
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Keurulainen, Antti; Westerlund, Isak Rafael; Keurulainen, Oskar; Howes, Andrew
Amortised Experimental Design and Parameter Estimation for User Models of Pointing

Published in:
CHI 2023 - Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems

DOI:
[10.1145/3544548.3581483](https://doi.org/10.1145/3544548.3581483)

Published: 19/04/2023

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Keurulainen, A., Westerlund, I. R., Keurulainen, O., & Howes, A. (2023). Amortised Experimental Design and Parameter Estimation for User Models of Pointing. In *CHI 2023 - Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* Article 772 ACM. <https://doi.org/10.1145/3544548.3581483>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



Amortised Experimental Design and Parameter Estimation for User Models of Pointing

Antti Keurulainen
antti.keurulainen@bitville.com
Aalto University
Espoo, Finland
Bitville Oy
Espoo, Finland

Oskar Keurulainen
oskar.keurulainen@bitville.com
Bitville Oy
Espoo, Finland

Isak Westerlund
isak.westerlund@bitville.com
Bitville Oy
Espoo, Finland

Andrew Howes
a.howes@bham.ac.uk
University of Birmingham
Birmingham, United Kingdom
Aalto University
Espoo, Finland

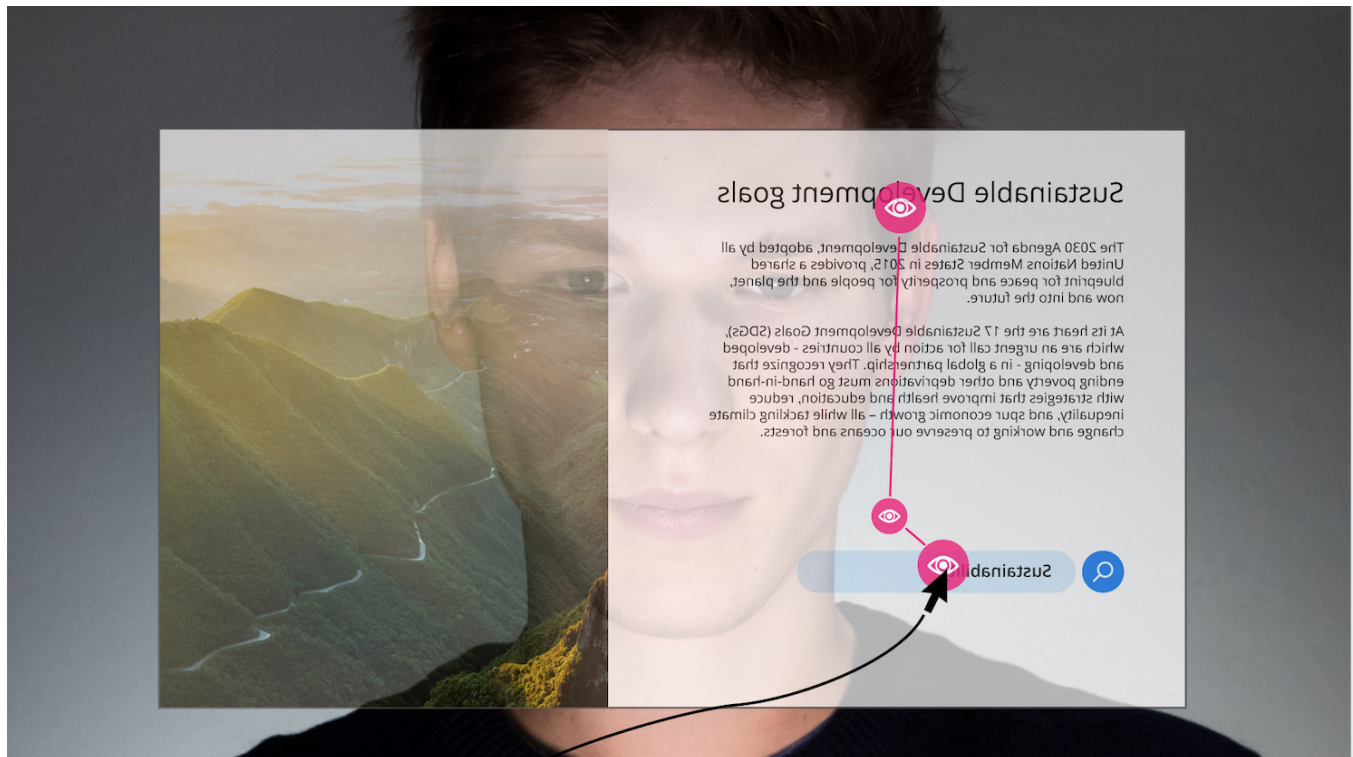


Figure 1: An intelligent interactive system views a user through sensors that include a keyboard and mouse and sometimes, camera, microphone and/or eye-tracker. What is the best way to infer a model of the user from these data? In the figure gaze path is represented in red, mouse movements in black and text entered in blue. A camera captures the user's gaze direction.

ABSTRACT

User models play an important role in interaction design, supporting automation of interaction design choices. In order to do so,



This work is licensed under a Creative Commons Attribution International 4.0 License.

CHI '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9421-5/23/04.
<https://doi.org/10.1145/3544548.3581483>

model parameters must be estimated from user data. While very large amounts of user data are sometimes required, recent research has shown how experiments can be designed so as to gather data and infer parameters as efficiently as possible, thereby minimising the data requirement. In the current article, we investigate a variant of these methods that amortises the computational cost of designing experiments by training a policy for choosing experimental designs with simulated participants. Our solution learns which experiments provide the most useful data for parameter estimation by interacting with in-silico agents sampled from the model space thereby using synthetic data rather than vast amounts of human data. The approach is demonstrated for three progressively complex models of pointing.

CCS CONCEPTS

• **Human-centered computing** → **HCI theory, concepts and models**; **User models**.

