# Supplementary Material for Predicting Preference Reversals via Gaussian Process Uncertainty Aversion

**Rikiya Takahashi**
IBM Research - Tokyo
Rikiya.Takahashi@gmail.com

**Tetsuro Morimura**
IBM Research - Tokyo
tetsuro@jp.ibm.com

## Abstract

This supplementary material provides an approach of optimizing the covariance function in the proposed choice model, as well as the objective evaluation function. While the isotropic RBF covariance function was sufficient for our datasets, optimization of the covariance function in the proposed discrete choice model is algorithmically possible. Based on an approximate Maximum A Posteriori estimation, a convex optimization algorithm for the covariance function is newly derived. The nice convexity stems from a quadratic form of the covariance function and an approximation of the log-determinant of a matrix by its trace.

## Quadratic Form of the Covariance Function and Optimizing Its Metric

Instead of the parametric covariance function $\frac{1}{\eta}\exp(-\frac{\xi}{2}\|\boldsymbol{x}'-\boldsymbol{x}\|^2)$, let us introduce a more flexible nonparametric form of the covariance function $K:\mathbb{R}^{d_X}\times\mathbb{R}^{d_X}\to\mathbb{R}$ as

$$K(\boldsymbol{x},\boldsymbol{x}')=\boldsymbol{\phi}(\boldsymbol{x})^\top\boldsymbol{\Omega}_\phi\boldsymbol{\phi}(\boldsymbol{x}'),$$

where $\mathbb{R}^{d_\phi\times d_\phi}\ni\boldsymbol{\Omega}_\phi\succ 0$ is a positive definite matrix that quantifies humans' intuitions for prioritizing features in comparing options. In other words, the similarity is defined as an inner product in the $d_\phi$-dimensional feature space, which is mapped by the function $\boldsymbol{\phi}$ and rotated by the matrix $\boldsymbol{\Omega}_\phi$.

Appearing in Proceedings of the $18^{th}$ International Conference on Artificial Intelligence and Statistics (AISTATS) 2015, San Diego, CA, USA. JMLR: W&CP volume 38. 

One direct way to optimize the covariance matrix is gradient ascending of the optimization objective, which is Eq. (7) in the main paper, with respect to $\boldsymbol{\Omega}_\phi$. By using the fact that $\boldsymbol{H}_i\triangleq(\boldsymbol{\Phi}_i\boldsymbol{\Omega}_\phi\boldsymbol{\Phi}_i^\top)(\boldsymbol{I}_{m[i]}+\boldsymbol{\Phi}_i\boldsymbol{\Omega}_\phi\boldsymbol{\Phi}_i^\top)^{-1}$, we can reach a local optimum of the metric matrix $\boldsymbol{\Omega}_\phi$.

Instead of such direct optimization, however, an approximated MAP estimation using a multi-task learning technique yields an efficient convex optimization of the matrix $\boldsymbol{\Omega}_\phi$. We can exploit the fact that the posterior mode, which is equivalent to the mean in GPR, is given by maximizing the sum of the log-likelihoods by the observation process $\mathcal{N}(\boldsymbol{\mu}_i,\sigma^2\boldsymbol{I}_{m[i]})$ and the subjective prior $\mathcal{N}(\boldsymbol{0}_{m[i]},\sigma^2\boldsymbol{K}(\boldsymbol{X}_i))$. We aim maximization of a joint log-likelihood, whose term related with context $i$ is given as a constant plus

$$\ell(\boldsymbol{u}_i^*,y_i)-\frac{m[i]}{2}\log\sigma^2-\frac{1}{2\sigma^2}\left\|\boldsymbol{u}_i^*-b\boldsymbol{1}_{m[i]}-\boldsymbol{\Phi}_i\boldsymbol{w}_\phi\right\|^2 -\frac{1}{2}\log|\boldsymbol{\Phi}_i\boldsymbol{\Omega}_\phi\boldsymbol{\Phi}_i^\top|-\frac{1}{2\sigma^2}\left\|\boldsymbol{u}_i^*\right\|^2_{\boldsymbol{\Phi}_i\boldsymbol{\Omega}_\phi\boldsymbol{\Phi}_i^\top}. \quad (1)$$

In (1), the term $\log|\boldsymbol{\Phi}_i\boldsymbol{\Omega}_\phi\boldsymbol{\Phi}_i^\top|$ mainly constrains the volume of the matrix $\boldsymbol{\Omega}_\phi$. Let us constrain the trace of the matrix $\boldsymbol{\Omega}_\phi$ instead of its log-determinant. As another modification, we replace the term $\|\boldsymbol{u}_i^*\|^2_{\boldsymbol{\Phi}_i\boldsymbol{\Omega}_\phi\boldsymbol{\Phi}_i^\top}$ by $\|\boldsymbol{\Phi}_i^\dagger\boldsymbol{u}_i^*\|^2_{\boldsymbol{\Omega}_\phi}$ with exploiting the Moore-Penrose pseudoinverse $(\cdot)^\dagger$.

