Preference Learning with Strength

Jonathan Downing, Steven Reece, Michael Osborne, Stephen Roberts

Department of Engineering Science, University of Oxford

Abstract. In order for machines to effectively collaborate with humans computers must be able to learn, and reason about, human preferences. Requesting a person to exhaustively assess their preferences regarding individual objects has been shown to be inaccurate and tiresome. We elicit preference automatically from data comprising human supplied pairwise comparisons, which invokes lower cognitive effort on the part of the person. We also seamlessly fold in the strength of preference, supplied by the human, to further improve the accuracy and efficiency of our algorithm. We take a Bayesian approach, exploiting the rich principles of probabilistic reasoning to reduce the number of comparisons undertaken by the human. In particular we use Bayesian non-parametrics, specifically Gaussian processes, to model a latent value function that allows us to infer global preference ordering from a small set of pairwise choices. We show that by integrating preference strength the accuracy of our approach improves by a factor of two.

1 Introduction

We live in a world where our interactions with technology are becoming more intimate and persistent. As this trend continues there is the risk that technologies that are 'mechanical' and transactional will fail to connect with the user. Learning the behaviour and preferences of the user is essential in mitigating this problem by reducing cognitive effort and increasing the efficiency of the technology. In order for machines to effectively collaborate with humans computers must be able to learn, and reason about, human preferences. This paper extends Bayesian preference learning [1] by incorporating strength of preference as a binary observation inspired by the economics [3] and Doyle's *Prospects for Preference* [2] in which he argues this as a more realistic representation of preference as opposed to treating all pairwise preferences equally. The challenge is to incorporate preference strength.

The problem is framed by presenting the user two objects and asking which the user prefers. Furthermore, we ask the user whether their preference is 'strong' or 'weak'. Depending on the application this can reflect either their strength in their choice (for example, "A is definitely much better than B") or their uncertainty in the choice ("Given the information I have, I think A is better than B"). The user is presented with multiple pairs of objects and their rankings are aggregated through our Gaussian process algorithm. The algorithm then enables us to infer pairwise relations not already observed and thus reduces the number of actual tasks performed by the user. Uncertainty in the user's preference and strength indicators are implicit in the user's response and these are represented in our model as a distribution over latent functions and learned from the data. This extends [1] in two ways: firstly, we introduce heteroscedastic latent functions that encode the strength in our belief about the preference ordering. Secondly, we introduce a likelihood function that encodes strong and weak preferences.

2 Mathematical Framework

The object A is preferred to B is expressed as a preference relation $A \succ B$. Furthermore, to express strong and weak preference we write $A \overset{s}{\succ} B$ and $A \overset{w}{\succ} B$, respectively. As per [1], preferences are represented as a latent function $f(x)$ over the objects, x. The latent function allows us to infer the global preference ordering from a small set of pairwise choices and is drawn from a Gaussian process to capture strong correlations between similar objects. The Gaussian process preserves the object preference ordering and $A \succ B$ if and only if $f(A) > f(B)$. Furthermore, when we consider the strength of the preference ordering then, when $A \overset{s}{\succ} B$ and $C \overset{w}{\succ} D$, then $f(A) - f(B) > f(C) - f(D)$. Both models are constructed from a set of n unique instances $x_i \in \mathbb{R}^d$ denoted as $\mathcal{X} = \{x_i : i = 1, ..., n\}$, a set of m observed pairwise preference relations on the instances, denoted, $\mathcal{D} = \{a_k \overset{s_k}{\succ} b_k : k = 1, ..., m\}$, where $a, b \in \mathcal{X}$ and $s \in \{\text{weak}, \text{strong}\}$.

Likelihood of Heteroscedastic Model: The likelihood model assumes independent pairwise preference observations. For a noiseless preference $A \succ B$ this consists of a step function over the domain of $f(A) - f(B)$, ensuring $f(A) - f(B) > 0$. Introducing Gaussian noise $\varepsilon \sim \mathcal{N}(0, \sigma^2)$ on the step function, i.e. $f(A) - f(B) > \varepsilon$. In our model we make the preference noise σ dependent on preference strength: $\mathcal{P}(a_k \overset{s_k}{\succ} b_k | f(a_k), f(a_k), s_k) = \Phi(z_k)$ where $z_k = \frac{f(a_k) - f(b_k)}{\sigma_{s_k}}$, $\Phi(z) = \int_{-\infty}^{z} \mathcal{N}(\delta; 0, 1) d\delta)$ and σ_{s_k} is drawn from a Gamma distribution.

Likelihood of Preferences Over Preferences Model: We introduce higher level preferences, also stated in the literature as *preference over preferences* [2]. A higher second level preference of $(A \succ B) \succ (C \succ D)$ is preserved in the latent function by $f(A) - f(B) > f(C) - f(D)$.

3 Results

We make a direct comparison of our methods against that of Chu et al [1]. We use the benchmark data set of the Boston Housing problem, which has been looked at extensively in the literature. Figure 1 shows the improvement achieved using preferences over preferences. This is in sharp contrast to the heteroscedastic model, which sits uniformly at the simple preference learning baseline. The left panel shows the variation of the Kendall Tau coefficient as the number of strong pairwise preferences increases. We note with interest that best performance is gained when the ratio of strong to weak preferences is approximately 0.5. The right panel shows the Kendall coefficient as the number of preferences increases.

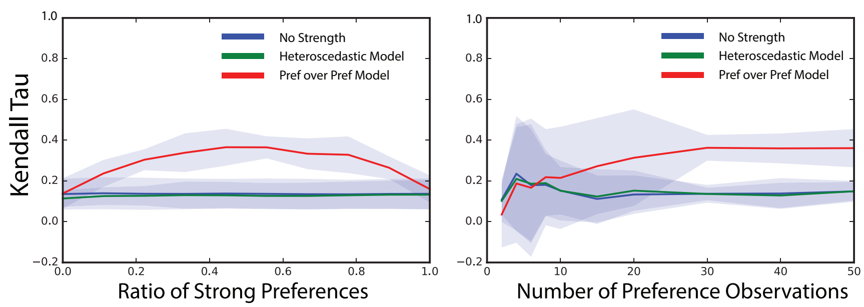

Fig. 1. Boston Housing data preference learning comparison.

References

[1] Wei Chu and Zoubin Ghahramani. "Preference Learning with Gaussian Processes". In: ICML '05. 2005, pp. 137–144.

[2] Jon Doyle. "Prospects for Preferences". en. In: *Computational Intelligence* 2 (May 2004), pp. 111–136.

[3] Veronika Köbberling. "Strength of preference and cardinal utility". en. In: *Economic Theory* 27.2 (), pp. 375–391.