

# A Semi-parametric Model for Decision Making in High-Dimensional Sensory Discrimination Tasks

Stephen Keeley<sup>1,2</sup>, Benjamin Letham<sup>2</sup>, Craig Sanders<sup>2</sup>, Chase Tymms<sup>2</sup>, Michael Shvartsman<sup>2</sup>

<sup>1</sup> Department of Natural Sciences, Fordham University, USA

<sup>2</sup> Meta

skeeley1@fordham.edu, {bletham,tymms,craigsanders,mshvarts}@meta.com

## Abstract

Psychometric functions typically characterize binary sensory decisions along a single stimulus dimension. However, real-life sensory tasks vary along a greater variety of dimensions (e.g. color, contrast and luminance for visual stimuli). Approaches to characterizing high-dimensional sensory spaces either require strong parametric assumptions about these additional contextual dimensions, or fail to leverage known properties of classical psychometric curves. We overcome both limitations by introducing a semi-parametric model of sensory discrimination that applies traditional psychophysical models along a stimulus intensity dimension, but puts Gaussian process (GP) priors on the parameters of these models with respect to the remaining dimensions. By combining the flexibility of the GP with the deep literature on parametric psychophysics, our semi-parametric models achieve good performance with much less data than baselines on both synthetic and real-world high-dimensional psychophysics datasets. We additionally show strong performance in a Bayesian active learning setting, and present a novel active learning paradigm for the semi-parametric model.

## Introduction

Understanding the mappings from physical stimuli to mental percepts is an important goal of perceptual neuroscience and psychophysics. A popular experimental technique is to measure behavioral responses to varying levels of a single stimulus feature, such as contrast of an image or volume of a sound, and generate a behavioral response curve. The average response probabilities are then often fit to a functional form that is sigmoidal in shape, for example using the probit, logit, or Weibull functions (Strasburger 2001). These parametric models tend to have specific interpretable parameters that are of use to practitioners, such as thresholds and slopes (Brand and Kollmeier 2002).

Improving the accuracy and sample efficiency of parameter estimation for these psychometric functions remains an area of active research (Schütt et al. 2016; Shen and Richards 2012). However, much of the existing work focuses on a single stimulus feature, ignoring the fact that stimuli continuously vary in important dimensions other than intensity, such

as color or pitch. To understand sensory sensitivity in these settings, univariate parameterized models require densely sampling the entire stimulus feature domain; i.e. creating a single psychometric function for each value of any non-intensity stimulus feature.

Recent work overcomes these limitations and extends the classical densely-sampled univariate psychophysical curve in terms of both modeling and experimental design. On the modeling front, recent work captures correlations in psychometric functions across stimulus dimensions, either using prespecified multidimensional parametric models (Watson 2017), or using nonparametric Gaussian processes (GPs) (Gardner et al. 2015a,b; Owen et al. 2021). For experimental design, dense sampling of the entire multidimensional input space has been replaced with efficient active learning schemes (Houlsby et al. 2011; Settles 2009; Gardner et al. 2015a,b; Watson 2017). Together, flexible psychometric models and multidimensional adaptive sampling methods have led to improvements in estimating psychometric tuning in multivariate stimulus settings (Gardner et al. 2015a; Owen et al. 2021; Letham et al. 2022).

However, existing high-dimensional psychometric models suffer from some limitations. The fully parametric approach, while interpretable, is constrained in practice, requiring an a priori model for all dimensions, with parameters learned using an inefficient grid-search (Watson 2017). On the other hand, the more flexible, nonparametric models that leverage the power of GPs are not regularized to have sensible tuning properties, such as positive monotonicity in stimulus intensity. Moreover, they do not afford an experimenter with interpretable parameters like those given in the classical univariate context (Gardner et al. 2015b; Owen et al. 2021).

To address these limitations, we propose a semi-parametric model of the psychometric field well-suited for high-dimensional stimuli. Over the intensity dimension, our psychometric model is governed by a characteristic function that has a sigmoidal shape with identifiable slope and offset parameters. Each of these parameters, however, is governed by a GP across context (non-intensity) stimulus dimensions, admitting a flexible characterization of the psychometric field in a high-dimensional continuous stimulus space. While the posterior of our model is intractable, we develop two approximations, each with different benefits. The first, the *full semi-parametric model*, uses a semi-parametric

variational posterior that factorizes over the slope and offset parameters of the sigmoid to learn the model. For the second, the *MVN approximate model*, we derive a new approximation for the elementwise product of the multivariate normal (MVN) distributions of the slope and offset GPs in our model. This approximation implicitly defines a new, single GP kernel specific to the psychophysics setting, and lets us perform inference with fewer variational parameters, maintaining a form that is adaptable to standard GP inference methods and active learning machinery. We evaluate our model on simulated and real data in up to 8 dimensions, and find that this semi-parametric approach not only provides interpretable results in high-dimensional stimulus settings, but also offers a faster and more accurate estimation procedure for the psychometric field.