KEYWORDS

user models, adaptive experiment design, parameter estimation, active inference, computational rationality

ACM Reference Format:

Antti Keurulainen, Isak Westerlund, Oskar Keurulainen, and Andrew Howes. 2023. Amortised Experimental Design and Parameter Estimation for User Models of Pointing. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3544548.3581483>

1 INTRODUCTION

User models take many forms in HCI, from simple lists of ‘psychological factors’ including, perhaps, personality variables and/or product preferences to cognitive models that simulate the processing of information in the mind. Some of the latter have focused on constraints imposed by the human perceptual and motor systems, others on the structure of human memory and others on control (what to do next). While these user models provide one of the theoretical anchors for the discipline, they are difficult to construct and difficult to fit to behaviour – sometimes requiring vast amounts of data from an expanding range of sensors (Figure 1). In the current article, we explore one particular approach to addressing this latter problem by automating model parameter estimation for three pointing tasks. The approach involves the design of an optimal sequence of experimental trials (or designs) that maximise the relevance of the available information.

In the statistics literature ‘optimal experimental design’ (OED) is the problem of choosing which experimental trial to do next so as to maximize some objective. For example, an HCI researcher may want to measure the effect of a new pointing device on a user’s movement accuracy, but how far away and how big should the movement target be on each successive trial so as to maximize the information gained from observing the user performing the task? We present an approach to efficiently solving this problem for estimating user model parameters in HCI.

We argue that the presented approach has the potential to enhance the contribution that user models make to HCI by providing interactive systems with the means to automatically and rapidly fit

user models to individual users and thereby personalise interaction so as to best fit the requirements of the individual. This capability is also important to cooperative/collaborative Artificial Intelligence (AI), that is the problem of how to get machines to work with people. Personalisation and collaboration are important objectives for HCI, in part, because they directly address the desire to design interaction for diverse users. While we do not investigate personalisation per se in the current article, we believe that user modeling is crucial to the future of personalisation and that this potential will only be fulfilled if the parameter estimation problem can be solved.

Fitting a user model to an individual was difficult in the early days of cognitive modeling, in part because models such as GOMS [6], were constructed manually. Production rules that mapped goals into actions were written by an analyst. Model constructing consisted of a painstaking and iterative process of protocol analysis, and production rule writing.

Since then significant advances have been made on automatic construction of models. Rather than hand-coding production rules, modern modelling techniques now automatically learn a *control policy* (task knowledge) using deep reinforcement learning. In particular, machine learning can be used to learn a model’s control policy through exploration of simulated interaction.¹ This approach has been successfully applied to the automatic construction of user models of menu search, decision making, gaze-based interaction, hierarchical control, and touchscreen typing [9, 11–13, 16, 18, 22, 37]. However, fitting the quantitative parameters of these models to individual humans still requires a significant contribution from the analyst.

One of the difficulties with fitting learning-based user models to human data is that the control policy must be trained to make predictions for each set of possible parameters. For example, in the gaze-based interaction model reported in [11] oculomotor noise and perceptual noise parameters are properties of an individual user. Both parameters introduce uncertainty in aimed movements to a target. The optimal control policy chooses an aim point for a target in accordance with these uncertainties. For example, when oculomotor noise is high then it makes sense to deliberately under-shoot the target and then make a corrective submovement. This is because on average this results in a lower overall movement time than overshooting (which takes longer) and correcting.

Unfortunately then, while automatically learning human-like control policies solves part of the user modelling problem, it does not solve the parameter estimation problem. User models typically have many parameters and the control policy, and therefore behaviour, is adapted to these parameter values. For example, it is known that the human control policy for pointing is adapted to noise both in the motor system and the visual system. In general, the parameter estimation problem is to find a set of values for model parameters such that the predicted behaviour of an individual user is as close as possible to the observed (measured) behaviour. It is quite often formalised as an optimisation problem; how to generate the best estimate of the parameters from the available data. Typically, the objective will be to find parameters that minimise the

¹The control policy is a function that maps observations into actions. It can be learned with a number of machine learning algorithms.

difference between the model behaviour and human behaviour, that is maximising fit.

Estimating parameters so as to minimise the discrepancy between model and human requires high quality data; a fact that, as we have said, has given rise to work on optimal experimental design (OED) [35]. This problem has been conceptualised as a problem of how to maximise expected information gain (EIG). In other words, the purpose of an experiment is to maximise the information that is gained about the parameter values that best fit the model to the human. Various approaches have been proposed to this problem. In one class of approaches, Bayesian Optimal Experimental Design (BOED), the idea is to choose experiments that maximise EIG between the prior probabilities of all possible parameters values and a posterior distribution that is conditioned on the expected observations [43]. BOED does not only give a point estimate of the best fitting parameter values but also a posterior probability that these (and all other parameter values) are the best fit.

In some approaches, the choice of experiment is ‘amortised’ meaning that a near-optimal policy for choosing experiments is computed before the deployment of the policy for parameter estimation. In HCI, the advantage of amortisation is that design of experiments for estimating user models can be computed without slowing interaction with the user. Another advantage of amortisation is that experimental design can be non-myopic. This means that, rather than maximising EIG for each individual experiment, instead it can be maximised for a whole sequence of experiments. One approach to amortisation involves defining the optimal experimental design problem as a reinforcement learning problem.

In the current article we propose a new approach to user model parameter estimation in HCI that takes advantage of the recent advances described above. The contribution is in offering a novel and practical method for estimating user model parameters through amortising the cost of choosing experiments. While the proposal is for a general method, we demonstrate its viability in this paper for pointing tasks. An overview of the approach is illustrated in Figure 2. The approach estimates the parameters of a user model for an individual human user (Phase 3), having previously computed an ‘ensemble’ user model for all possible parameter combinations (Phase 1) and then an optimal sequential experimental design policy (Phase 2). In phase 1, the ensemble model of the space of possible users is trained to perform the task for the distribution of possible parameter values and the distribution of possible task environments. The ensemble approach is an important recent advance in user modelling that is described in more detail below [27, 34]. In phase 2, the Analyst is trained to conduct the best sequence of experiments for determining the model parameters. Simulated users are randomly sampled given the parameter distribution and the Analyst learns to fit the model parameters to the simulated user. In phase 3, the trained Analyst conducts experiments on users and generates parameter fits.

Phase 1 amortises the cost of computing the implications of different parameter values for model behaviour. Phase 2 amortises the cost of computing a non-myopic sequential policy for choosing experiments and Phase 3 takes advantage of the computation conducted in Phases 1 and 2 in order to optimally gather data and estimate user model parameters without any user-noticeable interaction latency.

