

Mind Reading by Machine Learning: A Doubly Bayesian Method for Inferring Mental Representations – supplementary material –

Ferenc Huszár (fh277@eng.cam.ac.uk)

Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK

Uta Noppeney (uta.noppeney@tuebingen.mpi.de)

Max Planck Institute for Biological Cybernetics, Spemannstrasse 41, Tübingen 72076, Germany

Máté Lengyel (m.lengyel@eng.cam.ac.uk)

Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK

Quasi-ideal observer models

In order to be able to make inferences about subjective distributions based on subjects’ responses we need to construct a model of the process of generating these responses based on the information available to subjects and parameters that describe their behaviour (see Fig.1 in the main paper). Throughout the paper, we assume that the response variables r_i are discrete, although the methodology naturally extends to the case of continuous responses.

The first step of constructing an ideal observer model involves a Bayesian analysis of the task faced by subjects. A subject may entertain several alternative hypotheses, each corresponding to one of the R different available responses. The likelihood of hypothesis \mathcal{H}_j , $p(\mathcal{S}_i|\mathcal{H}_j, \mathcal{P}, \Theta_{\mathcal{O}})$, expresses the probability of seeing stimuli \mathcal{S}_i in a trial in which response j is correct. Although the precise form of this likelihood depends on the particular task, as described in detail in sections -, it always depends on the subjective distribution \mathcal{P} , and this dependence will allow us to invert the model and infer \mathcal{P} from behavioural data. Importantly, as also noted in the paper, we introduce a number of parameters, captured in $\Theta_{\mathcal{O}}$, to model different sources of suboptimal behaviour so that the resulting quasi-ideal observer models will be appropriate statistical models of subjects’ responding.

Given the likelihood of hypothesis j , its posterior probability can be obtained by applying Bayes’ rule:

$$\mathbb{P}[\mathcal{H}_j|\mathcal{S}_i, \mathcal{P}] = \frac{p(\mathcal{S}_i|\mathcal{H}_j, \mathcal{P}) \cdot \pi_j}{\sum_{k=1}^R p(\mathcal{S}_i|\mathcal{H}_k, \mathcal{P}) \cdot \pi_k} \quad (\text{SM1})$$

where $\mathbb{P}[\mathcal{H}_j] = \pi_j$ describes the subject’s prior preferences over hypotheses. The first source of sub-optimality we introduce is that, unlike in most ideal observer analyses, we assume that these prior probabilities are truly subjective and thus they need not reflect the veridical fraction of trial types in which their corresponding responses are correct.

An ideal observer bases its decisions on the posterior

probabilities of the hypotheses it is entertaining. Therefore, the second key modelling step is to specify the precise functional dependence between the probabilities with which different options are chosen and the posterior probabilities of the hypotheses associated with them. We apply the flexible and commonly used ?? softmax rule:

$$\mathbb{P}_{\text{softmax}}[j] = \frac{\mathbb{P}[\mathcal{H}_j|\mathcal{S}_i, \mathcal{P}]^{\beta}}{\sum_{k=1}^R \mathbb{P}[\mathcal{H}_k|\mathcal{S}_i, \mathcal{P}]^{\beta}} \quad (\text{SM2})$$

The rule is parameterised by the decision noise parameter β . We note that the softmax rule contains probability matching ? ($\beta = 1$) and maximum *a posteriori* decisions ($\beta \rightarrow \infty$) as special cases. For binary tasks ($R = 2$, $\pi_1 = \pi$, $\pi_2 = 1 - \pi$) considered in the paper this rule becomes a simple logistic sigmoidal function of the log-likelihood ratio:

$$\mathbb{P}_{\text{softmax}}[1] = 1 - \mathbb{P}_{\text{softmax}}[2] = \frac{1}{1 + e^{-\beta \log \frac{\pi}{1-\pi} \frac{p(\mathcal{S}_i|\mathcal{H}_1, \mathcal{P})}{p(\mathcal{S}_i|\mathcal{H}_2, \mathcal{P})}}} \quad (\text{SM3})$$

