

Towards autonomous decision-making: A probabilistic model for learning multi-user preferences

Markus Peters and Wolfgang Ketter

May 2013
ERS-2013-007-LIS

Department of Decision and Information Sciences
Rotterdam School of Management
Erasmus University
3062 PA Rotterdam, The Netherlands

Abstract

Information systems have revolutionized the provisioning of decision-relevant information, and decision support tools have improved human decisions in many domains. Autonomous decision-making, on the other hand, remains hampered by systems' inability to faithfully capture human preferences. We present a computational preference model that learns unobtrusively from limited data by pooling observations across like-minded users. Our model quantifies the certainty of its own predictions as input to autonomous decision-making tasks, and it infers probabilistic segments based on user choices in the process. We evaluate our model on real-world preference data collected on a commercial crowdsourcing platform, and we find that it outperforms both individual and population-level estimates in terms of predictive accuracy and the informativeness of its certainty estimates. Our work takes an important step toward systems that act autonomously on their users' behalf.

Keywords: Assistive technologies, Autonomous decision-making, Multi-task learning, Preferences, Software agents

Contents

1	Introduction	1
2	Related Work	2
2.1	Information Systems	2
2.2	Computer Science	3
2.3	Psychology	4
2.4	Marketing and Econometrics	4
3	Mixed-Effect Preference Model	4
3.1	Preliminaries	4
3.2	Modeling Single-User Preferences	6
3.3	Modeling Multi-User Preferences	8
4	Experimental Evaluation	12
5	Conclusions and Future Work	14
A	Probabilistic Inference in <i>MEP</i>	15
A.1	Initialization	16
A.2	Expectation Step	16
A.3	Maximization Step	17
A.3.1	Optimizing α (Q_1)	18
A.3.2	Optimizing \mathbf{s}, σ^2 (Q_2)	18
A.3.3	Optimizing θ^i (Q_3)	20
B	Experimental Details	21

1 Introduction

Enhancing individual and organizational performance through information technology is one of Information Systems’ (IS) fundamental promises – a promise on which the discipline has delivered in many important ways. Consumers can now choose among an unprecedented variety of affordable products and services online; and managers routinely make decisions based on rich, differentiated real-time information about their organizations and global markets, to name just two significant developments of recent years. But while information systems have become increasingly adept at provisioning decision-relevant information and supporting human decision-makers, *autonomous decision-making* remains elusive under all but the most highly structured circumstances.

Dynamic flight pricing and automated credit approvals are two prominent examples of processes that have been successfully codified to the point where business rule engines are now making most operative decisions quickly, cheaply, and reliably [Davenport and Harris, 2005]. But in less clearly structured settings – from booking the next vacation to trading in complex multi-echelon markets – computer-assisted human decision-making remains the gold standard. Delegating decision-making to an autonomous software agent in these settings is often not so much hampered by the agent’s inability to select a good decision, as it is by an inability to *determine what constitutes a good decision* in the user’s eyes to begin with. In many cases, the autonomous decision-making challenge amounts to a challenge in eliciting and faithfully representing human preferences using a computational model [Bichler et al., 2010].

To be useful, the preference model employed by an autonomous software agent must meet a broad range of requirements. Most importantly, it must respect the numerous systematic inconsistencies in human preferences that psychologists have discovered and carefully documented over the last decades [Slovic, 1995, Lichtenstein and Slovic, 2006]. The uncertainty that these inconsistencies give rise to are a crucial ingredient for autonomous decision-making: an agent should only make high-value decisions autonomously if past evidence suggests that its actions will be correct with high probability; otherwise, the agent should prompt its user for additional information. Allowing the *threshold* for user involvement to change over time is referred to as adjustable autonomy [Tambe et al., 2002]. Adjustable autonomy is an important determinant of users’ trust in autonomous agents, and a good preference model will provide estimates of its own uncertainty to facilitate it. Finally, autonomous agents cannot demand their user’s attention too frequently and the preference models they employ must therefore begin to make accurate predictions after observing limited training data. In an information systems setting, this means that a preference model should be (a) flexible enough to learn from multiple types of data (e.g., observed click streams and explicit critiques of alternatives by the user), and (b) pool data from like-minded users in a principled way [Ansari et al., 2000].

Our main contribution is *MEP*, a novel **Mixed-Effect Preference** model that meets these requirements. *MEP* is a semi-parametric Bayesian model that quantifies the uncertainty of its own predictions based on probabilistic principles. The Bayesian framework allows us to cleanly separate between evidence and inference and, as such, to integrate multiple types of preference information incrementally, an essential property in information system settings. The type of preference information that we focus on are pairwise comparisons of the form “User U prefers alternative A over B”, whereas much previous work has focussed on ratings [Adomavicius and Tuzhilin, 2005a].

While reasoning based on ratings is significantly easier, there are three important motivations for studying pairwise comparisons: First, pairwise comparisons are more informative in the sense that they do not *force* users into self-consistency. Whereas rating systems suppress, e.g., intransitivity, *MEP* embraces inconsistencies and lets them unfold in the form of uncertainties. Second, pairwise choice situations are cognitively easier on users and evoke a qualitative reasoning mode in contrast to the quantitative reasoning evoked by ratings, making them an interesting subject of study in their own right [Slovic, 1995]. And third, pairwise choices are

often readily derived from users’ actions in information systems which reduces the user’s burden compared to providing explicit training information.

MEP further reduces users’ burden by systematically pooling observations across like-minded users, thus reducing the amount of information required from each individual. In the process, it discovers naturally occurring segments of users based on their choices. Regular segmentation techniques proceed by first segmenting users heuristically (based on approximations such as demographics, or heuristic global search) and then building a separate model for each segment [Jiang and Tuzhilin, 2009], but user demographics and other statistics are inherently limited in their ability to truly reflect user preferences, e.g., [Reinecke and Bernstein, 2012]. In contrast, in *MEP* segments are *induced* by choices which gives rise to a principled new approach to user segmentation. Importantly, inference in our model remains computationally within reach of deterministic approximations, whereas previous work had to resort to slower sampling methods, e.g., [Ansari et al., 2000].

We demonstrate the qualities of *MEP* in experiments on a real-world dataset collected specifically for this study. Earlier work has measured pairwise preference learning performance mostly on converted rating or regression datasets. But these tasks are fundamentally different from pairwise comparisons and the resulting datasets are unlikely to reflect human choice behavior well. We acquired our data through Amazon Mechanical Turk, a commercial crowdsourcing platform, a practice that is gradually becoming more popular in psychology and the social sciences, but that we have so far not encountered in preference learning. As such, our data collection methodology is an interesting contribution in its own right.

The work we present here is related to an earlier study [Peters and Ketter, 2012] that was inspired by the same problem, but that used a more heuristic approach to preference modeling. In contrast, in this article we give a rigorous mathematical derivation of our mixed-effect preference model, complete with formal convergence guarantees. *MEP* now employs a fully probabilistic segmentation (“soft clustering”) whereas our previous work was limited to deterministic segment assignments.

The remainder of this article is structured as follows: In Section 2 we present related work in Computer Science, Econometrics, Information Systems, Marketing, and Psychology. *MEP* is introduced in Section 3 and experimentally evaluated in Section 4. We discuss our work, and conclude with directions for future research in Section 5.

2 Related Work

Human preferences have long been a subject of interest for scholars from a multitude of fields including Information Systems, Computer Science, Psychology, Marketing, and Econometrics. Much like our study is biased by our own Information Systems *weltanschauung*, these works have looked at human preferences from a diverse set of angles, from descriptive to normative. Before introducing our own preference model in Section 3, we first consider selected contributions made in each of these fields, and we discuss their relationship to our work. In the interest of space, we will not strive for comprehensiveness. Instead, the idea that we wish to convey is that preferences are too rich and faceted as a phenomenon to admit mono-disciplinary study; scholars from all aforementioned fields are now embracing this insight to develop models and theories that transcend the traditional boundaries of their disciplines.

2.1 Information Systems

Within IS, our work is most closely related to the literature on **recommender systems**, which are concerned with finding and recommending items of interest, e.g., in electronic commerce settings [Adomavicius and Tuzhilin, 2005a]. Our work shares the design science perspective [Hevner et al., 2004] that is prominent in this literature in the sense that we aim to design an artifact (a

preference model), the ultimate purpose of which is to be *useful*; we also adopt several dimensions of usefulness that are important in recommender systems. Specifically, our model must be computationally tractable, suitable for operation within real-world information systems, predict preferred alternatives accurately, and be able to operate on limited input data. In other regards, the problem we study is different and in some sense harder than the recommendation problem:

1. Our model targets *autonomous decision-making* settings instead of giving *recommendations*, a difference that is reflected in the probabilistic design of our model. Autonomous decisions are becoming increasingly important in fast-paced, data-intensive electronic marketplaces [Peters et al., 2013] and we contend that these settings will require a principled quantification of uncertainty that allows software agents to return control to their users in cases of doubt. Some previous recommender systems studies make use of probabilistic models in focussed application areas (e.g., click-stream mining [Manavoglu et al., 2003]), but these works are not easily generalized to the autonomous decision-making task we consider.
2. Most recommender systems operate based on *ratings* which have several drawbacks in the context of autonomous decision-making. We already pointed out that ratings force users into self-consistency, invoke quantitative reasoning patterns, and are potentially intrusive. Moreover, ratings provide a comparatively coarse image of user preferences, a problem that is acknowledged in the development of multi-criteria recommender systems [Manouselis and Costopoulou, 2007]. And finally, ratings are subjective in the sense that two decision-makers may share the same preferences but assign different ratings to alternatives [Jin et al., 2002].

A particularly interesting development in the recommender systems community is the increasing adoption of psychological notions like time variability, iterative refinement, or context dependencies [Palmisano et al., 2008, Adomavicius and Tuzhilin, 2011]. Our model contributes to this discussion through its ability to detect inconsistencies in preferences stated at different times or in different contexts.

Other related research includes work on personalization [Murthi and Sarkar, 2003, Adomavicius and Tuzhilin, 2005b] which studies methods for adapting products and services, often in digital environments where associated costs are low. For example, [Atahan and Sarkar, 2011] propose a Bayesian method for the accelerated learning of user profiles through dynamic adaptation of navigation options. Several other IS scholars have worked on methods that span across multiple academic disciplines. For example, [Huang et al., 2012] propose a discrete choice prediction method based on Machine Learning techniques, and [Roy et al., 2008] combine ideas from Multi-Criteria Decision Analysis (MCDA) and Operations Research to learn utility functions. These works are related to ours in isolated aspects, but they generally follow different objectives; to our knowledge, our study is the first to combine probabilistic estimates, systematic data pooling across multiple users, flexibility with respect to different types of preference data, and a principled approach to segmentation in one coherent method.

2.2 Computer Science

Researchers in **Artificial Intelligence** (AI) have a long-standing interest in preferences as a flexible alternative to the notion of goals, reviewed in [Brafman and Domshlak, 2009]. It may not be clear from the outset whether a certain *goal* is attainable, but using a *preference relation*, an intelligent agent can work towards the most favorable outcome that is attainable under given circumstances. AI researchers realized the importance of preferences for autonomous decision-making agents early on [Maes, 1994, Brachman, 2002], and they proposed a series of innovative representations for preference relations and associated reasoning schemes, e.g., [Boutilier et al., 2004]. However, most of these representation are based on a markedly normative understanding

of human preferences. That is, they make strong rationality assumptions and reasoning consequently becomes difficult when they are confronted with actual human preferences, e.g., [Chevalleyre et al., 2011].

