Efficient Preference Learning with Pairwise Continuous Observations and Gaussian Processes

Bjørn Sand Jensen[1], Jens Brehm Nielsen[1,2] & Jan Larsen[1]

[1]{bjje,jenb,jl}@imm.dtu.dk, DTU Informatics, Technical University of Denmark [2]jeb@widex.com, Widex A/S

Abstract

Human preferences can effectively be elicited using pairwise comparisons and in this paper current state-of-the-art based on binary decisions is extended by a new paradigm which allows subjects to convey their degree of preference as a continuous but bounded response. For this purpose, a novel Beta-type likelihood is proposed and applied in a Bayesian regression framework using Gaussian Process priors. Posterior estimation and inference is performed using a Laplace approximation. The potential of the paradigm is demonstrated and discussed in terms of learning rates and robustness by evaluating the predictive performance under various noise conditions on a synthetic dataset. It is demonstrated that the learning rate of the novel paradigm is not only faster under ideal conditions, where continuous responses are naturally more informative than binary decisions, but also under adverse conditions where it seemingly preserves the robustness of the binary paradigm, suggesting that the new paradigm is robust to human inconsistency.

Methods

Introduction

According to Lockhead [1] every aspect of human perception is relative. Formal treatment of relative aspects goes back to the ideas of Thurnstone [2] and was revisited by Chu *et al.* [3] who formulated a Bayesian approach to preference learning with Gaussian Processes, which has initiated various applications and studies [4], [5] and [6].

To accelerate the preference learning task, we formulate a likelihood model that describes not only which of two presented options that is preferred, but additionally the degree to which the prevailing option is preferred as a continues but bounded response.

Pairwise setup

We consider n distinct inputs $x_i \in \mathcal{X}$ denoted $\mathcal{X} = \{x_i | i = 1, ..., n\}$ and a set of m responses y_k on pairwise comparisons between any two inputs u_k and v_k in $\mathcal{X}$, denoted by

$$\mathcal{Y} = \{(y_k; u_k, v_k) | k = 1, ..., m\}.$$

We consider two types of responses

- **binary** (BR):
  $$y_k = d_k, d_k \in \{-1, 1\}$$
- **continuous & bounded** (CBR):
  $$y_k = \pi_k, \pi_k \in]0, 1[$$

Model

The preference for each input x_i is modeled by a latent function $f : \mathcal{X} \to \mathbb{R}$ defining an internal but latent reference. Given a function f we define the likelihood functions for each of the two responses as

- **binary** (BR) - State-of-the-art [3]
  $$p(d_k | \mathbf{f}_k) = \Phi\left(d_k \frac{f(v_k) - f(u_k)}{\sqrt{2}\sigma}\right)$$
- **continuous & bounded** (CBR) - Novel
  $$p(\pi_k | \mathbf{f}_k) = \mathrm{Beta}\left(\pi_k | \nu\mu(\mathbf{f}_k), \nu(1 - \mu(\mathbf{f}_k))\right)$$
  $$\text{with } \mu(\mathbf{f}_k) = \Phi\left(\frac{f(v_k) - f(u_k)}{\sqrt{2}\sigma}\right)$$

where we have defined $\mathbf{f}_k = [f(u_k), f(v_k)]^\top$ as the vector of function values for the two compared options u_k and v_k.

Gaussian Process Prior

We model the latent function $f(x)$ with a zero-mean Gaussian Process (GP) [7] $f(x) \sim \mathcal{GP}(\mathbf{0}, k(\cdot, \cdot))$ with a *squared-exponential* (SE) covariance function $k_{SE}(x, x')$. A fundamental consequence of this formalism is that the GP defines a prior over functions $p(\mathbf{f}|\mathcal{X})$, which leads us directly to a formulation given Bayes rule

$$p(\mathbf{f}|\mathcal{Y}, \mathcal{X}) = \frac{p(\mathbf{f}|\mathcal{X}) \prod_{k=1}^{m} p(y_k|\mathbf{f}_k)}{p(\mathcal{Y}|\mathcal{X})}.$$

where $p(y_k|\mathbf{f}_k)$ is any of the two pairwise likelihood functions, $\mathbf{f} = [f(x_1), f(x_2), \ldots, f(x_n)]^\top$ and we have assumed that the likelihood factorizes over observations.

Exact inference is intractable for both likelihood functions. Instead, the Laplace approximation is used for inference, parameter optimization and prediction.

