Multi-Task Preference Learning with Gaussian Processes

Adriana Birlutiu and Perry Groot and Tom Heskes *

Radboud University Nijmegen - Intelligent Systems
Toernooiveld 1, 6525 ED Nijmegen - the Netherlands

Abstract. We present an EM-algorithm for the problem of learning user preferences with Gaussian processes in the context of multi-task learning. We validate our approach on an audiological data set and show that predictive results for sound quality perception of normal hearing and hearing-impaired subjects, in the context of pairwise comparison experiments, can be improved using the hierarchical model.

1   Introduction

Learning user preferences appears in many contexts. Consider, for example, the case in which the parameters of a medical device such as a hearing aid have to be tuned such as to adapt them optimally to a user's preferences. Typically, preferences are learned from the user by repeatedly asking questions, for example, doing listening experiments in the case of hearing aid fitting. A major obstacle, however, is that obtaining new observations is often a time consuming process and a burden on the subject participating.

One approach to this problem is to not consider the parameter estimation for one subject, but for *multiple* subjects for similar tasks such that different subjects can regularize each other by assuming that model parameters are drawn from a common hyperprior [1, 2]. Using responses from other subjects effectively leads to an informed prior such that less observations are needed to obtain a good indication of the user's preferences.

In this paper, we extend earlier work on multi-task regression with Gaussian processes [3] to the case of multi-task learning of users' preferences. We demonstrate the usefulness of our model on an audiological data set collected by Arehart et. al. [4]. We show that the process of learning users' preferences can be significantly improved by using a *hierarchical non-parametric* model based on Gaussian processes.

The rest of this paper is organized as follows. Section 2 describes the probabalistic choice model used for preference learning. Section 3 introduces the Gaussian process framework we use for representing utility functions. Section 4 introduces the hierarchical extension of our Bayesian framework and describes the Expectation Maximization algorithm for learning a hierarchical prior. Section 5 reports experimental results with the hierarchical model for learning user preferences in the context of listening experiments. Section 6 presents our conclusions and directions for future work.

2 Probabilistic Choice Models

Let $X = \{\boldsymbol{x}^1, \ldots, \boldsymbol{x}^N\}$ be a set of N distinct inputs with $\boldsymbol{x}^i \in \mathbb{R}^d$, i.e., every input is represented by some features. Let $\mathcal{D}^j$ be a set of N^j observed preference comparisons over instances in X, corresponding to subject j,

$$\mathcal{D}^j = \{(\boldsymbol{x}^{i1}, \ldots, \boldsymbol{x}^{iK}, k) \,|\, 1 \leq i \leq N^j, \boldsymbol{x}^{i\cdot} \in X, k \in \{1, \ldots, K\}\}$$

where k means that alternative $\boldsymbol{x}^{ik}$ is preferred from the K inputs presented to subject j.

A standard modeling assumption [5, 6] is that the subject's decision in such type of forced-choice comparisons follows a probabilistic model defined as follows

$$P(k; \boldsymbol{x}^{i1}, \ldots, \boldsymbol{x}^{iK}, \boldsymbol{\theta}^j) = \frac{\exp\left[U(\boldsymbol{x}^{ik}, \boldsymbol{\theta}^j)\right]}{Z(\boldsymbol{\theta}^j)}, \tag{1}$$

with Z a normalization constant

$$Z(\boldsymbol{\theta}^j) \equiv \sum_{k=1}^{K} \exp\left[U(\boldsymbol{x}^{ik}, \boldsymbol{\theta}^j)\right].$$

$\boldsymbol{\theta}^j$ is a vector of parameters specific to subject j, and $U(\boldsymbol{x}^{ik}, \boldsymbol{\theta}^j)$ represents a utility function capturing the preference of subject j for option k. Equation (1) gives the probability that subject j prefers alternative k when given the inputs: $\boldsymbol{x}^{i1}, \ldots, \boldsymbol{x}^{iK}$. In this study, we restrict the model to pairwise-comparisons, i.e., $K = 2$, which is also known in the literature as the Bradley-Terry model [7].

3 Gaussian Processes

The probabilistic choice model used for learning preferences, defined above, can be reformulated in terms of Gaussian processes (GPs) [8]. The GP formalism allows for nonlinear utility functions and is nonparametric.