Machine Learning (ML), on the other hand, takes a data-driven, predictive approach to preferences, and researchers in the field have proposed models based on a broad variety of learning techniques [Fürnkranz and Hüllermeier, 2011]. Of particular interest to us is work on probabilistic preference models that yield uncertainty estimates based on noisy preference observations. [Chu and Ghahramani, 2005b] were the first to model preference learning as a problem in probabilistic ordinal regression, and their work has spawned a series of innovative extensions that, for example, incorporate relational information [Kersting and Xu, 2009, Xu et al., 2010] or improve scalability [Nielsen et al., 2011]. What is absent from these studies, however, is the ability to pool preferences across like-minded users to learn from limited information, or multi-task learning in ML terminology [Caruana, 1997]. [Bonilla et al., 2010] propose that multi-task preference learning can be realized by incorporating user information into preference statements, [Houlsby et al., 2012] propose a mixture of Gaussian processes model, and [Birlutiu et al., 2012] couple several single-user tasks through a hierarchical prior. But these approaches are either not applicable to our scenario (because user information may not be available), or they lack the segmentation based on stated preference that we aim for.

2.3 Psychology

Scholars in Psychology have primarily taken an explanatory interest in human preferences. One of the key achievements of the discipline has been the establishment of positive theories of human preferences [Tversky and Kahneman, 1992, Slovic, 1995, Thaler, 2000, Lichtenstein and Slovic, 2006] next to the normative theories from Economics [von Neumann et al., 2007]. These positive theories attest that human preferences are inconsistent, constructed as needed, and that the construction process depends on situational framing and environmental factors, among other influences. The probabilistic nature of our model acknowledges such inconsistencies and lets them unfold in the form of uncertainties. They are an important output of our model: an autonomous decision-making agent can use them to return control to its user when uncertain about a high-value decision.

2.4 Marketing and Econometrics

Finally, our work is related to the Marketing and Econometrics literature where **preference measurement** methods such as conjoint analysis, logit/probit models, and several other discrete choice prediction techniques [Greene, 2012] were pioneered. Early preference measurement was limited to population-level estimates, but more recent techniques accommodate heterogeneity across consumer segments [Allenby and Rossi, 1998, Evgeniou et al., 2007]. The primary target audience of these models are human decision makers, and their outputs are interpretable coefficients derived from single passes over full datasets. Not of concern in these methods are incremental updates, scalability, and other practical issues that arise when moving from passive preference measurements to active, autonomous decision-making [Netzer et al., 2008].

3 Mixed-Effect Preference Model

We now turn to deriving our mixed-effect preference model *MEP* in three steps: Section 3.1 introduces the required notation and our modeling objective. Section 3.2 gives a brief overview of Gaussian processes and their use in preference modeling for the case of a single user. And Section 3.3 extends the model to the multi user case, the core contribution of our work. We will focus on the intuitions behind our model in the main text and relegate most of the mathematical details to the appendix.

3.1 Preliminaries

Let $\mathbf{U} = \{u_1, u_2, \dots, u_{n_u}\}$ denote the set of **users** and $\mathbf{X} = \{x_1, x_2, \dots, x_{n_x} | x_i \in \mathbb{R}^{d_x}\}$ the set of **instances**, the objects or actions over which preferences are formed, where each instance is described by d_x real-valued attributes. For a given user $u \in \mathbf{U}$, the set of pairwise **choices** can be written as

$$\mathbf{D}_u = \left\{ (u, x_{u,j}^+, x_{u,j}^-) | 1 \leq j \leq n_{\mathbf{D}_u}, x_{u,j}^+, x_{u,j}^- \in \mathbf{X} \right\}$$

where (u, x^+, x^-) means that user u prefers instance x^+ over x^- , or $x^+ \succ_u x^-$ for short. Commonly, the goal in preference learning is to learn for each user u a total order \succeq_u over instances. Instead of operating directly on these orders \succeq_u , we consider latent **utility functions** $f_u : \mathbf{X} \rightarrow \mathbb{R}$ that assign a real number to each instance. The learning task then becomes finding functions f_u such that

$$f_u(x_{i_1}) \geq f_u(x_{i_2}) \geq \dots \geq f_u(x_{i_n}) \Leftrightarrow x_{i_1} \succeq_u x_{i_2} \succeq_u \dots \succeq_u x_{i_n}$$

at least in approximation. We summarize all mathematical notation used throughout this article in Table 1.

We argued above that human preferences are subject to numerous inconsistencies. From a modeling perspective, this entails that

1. in absence of prior information about the shape of the f_u , the population that these functions are selected from must be as broad as possible, so as to not exclude potentially interesting candidates.
2. we need to introduce *distributions* over functions that properly capture the ensuing uncertainty rather than point estimates. Specifically, given instances \mathbf{X} and choices \mathbf{D}_u we are interested in distributions $P(f_u | \mathbf{D}_u, \mathbf{X})$ over functions that quantify how likely it is that f_u is the true utility function that generated the observed choices.

Both these requirements are met by **Gaussian processes** (GPs), a type of stochastic process commonly used for representing distributions over functions in Bayesian modeling [Rasmussen and Williams, 2006, MacKay, 1998]. GPs are sets of random variables, any finite subset of which has a joint multivariate Gaussian distribution. In our case, these random variables are utility function evaluations at sets of instances; by always conditioning on finite sets of instances, a seemingly complicated inference problem over function distributions is reduced to manipulating multivariate Gaussians in accordance with the observed data. What is unclear, is how this manipulation should proceed given choices \mathbf{D}_u . Somewhat more formally, we are interested in computing posterior distributions

$$P(f_u | \mathbf{D}_u, \mathbf{X}) \text{ where } f_u \sim \mathcal{GP}(\mu_u, \mathbf{K}_u), u \in \mathbf{U}$$

over utility functions f_u given the data. As usual for multivariate Gaussians, the posterior distributions are fully characterized by their means μ_u and covariance matrices \mathbf{K}_u , but it is unclear how qualitative and possibly inconsistent constraints such as *the function value at instance x^+ should be higher than at x^-* (if the user stated that $x^+ \succ_u x^-$) should be incorporated into these parameters.

3.2 Modeling Single-User Preferences

Consider the choices \mathbf{D}_u of a single user u . Following [Chu and Ghahramani, 2005a,b] we model the preference learning task for user u as a problem in ordinal regression. In the Bayesian framework, this problem amounts to finding the Gaussian process posterior $P(f_u | \mathbf{D}_u, \mathbf{X}) \propto P(\mathbf{D}_u | f_u, \mathbf{X}) \times P(f_u)$ where $P(f_u)$ is a GP prior and $P(\mathbf{D}_u | f_u, \mathbf{X})$ denotes the likelihood of actually observing the choices in \mathbf{D}_u when the true utility function is f_u . Key contributions

Table 1: Summary of mathematical notation

Symbol	Interpretation
α	Segment probabilities for n_s segments, see also z_u
$\gamma_{u,k} = E[z_u = k \mathbf{D}_u, \mathbf{X}_u, \Theta^{t-1}]$	Fractional membership of user u in segment k
$\mathbf{D}_u = \{(u, x_{u,j}^+, x_{u,j}^-)\}$	Stated pairwise choices
$f_u = i_u + s_{z_u}$	Sum process for user u , sum of individual-level process i_u and most likely segment-level process s_{z_u}
$g \sim \mathcal{GP}(\mu, \mathbf{K})$	g follows a Gaussian process with mean μ and covariance \mathbf{K}
$i_u, \{i_u\}$	Individual-level process i_u for user u , and the set of all individual-level processes, respectively
\mathbf{i}	Evaluation of all individual-level processes $\{i_u\}$ at all instances \mathbf{X}
$N(x)$	Probability density function (PDF) of a standard Gaussian
$\Phi(x)$	Cumulative distribution function (CDF) of a standard Gaussian
Q^t, Q^{t-1}	Expectation of complete-data log likelihood at EM iteration t and $t - 1$, respectively
σ^2	Observational noise variance
$s_k, \{s_k\}$	Segment-level process s_k for segment k , and the set of all segment-level processes, respectively
\mathbf{s}	Evaluation of all segment-level processes $\{s_k\}$ at all instances \mathbf{X}
θ_u^i	Individual-level kernel hyperparameters for user u
θ^s	Segment-level kernel hyperparameters
$\Theta = \{\alpha, \mathbf{s}, \sigma^2, \theta_u^i\}, \Theta^{t-1}$	Collection of all EM parameters at iteration t and $t - 1$, respectively
$\mathbf{U} = \{u_1, u_2, \dots, u_{n_u}\}$	Users
$\mathbf{X} = \{x_1, x_2, \dots, x_{n_x} x_i \in \mathbb{R}^{d_x}\}$	Instances; alternatives or actions over which preferences are formed
x^+, x^-	Some preferred and dispreferred instance, respectively
$x^+ \succ_u x^-$	User u prefers x^+ over x^-
$w_i, W_{i,j}$	Gauss Hermite quadrature weights in one and two dimensions, respectively
$z_u \sim \text{Multi}(\alpha)$	Probabilistic segment assignment of user u to one of n_s segments
\mathbf{z}	Evaluation of all z_u




Figure 1: *GPPREF* result for a single user; the thick black line depicts the posterior mean $E[f_u|\mathbf{D}_u, \mathbf{X}]$ of the GP; dashed lines represent pairwise choices (training examples, right axis). For example, the leftmost dashed line states that the user preferred a 5.25" floppy disk (upper end of the line) to a punchcard (lower end); the dashed lines' thickness indicates the certainty with which the trained model reproduces the training examples. All training examples are reconstructed correctly, but the regression indicates lower certainty for the unusual training example on the far right.

of [Chu and Ghahramani, 2005b] are the definition of a plausible likelihood function for pairwise choice data, and the description of a tractable procedure for approximating the posterior $P(f_u|\mathbf{D}_u, \mathbf{X})$ as a multivariate Gaussian. One strength of the Bayesian framework is the clean separation between evidence and beliefs and, as such, other types of choice data (e.g., n-ary discrete choices) can in principle also be accommodated. We refer the interested reader to their work for technical details; in the following, we abbreviate this standard form of Gaussian process preference regression as *GPPREF*.

Figure 1 gives an example of applying *GPPREF* to pairwise choices among storage technologies from a single user. The thick black line depicts the mean of the GP posterior, $\mu_u = E[f_u|\mathbf{D}_u, \mathbf{X}]$. Note, that all training examples are correctly reproduced by this line. For example, the leftmost training example ($5.25'' \succ_u \text{Punchcard}$) is correctly reproduced because the posterior mean at $5.25''$ is greater than at Punchcard . A key advantage of *GPPREF* for our purposes is its principled approach to estimating predictive uncertainties. In the figure, the thickness of the choice lines captures the certainty with which they are reproduced by the trained model. Interestingly, most choices are reproduced with high certainty except for the unusual example on the far right, which conflicts with the user's general preference for more advanced storage technology. In an autonomous decision-making context, an agent could use this information to prompt the user to either revise or confirm the conflicting choice $\text{DVD-R} \succ_u \text{Quantum}$. Despite its advantages, *GPPREF* is limited when it comes to typical information systems scenarios with multiple users, because it estimates only a single posterior distribution over utility functions. To resolve this, one of two naïve approaches can be followed:

1. Infer **one utility function for all users** simultaneously. This approach works well as long as all users share similar preferences. It fails, however, in the more probable case of heterogeneity among users. Figure 2 illustrates this situation for the two user case.
2. Infer **one utility function for each user**. This approach alleviates the problem of conflicting preferences, but it only works if enough choices have been observed for *each user*. In other words, estimating one utility function per user misses out on an important opportunity to borrow statistical strength from like-minded users.