One confusion that arises about our approach is that once we have trained a simulator model for all possible parameter values (the ‘ensemble’ user model) then there is nothing else to learn. The confusion is resolved once it is realised that the ‘ensemble’ is a model of the distribution of all possible users (technically, a distribution over the parameters of the user model) and all possible tasks within the defined space. The assumption is that if we know the exact parameter values for a specific real-world user, then this model will accurately describe their behaviour, but the challenge is that, given a real-world user, we do not know the best parameters. We need the Analyst to design experiments and gather the data from the user in order to find the best parameters in the ensemble model.

For training the Analyst, we generated hypothetical users by sampling parameters from the ensemble distribution, and exposed these known parameter values to the Analyst as a training signal, but there was no expectation that any individual real-world user will have the same parameter values. At test time, the Analyst had no knowledge of the user parameters (and the sets of training-time and test-time users are disjoint), and these parameters were inferred from the designed experiments.

Another confusion is that it might seem that there is no point in learning about an individual user when Analyst already has a model of the distribution of all possible users. But, this model can only simulate an *specific* user if it knows the correct parameters for that user. The point of Analyst is to infer the parameters of a user with unknown parameters (e.g. a new human user whose behaviour has not been observed before). Analyst needs to infer the best values of those parameters, and only then can the user model simulate that particular user. Further, it is the capacity to simulate a specific user that promises a role for Analyst in interaction personalisation and cooperative Artificial Intelligence.

In what follows, we review the existing literature, formally define our approach, test it on two abstract tasks chosen so as to demonstrate the generality of the approach, and then report three studies of parameter estimation for user models of pointing. In the first of the studies, we demonstrate the effectiveness of the approach for estimating the parameters of a pointing user model from mouse click data. The user model has a single parameter for movement noise that gives rise to Fitts’s Law like behaviour through a learned policy that generates multiple submovements to achieve point and click goals.

In the second study, we extend the approach to simulated eye-movement data for gaze-based pointing. These data consist of variable length sequences of interleaved saccades and fixations. In this gaze-based model there are two parameters, one for oculomotor noise and the other for perceptual noise. This gives rise to a potential identifiability problem.

In the third study, we explore the capacity of the approach to not only identify perceptual/motor noise parameters but to also determine user preferences. Here the gaze-based model is applied to an interface in which pointing is achieved by looking at targets and pressing a button (a key on the keyboard). Because an experimental trial can be terminated at any time by pressing the button it gives rise to a speed/accuracy trade-off and the user’s policy is optimised for their preference, or otherwise, for accuracy over speed. In this

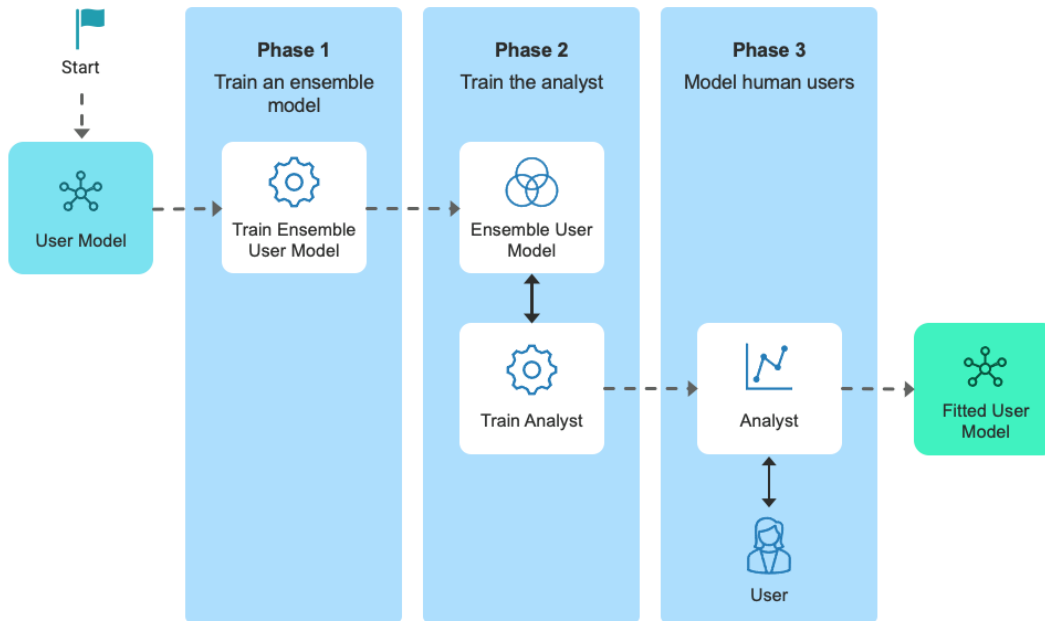


Figure 2: Our approach takes a user model as input. This user model has prior distributions over parameters θ but has not been fitted to individual human behaviour. In Phase 1, an Ensemble Cognitive Model (ECM) is trained to perform the task for the distribution of possible parameter values. In Phase 2, the cognitive model Analyst is trained to conduct the best sequence of experiments for determining the model parameters. In Phase 3, the trained Analyst is deployed with users and a fitted model is generated as output.

study both performance time and errors are used in the estimation of model parameters.

2 BACKGROUND

2.1 User models

User models – computable representations of psychological constraints on interaction – have been influential in HCI since its inception. GOMS, a rule-based model for representing hierarchical task knowledge, provides a formalism for HCI researchers to conduct detailed task analyses [6]. The Model Human Processor (MHP), a theory of the temporal properties of cognitive resources, supported the prediction of task performance time. ACT-R and EPIC provide means to simulate cognition and action; ACT-R with a particular focus on human memory and EPIC on constraints imposed by perceptual/motor systems [1, 24, 49]. Fitts’s Law, a mathematical formulation of the relationship between task difficulty and movement time, became a particularly influential model of pointing [30, 48]. Computationally rational models, based on machine learning problems but with cognitive bounds, provided a means to automatically learn control policies. Rather than hand crafted production rules, computationally rational models derive predictions by learning a control policy that is bounded only by human like resource constraints [19, 28, 37].

Most of the approaches to user modelling described above remain actively and productively investigated but our focus here is on computationally rational models [37]. While it has a distinctive

approach to the automation of a control policy, it shares a need for new approaches to parameter estimation. In what follows we look in detail at three computationally rational models in order to further understand the parameter estimation problem.

