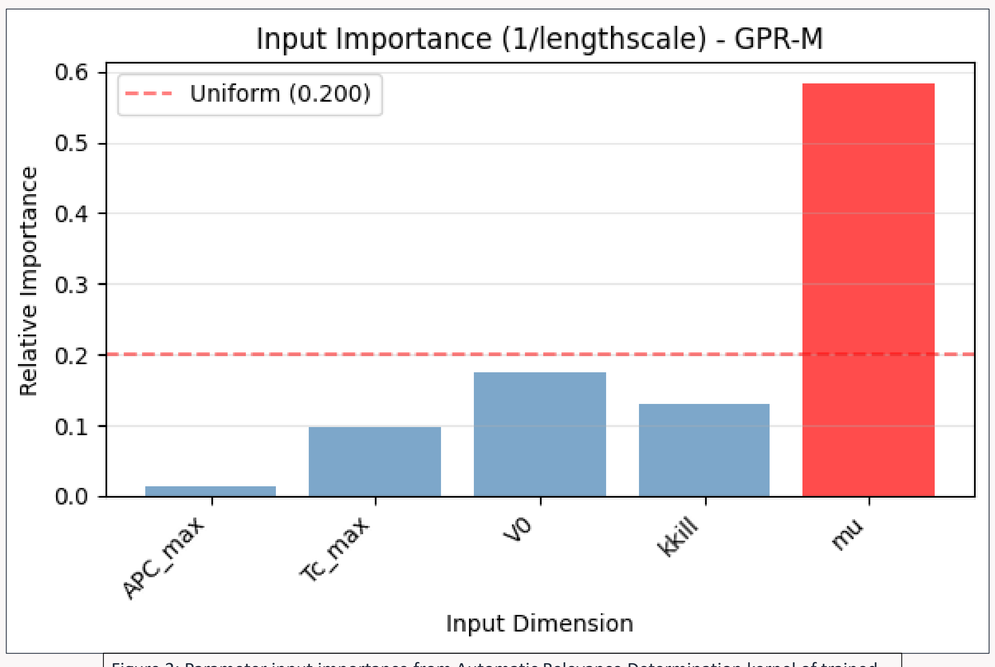

Figure 2: Parameter input importance from Automatic Relevance Determination kernel of trained GPR-M model. Higher values indicate stronger influence on distance to constraint violation.

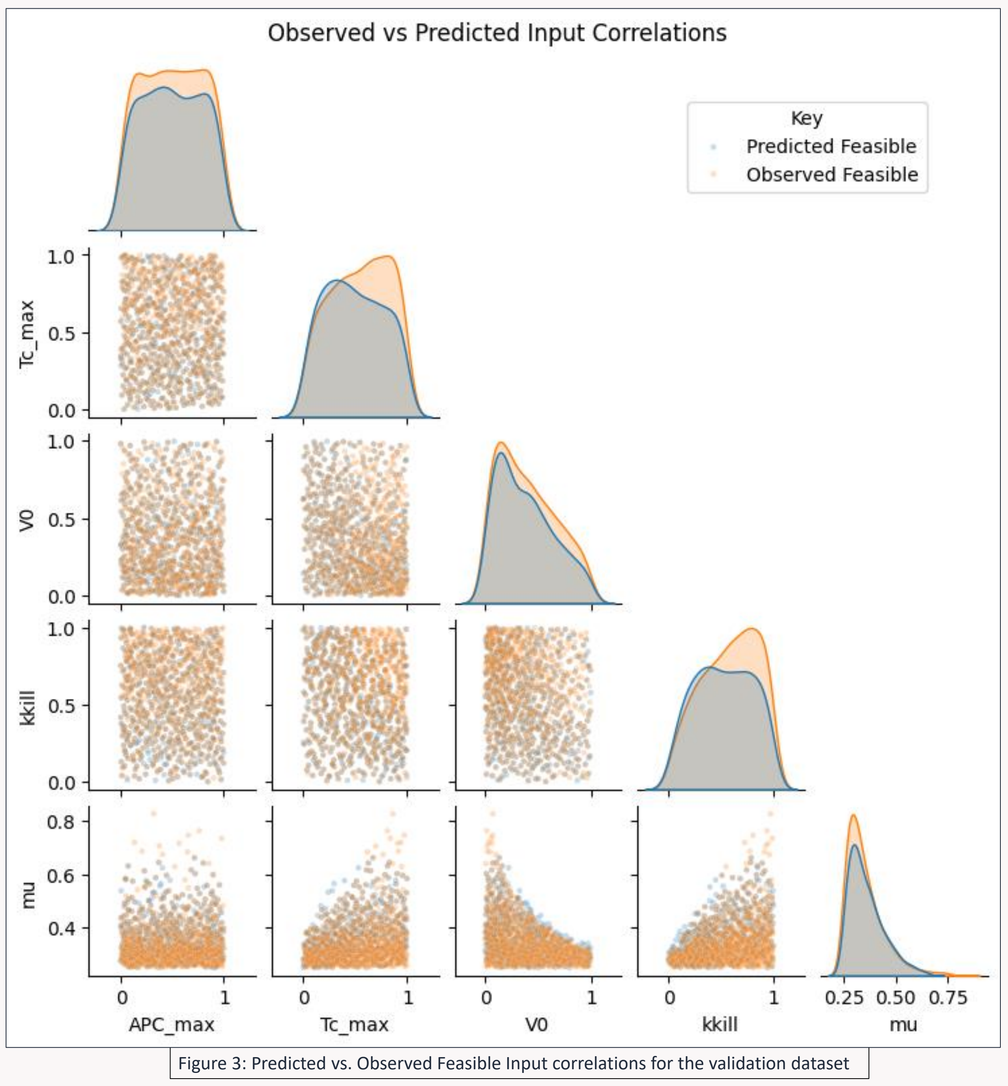

Figure 3: Predicted vs. Observed Feasible Input correlations for the validation dataset

References

[1] Reddyhoff D, Matteson A and Kierzek A. "Efficient Generation of Plausible Virtual Populations Using Machine Learning Surrogate Classification Models." (2026). doi: 10.70534/IKAY9510
[2] Lazarou, G et al. “Integration of Omics Data Sources to Inform Mechanistic Modeling of Immune-Oncology Therapies: A Tutorial for Clinical Pharmacologists.” Clinical pharmacology and therapeutics vol. 107,4 (2020): 858-870. doi:10.1002/cpt.1786
[3] Allen, R J et al. “Efficient Generation and Selection of Virtual Populations in Quantitative Systems Pharmacology Models.” CPT: Pharmacometrics & Systems Pharmacology vol. 5,3 (2016): 140-6. doi:10.1002/psp4.12063