For our purposes, both solutions are unsatisfactory since we are looking to combine high degrees




Figure 2: *GPPREF* result for two users with conflicting preferences; even though many training examples are reconstructed correctly, the regression indicates low certainty across all training examples, and pronounced difficulties in finding a utility function that generalizes across both users.

of certainty and accuracy with minimal intrusiveness. [Jiang and Tuzhilin, 2006] have shown that, for a broad range of models, the best results in this regard are to be expected *between* individual-level and population models, at the level of user **segments**. While it is possible to segment the user population based on, e.g., demographics and then use *GPPREF* on the resulting segments, this solution is unsatisfactory for both theoretical and practical reasons. On a theoretical level, demographics are likely too coarse to be a good approximation of users’ like-mindedness. It is therefore unclear whether segments derived in this way will lead to the intended gain in statistical strength. From a practical perspective, demographics may simply not be available, e.g., in online settings. A better solution would be to *induce* segments based on users’ choices, but it is unclear how such an induction could proceed using pairwise choices or similarly unwieldy types of data.

3.3 Modeling Multi-User Preferences

Our mixed-effect preference model *MEP* facilitates probabilistic preference estimation at the segment level and it segments users based on their choices in the process. The general idea behind *MEP* is that the latent utility function f_u for each user u is partly explained by the effects of the user’s association with one of n_s segments (segment-level effects) and partly by idiosyncratic effects of the user himself (individual-level effects). More formally

$$f_u(x) = s_{z_u}(x) + i_u(x), \text{ for } u \in \mathbf{U} \quad (1)$$

where $z_u \in \{1, \dots, n_s\}$ denotes the membership of user u in one of the n_s segments, and the $\{s_{z_u}\}$ and $\{i_u\}$ are Gaussian processes which capture the segment-level and individual-level effects, respectively.

While this formulation may look superficially similar to discrete choice models, it is important to note that the $\{f_u\}$, $\{s_{z_u}\}$, and $\{i_u\}$ in our notation are *stochastic processes*. Most discrete choice models assume that users’ latent utility functions come from a predetermined parametric class of functions (e.g., the class of linear functions), and they estimate the concrete parameters of the hypothesized utility function based on observed choices. In contrast, our model computes a full distribution over a broad range of utility functions without presupposing a certain function class. This formulation allows us to sidestep the question of the “true” utility function’s shape, and it yields predictive uncertainty estimates that are crucial in an autonomous decision-making context.

Figure 3 shows the probabilistic graphical model for *MEP*. Given instances \mathbf{X} and choices \mathbf{D}_u , our goal is to estimate a probabilistic model that describes how the \mathbf{D}_u are generated for each user $u \in \mathbf{U}$. Working backwards from the \mathbf{D}_u , this generative process proceeds as follows:

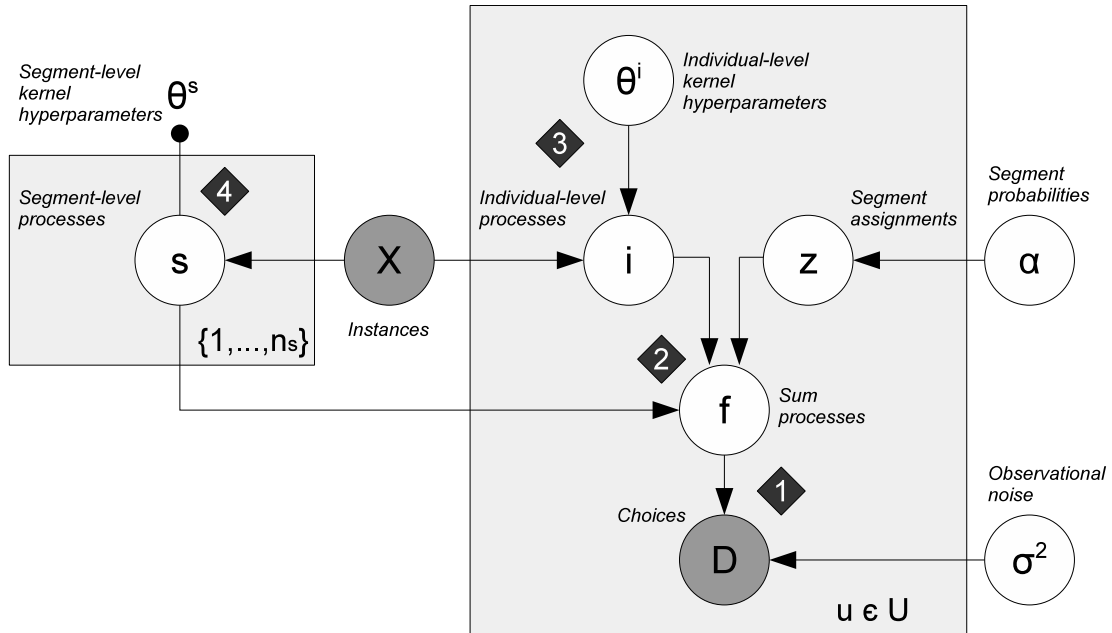


Figure 3: Probabilistic graphical model for *MEP*; circles indicate random variables; dark coloring indicates observed data; arrows denote probabilistic dependence; shaded *plates* are instantiated for each segment or user, respectively

- ❖ Observed choices in \mathbf{D}_u are generated by utility functions $f_u \sim \mathcal{GP}$. In an ideal world, a user's choices could simply be read off any concrete utility function drawn from $P(f_u|\cdot)$. But in reality, choices will be distorted by behavioral influences and by several types of measurement errors (e.g., when users accidentally make unintended choices). These deviations are captured in the noise parameter σ^2 .
- ❖ Utility functions f_u are generated by summing an individual-level effect i_u and a segment-level effect s_{z_u} . We model the membership of a user u in segment k as a multinomial distribution with n_s segments and segment probabilities α , that is, $z_u \sim \text{Multi}(\alpha)$. The number of segments is a user-defined input to our model. Alternatively, it could be modeled probabilistically as a Dirichlet process [Ghahramani, 2012] at the expense of additional computational cost. But, as we will show in our experiments, *MEP* is reasonably robust to mis-specification of n_s and we therefore prefer this simpler formulation.
- ❖ Individual-level effects $i_u \sim \mathcal{GP}(\mu_u, \mathbf{K}_u^i)$ are generated by GPs with kernel functions $\mathbf{K}_u^i(\theta_u^i)$ that describe the processes' covariance structure. Previous researchers have successfully used Gaussian kernels for preference estimation, e.g., [Birlutiu et al., 2012], and we follow this choice in our experiments below, while noting that our model can accommodate any other kind of kernel as well. The kernel functions at the individual level are parametrized with parameters θ_u^i that are estimated from the data.
- ❖ Segment-level effects $s_k \sim \mathcal{GP}(\mu_k, \mathbf{K}^s)$ are generated by GPs with a common kernel function $\mathbf{K}^s(\theta^s)$, the parameters of which are a user-defined input to the model. Generally speaking, the segment-level parameters will be chosen such that the s_k are more adaptive since they are estimated based on all of a segment's users, whereas the θ_u^i are estimated

to emphasize each user’s idiosyncrasies. The fixed choice of θ^s allows us to simplify the inference task; below, we will only consider the maximum a-posteriori (MAP) estimates for the s_k while maintaining full distributions over the i_u . The Gaussian process formulation can in this case be thought of as a regularization on the realizations of the s_k [Lu et al., 2008, Wang et al., 2010].

In sum, we are given instances $\mathbf{X} = \{\mathbf{X}_u\}$, choices $\mathbf{D} = \{\mathbf{D}_u\}$, the number of segments n_s , and the segment-level kernel function \mathbf{K}^s with its parameters θ^s . What we are interested in finding is the maximum likelihood estimate Θ^* for $\Theta = \{\alpha, \mathbf{s}, \sigma^2, \theta_u^i\}$ where $\mathbf{s} = \{s_k(\mathbf{X})\}$ denotes a random vector containing the evaluation of the segment-level GPs s_k at all instances $x \in \mathbf{X}$. More formally we aim to find

$$\Theta^* = \arg \max_{\Theta} P(\mathbf{D}|\mathbf{X}, \Theta)P(\mathbf{s}|\theta^s) \quad (2)$$

Once we have estimated Θ^* , we can answer a number of interesting queries, such as

- What is the probability that user u prefers instance x^+ to instance x^- ? Here, x^+ and x^- may or may not have been referenced in earlier choices by the user. In fact, x^+ and x^- may be entirely new in the sense that *no user* has made any earlier statement about them.
- What is the probability that user u belongs to segment k ; or what is the most likely association between users and segments? Which users are borderline cases between two or more segments?
- What are the commonalities of all users that are likely associated with segment k ?

Input: Instances $\mathbf{X} = \{X_u\}$, choices $\mathbf{D} = \{\mathbf{D}_u\}$, number of segments n_s , segment-level kernel parameters θ^s

Output: Segment-level effects \mathbf{s} , individual-level kernel parameters θ_u^i , segment assignments z_u , segment probabilities α , observational noise parameter σ^2

$Q^t = Q^{t-1} = -\infty$;

Initialize Θ ; $\Theta^{t-1} = \Theta$;

/* Section A.1 */

repeat

$Q(\Theta, \Theta^{t-1}) = E_{\{\mathbf{i}, \mathbf{z}|\mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} \log [P(\mathbf{D}, \mathbf{i}, \mathbf{z}|\mathbf{X}, \Theta) \times P(\{s_k\}|\theta^s)]$; /* E-Step, Section A.2 */

$\Theta^* = \arg \max_x Q(x, \Theta^{t-1})$;

/* M-Step, Section A.3 */

$Q^{t-1} = Q^t$; $Q^t = Q(\Theta^*, \Theta^{t-1})$;

$\Theta^{t-1} = \Theta$; $\Theta = \Theta^*$;

until $Q^t \approx Q^{t-1}$;

Algorithm 1: Mixed-Effect Preference Model

The direct optimization of the objective function in Equation (2) is intractable and we therefore resort to the Expectation Maximization (EM) algorithm [Dempster et al., 1977]. The hidden variables in our model are $\mathbf{z} = \{z_u\}$ and $\mathbf{i} = \{i_u(\mathbf{X}_u)\}$ where $i_u(\mathbf{X}_u)$ denotes a random vector containing the evaluation of the GP i_u at all instance $x \in \mathbf{X}_u$ that user u has stated something about. The EM algorithm estimates the desired values in an iterative fashion where each step is guaranteed to improve the objective function, and the algorithm is guaranteed to converge to a local optimum. Algorithm 1 shows an outline of this process; full technical details are provided in Appendix A.

