GaSPing for Utility

Mengyang Gu,[†] Debarun Bhattacharjya,[§] Dharmashankar Subramanian[§]
[†]Department of Statistics and Applied Probability, University of California, Santa Barbara, CA, USA
[§]Research AI, IBM T. J. Watson Research Center, Yorktown Heights, NY, USA
mengyang@pstat.ucsb.edu, {debarunb, dharmash}@us.ibm.com

Abstract

High-consequence decisions often require a detailed investigation of a decision maker's preferences, as represented by a utility function. Inferring a decision maker's utility function through assessments typically involves an elicitation phase where the decision maker responds to a series of elicitation queries, followed by an estimation phase where the state-of-the-art for direct elicitation approaches in practice is to either fit responses to a parametric form or perform linear interpolation. We introduce a Bayesian nonparametric method involving Gaussian stochastic processes for estimating a utility function from direct elicitation responses. Advantages include the flexibility to fit a large class of functions, favorable theoretical properties, and a fully probabilistic view of the decision maker's preference properties including risk attitude. Through extensive simulation experiments as well as two real datasets from management science, we demonstrate that the proposed approach results in better function fitting.

Introduction & Related Work

Making decisions under uncertainty using decision theory requires that beliefs about uncertainties be represented by probabilities and preferences over outcomes be summarized by utilities. It is often easier in practice to justify learning probabilities from historical data, which is reasonable when the decision maker believes that the future will resemble the past, than it is to learn preferences using other decision makers' choices. This is particularly true for high-consequence decisions such as a potentially life-changing medical decision, a new product launch, a government policy decision involving numerous stakeholders, etc. In these situations, a decision maker's preferences should be represented through their *utility function* either over a single attribute (such as monetary units) or over multiple attributes, depending on the decision situation under consideration.

The vast literature on assessing utility functions describes schemes where the decision maker responds to elicitation queries; these responses are then subsequently used to estimate the decision maker's utility function. We distinguish between these phases and refer to them as *elicitation* and *estimation* respectively. Posing elicitation questions and then estimating the functional form that best represents the decision maker's preferences often go hand in hand.

The elicitation phase could of course be conducted in various ways. We distinguish between *direct* and *indirect* approaches. In the former, elicitation directly results in outcomes with their utilities, whereas in the latter, utilities are typically inferred through choices – the most popular approach across domains (including AI) is to use pairwise comparisons between alternatives. Both these approaches have their respective drawbacks. The behavioral literature demonstrates that people are plagued with cognitive biases while responding to questions that aim to assess their preferences, and that they *construct* their preferences according to the situational context, including the method of elicitation, and often in an inconsistent manner (Lichtenstein and Slovic 2006). Systematic biases from choice inconsistencies arising from indirect elicitation approaches are well known in this literature (Fischer, Jia, and Luce 2000; Bleichrodt, Pinto, and Wakker 2001).

In this paper, we contribute to the literature on estimation by introducing a Bayesian nonparametric approach for modeling utility functions from direct elicitation. Specifically, we demonstrate the advantages of using a Gaussian stochastic process (henceforth GaSP) over the state-of-the-art in this space, where parametric methods and the nonparametric method of linear interpolation are most popular. We show that linear interpolation corresponds to a subclass of our GaSP model using a specific covariance function in which the process turns out to be a *Wiener process* (Mörters and Peres 2010); moreover, parametric methods can be incorporated into the GaSP model through the mean function.

The GaSP model is popular in domains such as nonlinear regression and classification in machine learning (Rasmussen and Williams 2006), spatial statistics (Gelfand et al. 2010), computer model emulation and calibration (Sacks et al. 1989). It has also been applied to preference modeling. For instance, Chu and Ghahramani (2005), Birlutiu, Groot, and Heskes (2009) and Bonilla, Guo, and Sanner (2010) use GaSP models for preference learning in the indirect elicitation case involving pairwise comparisons.

There is substantial AI literature on Bayesian models for preference elicitation in general, where uncertainty in utilities is typically exploited for the purpose of adaptive elicitation for specific decisions (Jimison et al. 1992;

Poh and Horvitz 1993; Chajewska, Koller, and Parr 2000; Boutilier 2002; Bhattacharjya and Kephart 2014). In contrast, our work is based on direct elicitation schemes, where the observations are assessed tuples (see Definition 2 for details). Direct elicitation is more popular in areas such as operations research, behavioral economics and management science, see e.g. Farquhar (1984). To the best of our knowledge, we are not aware of GaSP being deployed for such schemes, and it appears that direct elicitation has not been actively explored in the AI community, barring early work such as Chajewska, Koller, and Parr (2000).

Contributions. This work can be viewed as an application of AI/ML to an important problem in management science – utility elicitation & estimation. A fundamental challenge is that the number of assessed queries is often very small, since elicitation burden increases significantly and responses become less reliable when more questions are posed. Here we propose a robust estimation approach using GaSP models that quantify uncertainty associated with the utility elicitation process. While the application of GaSP to the problem is relatively straightforward, our novel contributions include showing that our method: 1) effectively handles limited observations through use of the Matérn kernel and robust posterior mode estimation, 2) generalizes linear interpolation, which is most widely used in practice, 3) more accurately assesses risk attitudes of decision makers, 4) reduces predictive error in comparison with baseline approaches, including a new strong baseline, and 5) is also applicable for multi-attribute problems. We conduct experiments with synthetic data, and perhaps more importantly, with real-world data from the management science literature.