We define a GP over the utility function for subject j, by assuming that the utility values are drawn from a multivariate Gaussian distribution, i.e.,

$$\{U^j(\boldsymbol{x}^1), \ldots, U^j(\boldsymbol{x}^N)\} \sim \mathcal{N}(\boldsymbol{\mu}_U, \boldsymbol{K}).$$

The covariance matrix $\boldsymbol{K}$ can be specified by a symmetric positive definite kernel function κ by setting $K_{ij} = \kappa(\boldsymbol{x}^i, \boldsymbol{x}^j)$. Examples for κ are the linear kernel and the Gaussian kernel defined as

$$\kappa_{\text{Linear}}(\boldsymbol{x}, \boldsymbol{y}) = \sum_{i=1}^{d} x_i y_i,$$

$$\kappa_{\text{Gauss}}(\boldsymbol{x}, \boldsymbol{y}) = \exp\left[-\frac{\ell}{2} \sum_{i=1}^{d} (x_i - y_i)^2\right],$$

where ℓ is a length-scale parameter.

A similar approach using GPs for learning preferences is introduced in [9].

4 Multi-Task Gaussian Processes

4.1 General Formulation

In many scenarios, data used for learning preferences is available from a group of subjects. In order to optimize the process of learning the utility function for a *new* subject, we make use of the data available from the *other* subjects for which the preferences were already learned. To implement this idea, we consider the estimation of *M related functions* $U^j, j = 1, \ldots, M$ for M different subjects, using a hierarchical Bayesian model, allowing individual models to interact and regularize each other.

Following [3], we assume that in the multi-task setting, the prior $P(\boldsymbol{\mu}_U, \boldsymbol{K})$ over the utility functions is the same for each subject, and it is drawn from a normal-inverse-Wishart distribution

$$P(\boldsymbol{\mu}_U, \boldsymbol{K}) = \mathcal{N}(\boldsymbol{\mu}|0, \frac{1}{\pi}\boldsymbol{K})\mathcal{IW}(\boldsymbol{K}|\tau, \boldsymbol{\kappa}) \,, \tag{2}$$

with $\boldsymbol{\kappa}$ the so-called *base kernel*. Possible choices for $\boldsymbol{\kappa}$ are the Gaussian kernel κ_{Gauss} or the linear kernel κ_{Linear} defined above.

4.2 Alternative Formulation

Inspired by the approach of [3], we derive an equivalent representation for the multi-task GPs. In this representation, for each subject j we have a vector of parameters $\boldsymbol{\alpha}^j$ (with dimension equal to N, the number of distinct input points), which captures in a compact form the information collected from the data set related to subject j. An *inductive utility function* U^j, for subject j can be defined for an unseen input $\boldsymbol{x}$ as follows:

$$U^j(\boldsymbol{x}) = \sum_{i=1}^{N} \alpha_i^j \kappa(\boldsymbol{x}, \boldsymbol{x}^i) = U(\boldsymbol{x}, \boldsymbol{\alpha}^j) \,,$$

where $\boldsymbol{x}^i \in X$ and κ is the *base kernel* defined above.[1] The dual representation is a consequence of the *representer theorem* [8].

It follows from Equation (2) that the vectors of parameters $\boldsymbol{\alpha}^j$ are sampled from a hierarchical prior distribution. In order to learn this prior, we couple the inference tasks of all the subjects. We set $P(\boldsymbol{\alpha}^j) = \mathcal{N}(\boldsymbol{\alpha}^j|\boldsymbol{\mu}_\alpha, \boldsymbol{C})$ a Gaussian prior with the same $\boldsymbol{\mu}_\alpha$ and $\boldsymbol{C}$ for every subject j, where $\boldsymbol{\mu}_\alpha$ and $\boldsymbol{C}$ are sampled once from a normal-inverse-Wishart distribution (with scale matrix $\boldsymbol{\kappa}^{-1}$)

$$P(\boldsymbol{\mu}_\alpha, \boldsymbol{C}) = \mathcal{N}(\boldsymbol{\mu}_\alpha|0, \frac{1}{\pi}\boldsymbol{C})\mathcal{IW}(\boldsymbol{C}|\tau, \boldsymbol{\kappa}^{-1}) \,.$$

The posterior distribution over $\boldsymbol{\alpha}^j$, that results from the hierarchical prior and all the data available from subject j, is assumed to be close to a Gaussian,

---

[1] Note that we use this representation of the utility function in the probabilistic choice model defined in Equation (1).