We limit the discussion of inference in *MEP* here to an example. Figure 4 shows two *MEP* iterations for three users (columns) and two segments. The choices of each user (training examples) are depicted in the top row. For example, the first user prefers an instance with attribute

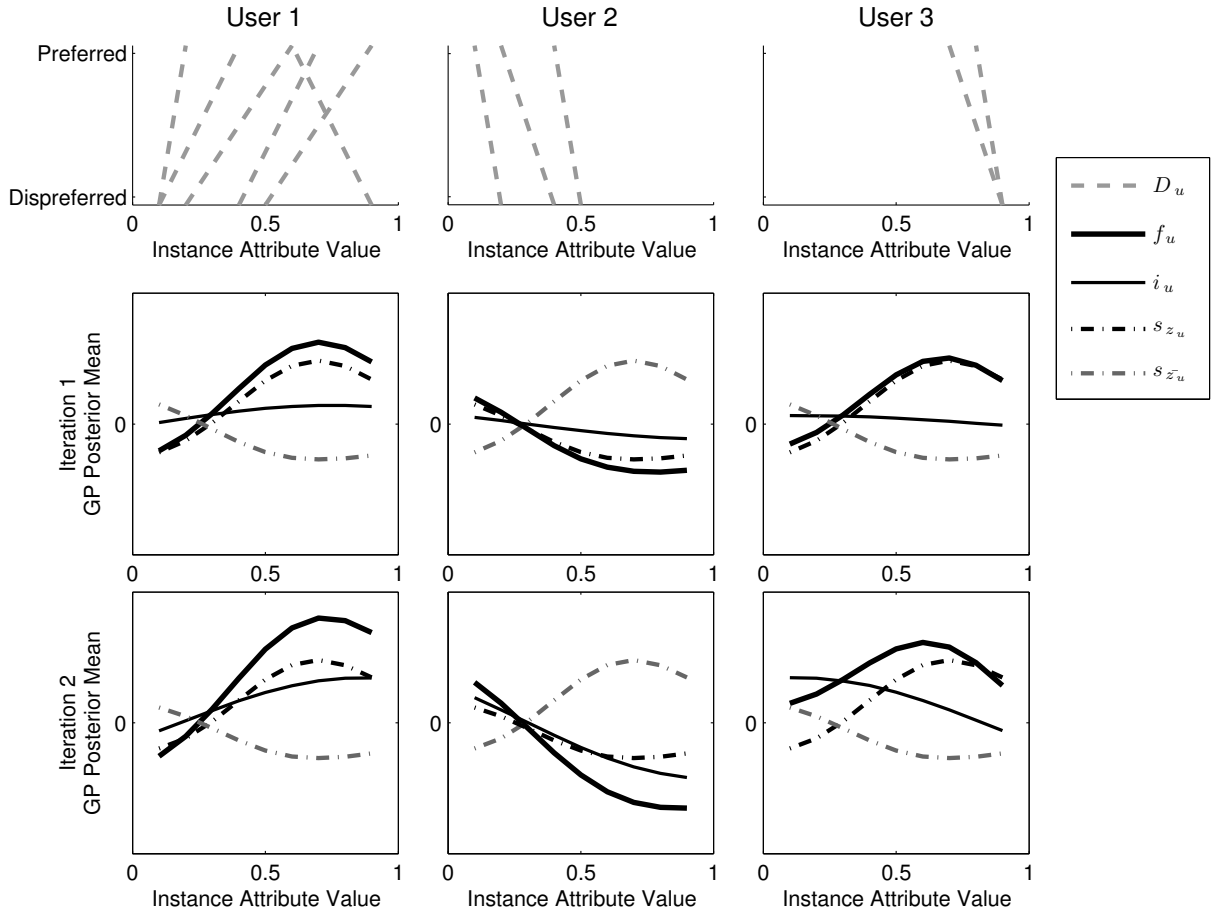


Figure 4: *MEP* sample run for three users and two segments; each column represents one user; top row shows users' pairwise choices, remaining rows show the first two iterations of the algorithm. Dash-dotted lines represent posterior means of the segment-level effects $E[s_k|\mathbf{D}, \mathbf{X}]$ for two segments, where the highest-probability segment for each user is drawn in black. Thin black lines represent the posterior means of the individual-level effects $E[i_u|\mathbf{D}_u, \mathbf{X}]$; thick black lines represent the sum of the individual-level and the most likely segment-level means, see Equation (1).

value 0.2 over an instance with attribute value 0.1, which gives rise to the leftmost dashed line in the top-left panel in Figure 4. The middle and bottom rows of the figure visualize the first two iterations on these input data. There, the dash-dotted lines depict the two segment-level effects s_k . We replicated these effects across all plots in each row and colored the most likely segment assignment for the respective user black. The thin black lines represent the individual-level effects i_u for each user, and the thick black lines additionally include the most probable segment-level effect $f_u = i_u + s_{z_u}$.¹

Consider the first iteration in the middle row of Figure 4. The overall utility f_1 for u_1 is comparable to the result of the single-user regression from Figure 1, but unlike the single-user case much of its general shape has now been absorbed by the segment-level effect. The individual-level effect i_1 only serves to slightly increase the steepness at low attribute values. Similarly, for u_2 much of the downward-sloping tendency suggested by her choices is absorbed by the second segment-level effect. For u_3 finally, only very limited observations are available. In deciding which segment to assign the user to, *MEP* opts for the same segment as u_1 based on the observation that both users share the same preferences with respect to high attribute values. The individual-level effect i_3 amplifies the (more conclusive) evidence for u_3 in that region and adds further to the downward sloping trend towards the right of the plot.

Such generalizations from the individual to the segment level are the driver behind *MEP*'s ability to make inferences from limited observations per user: by inferring that u_3 's preferences so far have been similar to u_1 , it can make more informed statements about u_3 's preferences in the low attribute value region that the users has not stated anything about. Importantly, the *MEP* estimate acknowledges the speculative nature of this generalization: while the segment assignment probabilities (not shown in the figure) are close to certainty for u_1 (89%) and u_2 (84%), *MEP* only deems it *more likely* for u_3 to belong to the same segment as u_1 (62%) than to the other segment (38%, plausible based on u_2 's general preference for lower attribute values). In an autonomous decision-making context, such certainty information is crucial for deciding whether a user should be prompted for additional information or not.

The second iteration in the bottom row of Figure 4 only serves to refine this estimate. During this iteration, the algorithm finds a slightly better allocation of the effects between the individual and the segment level. Specifically, the hump shape in s_1 remains while more of the general upward trend for u_1 and downward trend for u_3 are now captured in the individual-level effects. Similarly, more of u_2 's behavior is now captured by her individual-level effect i_2 as she is the only high-impact member of that segment. This behavior is typical of *MEP*, and we found during our experiments that it usually took three or less EM iterations to arrive at such high-quality estimates.

4 Experimental Evaluation

Given the fragile, constructed nature of human preferences, one can question the evaluation of preference models on converted ranking, rating, or regression datasets as is frequently the case in preference learning studies. For example, ranking tasks evoke a quantitative mode of reasoning in human decision-makers [Slovic, 1995] which is markedly different from the qualitative reasoning mode evoked in pairwise choices.

For this study, we collected a dedicated set of pairwise choice data in accordance with the focus of our model. Our data was collected on Amazon Mechanical Turk (MTurk, <http://www.mturk.com>), a crowdsourcing platform that has recently attracted attention among behavioral scientists [Buhrmester et al., 2011, Mason and Suri, 2012] but that has so far not been used

¹To keep the text uncluttered, we omit the expectation operators and the conditioning of the processes in this example. That is, we refer to the posterior mean $E[s_k|\mathbf{D}, \mathbf{X}]$ of the segment-level processes as s_k , and to the posterior mean $E[i_u|\mathbf{D}_u, \mathbf{X}]$ of the individual-level processes as i_u . The reader should bear in mind, however, that the s_k and i_u actually refer to stochastic processes instead of regular functions.

for the collection of preference data. Several scholars have studied the demographics of MTurk workers [Ipeirotis, 2010] and have proposed guidelines for assuring the quality of data collected through MTurk tasks [Paolacci et al., 2010]. These studies give reason to believe that (1) MTurk data can be of equal or better quality than data selected through channels such as student surveys, (2) MTurk workers are highly diverse (increasing external validity), and (3) the unsupervised nature of MTurk tasks may reduce the risk of experimenter bias (increasing internal validity), all if proper precautions are taken against distractions and random responses.

Our 80 adult American participants were asked to state their preferences in ten pairwise choices between electricity tariffs from the Texan retail electricity market. As a commodity, electricity is exhaustively described using a few quantifiable attributes (price per kilowatt-hour, renewable content, etc.) and unaffected by appearance or workmanship, which makes it a good object for online preference experiments. Each tariff pair was selected at random from a total of 30 tariffs. We also asked participants to answer ten questions on their demographics and electricity consumption behavior, some of which were attention checkers for which the correct answer had to be consistent with an answer given on another question. We restricted the final dataset to 58 participants who had taken three minutes or more to answer all our questions, and had passed all attention checks. Further details on the data collection process are given in Appendix B. Our interest in the experimental evaluation was threefold:

1. **Predictive accuracy** is a key quality indicator for preference models. We measured how well trained *MEP* instances performed on held-out test data compared to *GPPREF* (population level and individual level) for different training set sizes.
2. The second important indicator for the autonomous decision-making task we consider is **predictive certainty**. Note, that predictive certainty is also the key differentiator of our probabilistic approach. Certainty indications should be highly informative in the sense that the model should indicate high certainty on predictions that turn out to be correct and low certainty otherwise.
3. And finally, we were interested in *MEP*'s sensitivity to the user-defined number of segments n_s , and the segment-level kernel hyperparameter θ^s .

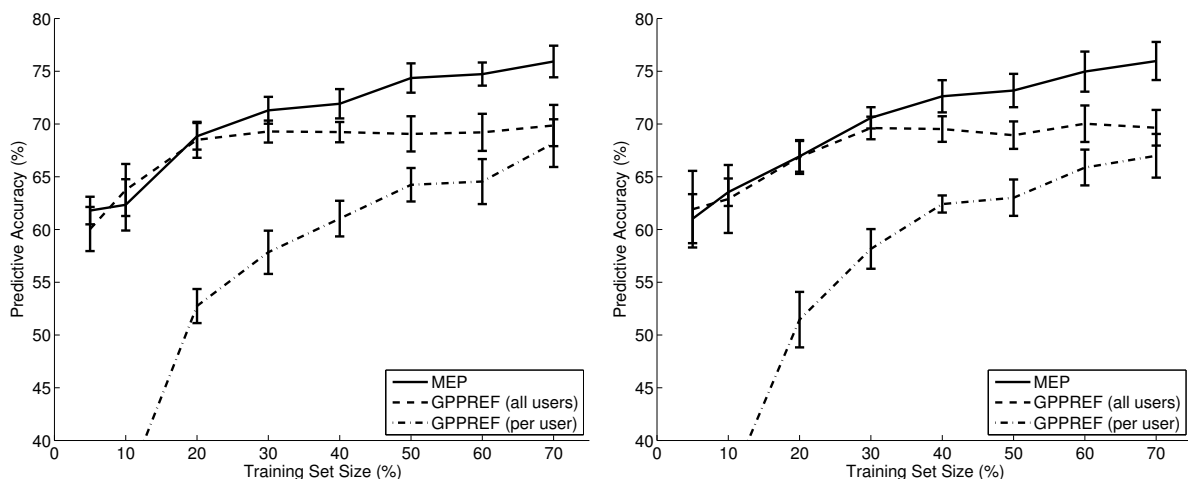


Figure 5: Predictive accuracy for *MEP* versus *GPPREF*; two segments (left) versus five segments (right); error bars show confidence at the 95% level based on 10-fold cross validation; *MEP* performs similar to *GPPREF* (all users) for limited observations and significantly better from about 3-4 pairwise choices per user; *GPPREF* (per user) requires significant input before predicting well and performs worse than random for limited observations