Given that Bayesian active learning also plays an important role in high dimensional psychophysical tuning estimation (Gardner et al. 2015a; Owen et al. 2021; Letham et al. 2022), we conclude our model evaluation showing semi-parametric performance under a variety of existing active learning objectives. We find that our full semi-parametric model shows strong performance under a variety of these objectives, and our MVN approximate model allows for analytic acquisition functions, including recently-developed look-ahead approaches (Letham et al. 2022). Finally, we introduce a novel threshold-based acquisition function for use with our full semi-parametric model that shows strong active learning performance in a high-dimensional experiment.

### The Semi-parametric Psychophysical Model

We consider data of the form  $D = \{\mathbf{x}_n, y_n\}_{n=1}^N$  where  $y_n \in \{0, 1\}$  are participant responses and  $\mathbf{x}_n = (x_{i_n}, \mathbf{x}_{s_n})$  describe stimulus configurations. We separate stimulus context ( $s$ ) dimensions from the intensity ( $i$ ) dimension. Throughout the manuscript we only consider 1-dimensional intensity, so that  $x_{i_n}$  is a scalar. We assume that intensity predicts responses according to a standard psychophysical parametric model of the form  $p = \sigma(k(x_i + c))$  for some slope  $k$ , offset  $c$ , and sigmoid link  $\sigma: \mathbb{R} \rightarrow [0, 1]$  and probability of detection  $p$ . Our framework can be flexibly adapted to incorporate any sigmoid from the literature, including the probit (as standard in GP classification), logistic (common elsewhere in machine learning), Weibull or Gumbell CDF (common in psychophysics, (Strasburger 2001; May and Solomon 2013)), or other common probability mappings. We can additionally modify these sigmoids, for example shifting or scaling them to manage psychophysics experiments where participants are asked to discriminate between stimuli, and the response probability is lower-bounded above 0. Common experimental paradigms can produce a lower bound response probability at 0.5 (for two-alternative decisions) or 0.25 (for four-alternative ‘odd one out’ trials).

We place independent GP priors on the slope and offset governed by the context dimensions, which implies that each of  $\mathbf{f}_k = [k(\mathbf{x}_{s_1}), k(\mathbf{x}_{s_2}), \dots, k(\mathbf{x}_{s_N})]$  and  $\mathbf{f}_c = [c(\mathbf{x}_{s_1}), c(\mathbf{x}_{s_2}), \dots, c(\mathbf{x}_{s_N})]$  is Gaussian distributed:

$$\mathbf{f}_k \sim \mathcal{N}(m, \Sigma_k), \mathbf{f}_c \sim \mathcal{N}(0, \Sigma_c), \quad (1)$$

where  $\Sigma_k, \Sigma_c$  are  $N \times N$  covariance matrices whose

$(n, n')$ th entry is given by a kernel function  $\kappa(\mathbf{x}_{s_n}, \mathbf{x}_{s_{n'}})$  and  $N$  is the total number of stimuli sampled. We use the standard radial basis function (RBF) kernel with independent hyperparameters for the slope and offset GP kernels. Each kernel is governed by its own hyperparameters  $\theta_c$  and  $\theta_k$  for the offset and slope GPs, respectively. Finally,  $m$  is a positive constant to center the prior distribution of slope values at some positive number (we use  $m = 2$  for all experiments; see appendix for additional detail and evaluation).

In this formulation, we can write the joint distribution of latent stimulus values as:

$$\mathbf{z} = \mathbf{f}_k \circ (\mathbf{f}_c + \mathbf{x}_i), \quad (2)$$

where  $\circ$  denotes the Hadamard (elementwise) product and  $\mathbf{x}_i = (x_{i_1}, x_{i_2}, \dots, x_{i_N})$ . The likelihood of a set of observations is given by a set of independent Bernoulli distributions with probabilities equal to  $\mathbf{z}$  taken through the sigmoid link. To make sense of this mapping, consider a joint draw of  $N$  slopes  $\mathbf{f}_k \in \mathbb{R}^N$  and intercepts  $\mathbf{f}_c \in \mathbb{R}^N$ . Suppressing the dependence of  $k$  and  $c$  on  $\mathbf{x}_s$ , the resultant transformed collection of variables has each  $z_n = k_n c_n + k_n x_{i_n}$ , producing  $\mathbf{z} = [k_1 c_1 + k_1 x_{i_1}, k_2 c_2 + k_2 x_{i_2}, \dots, k_N c_N + k_N x_{i_N}]$ . Note that this is just a multivariate extension of the standard  $k(x + c)$  input into the link, written jointly to accommodate the GP prior on  $k$  and  $c$ . A graphical depiction of the semi-parametric model is given in Fig. 1A.

Fig. 1B shows prior samples of the probability of detection,  $p$ , plotted along the intensity dimension, for three different models. The first is an unconstrained GP model (GP-RBF) which treats all dimensions equally (Owen et al. 2021), and the second has an additive GP kernel that is linear in the intensity dimension and RBF in context dimensions (GP-linear) (Gardner et al. 2015b). The third (Fig. 1B, Right) is our proposed semi-parametric model. These probability samples for each model are drawn from the prior and illustrate the hypothesis space of the considered psychometric functions. The GP-RBF model allows for tuning that is not sigmoidal (or even monotonic), and the GP-linear model restricts the shape of the sigmoid while still permitting negative slopes. In contrast, our semi-parametric model shows classical psychometric curves along the intensity dimension and is free to vary independently in its offset  $c$  and slope  $k$ .