$\mathcal{N}(\boldsymbol{\alpha}^j | \hat{\boldsymbol{\alpha}}^j, \hat{\boldsymbol{C}}_{\alpha^j})$. There are several alternatives to approximate the posterior distribution to a Gaussian: deterministic methods for approximate inference (e.g., Laplace's method [10], Expectation propagation [11]) or methods based on sampling.

EM algorithm

The hierarchical prior is obtained by maximizing the penalized loglikelihood of all data. This optimization is performed by applying the Expectation Maximization algorithm [1, 3], which reduces in our case to the iteration, until convergence, of the following two steps.

E-step: For each subject j, estimate the statistics (mean $\hat{\boldsymbol{\alpha}}^j$ and covariance matrix $\hat{\boldsymbol{C}}_{\alpha^j}$) of the posterior distribution over $\boldsymbol{\alpha}^j$, given the current estimates, $\boldsymbol{\mu}_\alpha$ and $\boldsymbol{C}$, of the hierarchical prior.

M-step: Re-estimate the parameters of the hierarchical prior:

$$\boldsymbol{\mu}_\alpha = \frac{1}{\pi + M} \sum_{j=1}^{M} \hat{\boldsymbol{\alpha}}^j$$

$$\boldsymbol{C} = \frac{1}{\tau + M} \left[\pi \boldsymbol{\mu}_\alpha \boldsymbol{\mu}_\alpha^T + \boldsymbol{\kappa}^{-1} + \sum_{j=1}^{M} \hat{\boldsymbol{C}}_{\alpha^j} + \sum_{j=1}^{M} (\hat{\boldsymbol{\alpha}}^j - \boldsymbol{\mu}_\alpha)(\hat{\boldsymbol{\alpha}}^j - \boldsymbol{\mu}_\alpha)^T \right] .$$

5 Experiments

We validate our approach of hierarchical preference learning on an audiological data set containing listening experiments, described in [4]. The data set consists of 576 pairwise comparisons per subject for 14 normal-hearing and 18 hearing-impaired subjects. Each listening experiment is of the form $(\boldsymbol{x}_1, \boldsymbol{x}_2, k)$, where $k = \{1, 2\}$ denotes whether sound sample $\boldsymbol{x}_1$ or $\boldsymbol{x}_2$ was preferred by the subject, respectively.

We used this data set in order to test empirically whether, in the context of learning preferences in the GP framework, the preferences of a *new subject* can be learned faster, by using the data available from a *group of subjects* for which the preferences were already learned. We compared the performances obtained using the hierarchical prior versus a flat prior which assumes no information about subject's preferences. In a simulation, the jth subject was left out, and the EM algorithm described in the previous section was used to gather data from the rest of the subjects in a probability distribution over $\boldsymbol{\alpha}^j$, which was used as the starting prior for the left-out subject. The data set for the left-out subject, was split into training (used for learning preferences) and testing (the accuracy of the predictions on the test data was used as a measure of how much we learned about subject's preferences). For each subject, we averaged the results using 10-fold cross-validation. Furthermore, the results were averaged within each group of subjects, i.e., normal-hearing and hearing-impaired subjects.
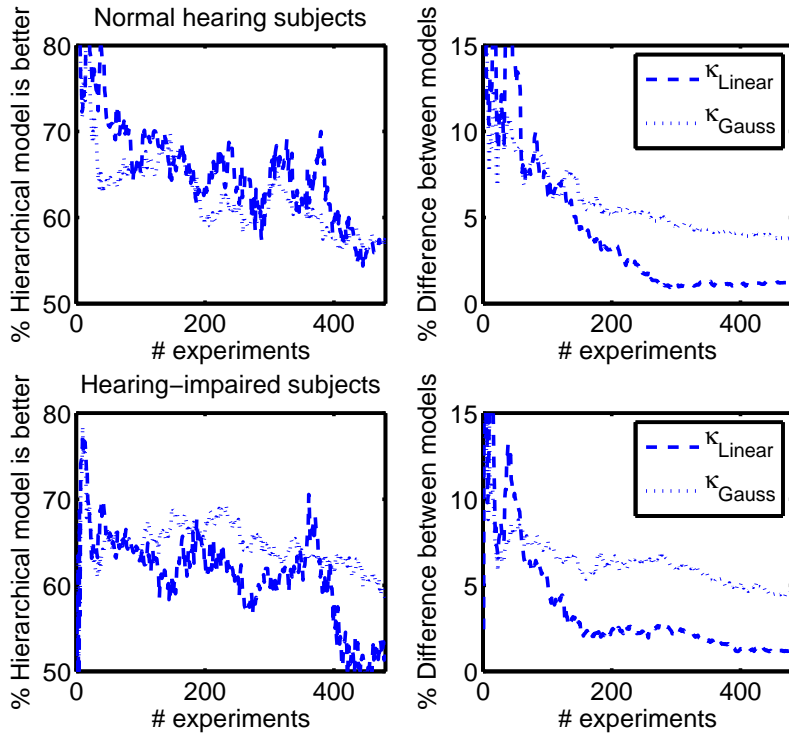
Fig. 1: Left: percentage of the number of times the prediction accuracy using the learned prior is better than the prediction accuracy with a flat prior. Right: percentage of the number of predictions on which the two models (with the learned and with a flat prior) disagree. For the Gaussian kernel we set $\ell = 1$; the results are rather insensitive to the specific choice for this parameter. Top and bottom rows refer to experiments on the data set from normal-hearing and hearing-impaired subjects, respectively.