Figure 5 shows predictive accuracies on test data for *MEP* and *GPPREF* when *MEP* is initialized to learn two segments (left) or five segments (right). The horizontal axis shows the percentage of pairwise choices used to train the model; all remaining choices are held out for testing. The 40% point, for examples, shows runs where four out of ten choices per user (on average) are used to train the model. It is unsurprising that predictive accuracy increases as more training data become available. But *MEP* combines reasonable accuracy for cases with very limited data (where at best a general tendency can be estimated) with substantial accuracy growth as more choices are observed. A single *GPPREF* model over all users, on the other hand, levels out for denser datasets because the increasing detail of these heterogeneous data cannot be captured in a single utility function. And estimating separate *GPPREF* models per user delivers unsatisfactory performance for all but the densest training sets, as these localized models do not generalize across users. In fact, the per-user estimates perform worse than random guessing for less than three observations due to strong overfitting. It is interesting to note that all aforementioned results are identical for the two and the five segment case, except for small differences in confidence interval widths.²

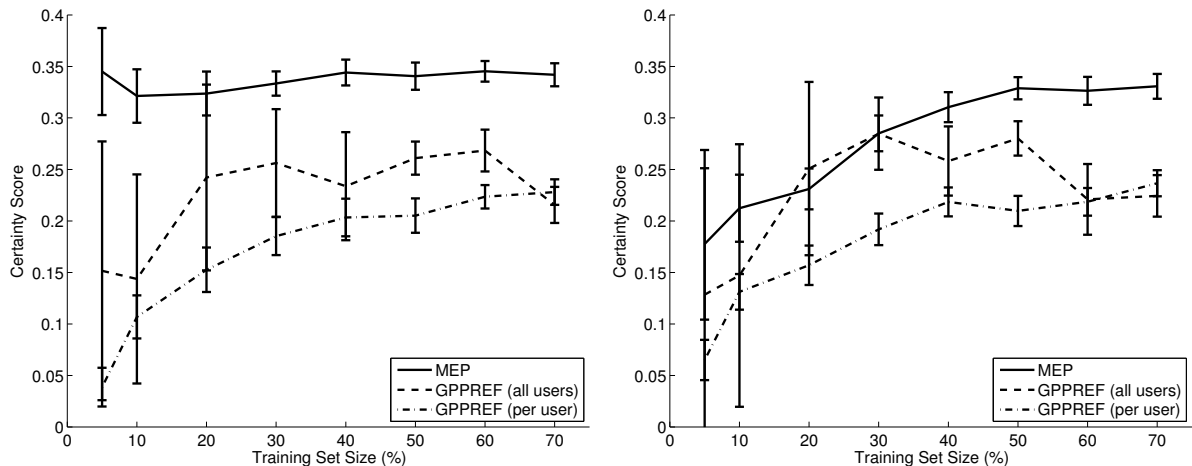


Figure 6: Certainty scores for *MEP* versus *GPPREF*; two segments (left) versus five segments (right); error bars show confidence at the 95% level based on 10-fold cross validation; certainties are strongly affected when misestimating segment numbers for limited observations, but stabilize at about 3-4 pairwise preferences per user

Figure 6 shows similar graphs for the **certainty scores** of the different models, where higher scores are better.³ *MEP*'s certainty estimates are highly informative in the sense that it only indicates certainty on predictions that later turn out to be correct. Interestingly, certainty scores for very limited data are more strongly affected by the of choice segments n_s than predictive accuracies. This effect is, however, simply caused by distributing very limited evidence across many segments. Beginning at about four to five training samples per user, this choice becomes inconsequential and *MEP*'s certainty estimates are again significantly more informative than those given by other models.

²We also examined *MEP*'s robustness with respect to different segment-level kernel parameters θ^s and found the same degree of robustness. We omit details in the interest of brevity.

³Our certainty score is designed to capture how informative the predictive certainty output of a model is. For each test sample, we determine how far the model's certainty indication lies from the indifference point (50%). For predictions that later turn out to be correct, the certainty indicated by the model should have been high (ideally 100%), whereas for predictions that turn out to be incorrect the opposite should be the case (ideally 0%). The certainty score is the average of these contributions over all test samples, and it therefore ranges from -0.5 (worst) to +0.5 (best).

5 Conclusions and Future Work

Models of human preferences will play an important role in the personalized, autonomous information systems of the future. Useful preference models have to learn from multiple types of observations; make inferences based on limited, unobtrusively collected data; and yield confidence estimates so that autonomous agents can return control to their users when unsure about high-value decisions. We are gradually beginning to understand how preferences are formed based on insights from Psychology and Marketing. But the best design for a computational model that captures such inherently fragile preferences while satisfying the aforementioned requirements remains to be found.

Our preference model *MEP* is the first to combine probabilistic estimates, systematic data pooling across multiple users, flexibility with respect to different types of preference data, and a principled approach to segmentation in one coherent method. We demonstrated *MEP*'s ability to deliver more accurate predictions and more informative certainty indications than previous probabilistic models for pairwise choices. Importantly, we obtained our results on real-world preference data instead of converted rating datasets with artificially high internal consistency. This mirrors our intention to embrace inconsistencies, and to let them unfold in the form of uncertainties that can guide autonomous agents. We acquired this data through Amazon Mechanical Turk, a commercial crowdsourcing platform, a practice which we have not previously encountered in preference learning. Our work confirms the finding of [Jiang and Tuzhilin, 2006] that the best predictive performances can be expected at the segment-level for the case of probabilistic preference models.

The work presented in this article lays the foundation for a number of interesting applications and extensions:

1. The clean separation between the internal GP representation and the observations inherited from *GPPREF* allows us to learn from other types of information besides pairwise choices. For example, choices between more than two alternatives, and contextual and temporal information could all be used to accelerate learning and improve predictive certainty.
2. The principled uncertainty estimates produced by *MEP* can be used for active learning schemes that determine the highest-value question to ask when eliciting preferences interactively, thereby reducing the user's burden, e.g., [Birlutiu et al., 2012, Saar-Tsechansky et al., 2009, Houlisby et al., 2012]. Uncertainty estimates can also serve to locate inconsistencies in user preferences over time. Such inconsistencies can point to regime shifts [Ketter et al., 2012] caused by changing tastes or more mundane causes like multiple users working from one account [Murthi and Sarkar, 2003]. We speculate that our data-driven approach to segmentation can help to identify and resolve such situations.
3. One disadvantage inherent in naïve Gaussian process inference is its cubic complexity in the number of instances (on which the user made choices, $|\mathbf{X}_u|$). To make *MEP* practically applicable we are evaluating approximation schemes that overcome this deficit. Several approaches have been proposed, e.g., [Snelson and Ghahramani, 2006, Quinonero-Candela et al., 2007, Qi et al., 2010], and we are evaluating the speed versus quality tradeoffs they entail for our application.

Moreover, behavioral quality indicators, such as users' trust in the model's predictions, should be investigated as starting point for further improvements in user preference models, e.g., [Häubl and Trifts, 2000, Häubl and Murray, 2003, Lee and Benbasat, 2011].

Our work contributes to the foundations for next-generation information systems, and it constitutes an important step towards systems that act autonomously on their users' behalf.

Acknowledgements The authors would like to thank Perry Groot and Tom Heskes for insightful discussions on probabilistic preference modeling, and for sharing their implementation of Gaussian process regression with preference likelihoods. We are also grateful for constructive comments on an earlier version of this work from four anonymous WITS 2012 reviewers. Their comments inspired, among other things, the rigorous mathematical derivation that we give for our algorithm in this paper.

APPENDIX

A Probabilistic Inference in *MEP*

We now derive the technical details required for probabilistic inference in *MEP*. Recall from Algorithm 1 in Section 3.3 that we are given as inputs

- instances $\mathbf{X} = \{\mathbf{X}_u\}$,
- choices $\mathbf{D} = \{\mathbf{D}_u\}$,
- the number of segments n_s , and
- the segment-level kernel function \mathbf{K}^s with its parameters θ^s .

What we seek to infer is the maximum likelihood estimate Θ^* for $\Theta = \{\alpha, \mathbf{s}, \sigma^2, \theta_u^i\}$ which we restate here for convenience

$$\Theta^* = \arg \max_{\Theta} P(\mathbf{D}|\mathbf{X}, \Theta)P(\mathbf{s}|\theta^s) \quad (3)$$

There is no analytical solution for this optimization problem and we instead resort to the Expectation Maximization (EM) algorithm [Dempster et al., 1977] with hidden variables $\mathbf{z} = \{z_u\}$ and $\mathbf{i} = \{i_u(\mathbf{X}_u)\}$. Our presentation proceeds in three steps: A.1 discusses the initialization of Θ , A.2 derives the complete-data log-likelihood $Q(\Theta, \Theta^{t-1})$ as required in the expectation step of Algorithm 1, and A.3 explains how to optimize $Q(\Theta, \Theta^{t-1})$ with respect to Θ in the maximization step.

A.1 Initialization

We follow one of the following two approaches in initializing Θ :

1. When no additional data on individual users is available, we first draw n_s function realizations for s_k from a GP prior with given parameter θ^s . We then determine, for each user, the likelihood that his training observations were generated by each of the s_k and initialize α as the weighted segment probabilities over all users.
2. In the case where additional demographics on users are available, it can be leveraged for model initialization as follows: We perform a regular user segmentation by clustering users along their demographic data with an arbitrary clustering method for metric data. We then run a standard *GPPREF* regression (including hyperparameter estimation) on the resulting segments and initialize \mathbf{s} using the estimated posterior means for each segment, and σ^2, θ_u^i from the hyperparameter estimates. The segment probabilities α are initialized as in case (1).

A.2 Expectation Step

In the expectation step, we compute the expectation of the complete-data log-likelihood with respect to the hidden variables \mathbf{z} and \mathbf{i} , or formally

$$Q(\Theta, \Theta^{t-1}) = E_{\{\mathbf{i}, \mathbf{z} | \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} \log [P(\mathbf{D}, \mathbf{i}, \mathbf{z} | \mathbf{X}, \Theta) \times P(\{s_k\} | \theta^s)]$$

Our derivation is similar to that in [Wang et al., 2010] with the exception that they consider the simpler case of real-valued observations distorted by Gaussian noise, whereas in *MEP* we adopt the more unwieldy noise formulation for pairwise choices from [Chu and Ghahramani, 2005b]. First, note that

$$P(\mathbf{i}, \mathbf{z} | \mathbf{X}, \mathbf{D}, \Theta^{t-1}) = \prod_u P(i_u, z_u | \mathbf{X}, \mathbf{D}, \Theta^{t-1})$$

and we have that

$$P(i_u, z_u | \mathbf{X}, \mathbf{D}, \Theta^{t-1}) = P(i_u | z_u, \mathbf{X}_u, \mathbf{D}_u, \Theta^{t-1}) \times P(z_u | \mathbf{X}_u, \mathbf{D}_u, \Theta^{t-1}) \quad (4)$$

Let us consider each of the terms on the right-hand side of Equation (4) separately. The first term amounts to evaluating the posterior of a Gaussian process preference regression with a known segment-level effect s_{z_u} . In other words, we perform a standard *GPPREF* regression where the likelihood terms now include the segment-level effect s_{z_u}

$$P(\mathbf{D}_u | z_u, \mathbf{X}_u, \Theta^{t-1}) = \prod_{j=1}^{n_{\mathbf{D}_u}} P(x_{u,j}^+ \succ_u x_{u,j}^- | (i_u + s_{z_u})(x_{u,j}^+), (i_u + s_{z_u})(x_{u,j}^-)) \quad (5)$$