Our semi-parametric model is easily adaptable to use any link function with interpretable slope and intercept values, as well as scaled and/or shifted variants thereof. In the present contribution we consider the logistic and Probit mapping, as well as the so-called Weibull function which is the cdf of a left Gumbel distribution in log-space (Strasburger 2001; May and Solomon 2013). We also consider a ‘floor link’ whose minimum is set to the known lower bound of the response probability in the given task. For our results on simulated functions, we choose the link that performed best with quasi-random sampling for each model<sup>1</sup>. We show results with all links on human psychophysical data and include results from all links and floors in the appendix.

<sup>1</sup>This is the Gumbel-link with floor for the semi-parametric model for all except for GlobalMI acquisition, which requires a probit link with a floor of 0. For all other models this is the Probit link with a floor of 0.

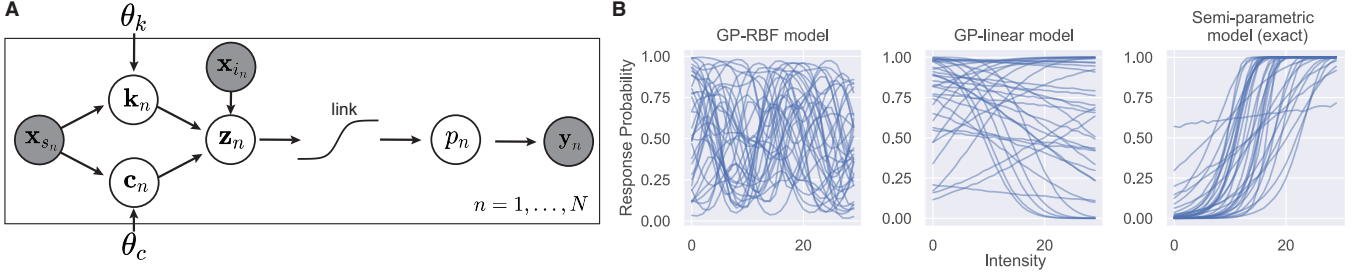


Figure 1: **A.** Graphical model of our semi-parametric approach. **B.** Samples from the prior along the intensity dimension for three psychophysical models. The GP-RBF model’s prior contains unrealistic psychometric functions that are not monotonic. The previously published GP-linear model overly restricts the shape of the sigmoid due to the linear kernel in intensity, while admitting functions with negative slopes. Our new model’s prior, here with a Gumbel link, contains a more diverse set of realistic psychometric functions than either baseline variant.

## Inference for the Semi-parametric Psychophysical Model

The marginal likelihood for our model is:

$$P(\mathbf{y}|\mathbf{X}_s, \mathbf{x}_i, \theta_k, \theta_c) = \int p(\mathbf{y}|\mathbf{x}_i, \mathbf{f}_k, \mathbf{f}_c) p(\mathbf{f}_k|\mathbf{X}_s, \theta_k) p(\mathbf{f}_c|\mathbf{X}_s, \theta_c) d\mathbf{f}_k d\mathbf{f}_c, \quad (3)$$

where  $\mathbf{X}_s$  is a concatenation of  $(\mathbf{x}_{s,1}, \mathbf{x}_{s,2} \dots \mathbf{x}_{s,N})$  and  $\mathbf{y}$  is a concatenation of the  $(y_1, y_2 \dots y_n)$  observations. This likelihood is intractable, but we provide two distinct strategies for approximating it: one by factorized variational inference (VI), and the other by approximating the model itself, which lets us apply standard VI methods for GPs.

### Factorized Variational Inference for the Semi-parametric Model

We define MVN variational distributions  $q_k$  and  $q_c$  for the slope and offset to perform VI. We can write an evidence lower bound (ELBO) as follows:

$$\begin{aligned} \log p(\mathbf{y} | \mathbf{X}_s, \mathbf{x}_i) \geq \\ \mathcal{L} := & \mathbb{E}_{q_c(\mathbf{f}_c), q_k(\mathbf{f}_k)} [\log p(\mathbf{y} | \mathbf{f}_k, \mathbf{f}_c, \mathbf{x}_i)] \\ & - \text{KL}[q_k(\mathbf{f}_k) \| p(\mathbf{f}_k | \mathbf{X}_s)] \\ & - \text{KL}[q_c(\mathbf{f}_c) \| p(\mathbf{f}_c | \mathbf{X}_s)], \end{aligned}$$

where the two KL terms between MVNs are available in closed form. For the remaining term, since the slope and offset GPs are combined into a scalar latent term per observation, we can compute the expectation by one-dimensional Gauss-Hermite quadrature. Gradients of all terms are available by automatic differentiation, which lets us optimize this objective by standard methods (Hensman, Matthews, and Ghahramani 2015; Balandat et al. 2020).