Utility Elicitation

Consider outcomes $\mathbf{x}$ in a separable metric space $\mathcal{X}$ (e.g. $\mathbb{R}^n$). Debreu (1954) showed that preferences over uncertain outcomes in such a space are complete, transitive and continuous in $\mathcal{X}$ iff there exists a continuous utility function representation $U : \mathcal{X} \to \mathbb{R}$. Since people often provide inconsistent responses to queries, a decision maker's utility function U should perhaps be considered an approximate representation of their preferences; for this and other reasons we shall discuss shortly, we make the following distinction:

Definition 1 (Noise-free vs. noisy assessment) *When a decision maker answers all the preference elicitation questions consistently with the same underlying (and typically unknown) utility function, the assessment is said to be noise-free, otherwise it is noisy.*

The notion of a 'true' underlying utility function can be viewed as a theoretical construct and one that is often discussed in the literature. One could model the uncertainty in preference elicitation responses as a random response error to a systematic component or to treat the utility function as inherently stochastic. Making a noise-related distinction enables us to express different sources of uncertainty in the elicitation process. One source of uncertainty is prediction uncertainty, representing the system's uncertainty about the decision maker's utility at an unassessed $\mathbf{x}$. The second source of uncertainty depends on the decision maker – when they consistently answer questions with the same underlying utility function, the elicited utility $u(\mathbf{x})$ is identical to $U(\mathbf{x})$ at each assessed $\mathbf{x}$. We refer to an estimation method that agrees with the elicited utility at each assessed $\mathbf{x}$ as an *interpolator*. It is ideal for a method to be an interpolator for a consistent decision maker. Note that the second source of uncertainty only appears when decision makers do not answer questions consistently, leading to noisy assessments.

Preference elicitation for decision making under risk is typically conducted using *gambles*. Consider gamble $(\mathbf{x}_a, p; \mathbf{x}_b)$ that results in outcome $\mathbf{x}_a$ with probability p and outcome $\mathbf{x}_b$ with probability $1-p$, where $\mathbf{x}_a, \mathbf{x}_b \in \mathcal{X}$. In an elicitation query, the decision maker must evaluate two or more gambles presented to them, e.g. they may be asked to compare gamble $(\mathbf{x}_a, p; \mathbf{x}_b)$ with the degenerate gamble $(\mathbf{x}_c)$. If the assessment is noise-free and if the evaluation is done by expected utility theory (EUT), they prefer the first gamble iff $V_{EUT}(\mathbf{x}_a, p; \mathbf{x}_b) = pU(\mathbf{x}_a) + (1-p)U(\mathbf{x}_b) \geq U(\mathbf{x}_c)$, for some underlying utility function U (von Neumann and Morgenstern 1947). A subsequent estimation task must be performed to infer U from the responses.

There is significant empirical evidence from the descriptive literature on prospect theory (PT) demonstrating that people tend to overweight low probabilities and underweight high probabilities. Thus, under prospect theory, the gamble $(\mathbf{x}_a, p; \mathbf{x}_b)$ evaluates to $V_{PT}(\mathbf{x}_a, p; \mathbf{x}_b) = \omega(p)U(\mathbf{x}_a) + \omega(1-p)U(\mathbf{x}_b)$, where $\omega(.)$ is a probability weighting function and the reference point is assumed to be 0 (Kahneman and Tversky 1979; Tversky and Kahneman 1992). Prospect theory explains observed behavior such as loss aversion and diminishing sensitivity relative to the reference point.

Bleichrodt, Pinto, and Wakker (2001) and others argue that descriptive violations of expected utility bias utility elicitation; therefore the design of elicitation schemes and estimation of utility functions should be conducted based on PT rather than EUT, resulting in more accurate estimates of U. We are sympathetic to this view but as we explain in the next section, our research goal is purely that of estimating U, provided elicitation responses in the following form:

Definition 2 (Assessed tuples) *Assessed tuples are assessments of points on the utility function,* $(\mathbf{x}_i, u(\mathbf{x}_i))$ *for* $i = 1, \ldots, n$*,* $\mathbf{x}_i \in \mathcal{X}$*.*

We propose a GaSP approach for estimating a utility function given assessed tuples as input. Our method can thus be applied to different data sets and is agnostic to the underlying theory. In fact, we demonstrate GaSP estimation using elicitation schemes based on both EUT and PT. Furthermore, it is also agnostic to the generative structure of randomness in the case of probabilistic preferences, i.e. any stochastic utility model. In this work, we consider both noise-free and noisy assessments. Note that for noise-free assessments, assessed tuples $(\mathbf{x}_i, u(\mathbf{x}_i)) = (\mathbf{x}_i, U(\mathbf{x}_i))$ for an underlying utility function U.

Utility Estimation

The goal of any estimation task associated with preference elicitation is to use responses to the elicitation queries to infer the decision maker's utility function $U(\mathbf{x})$. Here we discuss parametric estimation and linear interpolation, highlighting their limitations, along with a novel baseline.

Parametric Estimation

Parametric estimation is a popular approach for estimating utility functions involving a single attribute x (Eliashberg and Hauser 1985; Kirkwood 2004). The most common parametric forms are the exponential and power functions. The exponential utility function follows $a - b\,\text{sgn}(\rho)\,\exp\left(-x/\rho\right)$, where a and $b > 0$ are constants and $\text{sgn}(\rho)$ is the sign of the risk tolerance parameter $\rho \neq \infty$. The power utility function is of the form $a + b\ \text{sgn}(\alpha)\ \text{sgn}(x)\ |x|^{\alpha}$, with constants a and $b > 0$, where $\text{sgn}(\alpha)$ and $\text{sgn}(x)$ are the signs of $\alpha \neq 0$ and x. The limiting cases for the exponential and power functions are the linear and logarithmic functions; these two families of functions are the only ones that satisfy constant risk aversion and constant relative risk aversion respectively (Pratt 1964). It is not hard to see that a parametric method will not be an interpolator unless the underlying utility function follows the parametric class being used.

Linear Interpolation

An alternate approach that is popular in the empirical literature on estimating single-attribute utility functions is that of piece-wise linear interpolation across assessed tuples (Abdellaoui 2000; Abdellaoui, Bleichrodt, and Paraschiv 2007). Such an approach is essentially equivalent to the *predictive mean* of the extended Wiener process, defined as a stochastic process W_t with independent, normally distributed increments $W_t - W_s \sim \mathcal{N}(0, t-s)$ for $t \geq s \geq 0$ with continuous sample paths (Karlin 1975). A Wiener process is typically defined to have initial value $W_0 = 0$ but we may relax this assumption. The predictive distribution is formalized in the following lemma.

