

Quantifying Uncertainty in Health Decision Making using Efficient High-Dimensional Gaussian Process Regression

Anna Heath, Ioanna Manolopoulou & Gianluca Baio



Motivation

Health care decisions are made under a strict decision theoretic framework. Parameters in these models are concerned with the prevalence of a disease or the length of time spent in hospital. These parameters are used to quantify the costs and effects of different treatment options and are then combined into a utility function that gives the *net benefit*, NB_t , of the different treatments. The adopted treatment has the highest expected net benefit.

Due to the cost of implementing health care interventions, health decision makers are concerned with the impact of parameter uncertainty on the decision. A decision theoretic measure of this uncertainty is known as the Expected Value of Partial Perfect Information (EVPPI). This measure quantifies the expected value of learning a subset of the parameter space ϕ *perfectly*, i.e. resolving all the uncertainty about this parameter set.

The main barrier to the widespread implementation of this method is computational as it can take days to calculate an estimate for the EVPPI. However, recent developments [4] have demonstrated that the problem can be transformed into a non-parametric regression problem. Specifically, in some cases, a Gaussian Process regression problem. This speeds up the computation time but when ϕ is high-dimensional the computation time is still infeasible for use in health economic evaluations.

Gaussian Process Regression

Gaussian Process regression models the outputs NB_t as a multivariate normal dependent on some inputs ϕ . The mean vector and covariance matrix of this normal are calculated dependent on ϕ . Most importantly the covariance matrix is dependent on the distance between the inputs ϕ and some hyperparameters ζ . Estimating these hyperparameters using standard methods from the data is computationally costly – so we developed a fast approximation method for this matrix using ideas from Spatial Statistics.

For a two-dimensional vector (ϕ_1, ϕ_2) , modelling as a Gaussian Process is equivalent to estimating β , ζ and σ , where

$$\begin{pmatrix} NB_t(\theta_1) \\ NB_t(\theta_2) \\ \vdots \\ NB_t(\theta_S) \end{pmatrix} \sim \text{Normal} \left(\begin{pmatrix} 1 & \phi_1^1 & \phi_2^1 \\ 1 & \phi_1^2 & \phi_2^2 \\ \vdots & \vdots & \vdots \\ 1 & \phi_1^S & \phi_2^S \end{pmatrix} \beta, \mathcal{C}(\zeta) + \sigma^2 \mathbf{I} \right)$$

and \mathcal{C} is a function of ζ that calculates the covariance matrix.

Covariance Matrix

It can be shown that a Gaussian Process with a *Matérn* covariance function is a solution to a Stochastic Partial Differential Equation:

$$(\kappa^2 - \Delta)^{\frac{\alpha}{2}} \tau f(\phi) = \mathcal{W}(\phi)$$

where Δ is the Laplacian and $\mathcal{W}(\phi)$ is Gaussian white noise.

Therefore, we can transform the estimation of the covariance matrix into the estimation of an SPDE for which there is a wealth of literature [2]. Specifically, the *finite element method* allows us to transform the estimation of the hyperparameters into the estimation of a set of parameters ω which can be estimated using a sparse precision matrix. This allows for fast computation, specifically using the INLA [3] algorithm for fast approximate Bayesian inference.

To use INLA we consider a specific representation of Gaussian Process regression:

$$NB_t | \omega, \beta, \sigma \sim N(\mathbf{H}\beta + \omega, \sigma^2 \mathbf{I})$$

and to extend this method to higher dimensions, we estimate ω based on a bi-dimensional projection from the full parameter space ϕ [1].

