Multi-Task Preference Learning with Gaussian Processes[1]

Adriana Birlutiu   Perry Groot   Tom Heskes

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

1 Introduction

Recently, there has been an increasing interest in learning preferences. In this paper, we consider preference learning in the context of having obtained preference data from *multiple* subjects in similar tasks. This setup is interesting as in many real-world application domains data for a specific single scenario is scarce, but data is already available from similar scenarios. The data from different subjects can be used to regularize individual user models by assuming that model parameters are drawn from a common hyperprior. We extend earlier work on multi-task regression with Gaussian processes to the case of multi-task learning of subjects's preferences. We demonstrate the usefulness of our model on an audiological data set. We show that the process of learning subject's preferences can be significantly improved by using a *hierarchical non-parametric* model based on Gaussian processes.

2 Multi-Task Framework

Let $X = \{\boldsymbol{x}^1, \ldots, \boldsymbol{x}^N | x^i \in \mathbb{R}^d\}$ be a set of N distinct inputs. 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 assumption is that the subject's decision in such forced-choice comparisons follows a probabilistic model

$$P(k; \boldsymbol{x}^{i1}, \ldots, \boldsymbol{x}^{iK}, \boldsymbol{\theta}^j) = \exp\left[U(\boldsymbol{x}^{ik}, \boldsymbol{\theta}^j)\right] / Z(\boldsymbol{\theta}^j), \qquad Z(\boldsymbol{\theta}^j) \equiv \sum_{k=1}^{K} \exp\left[U(\boldsymbol{x}^{ik}, \boldsymbol{\theta}^j)\right].$$

with Z a normalization constant, $\boldsymbol{\theta}^j$ a vector of parameters specific to subject j, and $U(\boldsymbol{x}^{ik}, \boldsymbol{\theta}^j)$ a utility function capturing the preference of subject j for option k.

We define a Gaussian process 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 $\boldsymbol{K}$ is specified by a symmetric positive definite kernel function κ, $K_{ij} = \kappa(\boldsymbol{x}^i, \boldsymbol{x}^j)$. As a consequence of the *representer theorem*, the utility function U^j has a dual representation

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

[1]Published in the Proceedings of the 17th European Symposium on Artificial Neural Networks (ESANN), pages 123-128, 2009.

The parameters $\boldsymbol{\alpha}^j$ are sampled from a hierarchical prior distribution. To learn this prior we couple the tasks of all subjects and set $P(\boldsymbol{\alpha}^j) = \mathcal{N}(\boldsymbol{\alpha}^j | \boldsymbol{\mu}_\alpha, \boldsymbol{C}_\alpha)$ a Gaussian prior with the same $\boldsymbol{\mu}_\alpha$ and $\boldsymbol{C}_\alpha$ for every subject j, where $\boldsymbol{\mu}_\alpha$ and $\boldsymbol{C}_\alpha$ are sampled once from a normal-inverse-Wishart distribution (with scale matrix $\boldsymbol{\kappa}^{-1}$). The hierarchical prior is obtained by maximizing the penalized loglikelihood of all data. This optimization is performed by applying the Expectation Maximization algorithm, which reduces in our case to the iteration, until convergence, of the following two steps.

E-step: For each subject j, estimate the sufficient 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 $\boldsymbol{C}_\alpha$, of the hierarchical prior.

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

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

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

3 Experiments

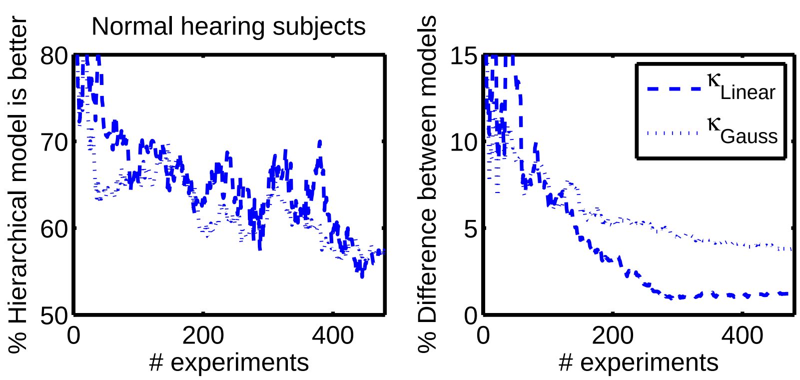

We validated our approach on an audiological data set of 14 normal-hearing subjects using two different kernels.[2] Each listening experiment consisted of a preference response for one of two sounds presented. The results are shown above. Left: percentage of the number of times the prediction accuracy 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. 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.

4 Conclusions and Future Work

We introduced a hierarchical modelling approach for learning related functions of multiple subjects performing similar tasks using Gaussian processes. The hierarchical model with a prior learned from other subjects significantly outperformed the model with a flat prior.

We are interested in further improvements of the hierarchical model. In particular, combining the hierarchical model with active learning and automatic clustering techniques to automatically group subjects with similar behaviour.

[2]The paper published also contains results on an audiological data set of 18 hearing-impaired subjects.