Lemma 1 *(Theorem 2.1 in (Karlin 1975)) Assume W_t, $t \in \mathcal{T}$ follows a Wiener process. Assume we have observations $W_{t_1}, ..., W_{t_n}$ with $0 < t_1 < t_2 < ... < t_n$. For any $t_i \leq t_* \leq t_{i+1}$, for any $1 \leq i < n$ and $i \in \mathbb{N}$, the predictive distribution of W_{t_*} given $W_{t_1}, ..., W_{t_n}$ is $W_{t_*}|W_{t_1}, ..., W_{t_n} \sim \mathcal{N}(\mu_*, V_*)$, where $\mu_* = \frac{(t_{i+1}-t_*)W_{t_i} + (t_* - t_i)W_{t_{i+1}}}{t_{i+1}-t_i}$ and $V_* = \frac{(t_{i+1}-t_*)(t_*-t_i)}{t_{i+1}-t_i}$.*

Lemma 1 states that if the utility function is modeled as a Wiener process for a single attribute t, the posterior mean $\mu_* = E[W_{t_*}|W_{t_1}, ..., W_{t_n}]$ for assessing the utility at a point t_* (that has not been assessed) is equivalent to linear interpolation between two nearby assessed tuples. Indeed, a Wiener process is a special case of GaSP (which we will formally define in the next section) with initial value $W_0 = 0$, mean zero and covariance function $Cov(W_s, W_t) = min(s, t)$ at any $s, t \geq 0$, and continuous sample path. However, a Wiener process is not differentiable everywhere and consequently using the posterior mean (i.e. linear interpolation) poses problems for estimating the risk aversion coefficient, since it relies on computing derivatives.

Example 1 (Linear interpolation and overconfidence) *Figure 1 displays the function $y = 3sin(5\pi t) + cos(7\pi t)$ (treated as unknown) with assessments on 12 equally spaced points in $[0, 1]$. The left panel shows the prediction (blue curves) by the extended Wiener process (for which we do not assume $W_0 = 0$). Not only does the prediction show discrepancy at places where the derivatives of the function change, it is also clearly overconfident as the 95% confidence interval covers regions that are a lot smaller than the nominal 95%. In comparison, the right panel is the prediction by the method we propose with the same 12 assessed points, using the default setting in the RobustGaSP R Package (Gu, Palomo, and Berger 2019). While this example illustrates a shortcoming of linear interpolation using a generic function that is measured at a finite number of values in its domain, the above weakness is relevant to our objective of estimating a utility function using experimentally assessed tuples.*

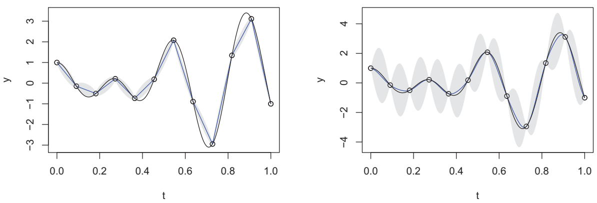

Figure 1: Interpolation of the function $y = 3sin(5\pi t) + cos(7\pi t)$ plotted as black curves with 12 assessments equally spaced in $[0, 1]$ (black dots). Predictions by the extended Wiener process (left panel) and the GaSP model (right panel) are plotted as blue curves. 95% predictive confidence intervals are shown by the grey area.

Another limitation of a Wiener process estimation approach is that it is only defined in a one-dimensional domain and thus is limited to single-attribute utility function estimation. Although there is some literature on interpolation in multi-attribute problems, e.g. Bell (1979), it is challenging due to the curse of dimensionality. A more general GaSP approach with suitable covariance functions built on the space of multiple attributes may be more suitable.

Quantile-Parameterized Distributions

Quantile-parameterized distributions (QPD) have recently been introduced for modeling uncertainties (Keelin and Powley 2011). This approach characterizes a continuous probability distribution based on a number of assessed quantile/probability pairs. Although QPDs have not been discussed in the context of utility elicitation, it is straightforward to apply the approach conceptually, since assessed tuples are analogous to quantile/probability pairs. Denoting these assessed tuples as (x_i, u_i) for $i = 1, .., n$, where u_i is the assessed utility scaled from $[0, 1]$, the inverse CDF of a QPD takes the form:

$$F^{-1}(u) = \begin{cases} L_0 & u = 0 \\ \sum_{i=1}^{q} a_i g_i(u) & 0 < u < 1 \\ L_1 & u = 1 \end{cases}$$

with left handed limit $L_0 = \lim_{y \to 0^+} F^{-1}(u)$, right handed limit $L_1 = \lim_{y \to 1^-} F^{-1}(u)$, and $g_i(\cdot)$ referring to basis functions for $i = 1, ..., q$, $q \leq n$. When the domain of x is a real line, a simple choice is the Q-normal distribution, where the basis functions are $g_1(u) = 1$, $g_2(u) = \Phi^{-1}(u)$, $g_3(u) = y\Phi^{-1}(u)$, $g_4(u) = u$, with Φ being the normal CDF. Note that the results depend entirely on the choice of basis functions; further research is required to explore suitable models for utility elicitation. Here we consider QPDs solely as another benchmark for comparison, since they provide some flexibility to model a curve ranging from $[0, 1]$.

Estimation with GaSP

We introduce a Bayesian nonparametric method that takes assessed tuples as training data input, regardless of the underlying theory and assumptions used to derive them, and provide an estimated utility function $\hat{u}(\mathbf{x})$, where $\mathbf{x}$ could either be a single attribute or multiple attributes.