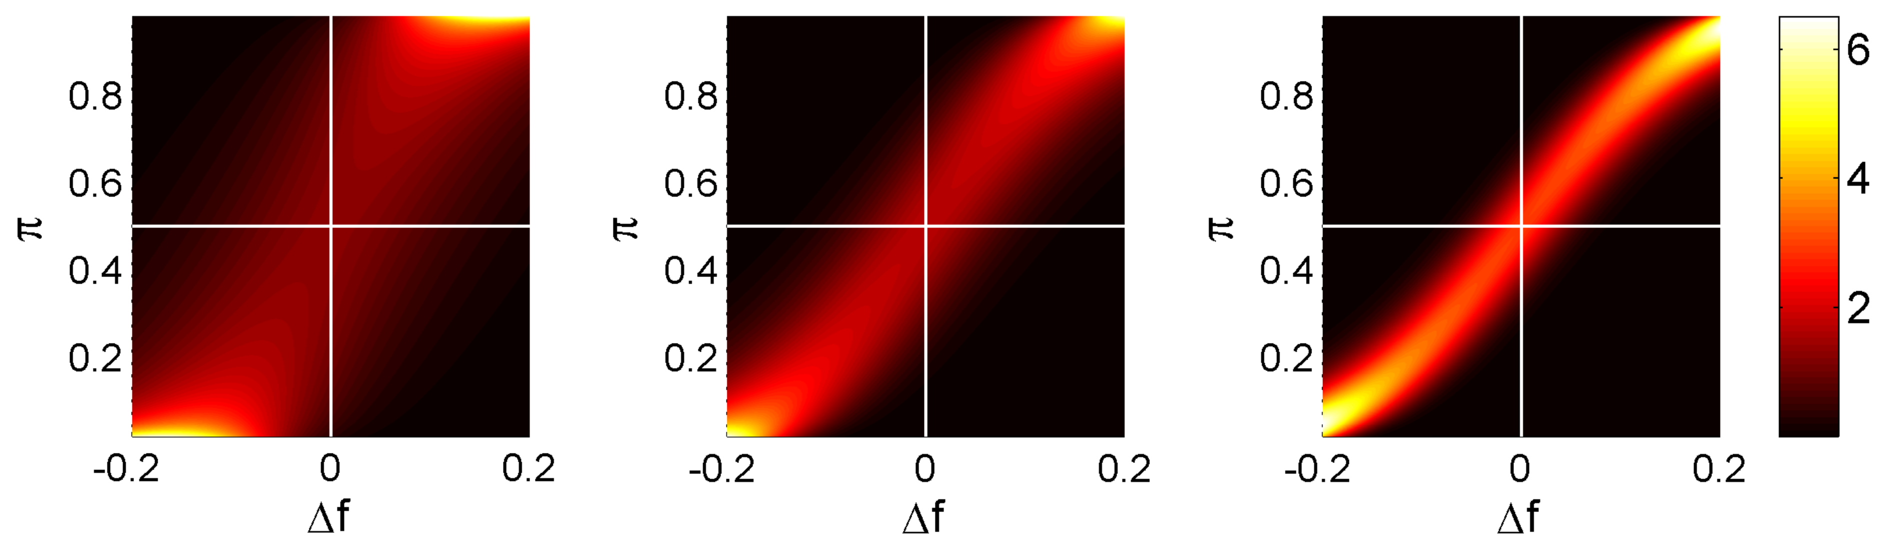

Fig. 1: *Illustration of the proposed likelihood with $p(\pi_k|\mathbf{f}_k)$ shown as a color level. The likelihood parameters are $\sigma = 0.1$ and left: $\nu = 3$, middle: $\nu = 10$ and right: $\nu = 30$*