### Approximation to the Semi-parametric Model

While the inference approach above is tractable, its key disadvantage is that the prior on the latent  $\mathbf{z}$  is no longer a Gaussian Process prior. This limits our ability to use standard black box GP variational inference for this model, as

well as other tooling that relies on an MVN prior. Furthermore, since we have MVN posteriors on both slope and offset, the variational approximation has twice as many parameters, which potentially slows down inference (of interest for human-in-the-loop applications). To address these disadvantages, we derive an approximation to our model that avoids the need to use a variational approximation for both the slope and offset GP function values. Specifically, we directly approximate the latent function  $\mathbf{z}$  with an MVN using moment-matching.

We are interested in a MVN approximation to the latent function:

$$Z = F_k \circ (F_c + \mathbf{x}_i).$$

Here,  $F_k, F_c$  are random variables distributed according to the Gaussians in (1), and  $Z$  is now also a random variable (in contrast to  $\mathbf{f}_k, \mathbf{f}_c, \mathbf{z}$ , indicating realizations of the random variables). For convenience of derivation, let  $\tilde{F}_c = F_c + \mathbf{x}_i$ . The GP prior on  $F_c$  implies that  $\tilde{F}_c \sim \mathcal{N}(\mathbf{x}_i, \Sigma_c)$ . We can compute the mean of  $Z$  in terms of the slope and offset distributions:

$$\mathbb{E}[Z] = \mathbb{E}[F_k \circ \tilde{F}_c] = m\mathbf{x}_i.$$

We can similarly compute the covariance of the latent  $Z$  in terms of the slope and offset. For convenience, denote  $Y_k = F_k - m$  and  $Y_c = \tilde{F}_c - \mathbf{x}_i$ . We have that

$$\begin{aligned} Z &= (Y_k + m) \circ (Y_c + \mathbf{x}_i) \\ &= m\mathbf{x}_i + mY_c + \mathbf{x}_i \circ Y_k + Y_k \circ Y_c, \end{aligned} \quad (4)$$

with  $Y_c \sim \mathcal{N}(0, \Sigma_c)$  and  $Y_k \sim \mathcal{N}(0, \Sigma_k)$ . The covariance of the cross-term can be computed as

$$\begin{aligned} \text{Cov}[Y_k \circ Y_c]_{i,j} &= \mathbb{E}[(Y_k \circ Y_c)_i (Y_k \circ Y_c)_j] \\ &= \mathbb{E}[(Y_{k,i} Y_{c,i}) (Y_{k,j} Y_{c,j})] \\ &= \mathbb{E}[Y_{k,i} Y_{k,j}] \mathbb{E}[Y_{c,i} Y_{c,j}] \\ &= (\Sigma_k)_{i,j} (\Sigma_c)_{i,j}, \end{aligned} \quad (5)$$

where (5) uses that  $Y_k$  and  $Y_c$  have 0 mean, and (6) uses their independence. Thus,  $\text{Cov}[Y_k \circ Y_c] = \Sigma_k \circ \Sigma_c$ .

Applying this result to (4) we can compute the covariance for the latent function:

$$\begin{aligned} \text{Cov}[Z] &= m^2 \Sigma_c + \mathbf{x}_i \mathbf{x}_i^T \circ \Sigma_k + \Sigma_c \circ \Sigma_k \\ &= m^2 \Sigma_c + (\Sigma_c + \mathbf{x}_i \mathbf{x}_i^T) \circ \Sigma_k. \end{aligned}$$

The primary benefit of this approximation is that rather than learn the full semi-parametric latent function  $Z$ , which includes parameterizing slopes,  $\mathbf{f}_k$ , and offsets,  $\mathbf{f}_c$ , we simply infer values for the moment-matched approximate latent

$$\tilde{Z} \sim \mathcal{N}(m\mathbf{x}_i, m^2\Sigma_c + (\Sigma_c + \mathbf{x}_i\mathbf{x}_i^T) \circ \Sigma_k). \quad (7)$$

In this formulation, the latent function is effectively a GP with a novel kernel function specific to the psychophysics problem, and we can thus apply standard methods for model fitting, as well as for using the model in active learning as we will see below.

**Approximate Normality of the Latent Function** The accuracy of the moment-matched MVN in (7) will depend on how close the true posterior for the latent function  $Z$  is to being normally distributed. We show here the conditions under which  $Z$  is Gaussian (Pinelis 2018).

Let  $\Sigma = m^2\Sigma_c + \mathbf{x}_i\mathbf{x}_i^T \circ \Sigma_k$ . Then,

$$\Sigma^{-1/2}Z = W + \Sigma^{-1/2}(Y_k \circ Y_c),$$

where  $W \sim \mathcal{N}(0, 1)$ . Note that

$$\begin{aligned} \mathbb{E}\|\Sigma^{-1/2}(Y_k \circ Y_c)\|^2 &\leq \left\|\Sigma^{-1/2}\right\|^2 \mathbb{E}\|Y_k \circ Y_c\|^2 \\ &= \left\|\Sigma^{-1}\right\| \text{tr}(\Sigma_k \circ \Sigma_c). \end{aligned}$$

Thus, as  $\left\|\Sigma^{-1}\right\| \text{tr}(\Sigma_k \circ \Sigma_c) \rightarrow 0$ ,  $Z$  is Gaussian with mean 0 and covariance  $\Sigma$ . This means that when the variances of the slope and offset values,  $k$  and  $c$ , are low, the Hadamard product of the GPs is approximately Gaussian, and there will be little loss to using the Hadamard approximation model. Empirically, we later show that the Hadamard model incurs a modest performance loss relative to the full semi-parametric model, but (as shown in the appendix) the Hadamard model is much less able to benefit from psychophysics-specific modifications to the link function.

ere about cholesky and specifics of inference....