Model Formulation

To set notation, let $\mathbf{x} = (x_1, ..., x_p)^T$ be a vector of p different attributes and let $u(\mathbf{x})$ be the utility evaluated at $\mathbf{x}$. Let us consider a random utility function modeled in a general regression way with the form $u(\mathbf{x}) = m(\mathbf{x}) + z(\mathbf{x})$. $m(\mathbf{x})$ is the *mean function*, modeled as:

$$E[u(\mathbf{x})] = m(\mathbf{x}) = \mathbf{h}(\mathbf{x})\boldsymbol{\theta} = \sum_{j=1}^{q} h_j(\mathbf{x})\theta_j, \qquad (1)$$

where $\mathbf{h}(\mathbf{x})$ is assumed to be a q dimensional domain dependent basis function for any $\mathbf{x} \in \mathcal{X}$, with unknown regression parameters θ_j for each basis function $h_j(\mathbf{x})$. $h_j(\mathbf{x})$ could be chosen, e.g., as a particular parametric form or as a polynomial function in $\mathbf{x}$. For the additive residual term, instead of taking $z(\mathbf{x})$ as independent measurement errors as in Eliashberg and Hauser (1985), we model $z(\cdot)$ as a stationary GaSP:

$$z(\cdot) \sim GaSP(0,\, \sigma^2 c(\cdot, \cdot))\,, \qquad (2)$$

with variance σ^2 and pair-wise correlation function $c(\cdot, \cdot)$. In return, the joint distribution of any n inputs $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\} \in \mathscr{X}$, follows a multivariate normal distribution,

$$\left((z(\mathbf{x}_1), \ldots, z(\mathbf{x}_n))^T \mid (\sigma^2,\, \mathbf{C})\right) \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{C})\,, \qquad (3)$$

i.e. a normal distribution that is conditional on the unknown variance σ^2 and the Gram (correlation) matrix $\mathbf{C}$ (Rasmussen and Williams 2006) whose (i, j) element is $c\,(\mathbf{x}_i, \mathbf{x}_j)$. By definition, the covariance of the utility is:

$$\mathrm{Cov}(u(\mathbf{x}_a), u(\mathbf{x}_b)) = \sigma^2 c(\mathbf{x}_a, \mathbf{x}_b),$$

for any $\mathbf{x}_a, \mathbf{x}_b \in \mathcal{X}$. If $p = 1$ (i.e. single attribute), $c(x_a, x_b) = min(x_a, x_b)$ on input domain $[0, +\infty)$ and initial value 0, GaSP becomes a Wiener Process . To extend the definition to the case when $p > 1$, the isotropic assumption is sometimes made for modeling a spatial process (Gelfand et al. 2010), meaning that the correlation function $c(\mathbf{x}_a, \mathbf{x}_b)$ is a function of $||\mathbf{x}_a - \mathbf{x}_b||$ where $|| \cdot ||$ is the Euclidean distance. However, the domain of attributes typically varies on

	$c_l(d_l)$
Power Exponential	$\exp\{-(d_l/\gamma_l)^{\nu_l}\}$, $\nu_l \in (0, 2]$
Spherical	$\left(1 - \frac{3}{2}\left(\frac{d_l}{\gamma_l}\right) + \frac{1}{2}\left(\frac{d_l}{\gamma_l}\right)^3\right)\mathbf{1}_{[d_l/\gamma_l \leq 1]}$
Rational Quadratic	$\left(1 + \left(\frac{d_l}{\gamma_l}\right)^2\right)^{-\nu_l}$, $\nu_l \in (0, +\infty)$
Matérn	$\frac{1}{2^{\nu_l - 1}\Gamma(\nu_l)}\left(\frac{d_l}{\gamma_l}\right)^{\nu_l}\mathcal{K}_{\nu_l}\left(\frac{d_l}{\gamma_l}\right)$, $\nu_l \in (0, +\infty)$

Table 1: Popular choices of correlation functions, when $c_l(x_{al}, x_{bl}) = c_l(d_l)$ with $d_l = |x_{al} - x_{bl}|$ for $l = 1, ..., p$. Here ν_l is the roughness parameter, γ_l is the range parameter, $\Gamma(\cdot)$ is the gamma function and $\mathcal{K}_{\nu_l}(\cdot)$ is the modified Bessel function of second kind of order ν_l.

completely different scales (e.g. between the price and comfort of a car), so the effect of the attributes on the correlations will be highly variable. Consequently, the assumption of isotropy may not be reasonable. Instead, the product correlation function is often assumed:

$$c(\mathbf{x}_a,\, \mathbf{x}_b) = \prod_{l=1}^{p} c_l(x_{al}, x_{bl}), \qquad (4)$$

where $c_l(\cdot,\, \cdot)$ is a one-dimensional correlation function for the l^{th} attribute. We list several frequently used correlation functions in Table 1. The difference between the above product correlation function and the isotropic assumption is that for the former, there are parameter(s) in each correlation $c_l(\cdot,\, \cdot)$ that can control the size of correlation and smoothness of the utility function on this attribute (which could be learned from the data), whereas the isotropic correlation is a function of the Euclidean distance between two attributes.

The power exponential covariance and the Matérn covariance have been used in many applications. When $\nu_l = (2k + 1)/2$ where $k \in \mathbb{N}$, Matérn correlation has a closed form. For example, when the roughness parameter $\nu_l = 5/2$, the Matérn correlation is:

$$c_{Mat}(d_l) = \left(1 + \frac{\sqrt{5}d_l}{\gamma_l} + \frac{5d_l^2}{3\gamma_l^2}\right) exp\left(-\frac{\sqrt{5}d_l}{\gamma_l}\right), \quad (5)$$

where $d_l = |x_{al} - x_{bl}|$. As the roughness parameter (ν_l) is fixed at a chosen value, we only need to estimate the range parameters (γ_l) in the correlation function. The sample path of GaSP with the Matérn correlation defined in equation (5) is twice mean square differentiable (Rasmussen and Williams 2006), allowing one to infer the risk attitude using twice derivatives of the GaSP as discussed later. We found the GaSP model with the Matérn correlation in equation (5) performs well in both simulated and real studies of utility elicitation, but we do not preclude the use of other correlation functions with suitable differentiable results in future applications.