It is also possible, and desirable, to model other sources of sub-optimality in responding. For instance, a subject may misunderstand the instructions or become bored and start pressing the same button irrespective of the stimuli. To account for this, we apply a probabilistic mixture of an ideal and a ‘misbehaved’ subject, parameterised by the ‘attention’ parameter κ . Actual response probabilities thus take the following form:

$$\mathbb{P}[r_i = j|\mathcal{S}_i, \mathcal{P}, \Theta_{\mathcal{O}}] = \frac{1 - \kappa}{2} + \kappa \mathbb{P}_{\text{softmax}}[j],$$

A final source of stochasticity we consider in subjects’ behaviour is observation noise (or, more precisely, the subject’s internal model of her observation noise), describing the way ‘true’ stimulus feature values, s^* , controlled by the experimenter are distorted in the subject’s perception to yield ‘effective’ features, s . We model observation noise as a Gaussian distribution around the

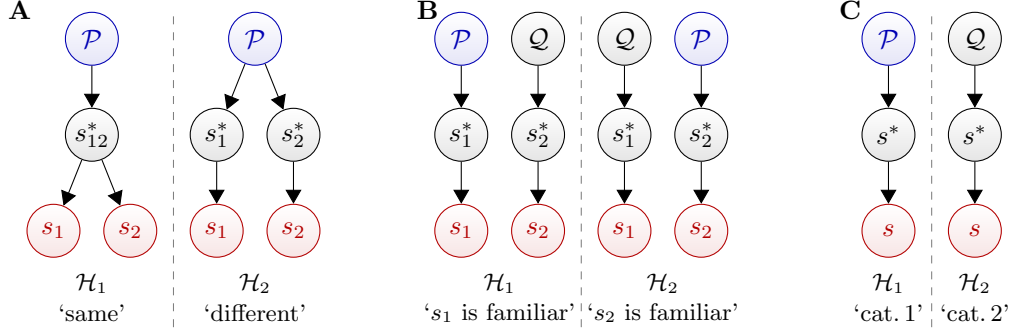


Figure SM1: Graphical models describing alternative hypotheses in the psychophysical tasks treated in the paper. **A**, Alternative hypotheses in a trial of the one-back discrimination task. Under \mathcal{H}_1 (*left*), the two observed stimuli, s_1 and s_2 are noisy observations of the same underlying stimulus s_{12}^* that was drawn from \mathcal{P} , whereas \mathcal{H}_2 (*right*) holds that s_1 and s_2 are noisy observations of two different stimuli s_1^* and s_2^* , that are independent draws from \mathcal{P} . **B**, Alternative hypotheses in a trial of the 2AFC task. The subject is presented with two stimuli s_1 and s_2 . Under \mathcal{H}_1 (*left*), s_1 is familiar, i.e. it is a noisy observation of a stimulus drawn from \mathcal{P} , whereas s_2 is a noisy observation of another stimulus drawn from some alternative distribution Q . Under \mathcal{H}_2 (*right*), s_2 is the familiar stimulus. **C**, Alternative hypotheses in a trial of the binary classification task. The subject is presented with a stimulus s . Under \mathcal{H}_1 (*left*), s is a noisy observation from category-conditional distribution \mathcal{P} . Under \mathcal{H}_2 (*right*), s is a noisy observation from Q .

true value with constant (i.e. stimulus-independent) covariance Σ_n :

$$p_n[s|s^*] = \mathcal{N}(s; s^*, \Sigma_N) \quad (\text{SM4})$$

Of course, it is possible to incorporate more complex observation noise models when it is particularly required. A possible extension would be to allow the covariance of the Gaussian noise to be stimulus-dependent.