### Active Learning with Semi-parametric GPs

Our semi-parametric model grants additional benefits for Bayesian active learning. Active learning methods define an acquisition function that prescribes the value of sampling a particular candidate point given the previously observed data and the current estimate of the model posterior. By optimizing this acquisition function, the next input ( $\mathbf{x}$ ) is chosen for sampling. A key goal in psychophysics is the estimation of psychometric detection thresholds, more formally known as a *level set estimation* (LSE) problem, i.e. finding the regions where the psychometric function is above or below some pre-determined threshold value  $r$ . In the general case of active learning for LSE, the location of the threshold is defined implicitly, and sampling strategies operate on an estimate of an arbitrary latent function.

With our semi-parametric model, the threshold is uniquely determined as a function of the context dimensions, and can be computed directly from the model posterior for  $k$  and  $c$  as:

$$x_i^r = \frac{\sigma^{-1}(r)}{k(\mathbf{x}_s)} - c(\mathbf{x}_s).$$

We can use samples from this posterior to compute several quantities of interest, for example the threshold posterior variance. In addition to being of interest to practitioners, reducing the posterior variance of the threshold is a natural, simple objective for active learning. We term it Threshold-BALV as it applies Bayesian Active Learning by Variance (Settles 2009) to the threshold posterior.

## Results

We showcase the benefits of the semi-parametric model in a few ways. First, we demonstrate performance on two synthetic psychometric test functions, where we show that our model can achieve good performance with less data than previously proposed baselines. Second, we evaluate performance on multiple real-world datasets, and show that our models outperform baselines in terms of predictive performance on unseen data. Finally, we demonstrate the compatibility of our models with active learning methods, again showing good performance with far less data than baselines, as well as competitive behavior of our novel Threshold-BALV acquisition function. We consider as baselines previous models used for flexible modeling and active learning for psychophysics, namely an otherwise-unconstrained GP-RBF model (Owen et al. 2021), and a GP-linear model with a linear kernel in the intensity dimension and an RBF kernel in the remaining dimensions (Schlittenlacher, Turner, and Moore 2018, 2020; Song, Garnett, and Barbour 2017; Gardner et al. 2015b).

### Two Dimensional Task

Before evaluating the semi-parametric model in a high dimensional setting, we first demonstrate performance in a simple 2-d psychometric test function, previously proposed in (Owen et al. 2021), and detailed in the appendix. This function has a monotonic (probit) probability of detection along an intensity dimension, and smoothly varies as a linear combination of sines and cosines in a second dimension.

We used a quasi-random Sobol sequence (Sobol' 1967) to select stimulus locations  $\mathbf{x}$  for our 2d test function.

Fig. 2A shows the prediction performance of the semi-parametric model on this function and Fig. 2C shows better estimation of the psychometric curve after 500 samples. For our evaluation metric we use the Brier score (Brier 1950), computed in expectation over the model's posterior. We use the Brier score because it assesses the calibration of the approximate posterior, and we use the expectation to account for the quality of posterior uncertainty estimation (we consider other metrics in the appendix). Our proposed model, with or without the MVN approximation, achieves low Brier scores much faster than baselines, though all models eventually achieve very good performance in this relatively simple test function. We will see next that the benefits are magnified in higher dimensions.

### Eight Dimensional Task

To simulate more a realistic, multidimensional sensory context we test the semi-parametric model using an 8-d psychometric test function previously reported by (Letham et al.