Importantly, under the Bayesian framework, other types of data and likelihoods can also be accommodated instead of pairwise choices, as long as a tractable procedure for approximating the GP posterior exists. For the second term on the right-hand side of Equation (4) we have by Bayes' rule that

$$P(z_u | \mathbf{X}_u, \mathbf{D}_u, \Theta^{t-1}) \propto P(z_u | \Theta^{t-1}) \times P(\mathbf{D}_u | z_u, \mathbf{X}_u, \Theta^{t-1})$$

where $P(z_u | \Theta^{t-1}) = \text{Multi}(\alpha^{t-1})$ is known, and $P(\mathbf{D}_u | z_u, \mathbf{X}_u, \Theta^{t-1})$ is the likelihood of observing the choices \mathbf{D}_u under the segment-level effect s_{z_u} . These latter likelihoods are already known from the *GPPREF* regression (Equation (5)) above. Combining these results, we have that

$$\begin{aligned} Q &= E_{\{\mathbf{z}, \mathbf{i} | \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(\mathbf{D}, \mathbf{i}, \mathbf{z} | \mathbf{X}, \Theta)] + \log P(\{s_k\} | \theta^s) \\ &= E \left[\log \prod_u \prod_k P(\mathbf{D}_u | i_u, z_u = k, \mathbf{X}, \Theta) P(i_u | z_u = k, \Theta) P(z_u = k | \Theta) \right] \\ &\quad + \log P(\{s_k\} | \theta^s) \end{aligned} \quad (6)$$

Denote the expectation term in Equation (6) as $\tilde{Q} = \tilde{Q}(\Theta, \Theta^{t-1})$.

$$\begin{aligned} \tilde{Q}(\Theta, \Theta^{t-1}) &= E_{\{\mathbf{z}, \mathbf{i} | \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(\mathbf{D}, \mathbf{i}, \mathbf{z} | \mathbf{X}, \Theta)] \\ &= E_{\{\mathbf{z} | \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} E_{\{\mathbf{i} | \mathbf{z}, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(\mathbf{D}, \mathbf{i}, \mathbf{z} | \mathbf{X}, \Theta)] \\ &= \sum_z P(\mathbf{z} | \mathbf{X}, \mathbf{D}, \Theta^{t-1}) \sum_u \sum_k \mathbb{1}[z_u = k] \times \\ &\quad \int dP(i_u | z_u = k) \log [P(\mathbf{D}_u | i_u, z_u = k, \mathbf{X}, \Theta) P(i_u | z_u = k, \Theta) P(z_u = k | \Theta)] \\ &= \sum_u \sum_k \underbrace{\left(\sum_z P(\mathbf{z} | \mathbf{X}, \mathbf{D}, \Theta^{t-1}) \mathbb{1}[z_u = k] \right)}_{\gamma_{u,k}} \times \int \dots \end{aligned}$$

Note that $(\sum_z P(\mathbf{z}|\mathbf{D}, \mathbf{X}, \Theta^{t-1}) \mathbb{1}[z_u = k]) = E[z_u = k|\mathbf{D}_u, \mathbf{X}_u, \Theta^{t-1}]$ can be interpreted as the fractional membership of user u in segment k and we will denote this term as $\gamma_{u,k}$ henceforth. Substituting these results back into Q we obtain

$$\begin{aligned} Q(\Theta, \Theta^{t-1}) &= \log P(\{s_k\}|\theta^s) + \tilde{Q}(\Theta, \Theta^{t-1}) \\ &= \log P(\{s_k\}|\theta^s) + \sum_u \sum_k \gamma_{u,k} \times \\ &\quad \left[E_{\{i_u|z_u=k, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(i_u|z_u = k, \Theta)] + \right. \\ &\quad E_{\{i_u|z_u=k, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(\mathbf{D}_u|i_u, z_u = k, \mathbf{X}, \Theta)] + \\ &\quad \left. E_{\{i_u|z_u=k, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(z_u = k|\Theta)] \right] \end{aligned} \quad (7)$$

A.3 Maximization Step

In the maximization step we aim to optimize

$$\Theta^* = \arg \max_{\Theta} Q(\Theta, \Theta^{t-1})$$

This problem can be decomposed into three separate optimization problems by separating the terms from Equation (7) into $Q = Q_1 + Q_2 + Q_3$ as follows:

$$Q_1(\Theta, \Theta^{t-1}) = \sum_u \sum_k \gamma_{u,k} \times E_{\{i_u|z_u=k, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(z_u = k|\Theta)] \quad (8)$$

$$\begin{aligned} Q_2(\Theta, \Theta^{t-1}) &= \log P(\{s_k\}|\theta^s) + \\ &\quad \sum_u \sum_k \gamma_{u,k} \times E_{\{i_u|z_u=k, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(\mathbf{D}_u|i_u, z_u = k, \mathbf{X}, \Theta)] \end{aligned} \quad (9)$$

$$Q_3(\Theta, \Theta^{t-1}) = \sum_u \sum_k \gamma_{u,k} \times E_{\{i_u|z_u=k, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(i_u|z_u = k, \Theta)] \quad (10)$$

We now address the optimization of each of these terms in turn.

A.3.1 Optimizing α (Q_1)

Note that $P(z_u = k|\Theta) = \alpha_k$ and, hence, optimizing Q_1 amounts to the following constrained optimization problem:

$$\max_{\alpha} \sum_u \sum_k \gamma_{u,k} \times \log \alpha_k \quad \text{s.t.} \quad \sum_k \alpha_k = 1$$

Computing the derivatives of the Lagrangian, we obtain the following system of equations

$$\frac{\partial}{\partial \alpha_i} \left[\sum_u \sum_k [\gamma_{u,k} \times \log \alpha_k] + \lambda \left(\sum_k \alpha_k - 1 \right) \right] = \left[\sum_u \frac{\gamma_{u,i}}{\alpha_i} + \lambda \right]_i \stackrel{!}{=} 0$$

The solution to which is $\alpha_i = -\frac{1}{\lambda} \sum_u \gamma_{u,i}$. We know that

$$1 \stackrel{!}{=} \sum_k \alpha_k = - \sum_k \left(\frac{1}{\lambda} \sum_u \gamma_{u,k} \right) = -\frac{1}{\lambda} \sum_k \sum_u \gamma_{u,k} = -\frac{1}{\lambda} \sum_u 1$$

which means that $\lambda = -n_u$ and

$$\alpha_i = \frac{1}{n_u} \sum_u \gamma_{u,i}$$

A.3.2 Optimizing \mathbf{s}, σ^2 (Q_2)

In the next step, we consider \mathbf{s} to be the $(n_s \times n_x)$ -elements column vector of segment-level effect values where each segment-level effect s_k is evaluated at all $x \in \mathbf{X}$. We denote as \mathbf{s}_k all elements of \mathbf{s} that pertain to the k -th segment, and as $\mathbf{s}_{k,i}$ the single element that pertains to instance x_i and segment k . Expanding the expectation term in Q_2 yields

$$\begin{aligned} & E_{\{i_u | z_u=k, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(\mathbf{D}_u | i_u, z_u = k, \mathbf{X}, \Theta)] \\ &= \int_{i_u} P(i_u | \cdot) \times \sum_{j=1}^{n_{\mathbf{D}_u}} \log P(x_{u,j}^+ \succ_u x_{u,j}^- | \mathbf{s}_k, \sigma^2) di_u \\ &= \int_{i_u} P(i_u | \cdot) \times \sum_{j=1}^{n_{\mathbf{D}_u}} \log \Phi \left(\frac{i_+ + s_+ - i_- - s_-}{\sqrt{2}\sigma} \middle| \mathbf{s}_k, \sigma^2 \right) di_u \end{aligned} \quad (11)$$

where Φ stands for the CDF of a standard Gaussian and the last step in Equation (11) uses the definition of the pairwise choice likelihood from [Chu and Ghahramani, 2005b]. The integral above has no closed-form solution but it can be approximated using Gauss-Hermite quadrature where we take $P(i_u | \cdot)$ to be the envelope function. Note, that Gauss-Hermite quadrature in general would require the computation of n_x nested sums, where n_x is the number of instances or, alternatively, the number of elements in i_u . However, by realizing that (a) each element of \mathbf{D}_u makes use of only two components of i_u (namely its evaluation i_+ at the preferred and i_- at the dispreferred instance), and that (b) the marginal of i_u with respect to these two dimensions is a bivariate Gaussian we can rewrite the integral as

$$\sum_{j=1}^{n_{\mathbf{D}_u}} \int_{d(i_+, i_-)} P(i_+, i_- | \cdot) \times \log \Phi \left(\frac{i_+ + s_+ - i_- - s_-}{\sqrt{2}\sigma} \middle| \mathbf{s}_k, \sigma^2 \right) d(i_+, i_-) \quad (12)$$

where the integral is now an integral over the density of a bivariate Gaussian. One last correction relates to the correlation structure of the (i_+, i_-) . In general, the (i_+, i_-) are jointly Gaussian with mean (μ_+, μ_-) a correlation that is reflected in the posterior covariance matrix of their Gaussian process i_u . Using the transform

- $t_+ = \frac{i_+ - \mu_+}{\sigma_+} \Rightarrow i_+(t_+) = t_+ \sigma_+ + \mu_+$
- $t_- = \frac{i_- - \mu_- - \rho i_+}{\sigma_- \sqrt{1 - \rho^2}} \Rightarrow i_-(t_+, t_-) = t_- \sigma_- \sqrt{1 - \rho^2} + \rho(t_+ \sigma_+ + \mu_+) + \mu_-$

where σ_+^2 , σ_-^2 , and ρ are the variances of i_+ and i_- and their correlation, respectively, we obtain two uncorrelated and independent Gaussians. Following [Varin et al., 2005] we rewrite integral (12) as

$$\begin{aligned} & \frac{1}{2\pi} \sum_{j=1}^{n_{\mathbf{D}_u}} \int_{t_-} \int_{t_+} \log \Phi \left(\frac{i_+(t_+) + s_+ - i_-(t_-) - s_-}{\sqrt{2}\sigma} \middle| \mathbf{s}_k, \sigma^2 \right) e^{-\frac{t_+^2}{2}} e^{-\frac{t_-^2}{2}} dt_+ dt_- \\ & \approx \frac{1}{2\pi} \sum_{j=1}^{n_{\mathbf{D}_u}} \sum_{m_1=1}^M \sum_{m_2=1}^M \log \Phi \left(\underbrace{\frac{i_+(h_{m_1}) + s_+ - i_-(h_{m_2}) - s_-}{\sqrt{2}\sigma}}_{=:v_j} \middle| \mathbf{s}_k, \sigma^2 \right) \times \\ & \quad \underbrace{\mathbb{1}[w_{m_1} w_{m_2} > \epsilon] \times w_{m_1} w_{m_2}}_{:=W_{m_1, m_2}} \end{aligned}$$

where the h are the quadrature nodes, w the corresponding quadrature weights, and M the number of quadrature terms used. As many of the products $w_{m_1} w_{m_2}$ will be close to zero, it

is common to use some form of pruning [Jaeckel, 2005] to further reduce the computational requirements of the approximation. Above we used the indicator function to include only products greater than a fixed given value ϵ . Towards the optimization of Q_2 , we can now compute its derivatives with respect to \mathbf{s} and σ . We first consider the derivative w.r.t. $\mathbf{s}_{k,i}$, the i -th component of the k -th segment-level effect:⁴

$$\begin{aligned} \frac{\partial Q_2(\Theta, \Theta^{t-1})}{\partial s_{k,i}} &\approx \frac{\partial \log P(\mathbf{s}|\theta^s)}{\partial \mathbf{s}_{k,i}} + \\ &\quad \frac{1}{2\pi} \frac{\partial}{\partial \mathbf{s}_{k,i}} \left[\sum_u \sum_k \gamma_{u,k} \times \sum_{j=1}^{n_{\mathbf{D}_u}} \sum_{m_1, m_2=1}^M \log \Phi(v_j) \times W_{m_1, m_2} \right] \end{aligned} \quad (13)$$