One approach to computational rationality involves defining an interactive cognitive task as a Markovian problem and solving it using machine learning, usually reinforcement learning. For example, multi-attribute decision making can be defined as a Partially Observable Markov Decision Problem (POMDP) in which information about relevant attributes is gathered using saccadic eye movements. The reward function specifies a trade-off between speed and accuracy such that information is only gathered if the benefits to decision accuracy outweigh the temporal cost [13]. As a consequence, once an optimal control policy has been learned, it generates attribute-wise, rather than option-wise information gathering – much like humans. While automatically acquired optimal control policies address a major part of the user modelling problem, they leave open the question of how to set model parameters. In the case of multi-attribute decision making the predictions are only as good as the attribute weights that define the user preference function.

Similarly, gaze-based interaction can be defined as a POMDP in which partial observations are constrained by foveated vision and saccadic eye-movements by oculomotor noise [11]. The model of foveated vision imposes increasing localisation error with eccentricity of the target from the fovea. The solution to this POMDP is an optimal control policy that – again like humans – undershoots the

target in order to minimise movement time (because undershoots take less time than overshoots). Again, the capacity to automatically generate a control policy is an important advance but it leaves open the question of how the parameter values are set so as to model individual users. In the case of gaze-based interaction parameters include the perceptual and oculomotor noise weights as well as the saccade duration intercept and slope parameters.

An important aspect of user modeling is determining the prior distribution of possible parameter values. This distribution can be thought of as an hypothesis space that covers all of the possible behaviours of the population of users. It is constructed using knowledge that is available before performing experiments on individual humans. This prior parameterisation of the model induces an ensemble of possible user models, each of which can be expressed in the form of a particular POMDP. Importantly, this modelling framework is agnostic to the amount of prior information available, since the experimenter can specify arbitrary prior distributions for the parameters of the model. In the case of limited available prior knowledge, the simulator implementing the model can be initialised with non-informative priors, thus describing a diverse distribution of users with highly varying cognitive bounds and preferences. Once the prior distribution of possible user models is established then experiments on humans can be used to determine which user model (i.e. which parameter settings) best fits.

2.2 Adaptive Experimental Design in ML

In Bayesian approaches to experimental design the starting assumption is that the posterior probability of a parameter value given an experiment is proportional to the likelihood of the data times the prior $p(\theta|y, d) \propto p(y|\theta, d) \times p(\theta)$. The key question is how to choose experiments that maximise the utility of the data. In Bayesian Optimal Experimental Design (BOED) it is assumed that the objective is to choose experiments that maximise Expected Information Gain (EIG). BOED usually relies on a likelihood model $p(y|\theta, d)$ for predicting the probability of data y (the outcome of an experiment) given an experimental design d and parameter values θ . The objective is then to optimise EIG. EIG can be thought of as the *mutual information* between θ and y .

BOED methods have been successfully applied to the design optimisation problem in various settings [2, 8]. However, BOED can require computationally expensive calculations, such as updating the posterior or estimating the mutual information between the model parameters and experiment outcomes. As these calculations are needed between the time steps of the experiment, this approach becomes impractical for many real-life settings. More recent work has amortised the cost of experiment selection using pre-trained deep neural networks. As an example, Foster et. al [17] suggest a policy network, parameterised by a deep neural network, to produce informative experimental design values. In their approach, the loss function is based on calculating a lower bound of the mutual information instead of costly exact values. This work was extended in [21], in which likelihood functions can be unknown, thus expanding this approach to implicit models. Another line of work is based on using the mutual information as the main criteria for selecting design values, but using neural estimators on the mutual information or its lower bounds [4, 25, 26, 39]. Blau et al [5]

present a Reinforcement Learning formulation for design optimisation. They defined sequential experimental design as a Markov decision process (MDP), highlighting the strong exploration capability of RL-based methods. In their approach, the optimisation target is the lower bound of expected information gain (EIG), comparable to the DAD method in [17].

3 THEORY

We present the theory in two parts. In the first part, we describe a user model with parameters that must be estimated from data. The user model is an example of a class of simulation-based reinforcement learning model that has recently become influential in HCI [12, 13, 22, 37] but which because of their complexity currently lack adequate parameter estimation methods (though see [23, 34].

In the second part of the current section we introduce our proposed *Analyst*. Like the user model, Analyst is also an RL agent and care is needed not to cause confusion. Where for the user model, RL learns a policy that models human cognitive control knowledge, for the analyst, RL learns an policy for choosing experiments and inferring parameters.

3.1 User model

We extended a reinforcement learning model of gaze based target selection previously reported in [11]. The key assumption in the model is that the control of movement is computationally rational: that is, the saccade path and fixations are determined by an attempt to optimise some objective function (e.g. to minimise selection time) given the bounds imposed by the perceptual/motor system. The predicted eye movement strategies are therefore an adaptive consequence of the following constraints: (1) target eccentricity, (2) target size, (3) oculomotor saccade noise, (4) a target detection threshold, and (5) location and target size estimation noise in peripheral vision. The interaction between the target size, target eccentricity, signal-dependent oculomotor noise, eccentricity-dependent estimation noise and size- and eccentricity-dependent target detection results in a multi-step gaze-based selection process. Two-step selections are typical in humans but under some circumstances either one-step (for large targets) or 3-plus-steps (very small targets) can be observed [32].