Posterior Predictive Distribution

Denote assessed tuples as $\left(\mathbf{x}^{\mathcal{D}}, \mathbf{u}(\mathbf{x}^{\mathcal{D}})\right)$, where $\mathbf{x}^{\mathcal{D}} = \left\{\mathbf{x}_1^{\mathcal{D}}, \mathbf{x}_2^{\mathcal{D}}, \ldots, \mathbf{x}_n^{\mathcal{D}}\right\}$ are n points in the domain of the multi-attributes. As we are empirically limited by a small number of assessed utility values, we seek the *posterior predictive distribution* of the utility function $u^*(x^*)$ at any input

$x^* \in \mathcal{X}$ based on the assessed tuples $\left(\mathbf{x}^{\mathcal{D}}, \mathbf{u}(\mathbf{x}^{\mathcal{D}})\right)$. For simplicity, denote $\mathbf{u}^{\mathcal{D}} = \left(u(\mathbf{x}_1^{\mathcal{D}}), u(\mathbf{x}_2^{\mathcal{D}}), ..., u(\mathbf{x}_n^{\mathcal{D}})\right)^T$ as the assessed utility points in the design. As per the chosen GaSP model, the likelihood is a multivariate normal likelihood:

$$\mathcal{L}(\mathbf{u}^{\mathcal{D}}|\boldsymbol{\theta}, \sigma^2, \boldsymbol{\gamma}) = (2\pi\sigma^2)^{-n/2}|\mathbf{C}|^{-1/2} \times \exp\left\{-\frac{(\mathbf{u}^{\mathcal{D}} - \mathbf{h}(\mathbf{x}^{\mathcal{D}})\boldsymbol{\theta})^T\mathbf{C}^{-1}(\mathbf{u}^{\mathcal{D}} - \mathbf{h}(\mathbf{x}^{\mathcal{D}})\boldsymbol{\theta})}{2\sigma^2}\right\},$$

where $\mathbf{h}(\mathbf{x}^{\mathcal{D}})$ is the $n \times q$ basis design matrix with (i, j) element $h_j(\mathbf{x}_i^{\mathcal{D}})$. The model parameters for posterior estimation are the mean parameter $\boldsymbol{\theta} = (\theta_1, \theta_2, ..., \theta_q)^T$, variance parameter σ^2 and range parameters $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, ..., \gamma_p)^T$ in the correlation function in equation (5).

We take an objective Bayesian approach using the reference prior for the model parameters (Berger, De Oliveira, and Sansó 2001), and estimate $\boldsymbol{\gamma}$ by the maximum marginal posterior mode with robust parameterization (Gu, Wang, and Berger 2018; Gu 2018). We use robust estimation as the number of assessed tuples is typically small, a scenario where some routinely used methods (e.g. the maximum likelihood estimator) are unstable, leading to a large predictive error; see e.g. Figure 1 in Gu, Wang, and Berger (2018).

The predictive distribution of $u(\mathbf{x}^*)$ at a new point $\mathbf{x}^* \in \mathcal{X}$, given the assessed tuples and the estimated range parameter $\hat{\boldsymbol{\gamma}}$, is a t-distribution with $n - q$ degrees of freedom:

$$u(\mathbf{x}^*) \mid \mathbf{u}^{\mathcal{D}}, \hat{\gamma} \sim t(\hat{u}(\mathbf{x}^*), \hat{\sigma}^2 c^{**}, n - q). \tag{6}$$

The closed form expressions for $\hat{u}(\mathbf{x}^*)$, $\hat{\sigma}^2$ and c^{**} are given in Gu, Wang, and Berger (2018).

The predictive mean $\hat{u}(\mathbf{x}^*)$ will be used for estimation of a utility function at any $\mathbf{x}^*$. The predictive mean estimator at any $\mathbf{x}_i^{\mathcal{D}}$ is an *interpolator* meaning that $\hat{u}(\mathbf{x}_i^{\mathcal{D}}) = u(\mathbf{x}_i^{\mathcal{D}})$ (Gu and Berger 2016). For noisy assessments, we do not expect the prediction of GaSP to be exact at the assessed points. In this scenario, an independent noise can be added in the model by defining $\tilde{z}(\mathbf{x}) = z(\mathbf{x}) + \varepsilon$, where ε is independent white noise. The objective Bayesian inference for such a GaSP model is similar to the method discussed above (Ren, Sun, and He 2012).

Derivatives of the Utility Function

Differentiability is an important property of utility functions, e.g. the Arrow-Pratt measure of local risk aversion is defined as $\lambda(x) = -u''(x)/u'(x)$. Our proposal to use GaSP is helpful in this regard since the derivative processes are also GaSP when the covariance function is mean square differentiable (Rasmussen and Williams 2006).

For demonstration purposes, we derive the first and second order derivative processes of the Matérn class correlation with roughness parameter equal to 2.5[1]. The result can be easily extended to directional derivative processes with regard to each attribute in the multi-attribute case, i.e. $\frac{\partial u(\mathbf{x})}{\partial x_l}$, for $l = 1, ..., p$. The risk attitude can be assessed using the predictive distribution of the derivative processes. Unlike the estimation by linear interpolation (Abdellaoui, Bleichrodt, and Paraschiv 2007), our approach for estimating the risk attitude enables full assessment of the uncertainty.

[1]Please see the arXiv version for closed form expressions of the derivative processes: https://arxiv.org/abs/1807.10840

Synthetic Data Experiments

In this section, we explore practical ramifications through experiments with synthetic data. In preference elicitation, only a limited number of questions can be posed, thus the number of assessed tuples is typically small: $n = 7$ and 10 are considered herein. Out of sample mean squared error $MSE = \sum_{i=1}^{n^*} \{\hat{u}(x_i^*) - u(x_i^*)\}^2/n^*$ is utilized for comparison, where $x_i^* \in \mathcal{X}$ is the i^{th} equally spaced held-out point and $n^* = 1,001$ is used for testing throughout this section. We assume $U(x_{min}) = 0$ and $U(x_{max}) = 1$, where $x_{min} = 0$ and $x_{max} = 10^5$ are lower and upper bounds of $\mathcal{X}$ in simulated studies. Assuming ground truth of a power utility function, we compare the exponential function (Exp), linear interpolation (LI) and quantile-parameterized distribution (QPD) method with GaSP. The experiments were repeated with other functions, yielding similar results; these are omitted due to space limitation.

For the parametric methods and QPD method, we estimate the parameters with the minimum least squares error. For the QPD, we choose the basis function to be $g_1(u) = 1$, $g_2(u) = \Phi_*^{-1}(u)$, $g_3(u) = u\Phi_*^{-1}(u)$, $g_4(u) = u$. As the domain of x is $[0, 10^5]$, the usual Q-normal distribution is not a sensible choice. Instead, we let $\Phi_*^{-1}(u)$ be a normal distribution truncated at 0 and 10^5 centered at 0 with standard deviation 5×10^4. This seems to perform the best among all basis functions we explored.