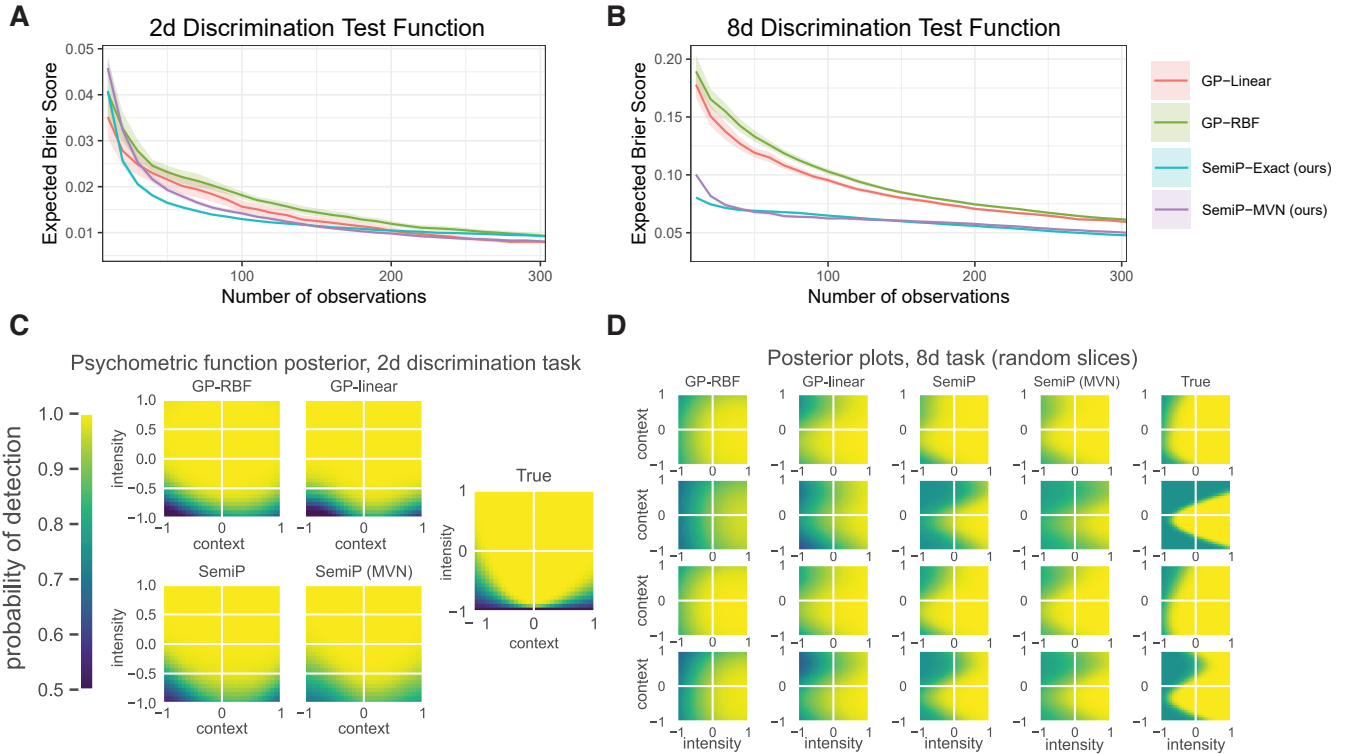


Figure 2: **A.** Performance on 2-d example as a function of number of quasi-random samples drawn, for all models. Curves are averages from 100 replications, with standard errors shaded. **B.** Same for 8-d test function. **C.** 2-d function estimation after 500 samples for all models. All models recover the test function with sufficient data. **D.** Inferred response probabilities in 8-d test function after 500 observations. Plots in each row have a randomly chosen dimension plotted along the y axis, and the intensity dimension along the x axis.

2022) and described in the supplement. This function retains monotonicity in detection probability along the intensity dimension in the same way as the 2d function, but detection probability is a smooth function of stimulus feature inputs in the other seven dimensions.

For this high dimensional test function, the semi-parametric model better estimates detection probability  $p$  in fewer samples than competing models, and the approximate MVN model achieves essentially identical performance as the full semi-parametric model. Fig 2 B shows the Brier score for each model for the first 300 samples. To get a sense of how well the models are actually able to estimate this high dimensional psychometric test function, we include 4 random 2d slices through the 8 dimensional space after training on 500 random samples (Fig 2 D).

## Results on Human Behavioral Data

To emphasize the generality and utility of our model, we evaluate performance on five participants from two real-world 6-dimensional datasets. All data are from visual psychophysical tasks. One dataset is a participant in a two-alternative forced choice (2AFC) task with 3000 trials from Letham et al. (2022), and the remaining are four participants performing a 4 alternative forced choice (4AFC) from Wuerger et al. (2020) whose trial counts range from  $\sim 200$ -

500 depending on the subject<sup>2</sup>. In the 2AFC task, the participant is presented with an animated circular Gabor patch, one half of which has been scrambled to resemble white noise. The scrambled side is selected at random. The stimulus varied along eight dimensions, six of which (contrast, background luminance, temporal and spatial frequency, size, and eccentricity) have data published. In the 4AFC task, a Gabor stimulus was presented in one of four quadrants of a screen, and participants were asked to select which quadrant contained the stimulus. This stimulus varied with size, orientation, frequency, and color. We used contrast as the intensity dimension for both tasks. For additional subject information and example stimuli, see supplementary materials. We run 15-fold cross-validation, train our model on 80% of the data and test on 20%, and report cross-validated log-likelihood. We see across all five subjects that some variant of the semi-parametric model consistently has superior cross-validated log-likelihood on held-out trials, regardless of whether we use the MVN or exact semi-parametric variant. We also note that no one link function consistently performs best—this inconsistency across choice of link is one reason the specific

<sup>2</sup>The former dataset is available at [https://github.com/facebookresearch/bernoulli\\_lse/tree/main/data](https://github.com/facebookresearch/bernoulli_lse/tree/main/data), and the latter at <https://www.repository.cam.ac.uk/handle/1810/304228>.




Figure 3: Mean cross-validated log-likelihood for 15-fold cross validation on an example subject from a 2AFC psychometric task, and four example subjects in a 4AFC task from (Watson 2017). Gray bars indicate standard error across folds.

parametric form of choice is still an area of active research. Nonetheless, we highlight that the semi-parametric model and MVN approximation are consistently strong performers.