We replace the matrix $\boldsymbol{\Omega}_\phi$ by $(\boldsymbol{A}_\phi/\eta)$ such that $\mathrm{Tr}(\boldsymbol{A}_\phi)=d_\phi$, and regularize the matrix $\boldsymbol{A}_\phi$ with a prior $p(\boldsymbol{A}_\phi)\propto\exp(-\frac{\delta\eta}{2\sigma^2}\mathrm{Tr}(\boldsymbol{A}_\phi^{-1}))$ by introducing a hyperparameter $\delta$. Given the mapping function $\boldsymbol{\phi}$ and the hyperparameters $(\sigma^2,\eta,c,\delta)$, an alternative optimization to approximate (1) is given as

$$\max_{\substack{(\boldsymbol{u}_i^*)_{i=1}^n, b\\ \boldsymbol{w}_\phi, \boldsymbol{A}_\phi}} \sum_{i=1}^n \Bigg[ \ell(\boldsymbol{u}_i^*, y_i) - \frac{1}{2\sigma^2} \left\| \boldsymbol{u}_i^* - b\boldsymbol{1}_{m[i]} - \boldsymbol{\Phi}_i \boldsymbol{w}_\phi \right\|^2 \frac{\eta}{2\sigma^2} \left\| \boldsymbol{\Phi}_i^\dagger \boldsymbol{u}_i^* \right\|_{\boldsymbol{A}_\phi}^2 \Bigg] - \frac{c}{2}\|\boldsymbol{w}_\phi\|^2 - \frac{\delta\eta}{2\sigma^2}\mathrm{Tr}(\boldsymbol{A}_\phi^{-1})$$
$$\text{subject to } \boldsymbol{A}_\phi \succ 0 \text{ and } \mathrm{Tr}(\boldsymbol{A}_\phi) = d_\phi. \quad (2)$$

The optimization (2) is jointly convex with respect to the variables $((\boldsymbol{u}_i^*)_{i=1}^n, b, \boldsymbol{w}_\phi, \boldsymbol{A}_\phi)$ (Zhang and Yeung, 2010), and hence gradient ascending consequently converges into the global optimum. We iterate Newton-Raphson updating in which Hessian matrix for each vector $\boldsymbol{u}_i^*$ is computed. The updating procedure for the matrix $\boldsymbol{A}_\phi$ is $\boldsymbol{A}_\phi \propto [\delta \boldsymbol{I}_{d_\phi} + \sum_{i=1}^n (\boldsymbol{\Phi}_i^\dagger \boldsymbol{u}_i^*)(\boldsymbol{\Phi}_i^\dagger \boldsymbol{u}_i^*)^\top]^{1/2}$, as derived in (Zhang and Yeung, 2010). Because iterative updating between the evaluations $(\boldsymbol{u}_i^*)_{i=1}^n$ and the parameters $(b, \boldsymbol{w}_\phi, \boldsymbol{A}_\phi)$ causes slow convergence, one should consider a warm start by setting $\boldsymbol{\Omega}_\phi = \boldsymbol{I}_{d_\phi}/\eta$ and initializing the parameters $(b, \boldsymbol{w}_\phi)$ and the evaluations $(\boldsymbol{u}_i^*)_{i=1}^n$.

Algorithm 1 summarizes the entire fitting algorithm involving the optimization of the metric. After setting $\boldsymbol{\Omega}_\phi = \boldsymbol{A}_\phi/\eta$, we should once again refit the values of $(b, \boldsymbol{w}_\phi)$ using the Optimization (7) in the main paper. This refitting provides better estimates than those produced only by Eq. (2), because the marginalization to compute the posterior mean compensates the errors induced by the trace-based MAP approximation.

Note that we could not confirm additional performance gains at least when we applied Algorithm 1 for our datasets. Gains by optimizing the covariance matrix would become clearer when the option attributes become more high-dimensional, while the option attributes of our datasets are only two- or three-dimensional.