Simulation Results

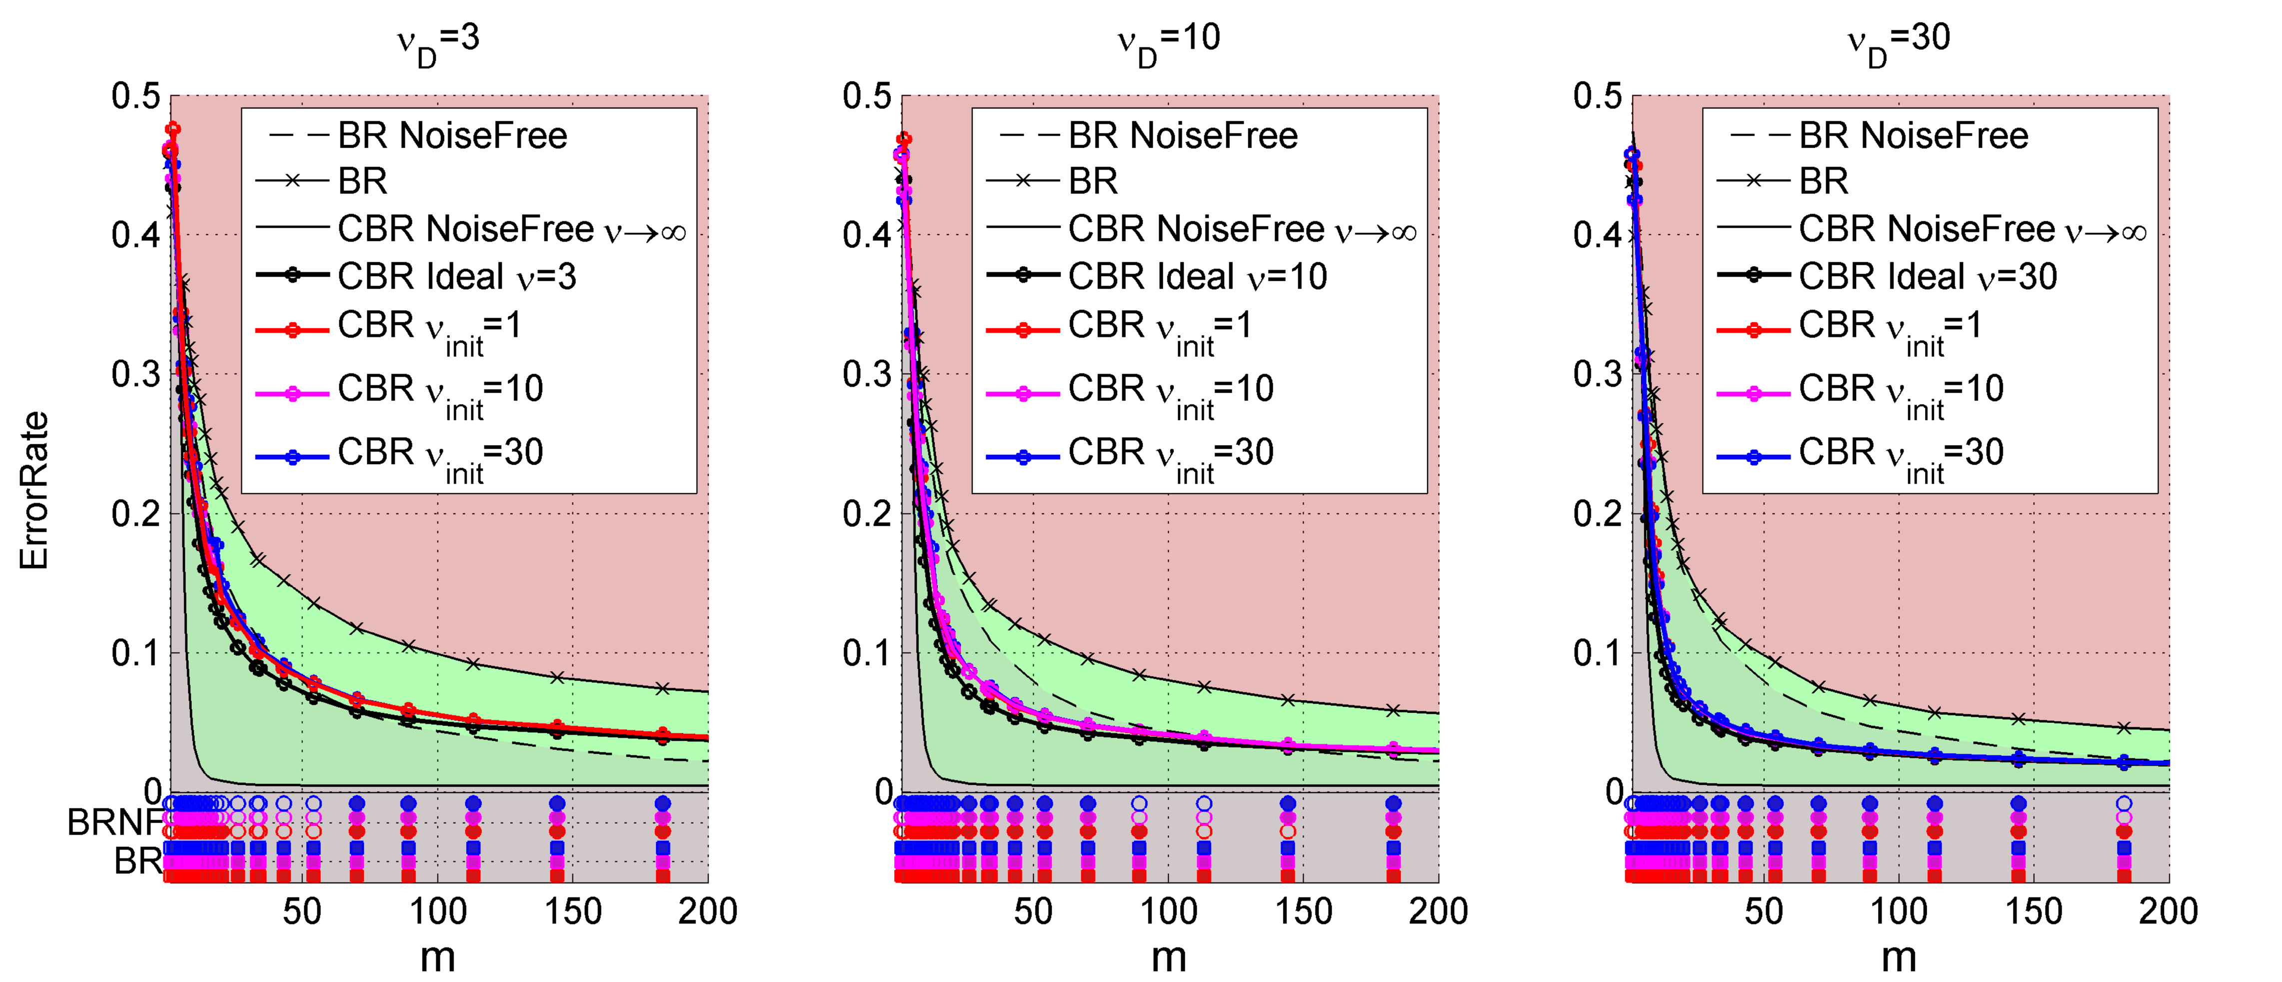

Fig. 2: *Mean error test rates (MER) as a function of the number of experiments over 100 different realizations of the training set generated with different ν_D. In the red and top green area MER are worse and better, respectively, than those obtained with the BR model on the noisy data. In the lower green area MER are also better than those obtained by the BR NoiseFree, and finally, the grey area corresponds to unrealistic MER better than those obtained with a CBR NoiseFree model with $\nu \to \infty$ evaluated with $\nu = 10^3$ on a noise-free data set. The six rows of markers indicate if the MER of the corresponding CBR model are significantly different from those resulting from the BR (squares) and from the BR NoiseFree (circles). If solid, the zero-hypothesis of the two means being equal is rejected at the 5% level using a paired t-test.*

Conclusions and Directions

The learning rates (Fig. 1) show that

- As expected, CBR outperforms BR under (near) ideal conditions.
- In noisy conditions, CBR outperforms corresponding BR.
- Actually, CBR shows similar or better performance in noisy conditions than BR in noise-free conditions.

This suggests that the CBR model

- **is robust to user inconsistency.**
- **can effectively exploit the additional information in continuous responses.**
- **reduces the number of required experiments/observations (faster learning rate).**

Future research includes

- Real world experiments, e.g., listening personalization.
- Suitable active learning criterion to increase the learning rate.