### Active Learning

Lastly, we evaluate the semi-parametric model’s estimation of the 8-d test function using a variety of common Bayesian active learning schemes, including the BALD acquisition function based on mutual information (Houlsby et al. 2011; Gardner et al. 2015a; Owen et al. 2021), the BALV active learning scheme based on posterior variance (Settles 2009), and GlobalMI, a global lookahead acquisition function based on threshold estimation (Letham et al. 2022).

The computation of the GlobalMI acquisition function requires an MVN posterior on the latent function  $\mathbf{z}$  and a probit link with a floor of 0, which means it can be applied directly with GP-RBF and GP-Linear, and to our MVN approximate model. To use GlobalMI with the full semi-parametric model, we switch the link to a probit with floor of 0 (from Gumbel with a floor set to chance), and apply the MVN approximation we derived above to the variational posteriors  $q_c(\mathbf{f}_c)$  and  $q_k(\mathbf{f}_k)$ . In this setting we can still use the full semi-parametric model for evaluation. In addition to these baselines, we include the ThresholdBALV acquisition function as described earlier for use with the semi-parametric model.

Fig. 4 shows performance of all models under the considered active learning schemes. Here, we are not using the standard Brier score as we did previously, which was over the full model posterior. Instead are using a Brier score on the sublevel-set (threshold) posterior, i.e. the Brier score on estimating the probability of  $p$  above and below a threshold of 0.75. We choose this performance metric as the ThresholdBALV and GlobalMI acquisition functions are based specifically on threshold estimation. Performance of these acquisition methods using other metrics are shown in the supplement. We see that the BALD and BALV baseline acquisition functions perform comparatively poorly, and this

is true irrespective of the model.

However, we see using threshold-based schemes (GlobalMI and ThresholdBALV) that the semi-parametric models perform well, especially during early acquisition. At the end of acquisition, GlobalMI acquisition in conjunction with the GP-RBF model performs marginally better than all other acquisition-based models. It is additionally important to note that quasi-random Sobol sampling for this 8-d function performs remarkably well against these active learning schemes, often as good or better than all acquisition functions tested. The unusual effectiveness of quasi-random sampling in this setting has been previously reported (Letham et al. 2022) and we see it with our models as well as with the GP-RBF and GP-linear baselines. Exploring performance of these acquisition functions compared to quasi-random sampling for high dimensional psychophysics is an interesting avenue for future work. Here, we simply wish to emphasize that the semi-parametric model and its MVN approximation are compatible and competitive under a variety of existing active learning schemes, and our proposed ThresholdBALV acquisition shows strong performance early on in sampling for an 8-d test function. For further evaluation of active learning for the semi-parametric models, see the supplement. boundaries, allowing for more informative samples.

### Conclusion

We have demonstrated that a semi-parametric model for psychometric field estimation based on a parameterized sigmoid function can be adapted to high-dimensional psychophysical contexts using GPs as a non-parametric constraint on the sigmoid parameters. This semi-parametric approach not only offers parameters with scientific interpretation in the context of discrimination behavior, but offers accuracy improvements for estimating high dimensional tuning curves compared to competing methods. It does the latter both by providing a better prior, and by enabling a new active learning objective based on the semi-parametric functional form. We further introduce a moment-matching approximation to

## Threshold active learning




Figure 4: Performance metrics for threshold estimation under a variety of common active learning paradigms.

our model that can be used as a psychophysics-specific GP prior, or to produce approximate MVN posteriors compatible with analytic acquisition functions. We evaluate our contributions relative to baselines on both synthetic and real data and show a number of performance gains, especially for smaller sample sizes that are important in real human-in-the-loop experiments.

**Limitations** First, the evaluations in this paper either use synthetic test functions or human visual psychophysical data on a per-subject basis, so we cannot say at this stage how the model will perform with data from other sensory modalities or for cross-participant prediction. Second, we focus our evaluation on the Brier score in expectation over the posterior. While taking other metrics in expectation over the posterior shows similar behavior (as we demonstrate in the appendix), focusing on the posterior mean only (as done in prior work) changes the story somewhat. In particular, if we only consider the posterior mean of the RBF model, its key deficiency of having an overly flexible hypothesis space is mitigated and its performance looks stronger. Third, while we demonstrate that our work is compatible with active learning methods (even ones that require an MVN posterior), we do not offer an exhaustive evaluation of active learning methods, and the benefits of our contributions for active learning appear to be focused on small sample sizes. In line with this, we focus our evaluation and narrative around performance with relatively small data (consistent with the goal of sample-efficient psychophysics) but it is likely that with much larger datasets, the GP-RBF model’s universality will let it match or outperform models with more restricted hypothesis spaces such as ours.

**Ethics Statement** Our work is primarily concerned with understanding low-level human perception, and as such carries relatively low risk of societal and ethical harm. Some risks include the misuse or de-anonymization of data, and overly broad or incorrect conclusions made based on data that is too limited, collected in a biased way, or based on misunderstanding or misusing the model. With respect to data misuse, we use only de-identified data that has been

previously published, where informed consent was obtained, and is of low sensitivity (it is behavioral responses to simple visual stimuli). With respect to overly broad conclusions, we keep our claims narrowly focused on the quality of the model, and do not provide new interpretations or conclusions related to the datasets we use for evaluation. Furthermore, we think the specificity of our model for the psychophysics problem domain makes it less likely to be applied (and misused) in other settings than more generic models. On a more positive note, increasing sample efficiency for psychophysics studies may improve the experience of human research participants, who can sometimes be required to participate in dozens of hours of data collection when traditional grid or staircase methods are used.