**Algorithm 1** Estimation for all of the model parameters

**input** Choices and features $(y_i \in \{1, \ldots m[i]\}, \boldsymbol{\Phi}_i \in \mathbb{R}^{m[i] \times d_\phi})_{i=1}^n$, hyperparameters $(\sigma^2, \eta, c, \delta)$

Initialize $(b, \boldsymbol{w}_\phi)$ by taking the values when $\boldsymbol{\Omega}_\phi := \boldsymbol{I}_{d_\phi}/\eta$.
**for** $i = 1$ **to** $n$ **do**
  $\boldsymbol{H}_i := (\boldsymbol{\Phi}_i \boldsymbol{\Phi}_i^\top)(\eta \boldsymbol{I}_{m[i]} + \boldsymbol{\Phi}_i \boldsymbol{\Phi}_i^\top)^{-1}$
  $\boldsymbol{u}_i^* := b\boldsymbol{H}_i \boldsymbol{1}_{m[i]} + \boldsymbol{H}_i \boldsymbol{\Phi}_i$
**end for**
$\boldsymbol{A}_\phi := \boldsymbol{I}_{d_\phi}$, $\boldsymbol{R} := c\sigma^2 \boldsymbol{I}_{d_\phi} + \sum_{i=1}^n \boldsymbol{\Phi}_i^\top \boldsymbol{\Phi}_i$
**repeat**
  $b^{\mathrm{N}} := 0, \boldsymbol{\rho} := \boldsymbol{0}_{d_\phi}, \boldsymbol{A}_\phi^{\mathrm{N}} := \delta \boldsymbol{I}_{d_\phi}$
  **for** $i = 1$ **to** $n$ **do**
    $\boldsymbol{g}_i := (\exp(u_{i1}^*), \ldots, \exp(u_{im[i]}^*))^\top / (\sum_{j=1}^{m[i]} \exp(u_{ij}^*))$
    $\boldsymbol{\Sigma}_i := \boldsymbol{\Phi}_i \boldsymbol{A}_\phi \boldsymbol{\Phi}_i^\top$
    $\boldsymbol{f}_i := \sigma^2 \boldsymbol{\Sigma}_i (\boldsymbol{e}_{y_i} - \boldsymbol{g}_i) - \boldsymbol{\Sigma}_i \left( \boldsymbol{u}_i^* - b\boldsymbol{1}_{m[i]} - \boldsymbol{\Phi}_i \boldsymbol{w}_\phi \right) - \eta \boldsymbol{u}_i^*$
    $\boldsymbol{F}_i := \sigma^2 \boldsymbol{\Sigma}_i \left( \boldsymbol{g}_i \boldsymbol{g}_i^\top - diag(\boldsymbol{g}_i) \right) - \boldsymbol{\Sigma}_i - \eta \boldsymbol{I}_{m[i]}$
    $\boldsymbol{u}_i^* := \boldsymbol{u}_i^* - \boldsymbol{F}_i^{-1} \boldsymbol{f}_i$
    $b^{\mathrm{N}} := b^{\mathrm{N}} + \boldsymbol{1}_{m[i]}^\top (\boldsymbol{u}_i^* - \boldsymbol{\Phi}_i \boldsymbol{w}_\phi)$
    $\boldsymbol{\rho} := \boldsymbol{\rho} + \boldsymbol{\Phi}_i^\top (\boldsymbol{u}_i^* - b\boldsymbol{1}_{m[i]})$
    $\boldsymbol{A}_\phi^{\mathrm{N}} := \boldsymbol{A}_\phi^{\mathrm{N}} + (\boldsymbol{\Phi}_i^\dagger \boldsymbol{u}_i^*)(\boldsymbol{\Phi}_i^\dagger \boldsymbol{u}_i^*)^\top$
  **end for**
  $b^{\mathrm{N}} := b^{\mathrm{N}} / (\sum_{i=1}^n m[i])$, $\boldsymbol{w}_\phi := \boldsymbol{R}^{-1} \boldsymbol{\rho}$
  $\boldsymbol{A}_\phi := (\boldsymbol{A}_\phi^{\mathrm{N}})^{1/2} / \mathrm{Tr}((\boldsymbol{A}_\phi^{\mathrm{N}})^{1/2})$
**until** Convergence
Refit $(b, \boldsymbol{w}_\phi)$ by applying $\boldsymbol{\Omega}_\phi := \boldsymbol{A}_\phi/\eta$ for Optimization (7)

**output** Parameters $(b, \boldsymbol{w}_\phi, \boldsymbol{\Omega}_\phi := \boldsymbol{A}_\phi/\eta)$

## References

Y. Zhang and D.-Y. Yeung. A convex formulation for learning task relationships in multi-task learning. In *Proc. the 26th Conference on Uncertainty in Artificial Intelligence (UAI 2010)*, pages 733–742, 2010.