Comparing Utility Function Estimates

First we assume that assessments are noise-free. Table 2 displays the out of sample MSE when the underlying utility function U is power. Since the power utility function is a subclass of models contained within the GaSP framework with mean basis $h(x) = x^\alpha$ and variance $\sigma = 0$, we choose the mean function of the GaSP to be misspecified by selecting an inconsistent mean basis $(x^{0.5})$, to highlight that GaSP performs well even in this scenario. GaSP outperforms the other methods and the discrepancy is usually several orders of magnitude less than that for parametric fitting, and it is usually ten to hundred times better than LI.

Method	$\alpha = 0.7$	$\alpha = 1.5$	$\alpha = 2.5$
Exp	5.6×10^{-4}	3.2×10^{-4}	6.9×10^{-4}
LI	9.5×10^{-6}	4.8×10^{-5}	3.8×10^{-4}
GaSP	$\mathbf{2.8 \times 10^{-7}}$	$\mathbf{5.5 \times 10^{-6}}$	$\mathbf{5.0 \times 10^{-6}}$
QPD	2.2×10^{-5}	2.8×10^{-4}	2.8×10^{-3}

Table 2: Out of sample MSE for noise-free assessments with $n = 7$. U is assumed to be power with different α.

Next we assume noisy assessments, specifically that utilities are random with additive noise:

$$u(x) = U(x) + \epsilon;\ \epsilon \mid \sigma_\epsilon^2 \sim \mathcal{N}(0, \sigma_\epsilon^2), \tag{7}$$

Method	$\alpha = 0.7$	$\alpha = 1.5$	$\alpha = 2.5$
Exp	5.6×10^{-4}	3.2×10^{-3}	6.9×10^{-4}
LI	1.8×10^{-5}	3.4×10^{-5}	1.9×10^{-4}
GaSP	$\mathbf{9.7 \times 10^{-6}}$	$\mathbf{1.7 \times 10^{-5}}$	$\mathbf{1.3 \times 10^{-4}}$
QPD	2.4×10^{-5}	2.1×10^{-4}	2.1×10^{-3}

Table 3: $AvgMSE$ for noisy assessments with $n = 10$ and when U is assumed to be power, with $\sigma_\epsilon = 0.005$.

where $U(x)$ is the underlying utility function and σ_ϵ^2 is the variance of the Gaussian noise. Since the assessed points are noisy, we simulated $N = 200$ experiements to average out the design effect, and calculate the average MSE: $AvgMSE = \frac{1}{N}\sum_{j=1}^{N} MSE_j$. Here the total held-out testing points for each case is thus $n^* \times N = 200, 200$.

Compared to the results of noise-free assessments, all methods have comparatively large MSE for noisy assessments, as shown in Table 3. Indeed, if the elicited utilities are dominated by noise, none of the methods works as well as the noise-free cases. Compared to all baselines, GaSP still has the smallest predictive error.

Comparing Derivatives

In Abdellaoui, Bleichrodt, and Paraschiv (2007) and Abdellaoui, Bleichrodt, and l'Haridon (2008), the curvature of utility functions is classified as either concave, convex or of mixed type, using LI or parametric fitting. Such classification is *global* and characterizes the dominant risk attitude implied by the assessed utility function throughout the entire domain. In the LI method, empirical derivatives of assessed utility points are utilized for estimation of curvature. Let P_i be an observed utility middle point between two neighboring elicited utility points P_{i-1} and P_{i+1}. Denote $S^-(P_i)$ and $S^+(P_i)$ as slopes of the straight line between P_i to P_{i-1} and P_i to P_{i+1} respectively. $\Delta S(P_i) = S^+(P_i) - S^-(P_i)$ is used as the estimate of convexity at point P_i. The utility function is typically estimated to be concave (convex) if more than $\approx 2/3$ elicited utility points are estimated to be concave (convex), otherwise it is denoted as a mixed type.

In the GaSP method, one can compute the posterior distribution of the second derivative for *any* point x^* in the domain as mentioned previously. When $\Pr(u''(x^*) \leq 0 | u(x_1), \ldots, u(x_n)) > 0.5$, the utility function is predicted to be concave at x^*, otherwise convex. In the following simulations, we compute the predictive distribution on $n^* = 10,000$ equally spaced inputs x_i^*, $i = 1, ..., n^*$, and use the proportion of points that are predicted to be concave/convex to predict the overall curvature of the function.

The results for estimating global concavity by LI and GaSP are shown in Figure 2. In the top row of Figure 2, the underlying utility functions are all concave. Due to the effect of noisy assessment, the proportion of concave points by LI is between $1/3$ to $2/3$ in most of the experiments, meaning that LI fails to identify the concavity of the functions, classifying them of mixed type instead. In the bottom row of Figure 2, when the underlying utility function is convex, LI fails to identify the convexity of the utility functions

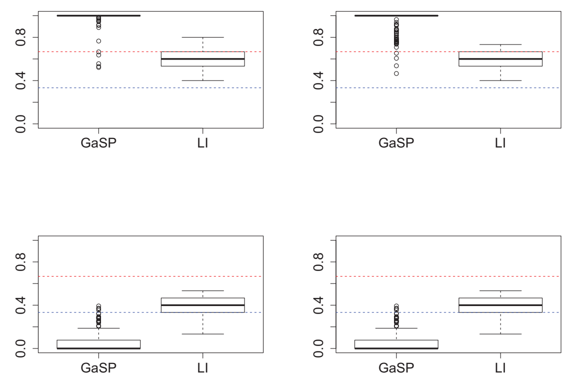

Figure 2: Boxplot of the proportion of points predicted as **concave** for $N = 500$ experiments. Upper figures show results for noisy assessment of concave exponential utilities with $\sigma_\epsilon = 0.01$, $\rho = 2/3$ (left) and $\sigma_\epsilon = 0.02$, $\rho = 1/3$ (right). Lower figures show results for noisy assessment of convex exponential utilities with $\sigma_\epsilon = 0.01$, $\rho = -1/2$ (left) and $\sigma_\epsilon = 0.02$, $\rho = -1/4$ (right). $n = 15$ for all.

and again classifies a majority of points as the mixed type. Compared to LI, GaSP predicts curvature more accurately. In the top row, the proportions of concavity points are almost all close to 1 across $N = 500$ experiments. In the bottom row, most points are predicted to be convex for a majority of experiments. Using $2/3$ as the threshold would correctly classify most of the utility functions using GaSP.