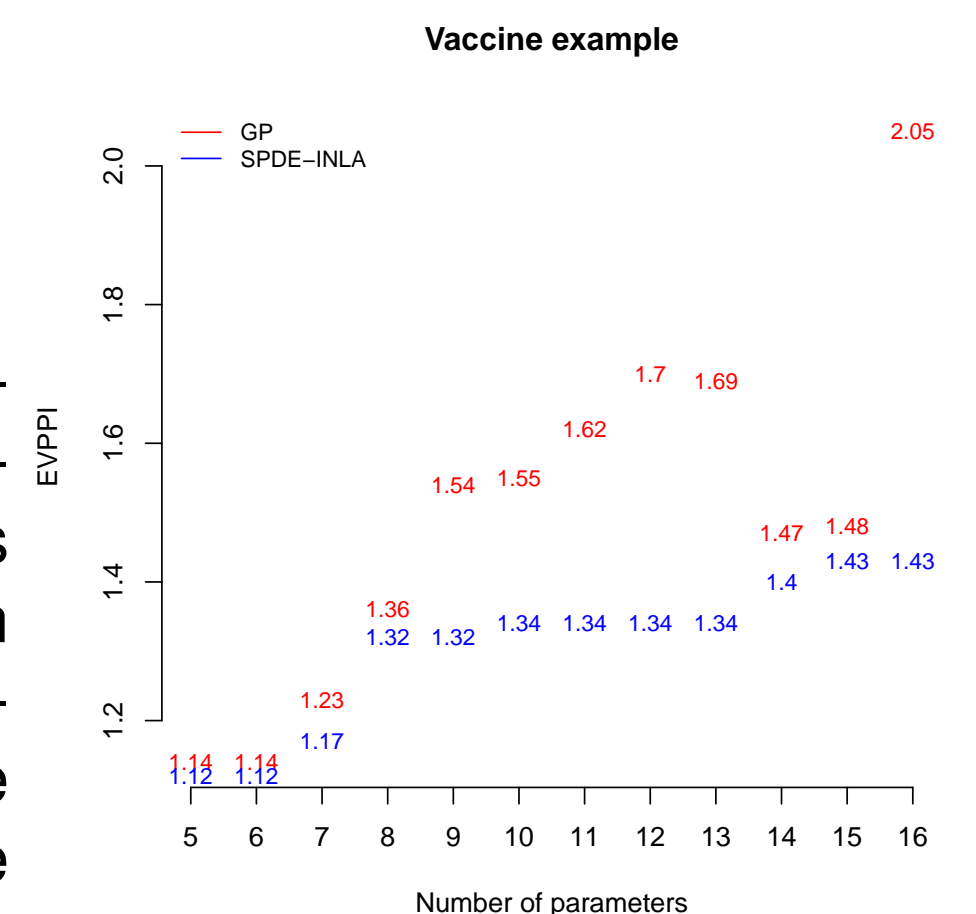
Results

This graphic shows the estimated EVPPI value for the standard Gaussian Process (GP) method and our SPDE-INLA method with projections.

The estimates should be:

- Positive
- Monotonic
- “Similar”

The difference between the two estimates is quite large at times. However, the SPDE-INLA method seems to estimate the EVPPI better in this example as monotonicity is respected. This gives evidence that the use of projections is valid and can be used to decrease computation time.



Computational Time

The computation time for our method is significantly shorter than the standard method, even in relatively low dimensions. Due to the use of projections the computation time required for our method is similar as the dimensionality increases, whereas, the standard method increases significantly with the dimensionality of ϕ .

This is key as the method can be used in complicated scenarios without significantly increasing the computational time.

Dimension of ϕ	Computation Time (s)	
	GP	SPDE INLA
5	24	9
6	46	9
7	222	9
8	128	9
9	252	8
10	198	11
11	776	8
12	264	11
13	660	13
14	695	12
15	910	11
16	559	13

Extensions

The flexibility of the INLA algorithm allows us to specify a more complicated linear predictor structure and therefore, without significantly increasing the computation time, model more complex relationships between ϕ and the net benefits. Second or third order interactions can be simply modelled along with other non-linear functions of the parameters.

Tests have shown that modelling these interactions increases the accuracy of the EVPPI and thus this flexibility is a very important extension to the SPDE INLA method for calculating the EVPPI.

References

- [1] A. Heath, I. Manolopoulou, and G. Baio. Efficient High-Dimensional Gaussian Process Regression to calculate the Expected Value of Partial Perfect Information in Health Economic Evaluations. Submitted.
- [2] F. Lindgren, H. Rue, and J. Lindström. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498, 2011.
- [3] H. Rue, S. Martino, and N. Chopin. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the royal statistical society: Series B (statistical methodology)*, 71(2):319–392, 2009.
- [4] M. Strong, J. Oakley, and A. Brennan. Estimating Multiparameter Partial Expected Value of Perfect Information from a Probabilistic Sensitivity Analysis Sample A Nonparametric Regression Approach. *Medical Decision Making*, 34(3):311–326, 2014.