The user model architecture is illustrated in Figure 3. The blue box contains processes (represented as white rectangles) that constitute a theory of human cognition in interaction with a ‘world’. Each trial begins with the controller choosing an ‘intent’ – a motor movement to an aim point (a location in the world). The chosen intent is implemented via a noisy ‘motor’ process. The motor process results in an ‘action’ which is the actual end point of the motor movement in the world. Additionally, the controller can perform a keypress that terminates the episode. Subsequently, a new stimulus is generated by the world which is the target location and width viewed from the new fixation. This stimulus is perceived by a foveated vision process that generates a noisy estimate of the target location and size (the ‘observation’). The observation provides evidence as to the location and size of the target. This evidence is optimally integrated in a Bayesian ‘memory’ process which outputs a ‘belief’. The memory is observed by the ‘control’ process and by a ‘utility’ function. The utility function generates a reward signal

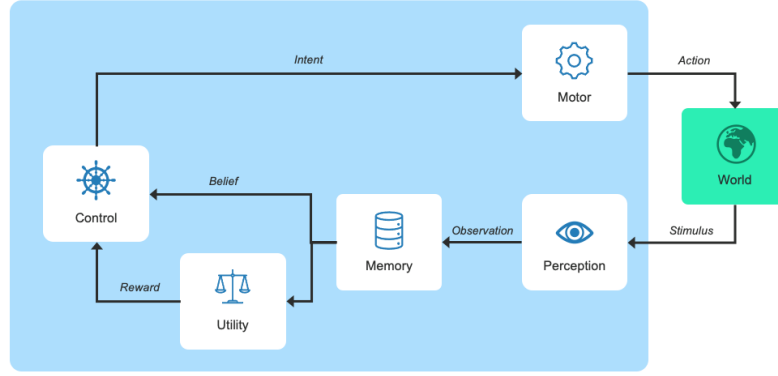


Figure 3: User model architecture. Human cognition is modelled as five processes in interaction with a World. The Motor process models movement noise. Perception models increasing visual noise with eccentricity from the fovea. Memory integrates multiple observations using Bayesian inference and Utility generates a reward signal that captures the trade-off between factors such as speed and accuracy.

that is used to train the controller. The controller's 'intent' are thereby conditioned on the belief.² The architecture implements an action-observation-reward cycle which repeats until the target is selected or the maximum step count for the episode is reached. Learning adjusts the mapping between the observation and the action so as to maximise the cumulative discounted rewards. The architecture can be described formally as a POMDP.

- **State space S :** At each time step t , the environment is in a state $s_t \in S$. A state represents a possible target position and width, and denotes as $s = (f_x, f_y, t_x, t_y, w)$ where (f_x, f_y) is the fixation location and (t_x, t_y) is the target location. For both, $x, y \in [-1, 1]$, with -1 and 1 being the edge of the display, and $w \in [0, 1]$.
- **Action space A :** An action, a_t , is taken at each time step t . On each of these steps the controller decides where to attempt to fixate next (the aim point). An aim point is denoted as a coordinate $a = (a_x, a_y)$ where $a_x, a_y \in [-1, 1]$.
- **Reward function $r(s, a)$:** At each time step t , a reward is generated by a utility function that models the preference utilities of a user. We assumed that users can trade speed for accuracy. Faster speeds are accompanied by more errors. The reward at time t is based on a linear gaze duration model $r(s_t, a_t) = -(\theta_a \times \text{Amplitude}(t) + \theta_b)$, where the slope θ_a and intercept θ_b are parameters of the user. If the user performs a keypress, the episode is terminated and a value $r_{max} \times \theta_{preference}$ is added to the final reward if the target is fixated (gaze is within target radius), otherwise the reward is $-r_{max} \times \theta_{preference}$ if the target is not fixated (an error).³ The parameter $\theta_{preference} \in [0, 1]$ describes the speed-accuracy preference of the user. Additionally, if the maximum amount of steps is reached without a keypress, a termination penalty is added to the final reward.

²Note, that the first action is selected before any observations are made. Therefore, while it is not conditioned on an observation, it may through training with the reward be informed by a prior expectation of the distribution of target locations.

³Here the utility is conditioned directly on the state (cf. Figure 3).

- **Transition function $T(s_{t+1}|s_t, a_t)$:** The environment switches to a new state according to a stochastic transition function. The target location remains unchanged but the fixation location changes according to the outcome of the action aim point. Aim points are corrupted by noise. Therefore, $T(s_{t+1}|s_t, a_t) = N((f_x, f_y)|(a_x, a_y), \sigma_{ocular}(t))$. The oculomotor noise is linearly dependent on the saccade distance (the amplitude) $\sigma_{ocular}(t) = \rho_{ocular} \times \text{Amplitude}(t)$.
- **Observation space O and observation function $o = f(s, a)$:** After taking the action (i.e. saccade to and fixate at a new position on the display), a new observation is received, which is a function of state and action $o_t = f(s_t, a_t)$. The observation of the target position is dependent on the true target location and width (state), and the current fixation location (action). Specifically, the spatial uncertainty of the target position (standard deviation) in peripheral vision is linearly dependent on the distance between the target and the current fixation position, i.e., eccentricity. We similarly assume a linear dependency between the uncertainty and the target size. Therefore, the perceived target position is $\tilde{i}_{x,t} \sim N(t_{x,t}, \sigma_o(t))$, $\tilde{i}_{y,t} \sim N(t_{y,t}, \sigma_o(t))$, where $\sigma_o(t) = \rho_{spatial} \times \text{eccentricity}(t) - \rho_w \times w + \rho_b$, $\rho_{spatial}$, ρ_w and ρ_b are parameters of the model. We also assume that the observed target width is corrupted by a Gaussian noise source $\tilde{w}_t \sim N(w_t, \sigma_w)$. Finally, a binary random variable $z_t \sim p(z|s_t, a_t)$, $z \in \{0, 1\}$, indicates whether the user detects the target. Thus the full observation at time t is the tuple $o_t = (\tilde{x}_t, \tilde{y}_t, \tilde{w}_t, a_t, z_t)$.
- **Discount rate $0 \leq \gamma < 1$.** The model receives a scalar reward at each time step, $r(s_t, a_t)$. The optimal strategy is the one that maximises the expected long-term sum of rewards: $E[\sum_{t=0}^T \gamma^t r(s_t, a_t)]$, given the constraints on the model defined above.