**Computational Load** With respect to computational load and environmental impact, the benchmarks were all carried out over the course of a few days ( $\approx 70$ -100 hours) on a single EC2 `c6i.metal` node, and the cross-validation folds were performed over a smaller node over a similar period of time. These hours are largely taken up by replications across seeds (for benchmarks) and folds (for cross-validation)—for practical usage, the models we use take seconds to estimate on a typical laptop, which makes them accessible for use by most practitioners and researchers.

**Code Release** The code needed to use our contribution in new experiments and benchmarks (models, likelihoods, acquisition function, etc) is available as part of AEPsych (<https://aepsych.org/>). The code needed to generate the specific benchmarks and figures in the paper is available at <https://github.com/facebookresearch/semi-parametric-psychophysics>.

## References

Balandat, M.; Karrer, B.; Jiang, D. R.; Daulton, S.; Letham, B.; Wilson, A. G.; and Bakshy, E. 2020. BoTorch: a framework for efficient Monte-Carlo Bayesian optimization. In *Advances in Neural Information Processing Systems 33*, NeurIPS.

- Brand, T.; and Kollmeier, B. 2002. Efficient adaptive procedures for threshold and concurrent slope estimates for psychophysics and speech intelligibility tests. *The Journal of the Acoustical Society of America*, 111(6): 2801–2810.
- Brier, G. W. 1950. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1): 1–3.
- Gardner, J. R.; Malkomes, G.; Garnett, R.; Weinberger, K. Q.; Barbour, D.; and Cunningham, J. P. 2015a. Bayesian active model selection with an application to automated audiometry. In *Advances in Neural Information Processing Systems 28*, NIPS, 2386–2394.
- Gardner, J. R.; Song, X. D.; Weinberger, K. Q.; Barbour, D.; and Cunningham, J. P. 2015b. Psychophysical detection testing with Bayesian active learning. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, UAI, 286–297.
- Hensman, J.; Matthews, A. G.; and Ghahramani, Z. 2015. Scalable variational Gaussian process classification. *Journal of Machine Learning Research*, 38: 351–360.
- Houlsby, N.; Huszár, F.; Ghahramani, Z.; and Lengyel, M. 2011. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*.
- Letham, B.; Guan, P.; Tymms, C.; Bakshy, E.; and Shvartsman, M. 2022. Look-ahead acquisition functions for Bernoulli level set estimation. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, AISTATS, 8493–8513.
- May, K. A.; and Solomon, J. A. 2013. Four theorems on the psychometric function. *PLOS One*, 8(10): e74815.
- Owen, L.; Browder, J.; Letham, B.; Stocek, G.; Tymms, C.; and Shvartsman, M. 2021. Adaptive nonparametric psychophysics. *arXiv preprint arXiv:2104.09549*.
- Pinelis, I. 2018. Normal approximation to the pointwise/Hadamard/Schur product of two multivariate Gaussian/normal random variables. <https://mathoverflow.net/q/293985>. Accessed: 2022-01-19.
- Schlittenlacher, J.; Turner, R. E.; and Moore, B. C. 2020. Application of Bayesian active learning to the estimation of auditory filter shapes using the notched-noise method. *Trends in Hearing*, 24: 1–13.
- Schlittenlacher, J.; Turner, R. E.; and Moore, B. C. J. 2018. Audiogram estimation using Bayesian active learning. *The Journal of the Acoustical Society of America*, 144(1): 421–430.
- Schütt, H. H.; Harmeling, S.; Macke, J. H.; and Wichmann, F. A. 2016. Painfree and accurate Bayesian estimation of psychometric functions for (potentially) overdispersed data. *Vision Research*, 122: 105–123.
- Settles, B. 2009. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin-Madison.
- Shen, Y.; and Richards, V. M. 2012. A maximum-likelihood procedure for estimating psychometric functions: thresholds, slopes, and lapses of attention. *The Journal of the Acoustical Society of America*, 132(2): 957–967.
- Sobol’, I. M. 1967. On the distribution of points in a cube and the approximate evaluation of integrals. *Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki*, 7(4): 784–802.
- Song, X. D.; Garnett, R.; and Barbour, D. L. 2017. Psychometric function estimation by probabilistic classification. *The Journal of the Acoustical Society of America*, 141(4): 2513–2525.
- Strasburger, H. 2001. Converting between measures of slope of the psychometric function. *Perception & Psychophysics*, 63(8): 1348–1355.
- Watson, A. B. 2017. QUEST+: A general multidimensional Bayesian adaptive psychometric method. *Journal of Vision*, 17(3): 10–10.
- Wuerger, S.; Ashraf, M.; Martinovic, J.; Kim, M.; Perez-Ortiz, M.; and Mantiuk, R. 2020. Measurements of spatio-chromatic contrast sensitivity up to 7000  $cd/m^2$  [Dataset].