Here, the first term is simply the i -th component of the derivative of a multivariate Gaussian log density:

$$\frac{\partial \log P(\mathbf{s}|\theta^s)}{\partial \mathbf{s}_k} = \frac{P(\mathbf{s}_k|\theta^s)}{P(\mathbf{s}_k|\theta^s)} \times -\Sigma_s^{-1}(\mathbf{s}_k - \vec{0}) = -\Sigma_s^{-1} \mathbf{s}_k$$

From the second term in Equation (13), summations over segments other than k disappear, as do choices in \mathbf{D}_u that do not involve instance i . Overall, we obtain

$$\begin{aligned} &\frac{1}{2\pi} \frac{\partial}{\partial \mathbf{s}_{k,i}} \sum_u \gamma_{u,k} \times \sum_{j \in \mathbf{D}_u|_i} \sum_{m_1, m_2=1}^M \log \Phi(v_j) \times W_{m_1, m_2} \\ &= \frac{1}{2\pi} \sum_u \gamma_{u,k} \times \sum_{j \in \mathbf{D}_u|_i} \sum_{m_1, m_2=1}^M \frac{p_j(\mathbf{s}_{k,i}) N(v_j)}{\sqrt{2\sigma} \Phi(v_j)} \times W_{m_1, m_2} \end{aligned}$$

where N represents the standard Gaussian density, and

$$p_j(\mathbf{s}_{k,i}) = \begin{cases} 1 & \text{if } x_i \text{ is the preferred instance in statement } j \\ -1 & \text{if } x_i \text{ is the dispreferred instance in statement } j \\ 0 & \text{otherwise} \end{cases}$$

While optimizing Q_2 w.r.t. σ we need to ensure that σ remains strictly positive. Instead of explicitly incorporating this constraint into the optimization, we optimize w.r.t. $\log(\sigma)$. The corresponding first derivative is

$$\begin{aligned} \frac{\partial Q_2(\Theta, \Theta^{t-1})}{\partial \log(\sigma)} &\approx \frac{1}{2\pi} \frac{\partial}{\partial \log(\sigma)} \sum_u \sum_k \gamma_{u,k} \times \sum_{j=1}^{n_{\mathbf{D}_u}} \sum_{m_1, m_2=1}^M \log \Phi(v_j) \times W_{m_1, m_2} \\ &= \frac{1}{2\pi} \sum_u \cdots \sum_{m_1, m_2=1}^M \frac{N(v_j) \times (-v_j)}{\Phi(v_j)} \times W_{m_1, m_2} \end{aligned}$$

To summarize:

$$Q_2(\Theta, \Theta^{t-1}) \approx \sum_k \log P(\mathbf{s}_k|\theta^s) + \sum_u \sum_k \frac{\gamma_{u,k}}{2\pi} \times \sum_{j=1}^{n_{\mathbf{D}_u}} \sum_{m_1, m_2=1}^M \log \Phi(v_j) \times W_{m_1, m_2} \quad (14)$$

$$\frac{\partial Q_2(\Theta, \Theta^{t-1})}{\partial \mathbf{s}_{k,i}} \approx (-\Sigma_s^{-1} \times \mathbf{s}_k)_i + \frac{1}{2\pi} \sum_u \gamma_{u,k} \times \sum_{j \in \mathbf{D}_u|_i} \sum_{m_1, m_2=1}^M \frac{p_j(\mathbf{s}_{k,i}) N(v_j)}{\sqrt{2\sigma} \Phi(v_j)} \times W_{m_1, m_2} \quad (15)$$

$$\frac{\partial Q_2(\Theta, \Theta^{t-1})}{\partial \log(\sigma)} \approx \frac{1}{2\pi} \sum_u \sum_k \gamma_{u,k} \times \sum_{j=1}^{n_{\mathbf{D}_u}} \sum_{m_1, m_2=1}^M \frac{N(v_j) \times (-v_j)}{\Phi(v_j)} \times W_{m_1, m_2} \quad (16)$$

⁴Recall that the $\mathbf{s}_{k,i}$ are regular (non-random) variables.

A.3.3 Optimizing θ^i (Q_3)

Lastly, we wish to find

$$\begin{aligned}\theta^{i*} &= \arg \max_{\theta^i} Q_3(\Theta, \Theta^{t-1}) \\ &= \arg \max_{\theta^i} \sum_u \sum_k \gamma_{u,k} \times E_{\{i_u|z_u=k, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(i_u|z_u = k, \Theta)]\end{aligned}\quad (17)$$

where $P(i_u|z_u = k, \Theta)$ is the approximate posterior of the Gaussian process preference regression from Equation (4). In other words, $i_u|z_u = k, \Theta$ is distributed $N(\mu_{u,k}, \Sigma_u)$ for some $\mu_{u,k} = E[i_u|z_u = k, \Theta]$ and $\Sigma_u = \Sigma_u(\theta^i)$. Expanding the Gaussian density in Equation (17) and collecting constant terms we obtain

$$\begin{aligned}& \sum_u \sum_k \gamma_{u,k} E_{\{i_u|z_u=k, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\log P(i_u|z_u = k, \Theta)] \\ &= C - \frac{1}{2} \sum_u \log |\Sigma_u| \\ & \quad - \frac{1}{2} \sum_u \sum_k \gamma_{u,k} E_{\{i_u|z_u=k, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [(x - \mu_{u,k})^T \Sigma_u^{-1} (x - \mu_{u,k})]\end{aligned}$$

Note, that the expectation is over the posterior, conditional on the previous parameters Θ^{t-1} . Since the $\mu_{u,k}$ are constants under this expectation, we have that $E[x - \mu_{u,k}] = E[x] =: \mu_{u,k}^{t-1}$ and $\text{Var}(x - \mu_{u,k}) = \text{Var}(x) =: C_u^{t-1}$. Recall that for the expectation of a quadratic form we have $E[v^T A v] = \text{tr}(A v) + m^T A m$ where $E[v] = m$ and $\text{Var}(v) = V$. Applying this result to $E_{\{i_u|z_u=k, \mathbf{X}, \mathbf{D}, \Theta^{t-1}\}} [\dots]$ with $v = (x - \mu_{u,k})$ and $A = \Sigma_u^{-1}$ we obtain $E[\dots] = \text{tr}[\Sigma_u^{-1} (C_u^{t-1} + \mu_{u,k}^{t-1} (\mu_{u,k}^{t-1})^T)]$, and hence the optimization task amounts to finding:

$$\begin{aligned}\theta^{i*} &= \arg \max_{\theta^i} Q_3(\Theta, \Theta^{t-1}) \\ &= \arg \max_{\theta^i} -\frac{1}{2} \sum_u \log |\Sigma_u(\theta_u^i)| - \frac{1}{2} \sum_u \sum_k \gamma_{u,k} \times \text{tr} \left[\Sigma_u(\theta_u^i)^{-1} \left(C_u^{t-1} + \mu_{u,k}^{t-1} (\mu_{u,k}^{t-1})^T \right) \right]\end{aligned}\quad (18)$$

Evaluation of this function is expensive. Computing the log determinant and solving the linear system of equations in Equation (18) require a Cholesky decomposition of each Σ_u which scales as $O(n_x^3)$ with the number of instances. In our experiments, we assumed that all users within a segment share the same kernel parameters θ^i , thereby reducing the *number* of evaluations of (18). Alternatively, the *cost* of evaluation can be reduced through approximations, e.g., [Snelson and Ghahramani, 2006].

B Experimental Details

We collected the dataset used in Section 4 using Amazon Mechanical Turk (MTurk, <http://www.mturk.com>), a commercial crowdsourcing platform. Eighty adult American participants were invited to fill in an academic survey about their electricity tariff preferences⁵ in exchange for a payment of \$0.30. All American MTurk workers could theoretically preview our survey through the MTurk platform, and 80 workers ultimately self-selected to participate. The survey consisted of three parts:

1. First, we reviewed basic electricity tariff concepts: fixed, variable, and indexed tariffs; and tariffs guaranteeing that a certain percentage of delivered electricity is produced from renewable sources.

⁵A tariff is sometimes also referred to as a plan.

Table 2: Example choice situation. Each participant in our experiment made ten choices.

Imagine having to choose between the following two tariffs for the household that you currently spend most of your time in. Which one would you prefer?

1. A Fixed tariff with 100 percent renewable energy content. Your monthly cost of electricity will be
 - 57.00\$ if you consume 500 kWh,
 - 106.00\$ if you consume 1000 kWh, and
 - 204.00\$ if you consume 2000 kWh

under this tariff. You pay your monthly electricity bill at the end of each month. A 12 months notice period applies before you can cancel this tariff.

2. A Variable tariff with 0 percent renewable energy content. The cost of electricity in the first month will be
 - 54.50\$ if you consume 500 kWh,
 - 101.00\$ if you consume 1000 kWh, and
 - 194.00\$ if you consume 2000 kWh

under this tariff. After the first month, the price of electricity may go up or down in accordance with the tariff's terms (and within legal bounds). You will have to pre-pay your monthly electricity bill at the beginning of the month. You can cancel your tariff anytime.

2. Next, participants were asked to make ten choices between pairs of tariffs, see Table 2 for an example. Each pair was randomly generated from a total of thirty tariffs that we kept constant throughout. The set of thirty, in turn, was randomly selected from over 230 tariffs offered in Austin, Texas in February 2013. We chose to work with a subset of all available tariffs to keep the number of required participants manageable. Texas has one of the most advanced competitive retail electricity markets in the United States and provides daily information on available tariffs, see <http://www.powertochoose.org>.
3. Finally we asked participants to answer ten questions on their demographics and electricity consumption behavior, some of which were attention checkers for which the correct answer had to be consistent with an answer given on another question.

Participants had a maximum of thirty minutes to fill out all questions, but could submit their results before that time. Participants could also withdraw, allowing another MTurk worker to fill out the survey instead. Next to the given answers, we recorded the time between self-selecting for participation and the submission of results. In pretests among colleagues, we had established that it took a quick reader at least three minutes to process all provided information. We therefore discarded surveys submitted before that time. As a further safeguard against random answers, we asked two pairs of attention check questions in the demographics section where the answers to one question depended on the answer of the other. We also discarded surveys where at least one of the attention check pairs was answered inconsistently, leaving us with a total of 58 surveys that met our quality standards.

References

- Gediminas Adomavicius and Alexander Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, pages 734–749, 2005a.
- Gediminas Adomavicius and Alexander Tuzhilin. Personalization technologies: a process-oriented perspective. *Communications of the ACM*, 48(10):83–90, 2005b.
- Gediminas Adomavicius and Alexander Tuzhilin. Context-aware recommender systems. In *Recommender Systems Handbook*, pages 217–253. Springer, 2011.
- Greg M. Allenby and Peter E. Rossi. Marketing models of consumer heterogeneity. *Journal of Econometrics*, 89(1):57–78, 1998.
- Asim Ansari, Skander Essegaier, and Rajeev Kohli. Internet recommendation systems. *Journal of Marketing Research*, 37(3):363–375, 2000.
- Pelin Atahan and Sumit Sarkar. Accelerated learning of user profiles. *Management Science*, 57(2):215–239, 2011.
- Martin Bichler, Alok Gupta, and Wolfgang Ketter. Designing smart markets. *Information Systems Research*, 21(4):688–699, 2010.
- Adriana Birlutiu, Perry Groot, and Tom Heskes. Efficiently learning the preferences of people. *Machine Learning*, pages 1–28, 2012.
- Edwin Bonilla, Shengbo Guo, and Scott Sanner. Gaussian process preference elicitation. *Neural Information Processing Systems*, 2010.
- Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos, and David Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, pages 135–191, 2004.