We made predictions for the outcomes of the experiments from the test data, using a model which resulted either by starting with a flat prior which assumes no information, or a model which uses the hierarchical prior as the starting prior for $\boldsymbol{\alpha}^j$. The plots on the right-hand side of Figure 1, give the percentage of predictions on which the two models (the one with the hierarchical and the one with the flat prior) disagree, with respect to the total number of predictions made; the dashed line refers to a GP with a linear kernel, the dotted line to a GP with a Gaussian kernel. The plots on the left-hand side of Figure 1, show the percentage of correct predictions made using the hierarchical prior, with respect to the number of predictions on which the two models disagree. Especially in

the beginning of the learning process, with few experiments, the model with a prior learned from the community of other subjects significantly outperforms the model with a flat prior.

6 Conclusions and Future Work

We have introduced a hierarchical modelling approach for learning related functions of multiple subjects performing similar tasks using Gaussian processes. A hierarchical prior was used from which model parameters were sampled in order to enforce a similar structure for the utility functions of each individual subject.

We are interested in further improvements of the model. Particularly, we plan to investigate how to select, in an active way, the most informative experiments in order to learn users' preferences. Furthermore, it might be interesting to automatically cluster, beforehand, the subjects into groups with similar behaviour; as in the current study we manually clustered the data set into two sets of normal-hearing and hearing-impaired subjects.

References

[1] A. Gelman, J.B. Carlin, H.S. Stern, and D.B. Rubin. *Bayesian Data Analysis, Second Edition*. Chapman & Hall/CRC, July 2003.

[2] B. Bakker and T. Heskes. Task clustering and gating for Bayesian multitask learning. *Journal of Machine Learning Research*, 4:83–99, 2003.

[3] K. Yu, V. Tresp, and A. Schwaighofer. Learning Gaussian Processes from Multiple Tasks. In *Proceedings of the 22nd International Conference on Machine Learning*, 2005.

[4] K.H. Arehart, J.M. Kates, C.A. Anderson, and L.O. Harvey Jr. Effects of noise and distortion on speech quality judgments in normal-hearing and hearing-impaired listeners. *J. Acoust. Soc. Am.*, 122(2):1150–1164, August 2007.

[5] M. Glickman and S. Jensen. Adaptive paired comparison design. *Journal of Statistical Planning and Inference*, 127:279–293, 2005.

[6] B. Kanninen. Optimal design for multinomial choice experiments. *Journal of Marketing Research*, 39:307–317, 2002.

[7] R.A. Bradley and M.E. Terry. Rank analysis of incomplete block designs: I, the method of paired comparisons. *Biometrika*, 1952.

[8] C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

[9] W. Chu and Z. Ghahramani. Preference Learning with Gaussian Processes. In *Proceedings of the 22nd International Conference on Machine Learning*, Bonn, Germany, 2005.

[10] D.J.C. Mackay. *Information Theory, Inference & Learning Algorithms*. Cambridge University Press, New York, NY, USA, 2002.

[11] T. P. Minka. *A Family of Algorithms for Approximate Bayesian Inference*. PhD thesis, M.I.T., 2001.