There are two main reasons why GaSP performs better. First, GaSP prediction of concavity of point x^* utilizes information from all assessed tuples rather than just the two neighboring points as in LI. Second, prediction of concavity by GaSP is averaged by many predictive samples of second derivatives ($n^* = 10,000$ chosen here) over the entire attribution domain rather than the limited number of n assessed tuples in LI. GaSP estimation is therefore better at analyzing preference properties like risk attitude.

Real Data Experiments

Single Attribute Dataset

Let us compare the parametric and nonparametric methods using a real dataset from Abdellaoui, Bleichrodt, and Paraschiv (2007), collected from a prospect theory based scheme. $k = 48$ people answered a series of questions about comparisons between risky gambles. Due to the effect of loss aversion, the range of the loss domain (negative outcomes) is assumed to be $[-1, 0]$ and the range of the gain domain (positive outcomes) is constructed to be $[0, 0.25]$, with 11 and 7 assessed tuples in each domain respectively. See Abdellaoui, Bleichrodt, and Paraschiv (2007) for details about the design and elicited scheme.

To test the predictive performance of different methods, we randomly sample $n^*_{loss} = 4$ and $n^*_{gain} = 3$ assessed tuples in the loss and gain domains for each person respec-

AvgMSE	GaSP	LI	Pow	Exp
Loss	$\mathbf{8.9 \times 10^{-4}}$	9.7×10^{-4}	1.5×10^{-3}	1.7×10^{-3}
Gain	$\mathbf{7.4 \times 10^{-5}}$	8.2×10^{-5}	1.1×10^{-4}	1.1×10^{-4}

Table 4: Average out of sample MSE for losses and gains using GaSP, linear interpolation (LI), power (Pow) and exponential (Exp) in the single attribute dataset. MSE is averaged over $n^*_{loss}kN = 96$K and $n^*_{gain}kN = 72$K respectively.

tively, saving them as the test set, while the remaining assessed tuples are used as the training set. We repeat this experiment $N = 500$ times, and fit each method for each experiment. We see in Table 4 that the average out of sample MSE for GaSP is similar to but better than LI. It is much smaller than those for the power and exponential functions. The parametric fit, albeit straightforward in interpretation, induces extra assumptions and may lead to poor estimates.

Multiple Attribute Dataset

To demonstrate the performance of the GaSP model for utility functions with multiple attributes, we study a real dataset with three attributes (Fischer, Jia, and Luce 2000). This experiment was conducted among 22 students at Duke about decisions pertaining to course selection involving three attributes. The attributes $\mathbf{x} = (x_1, ..., x_3)$ include the degree of interest, expected teaching quality and the average grade, each of which has 5 levels. The output $u(\mathbf{x})$ is the utility rating of a course given a set of attributes. Each volunteer is asked to rate the same 20 courses along the three attributes.

Fischer, Jia, and Luce (2000) proposed *RandMAU* as a sampling model to capture the stimulus properties of *attribute conflict* and *attribute extremity*. The model is intended to characterize potentially inconsistent responses to preference elicitation questions through a stochastic model of the assessed utility. It is defined as $u(x_1, ..., x_p) = \sum_{i=1}^{p} \omega_i u_i(x_i) + \omega \sum_{i=1}^{p} \sum_{j>i}^{p} \omega_i u_i(x_i) \omega_j u_j(x_j) + ... + \omega^{p-1} \prod_{i=1}^{p} \omega_i u_i(x_i)$, where $1 + \omega = \prod_{i=1}^{p}(1 + \omega\omega_i)$ and $\omega_i \sim_{ind} Beta(r_i, u_i)$. $p = 3$ is the number of attributes and $u_i(x_i) = \left[\frac{x_i - x_{i0}}{x_i^* - x_{i0}}\right]^{\alpha_i}$ with x_{i0} and x_i^* as the lower and upper limits for attribute x_i. This model has 7 parameters $(\omega_1, \omega_2, \omega_3, \omega, \alpha_1, \alpha_2, \alpha_3)$ but when $(\omega_1, \omega_2, \omega_3)$ are known, ω can be uniquely solved, which leaves the model with 6 degrees of freedom. The authors specify two ways of estimating parameters, namely the *corner point* and *nonlinear least squares* estimation methods. The first approach uses only 7 out of the 20 data points for each participant to fit the model while the second minimizes the squared error using all assessed tuples.

We compare the afore-mentioned two approaches with our proposed GaSP approach for modeling the average ratings (shown in Table 1 in Fischer, Jia, and Luce (2000)) based on the attributes. We compute the out of sample MSE and R^2 from our proposed GaSP estimation from the RandMAU-based methods of corner point and nonlinear least squares estimation, using the 13 data points that are not used in the corner points approach. Since the sample size is very small, each time we only leave a data point out and compute the MSE and R^2 for this point. The average out of sample MSE and R^2 are shown in Table 5.

	corner point	nonlinear least square	GaSP
MSE	0.0084	0.0056	**0.0017**
R^2	0.9087	0.9391	**0.9813**

Table 5: Out of sample performance using corner point, nonlinear least square and GaSP estimation in the multiple attribute dataset.

In Table 5, we find that the prediction by the GaSP model is several times better than the previous inference methods because it is flexible, while RandMAU is restrictive since each $u_i(x_i)$ is assumed to be the power utility function and thus cannot capture other shapes. For the RandMAU methods, we find that nonlinear least square estimates are better than corner point estimates in terms of MSE and R^2. This is because only 7 observations are used for prediction in corner point estimation, which is inefficient in estimation.