- Ronald J. Brachman. Systems that know what they're doing. *IEEE Intelligent Systems*, 17(6): 67–71, 2002.
- Ronen Brafman and Carmel Domshlak. Preference handling - an introductory tutorial. *AI Magazine*, 30(1):58–86, 2009.
- Michael Buhrmester, Tracy Kwang, and Samuel D. Gosling. Amazon's mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1):3–5, 2011.
- Rich Caruana. Multitask learning. *Machine Learning*, 28(1):41–75, 1997.
- Yann Chevaleyre, Frédéric Koriche, Jérôme Lang, Jérôme Mengin, and Bruno Zanuttini. Learning ordinal preferences on multiattribute domains: The case of CP-nets. In *Preference Learning*, pages 273–296. Springer, 2011.
- Wei Chu and Zoubin Ghahramani. Gaussian processes for ordinal regression. *The Journal of Machine Learning Research*, 6:1019–1041, 2005a.
- Wei Chu and Zoubin Ghahramani. Preference learning with Gaussian processes. In *Proceedings of the 22nd international conference on Machine Learning*, pages 137–144. ACM, 2005b.
- Thomas H Davenport and Jeanne G Harris. Automated decision making comes of age. *MIT Sloan Management Review*, 46(4):83, 2005.
- Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977.
- Theodoros Evgeniou, Massimiliano Pontil, and Olivier Toubia. A convex optimization approach to modeling consumer heterogeneity in conjoint estimation. *Marketing Science*, 26(6):805–818, 2007.
- Johannes Fürnkranz and Eyke Hüllermeier. *Preference learning*. Springer, 2011.
- Zoubin Ghahramani. Bayesian nonparametrics and the probabilistic approach to modelling. 2012.
- William Greene. *Econometric Analysis*. Prentice Hall, 7th edition, 2012.
- Gerald Häubl and Kyle B Murray. Preference construction and persistence in digital marketplaces: The role of electronic recommendation agents. *Journal of Consumer Psychology*, 13(1&2):75–91, 2003.
- Gerald Häubl and Valerie Trifts. Consumer decision making in online shopping environments: The effects of interactive decision aids. *Marketing Science*, 19(1):4–21, 2000.
- Alan R. Hevner, Salvatore T. March, Jinsoo Park, and Sudha Ram. Design Science in Information Systems research. *Management Information Systems Quarterly*, 28(1):75–105, 2004.
- Neil Houlsby, Jose M. Hernandez-Lobato, Ferenc Huszar, and Zoubin Ghahramani. Collaborative Gaussian processes for preference learning. *Advances in Neural Information Processing Systems*, 2012.
- Zan Huang, Huimin Zhao, and Dan Zhu. Two new prediction-driven approaches to discrete choice prediction. *ACM Transactions on Management Information Systems*, 3(2):9:1–9:32, July 2012. ISSN 2158-656X. doi: 10.1145/2229156.2229159.

- Panagiotis G. Ipeirotis. Analyzing the amazon mechanical turk marketplace. *XRDS*, 17(2): 16–21, December 2010. ISSN 1528-4972. doi: 10.1145/1869086.1869094. URL <http://doi.acm.org/10.1145/1869086.1869094>.
- Peter Jaeckel. A note on multivariate gauss-hermite quadrature, 2005. URL <http://www.pjaeckel.webspace.virginmedia.com/ANoteOnMultivariateGaussHermiteQuadrature.pdf>.
- Tianyi Jiang and Alexander Tuzhilin. Segmenting customers from population to individuals: Does 1-to-1 keep your customers forever? *IEEE Transactions on Knowledge and Data Engineering*, 18(10):1297–1311, 2006.
- Tianyi Jiang and Alexander Tuzhilin. Improving personalization solutions through optimal segmentation of customer bases. *IEEE Transactions on Knowledge and Data Engineering*, 21(3):305–320, 2009.
- Rong Jin, Luo Si, and ChengXiang Zhai. Preference-based graphic models for collaborative filtering. In *Proceedings of the Nineteenth conference on Uncertainty in Artificial Intelligence*, pages 329–336. Morgan Kaufmann Publishers Inc., 2002.
- Kristian Kersting and Zhao Xu. Learning preferences with hidden common cause relations. *Machine Learning and Knowledge Discovery in Databases*, pages 676–691, 2009.
- Wolfgang Ketter, John Collins, Maria Gini, Alok Gupta, and Paul Schrater. Real-time tactical and strategic sales management for intelligent agents guided by economic regimes. *Information Systems Research*, 23(4):1263–1283, December 2012. ISSN 1047-7047, 1526-5536.
- Young Eun Lee and Izak Benbasat. The influence of trade-off difficulty caused by preference elicitation methods on user acceptance of recommendation agents across loss and gain conditions. *Information Systems Research*, 22(4):867–884, 2011.
- Sarah Lichtenstein and Paul Slovic. *The construction of preference*. Cambridge University Press, 2006.
- Zhengdong Lu, Todd K. Leen, Yonghong Huang, and Deniz Erdogmus. A reproducing kernel Hilbert space framework for pairwise time series distances. In *Proceedings of the 25th international conference on Machine learning*, pages 624–631. ACM, 2008.
- David J.C. MacKay. Introduction to gaussian processes. *NATO ASI Series F Computer and Systems Sciences*, 168:133–166, 1998.
- Petty Maes. Agents that reduce work and information overload. *Communications of the ACM*, 37(7):30–40, 1994.
- Eren Manavoglu, Dmitry Pavlov, and C Lee Giles. Probabilistic user behavior models. In *Third IEEE International Conference on Data Mining (ICDM)*, pages 203–210. IEEE, 2003.
- Nikos Manouselis and Constantina Costopoulou. Analysis and classification of multi-criteria recommender systems. *World Wide Web*, 10(4):415–441, 2007.
- Winter Mason and Siddharth Suri. Conducting behavioral research on Amazon’s Mechanical Turk. *Behavior Research Methods*, 44:1–23, 2012. ISSN 1554-351X. URL <http://dx.doi.org/10.3758/s13428-011-0124-6>. 10.3758/s13428-011-0124-6.
- BPS Murthi and Sumit Sarkar. The role of the management sciences in research on personalization. *Management Science*, 49(10):1344–1362, 2003.

- Oded Netzer, Olivier Toubia, Eric T. Bradlow, Ely Dahan, Theodoros Evgeniou, Fred M. Feinberg, Eleanor M. Feit, Sam K. Hui, Joseph Johnson, John C. Liechty, et al. Beyond conjoint analysis: Advances in preference measurement. *Marketing Letters*, 19(3):337–354, 2008.
- Jens B. Nielsen, Björn S. Jensen, and Jan Larsen. On sparse multi-task gaussian process priors for music preference learning. In *NIPS Workshop on Choice Models and Preference Learning*, 2011.
- Cosimo Palmisano, Alexander Tuzhilin, and Michele Gorgoglione. Using context to improve predictive modeling of customers in personalization applications. *IEEE Transactions on Knowledge and Data Engineering*, 20(11):1535–1549, 2008.
- Gabriele Paolacci, Jesse Chandler, and Panagiotis Ipeirotis. Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5(5):411–419, 2010.
- Markus Peters and Wolfgang Ketter. Modeling sparse heterogeneous user preferences. In *Workshop on Information Systems and Technology*, Orlando, Florida, USA, December 2012.
- Markus Peters, Wolfgang Ketter, Maytal Saar-Tsechansky, and John E. Collins. A reinforcement-learning approach to autonomous decision-making in smart electricity markets. *Machine Learning*, 2013. doi: 10.1007/s10994-013-5340-0.
- Alan Qi, Ahmed Abdel-Gawad, and Thomas Minka. Sparse-posterior gaussian processes for general likelihoods. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence*, pages 450–457, 2010.
- Joaquin Quinonero-Candela, Carl Edward Rasmussen, and Christopher KI Williams. Approximation methods for gaussian process regression. *Large-scale kernel machines*, pages 203–224, 2007.
- Carl E. Rasmussen and Christopher Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Katharina Reinecke and Abraham Bernstein. Knowing what a user likes: A design science approach to interfaces that automatically adapt to culture. *Management Information Systems Quarterly*, 2012. Forthcoming.
- Asim Roy, Patrick Mackin, Jyrki Wallenius, James Corner, Mark Keith, Gregory Schymik, and Hina Arora. An interactive search method based on user preferences. *Decision Analysis*, 5(4):203–229, 2008.
- Maytal Saar-Tsechansky, Prem Melville, and Foster Provost. Active feature-value acquisition. *Management Science*, 55(4):664–684, 2009.
- Paul Slovic. The construction of preference. *American psychologist*, 50(5):364, 1995.
- Ed Snelson and Zoubin Ghahramani. Sparse Gaussian processes using pseudo-inputs, 2006.
- Milind Tambe, Paul Scerri, and David V Pynadath. Adjustable autonomy for the real world. *Journal of Artificial Intelligence Research*, 17(1):171–228, 2002.
- Richard H Thaler. From homo economicus to homo sapiens. *The Journal of Economic Perspectives*, 14(1):133–141, 2000.
- Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992.

- Cristiano Varin, Gudmund Høst, and Øivind Skare. Pairwise likelihood inference in spatial generalized linear mixed models. *Computational statistics & data analysis*, 49(4):1173–1191, 2005.
- John von Neumann, Oskar Morgenstern, Ariel Rubinstein, and Harold W. Kuhn. *Theory of games and economic behavior (Commemorative Edition)*. Princeton University Press, 2007.
- Yuyang Wang, Roni Khardon, and Pavlos Protopapas. Shift-invariant grouped multi-task learning for gaussian processes. *Machine Learning and Knowledge Discovery in Databases*, pages 418–434, 2010.
- Zhao Xu, Kristian Kersting, and Thorsten Joachims. Fast active exploration for link-based preference learning using gaussian processes. *Machine Learning and Knowledge Discovery in Databases*, pages 499–514, 2010.

ERIM Report Series <i>Research in Management</i>	
ERIM Report Series reference number	ERS-2013-007-LIS
Date of publication	2013-05-22
Version	22-05-2013
Number of pages	31
Persistent URL for paper	http://hdl.handle.net/1765/40144
Email address corresponding author	wketter@rsm.nl
Address	Erasmus Research Institute of Management (ERIM) RSM Erasmus University / Erasmus School of Economics Erasmus University Rotterdam PO Box 1738 3000 DR Rotterdam, The Netherlands Phone: +31104081182 Fax: +31104089640 Email: info@erim.eur.nl Internet: http://www.erim.eur.nl
Availability	The ERIM Report Series is distributed through the following platforms: RePub, the EUR institutional repository Social Science Research Network (SSRN) Research Papers in Economics (RePEc)
Classifications	The electronic versions of the papers in the ERIM Report Series contain bibliographic metadata from the following classification systems: Library of Congress Classification (LCC) Journal of Economic Literature (JEL) ACM Computing Classification System Inspec Classification Scheme (ICS)