3.1.1 Belief update. If the target is detected by the user (i.e. if $z_t = 1$ in the observation) the memory is updated by integrating

the current belief b_{t-1} and the new observation o_t using Bayes rule (a Kalman filter) [15, 42]. After taking an action (fixating at a location), the model receives noisy observations of the target location and width, which are sampled from a Gaussian distribution, $o_{t,k} \sim N(s_k, \sigma_{o,k}(t))$, with $k \in \{1, 2, 3\}$ (see Observation function above). We omit the index subscript k for clarity. At the time step 1, $b_1 = o_1$, $\sigma_b^2(t=1) = \sigma_o^2(t=1)$. The belief update from t to $t+1$ is shown in Equation (1) below.

$$\begin{aligned} b_{t+1} &= b_t + K_{t+1}[o_{t+1} - b_t] \\ \sigma_b^2(t+1) &= \sigma_b^2(t) - K_{t+1}\sigma_b^2(t) \\ K(t_{t+1}) &= \frac{\sigma_b(t)^2}{\sigma_b(t)^2 + \sigma_o(t+1)^2} \end{aligned} \quad (1)$$

3.1.2 Training. For each study, an ensemble model was trained, meaning that the parameters to be estimated were sampled from their prior distributions and given as additional inputs to the controller, thus distilling a population of user models covering the desired space of behaviours into a single set of neural network weights. The controller was implemented as a simple fully-connected feed-forward neural network and optimised with the proximal policy optimisation (PPO) algorithm [45]. The POMDP was solved for the restricted class of policies that conforms to the task bounds outlined above.

3.2 Analyst

As was the case for the user model, the environment of the Analyst is also formulated as a POMDP. The Analyst architecture is illustrated in Figure 4. The blue boxes form the reinforcement learning agent, and the green boxes represent the environment. The ‘user model’ represents the simulator of a synthetic user. By sampling user parameters from a specified prior distribution, and conditioning the user model on these parameters, a range of different behaviours can be simulated. The user model is also conditioned, and its behaviour affected, by the design values produced by the controller. The parameterised user model generates data given a specific experimental design. The Analyst tries various experimental designs to learn about the user model and thereby estimate its parameters.

The ‘memory’, which is considered as an internal process of the agent, stores the history of experimental designs and outcomes. The ‘control’ process includes the policy function that effectively maps the contents of the memory to a probability distribution over actions. The action contains the design for the next experiment as well as estimations of the user parameters with data collected from past experiments. The ‘discrepancy’ unit is used during the training of the Analyst to produce the reward signal for the RL agent. More formally, the Analyst POMDP is specified as follows:

- **State space S :** At each time step t , the environment is in a state $s_t \in S$. A state $s_t = (d_t, y_t, \theta_p)$ represents the tuple consisting of the design value d_t that was used to run the experiment at that time step, the experiment outcome y_t , and the user parameter vector θ_p .

- **Action space A :** An action, a_t , is taken at each time step t . The action $a_t = (d_{t+1}, \theta_e)$ tuple of the analyst includes designs d_{t+1} for the next experiment and parameter predictions θ_e based on the information gathered so far.
- **Reward function $r(s, a)$:** At each time step t , a reward is generated by a discrepancy function, which measures the similarity of predicted and true parameters of the user model as the negative L1 error $r(s_t, a_t) = -||\theta_p - \theta_e||_1$. The reward is directly influenced by the ability of the analyst to estimate parameters, but it is crucially also indirectly influenced by the analyst’s ability to design informative experiments.
- **Transition function $T(s_{t+1}|s_t, a_t)$:** The environment switches to a new state according to the transition function. The user parameter vector θ_p is sampled from the prior $p(\theta)$ at the beginning of the episode, and remains fixed until the end of the episode. At each time step t , the user parameters θ_p and the design d_{t+1} chosen by the analyst are used to run the simulator and produce an outcome $y_{t+1} \sim p(y|\theta_p, d_{t+1})$, which gives the new state $s_{t+1} = (d_{t+1}, y_{t+1}, \theta_p)$.
- **Observation space O and observation function $o = f(s, a)$:** At each time step t , the state is passed through an observation function $o_t = f(s_t, a_t) = (d_t, \tilde{y}_t)$ before given to the analyst, where \tilde{y}_t is a corrupted measurement of the true experiment outcome y_t . The user parameters θ_p are treated as a latent variable, they are included in the state but not the observation.

The partial observability of the environment motivates a policy that is conditioned on the full history of observations $o_{\leq t} = o_1 o_2 \dots o_t$. The analyst policy is a stochastic function that samples parameter predictions and designs for the next experiment conditioned on the observation history as $a_t \sim \pi^{analyst}(a_t|o_{\leq t})$. The objective of the analyst is to maximise the expectation of discounted return $E[\sum_{t=0}^M \gamma^t r(s_t, a_t)]$ where M is the number of experiments performed for a specific user and γ is the exponential discount rate. This objective is non-myopic since credit assignment for a particular experiment is performed based on the quality of all future parameter estimations.

3.2.1 Policy network. The studies reported below use two different architectures for representing the policy network. In our first study, the observation includes information of the movement time and final fixation, target location and target width. In this case, the policy network implementation is a multilayer perceptron (MLP) network followed by mean pooling across experiments, output layer and heads for action distribution and value estimation (see figure 5, left). In studies 2 and 3, the observation includes eye movement data and information about the target location and width. In these studies, the policy network architecture uses an inductive bias to support relational inference [3, 44] (see figure 5, right).

Considering the Studies 2 and 3 policy architecture, the output of the local pooling is an embedding of the observations over one episode,

$$e^l = \sum_{t=1}^T f_{enc}(c^{target}_t, c^{fixation}_t, c^{fixation}_{t+1}) \quad (2)$$

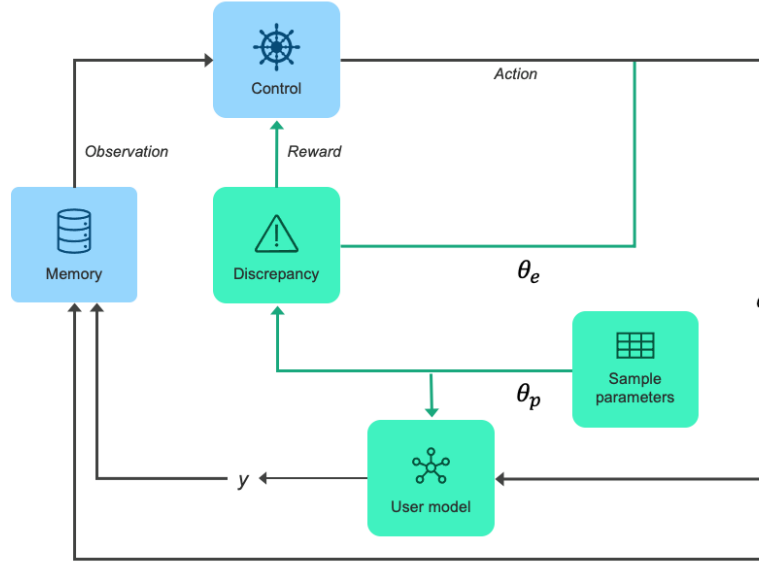


Figure 4: Analyst architecture. The problem for the reinforcement learner is to learn a sequential policy for the ‘control’ process by acting in an environment that consists of a sequence of user models with stochastically sampled parameters. Actions consist of a (d, θ_e) pair where d is a specification of the experiment and θ_e is an estimate of the model parameters. The ‘discrepancy’ function calculates a reward by comparing the estimated parameters to the true parameters θ_p . The ‘memory’ stores the sequence of experiment designs d and data y that have been conducted for the latest sampled parameters.