- More flexible Beta mean functions accounting for different user behaviors.

Acknowledgements

This work was supported in part by the IST Programme of the European Community, under the PASCAL2 Network of Excellence, IST-2007-216886. This publication only reflects the authors' views.

References

[1] G. R. Lockhead, "Absolute Judgments Are Relative: A Reinterpretation of Some Psychophysical Ideas.," *Review of General Psychology*, vol. 8, no. 4, pp. 265–272, 2004.

[2] L. L. Thurstone, "A law of comparative judgement.," *Psychological Review*, vol. 34, 1927.

[3] W. Chu and Z. Ghahramani, "Preference learning with Gaussian Processes," *ICML 2005 - Proceedings of the 22nd International Conference on Machine Learning*, pp. 137–144, 2005.

[4] P. Groot, T. Heskes, T. Dijkstra, and J. Kates, "Predicting preference judgments of individual normal and hearing-impaired listeners with Gaussian Processes," *IEEE Transactions on Audio, Sound, and Language Processing*, 2010.

[5] E. Bonilla, S. Guo, and S. Sanner, "Gaussian Process preference elicitation," in *Advances in Neural Information Processing Systems 23*, J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R.S. Zemel, and A. Culotta, Eds., pp. 262–270. 2010.

[6] W. Chu and Z. Ghahramani, "Extensions of Gaussian Processes for ranking: semi-supervised and active learning," in *Workshop Learning to Rank at Advances in Neural Information Processing Systems 18*, 2005.

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

DTU Informatics
Department of Informatics and Mathematical Modelling