In summary, the resulting quasi-ideal observer models have the following free parameters, $\Theta_{\mathcal{O}}$ (on top of $\Theta_{\mathcal{P}}$, characterising the subjective distribution \mathcal{P} , see section 2.2 of the main paper and section here): prior bias π , decision noise β , level of attention κ , and observation noise covariance matrix Σ_n . Importantly, these parameters are estimated from data rather than being fixed at some pre-defined values. Hyperparameters for priors over them are defined in section .

In the following we derive likelihood ratios (needed to compute binary choice probabilities based on eq. SM3) for the three binary task types considered in the paper.

One-back discrimination task

In a trial of a discrimination task the subject has to decide whether two stimuli presented to them are identical or not. In this paper we consider a special case of the discrimination task, the so-called one-back task in which stimuli are presented sequentially, and the subject has to discriminate the current stimulus from the previously presented one.

Our construction of the ideal observer model for the discrimination task relies on the Bayesian analysis of

similarity judgements put forward by Kemp et al.?. According to this analysis, the degree of subjective similarity of two objects is related to the probability of them having been generated by the same generative process. In our framework, the subject's belief about the process generating possible objects is captured by the subjective distribution \mathcal{P} . Based on this, we can derive the likelihood ratio of the two hypotheses (see Fig. SM1A for further explanation):

$$R = \frac{\int p_n(s_1|s_{12}^*) p_n(s_2|s_{12}^*) \mathcal{P}(s_{12}^*) ds_{12}^*}{\int p_n(s_1|s_1^*) \mathcal{P}(s_1^*) ds_1^* \cdot \int p_n(s_2|s_2^*) \mathcal{P}(s_2^*) ds_2^*} \quad (\text{SM5})$$

If the subjective distribution is a mixture of Gaussians, and the observation noise is Gaussian with constant covariance matrix, then this likelihood ratio can be computed analytically.

Two-alternative forced choice task

In each trial of the two alternative forced choice (2AFC) task, two stimuli are presented to the subject, who then needs to choose the one that appears more familiar, or the one that appears more to be a member of a particular category.

The two hypotheses that the subject has to choose between are explained in Figure SM1B. The likelihood ratio for this task is:

$$R = \frac{\int p_n(s_1|s_1^*) \mathcal{P}(s_1^*) ds_1^* \cdot \int p_n(s_2|s_2^*) \mathcal{Q}(s_2^*) ds_2^*}{\int p_n(s_1|s_1^*) \mathcal{Q}(s_1^*) ds_1^* \cdot \int p_n(s_2|s_2^*) \mathcal{P}(s_2^*) ds_2^*}, \quad (\text{SM6})$$

where \mathcal{Q} is the distribution of *non-familiar* objects. For convenience, we assume that the alternative distribution

\mathcal{Q} is approximately uniform, in which case the response probabilities will not depend on \mathcal{Q} , and reduce to standard formulæ also used in ???. Whether this assumption is reasonable depends largely on experimental conditions and on the exact phrasing of the instructions given to subjects. In principle, one could treat \mathcal{Q} as another parameter of the ideal observer and infer it in the same way as we infer \mathcal{P} .

Binary categorisation task

In each trial of the binary classification task the subject is presented with a single stimulus, and has to decide whether it belongs to category 1 or 2. For this task, we assume the subject entertains category-conditional subjective distributions \mathcal{P} and \mathcal{Q} over stimuli. Then, the likelihood ratio of hypotheses takes the following simple form (see Figure SM1C):

$$R = \frac{\int p_n(s|s^*) \mathcal{P}(s^*) ds^*}{\int p_n(s|s^*) \mathcal{Q}(s^*) ds^*} \quad (\text{SM7})$$

In the experiments with one-dimensional feature space, we consider a special case, when the subject is instructed that \mathcal{Q} is uniform, i. e. every stimulus is equally likely to be an exemplar of category 2. We also assume that – prior to the analysed part of the experiment – the subject had already learned the support and magnitude of the uniform distribution, and we can thus concentrate on inferring \mathcal{P} only.