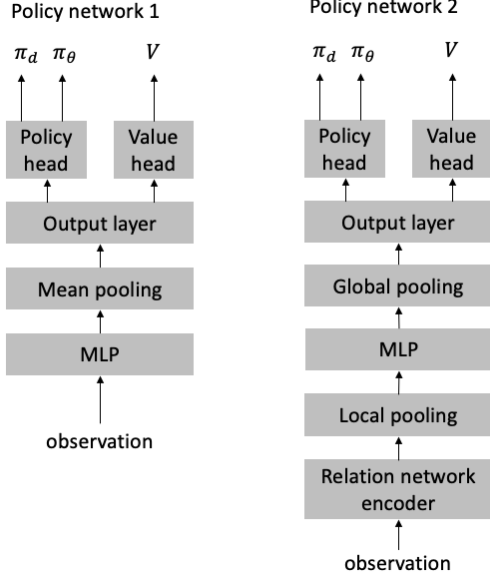


Figure 5: Two different policy network architectures. Study 1 uses a simpler policy network, whereas Studies 2 and 3 use a policy network architecture which implements relational reasoning.

where f_{enc} denotes a relation network encoder (in our case an MLP network), conditioned on information of the target location

c_t^{target} and information of the locations of two subsequent fixations $c_t^{fixation}$ and $c_{t+1}^{fixation}$ at time steps t and $t + 1$. The output of the global pooling is an embedding over M experiments

$$e^g = \sum_{i=1}^M g(e_i^l) \quad (3)$$

where M is the number of experiments and $g(\cdot)$ is an MLP network. The embedding e^g is fed to an fully connected output layer and to policy and value heads.

User models formulated as POMDPs, instead of closed form models, offer the possibility of capturing and simulating complicated and realistic human-like behaviour, but at the same time they raise some challenges. With such user models, we are often only able to draw samples from the simulator, without the possibility of evaluating the likelihood or differentiating through the simulator. Because of this, new methods are needed that fulfill this requirement. In addition to the requirement for non-differentiable user-models, other desired properties are amortisation, non-myopic designs, and adaptation to the experiment outcomes. Table 1 summarises the capabilities of various recent approaches for optimal experimental design. Our method is the most flexible as it is able to generate amortized, non-myopic and adaptive solutions in likelihood-free environments without the need to differentiate through the simulator.

3.3 Training

The policy network is trained by using both score-based and path-wise gradient estimators [33]. Since the design values are sampled

Table 1: Comparison of the capabilities of recent methods for optimal experimental design. * In MINEBED, a backup method is also described based on Bayesian optimisation, which does not require differentiable simulator.

Algorithm	Amortized	differentiable simulator not needed	non-myopic	Adaptive	likelihood-free
DAD [17]	✓	✓	✓	✓	✗
iDAD [21]	✓	✗	✓	✓	✓
Blau et al [5]	✓	✓	✓	✓	✗
MINEBED [25]	✗	✓*	✓	✓	✓
Valentin et al [47]	✗	✓	✓	✓	✓
Our method	✓	✓	✓	✓	✓

from the action distribution, it uses a score-based gradient estimator implemented with PPO [45]. Since the parameter estimations can be directly optimised with a loss function, extra parameter updates can be conducted with the pathwise gradient estimations thereby reducing variance.

The training method uses an exponential moving average (EMA) method in order to regularise the training process [46]. The parameters of a separate EMA network are updated based on the parameters from the policy network, and the previous values of the EMA network:

$$\phi'_t = \alpha \phi'_{t-1} + (1 - \alpha) \phi_t \quad (4)$$

where ϕ'_t is the EMA network parameter vector at time step t , ϕ_t is the parameter vector of the policy network, and α is the smoothing coefficient hyperparameter. The rewards for the policy network parameter updates are calculated by using parameter predictions from the EMA network, instead of the actual policy network. As a result, the policy parameter updates are less noisy and training is more efficient.

The training was conducted by using the Stable Baselines 3 (SB3) library [40] by implementing custom policy networks for the relation network, and by using the SB3 callback system for EMA implementation.

4 NON-MYOPIC AND ADAPTIVITY DEMONSTRATIONS

Our method produces amortised design strategies and parameter estimations in likelihood free settings without the need to differentiate through the simulator. As presented in Table 1, the existing methods are not applicable to our setting. To gain confidence for our method, its performance and range of applicability, we ran experiments on abstract tasks requiring non-myopic design strategies and adaptivity.

4.1 Non-myopic demonstration

We demonstrate first the capability of our method to produce non-myopic design selection strategies with an abstract task. We use a 1-dimensional Gaussian process to sample functions with a specified kernel. Depending on the selected kernel, nearby points are correlated. As a baseline for comparison to our approach, we use

a one step lookahead algorithm which minimises expected uncertainty reduction over the design space. This myopic algorithm will converge to non-optimal design strategy with two data points as the algorithm selects designs that are less uniformly distributed. In contrast, our non-myopic algorithm is able to take into account the total number of trials to be conducted and can design a whole series of experiments that lead to higher information gain overall. The figure 6 illustrates examples of design selections of both algorithms. In the left panel, the design is calculated with one-step lookahead so as to optimally reduce uncertainty. The mid-point design is chosen first in 'ignorance' of the fact that another experiment is to be conducted and as a consequence the overall distribution of experiments is not optimal. In the right panel, our non-myopic algorithm has successfully learned to choose designs that more evenly cover the design space, thereby gaining more information over the whole series of experiments and optimally reducing uncertainty overall.