Conclusions

We have presented a nonparametric Bayesian approach involving GaSP for inferring a decision maker's utility function using assessed tuples as training data, regardless of the choice of elicitation protocol and underlying theory. We describe theoretical benefits over parametric approaches around the desired property of interpolation for noise-free assessment. Unlike the linear interpolation approach, the proposed method guarantees the differentiability of the estimated utility function by choosing an appropriate correlation function, Our nonparametric estimation approach is flexible to fit a large class of functions, and the predictive distribution provides probabilistic quantification of the decision maker's preference properties, such as the risk attitude.

Simulated experiments confirm that the GaSP model has lower predictive error than parametric method, as well as the quantile parametrized distribution method and linear interpolation, even when the simulated data are corrupted with additive Gaussian noise. Our approach also has a smaller out-of-sample predictive error of estimating utility functions and risk attitudes using real data sets from the literature, which demonstrates the effectiveness of our method.

Utility functions have specific functional characteristics. For instance, a utility function is often assumed to be monotonic in its arguments. A potential avenue for future work is around constrained GaSPs, which would extend the proposed approach to respect monotonicity.

Acknowledgments

The authors thank three anonymous reviewers for useful feedback and suggestions for improving the paper.

References

Abdellaoui, M.; Bleichrodt, H.; and l'Haridon, O. 2008. A tractable method to measure utility and loss aversion under prospect theory. *Journal of Risk and Uncertainty* 36(3):245–266.

Abdellaoui, M.; Bleichrodt, H.; and Paraschiv, C. 2007. Loss aversion under prospect theory: A parameter-free measurement. *Management Science* 53(10):1659–1674.

Abdellaoui, M. 2000. Parameter-free elicitation of utility and probability weighting functions. *Management Science* 46(11):1497–1512.

Bell, D. E. 1979. Multiattribute utility functions: Decompositions using interpolation. *Management Science* 25(8):744–753.

Berger, J. O.; De Oliveira, V.; and Sansó, B. 2001. Objective Bayesian analysis of spatially correlated data. *Journal of the American Statistical Association* 96(456):1361–1374.

Bhattacharjya, D., and Kephart, J. O. 2014. Bayesian interactive decision support for multi-attribute problems with even swaps. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, 72–81. AUAI Press.

Birlutiu, A.; Groot, P.; and Heskes, T. 2009. Multi-task preference learning with Gaussian processes. In *Proceedings of the 17th European Symposium on Artificial Neural Networks*, 123–128.

Bleichrodt, H.; Pinto, J. L.; and Wakker, P. T. 2001. Making descriptive use of prospect theory to improve the prescriptive use of expected utility. *Management Science* 47(11):1498–1514.

Bonilla, E. V.; Guo, S.; and Sanner, S. 2010. Gaussian process preference elicitation. In *Advances in Neural Information Processing Systems*, 262–270.

Boutilier, C. 2002. A POMDP formulation of preference elicitation problems. In *Proceedings of the 18th National Conference on Artificial Intelligence*, 239–246.

Chajewska, U.; Koller, D.; and Parr, R. 2000. Making rational decisions using adaptive utility elicitation. In *Proceedings of the 17th National Conference on Artificial Intelligence*, 363–369.

Chu, W., and Ghahramani, Z. 2005. Preference learning with Gaussian processes. In *Proceedings of the 22nd International Conference on Machine Learning*, 137–144. ACM.

Debreu, G. 1954. Representation of a preference ordering by a numerical function. *Decision Processes* 3:159–165.

Eliashberg, J., and Hauser, J. R. 1985. A measurement error approach for modeling consumer risk preference. *Management Science* 31(1):1–25.

Farquhar, P. H. 1984. Utility assessment methods. *Management Science* 30(11):1283–1300.

Fischer, G. W.; Jia, J.; and Luce, M. F. 2000. Attribute conflict and preference uncertainty: The RandMAU model. *Management Science* 46(5):669–684.

Gelfand, A. E.; Diggle, P.; Guttorp, P.; and Fuentes, M. 2010. *Handbook of Spatial Statistics*. CRC Press.

Gu, M., and Berger, J. O. 2016. Parallel partial Gaussian process emulation for computer models with massive output. *The Annals of Applied Statistics* 10(3):1317–1347.

Gu, M.; Palomo, J.; and Berger, J. O. 2019. RobustGaSP: Robust Gaussian stochastic process emulation in R. *The R Journal* 11(1).

Gu, M.; Wang, X.; and Berger, J. O. 2018. Robust Gaussian stochastic process emulation. *Annals of Statistics* 46(6A):3038–3066.

Gu, M. 2018. Jointly robust prior for Gaussian stochastic process in emulation, calibration and variable selection. *Bayesian Analysis* 14(1).

Jimison, H. B.; Fagan, L. M.; Shachter, R. D.; and Shortliffe, E. H. 1992. Patient-specific explanation in models of chronic disease. *Artificial Intelligence in Medicine* 4(3):191–205.

Kahneman, D., and Tversky, A. 1979. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society* 263–291.

Karlin, S. 1975. *A First Course in Stochastic Processes*. Academic Press.

Keelin, T. W., and Powley, B. W. 2011. Quantile-parameterized distributions. *Decision Analysis* 8(3):206–219.

Kirkwood, C. W. 2004. Approximating risk aversion in decision analysis applications. *Decision Analysis* 1(1):51–67.

Lichtenstein, S., and Slovic, P. 2006. *The Construction of Preference*. Cambridge, U.K.: Cambridge University Press.

Mörters, P., and Peres, Y. 2010. *Brownian Motion*, volume 30. Cambridge University Press.

Poh, K. L., and Horvitz, E. J. 1993. Reasoning about the value of decision-model refinement: Methods and application. In *Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence*, 174–182. Morgan Kaufmann.

Pratt, J. W. 1964. Risk aversion in the small and in the large. *Econometrica* 32:122–136.

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

Ren, C.; Sun, D.; and He, C. 2012. Objective Bayesian analysis for a spatial model with nugget effects. *Journal of Statistical Planning and Inference* 142(7):1933–1946.

Sacks, J.; Welch, W. J.; Mitchell, T. J.; and Wynn, H. P. 1989. Design and analysis of computer experiments. *Statistical Science* 4(4):409—423.

Tversky, A., and Kahneman, D. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4):297–323.

von Neumann, J., and Morgenstern, O. 1947. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press, 2nd edition.