Prior hyperparameters

For ease of computation of derivatives, we chose simple exponential family distributions as priors over parameters $\Theta_{\mathcal{P}}$ and $\Theta_{\mathcal{O}}$.

$\Theta_{\mathcal{P}}$: In the one-dimensional experiments fish height (the only relevant stimulus feature) ranged from 0.3 to 8.35 cm. Accordingly, we set the prior over component means to be Gaussian with mean 4.16 cm and variance 1 cm². Component covariances were given a Wishart prior with mean 0.1 cm² and concentration 4. This way, training distributions corresponding to different conditions were approximately equally probable under the prior. We set the number of components to be 6.

In the second experiment with three-dimensional feature space, the stimulus features were normalised so that they took values between 0 and 1. To enforce subjective distributions to be concentrated in this effective stimulus set, we set the prior over component means to be Gaussian with mean $[0.5, 0.5, 0.5]^T$ and covariance 0.4I. We put a Wishart prior over component covariances with mean 0.1I and concentration 7. The prior average Kullback-Leibler divergences from the two training distributions were approximately equal. We set the number of components to be 8.

$\Theta_{\mathcal{O}}$: We used weak priors: Gamma with mean 1, concentration 3 for decision noise β ; Beta with mean 0.5 concentration 3 for the prior bias π ; Beta with mean 0.9 and concentration 3 for attention κ ; and Wisharts with means 0.02 and 0.02I, and concentrations 2 and 4 for the observation noise covariances Σ_n in the two experiments, respectively.

Predicting human responses

The problem of predicting human responses in simple binary decisions tasks can be interpreted as a problem of binary categorisation (mapping the feature values of a test stimulus, or of a pair of test stimuli, to one of two possible choices) . Figure SM2 shows the categorisation problems inherent in the three psychophysics tasks considered in the paper.

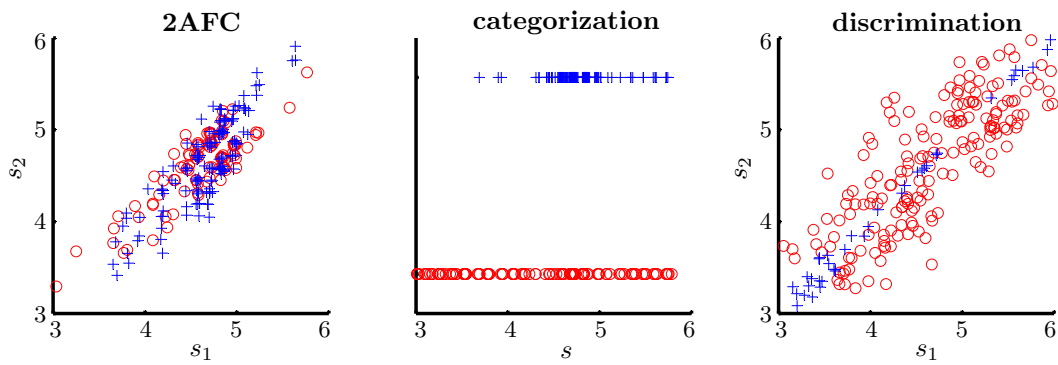


Figure SM2: Prediction of human responses viewed as a binary classification task. *Left:* in the 2AFC task, the subject is presented a pair of stimuli, and selects either s_1 (blue crosses) or s_2 (red circles) as being more familiar (human responses from the one-dimensional experiment). *Middle:* in the categorisation task, the subject is presented a single stimulus and classifies it as belonging to either category 1 (blue crosses) or 2 (red circles) (human responses from the one-dimensional experiment). *Right:* in the discrimination task, a pair of stimuli is presented in succession, and the subject decides whether they seem to be different (blue crosses) or identical (red circles) (simulated data).