The discrepancy function for the Analyst is the L2 distance between the true and estimated function. We measured the quality of the designs with two metrics, namely the L2 distance between the estimated and ground truth functions, and the reduction of variance as calculated with the integrated mean-squared error (IMSE) method as used in Chen et al. 2019 [10]. As the myopic algorithm is analytically calculated to reduce variance, the latter metric offers a reliable comparison to the baseline. The results are shown in table 2, indicating a clear benefit of the non-myopic method.

Table 2: Results of the non-myopic experiment. The Analyst learns a non-myopic design strategy that outperforms the optimal myopic design strategy. The mean and standard error was computed over an evaluation batch with 100 functions sampled from the prior.

Metrics	mean	standard error
Analyst discrepancy	0.0812	0.0089
Baseline discrepancy	0.1580	0.0204
Analyst IMSE	0.3909	0.1194
Baseline IMSE	0.7087	0.0

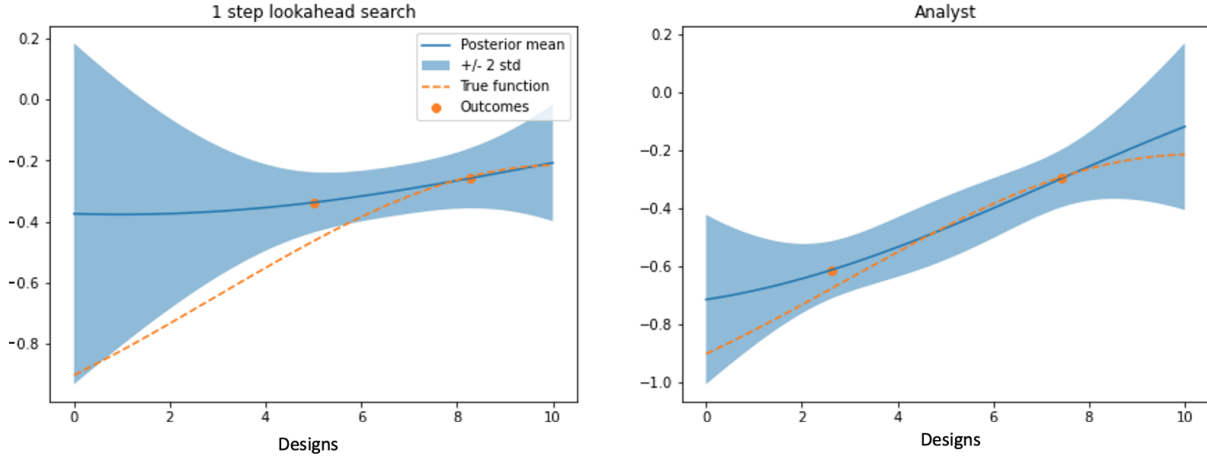


Figure 6: Examples of design selections by myopic and our non-myopic methods. Left panel: A myopic method has selected a design to reduce variance with optimal one-step lookahead. Right panel: Our method has learned a non-myopic design strategy and can select more informative design values as the designs are chosen to more evenly cover the whole space.

4.2 Adaptivity demonstration

The adaptivity is demonstrated in a setting where the task is to estimate a parameter that affects the positioning of a logistic sigmoid function on the x-axis. The data is generated by conducting Bernoulli trials where the probability $P(y_i = 1)$ is defined by a logistic function:

$$P(y_i = 1|\theta, d) = \frac{1}{1 + e^{-(d+\theta)}}$$

where y_i is the outcome of the i :th trial in the experiment, θ is the parameter to be estimated and d is the design value. To train a non-adaptive baseline, we masked the outcomes from the Analyst until the last time step of the episode. In contrast, when the outcomes of the previous trials during the episode are available, the Analyst can generate an adaptive strategy where design choices are affected by the previous outcomes during the episode. To compare Analyst to the baseline, we use the MSE of the estimated parameter. The table 3 indicates a clear benefit of the adaptive design strategy over the non-adaptive strategy. The adaptivity of the designs is illustrated in Figure 7 which shows how the Analyst design selections converge on the mid-point of the slope in the sigmoid function.

5 STUDY 1: ESTIMATION OF SUBMOVEMENT NOISE FROM SUMMARY DATA

Following Meyer’s law we assume that people make a series of submovements to achieve a goal [31]. In the first study, we apply our approach to a scenario where only summary statistics are provided to the analyst at the end of each episode. The summary statistics includes movement time from the beginning of the episode until the end of the episode, the target location, target width and the location of the final submovement. Although individual submovements are not observed, the goal is to infer the noise parameter affecting these movements. The movement time is calculated with the formula $mt = \sum_{i=1}^I (\theta_a x_i + \theta_b)$, where I is the number of user steps within the episode, x_i is the distance of the submovement during the current user step i , θ_a is the slope parameter, and θ_b is the intercept

Table 3: Results of the adaptivity demonstration. To provide a baseline, the Analyst trained without information of the previous outcomes during the episode learns a non-adaptive strategy. In contrast, when previous outcome information is available, the Analyst produces an adaptive strategy which beats both the non-adaptive baseline and a random baseline. The MSE mean and MSE standard error was computed over an evaluation batch with 10000 models sampled from the prior.

Method	MSE mean	MSE standard error
Adaptive Analyst	2.018	0.034
Non-adaptive baseline	6.265	0.085
Random design baseline	14.434	0.234

parameter of the movement time model. In Study 1, (though not in subsequent studies) both of these movement time parameters are fixed. The goal in Study 1 is to estimate the movement noise parameter.

5.1 Results

The larger the movement noise, the further the actual movements can be from the intended aim points. As a result, with large noise values the user model requires more steps to reach the target. In order to verify that the user model has adapted to the various movement noise values, we plot the number of steps that the user model needs to reach the target on average. Figure 8a shows that the fully trained user model require more steps to reach the target when the noise levels are higher. As the user model’s behaviour is impacted by the parameter value, it should be possible to train the Analyst to infer the movement noise from behavioural observations.

We tested the ability of the Analyst to estimate parameters and the results are reported in Figure 9. In each of the panel’s (a), (b) and (c) the true value of the movement noise is plotted against the estimated value. A linear regression fit was performed to all three

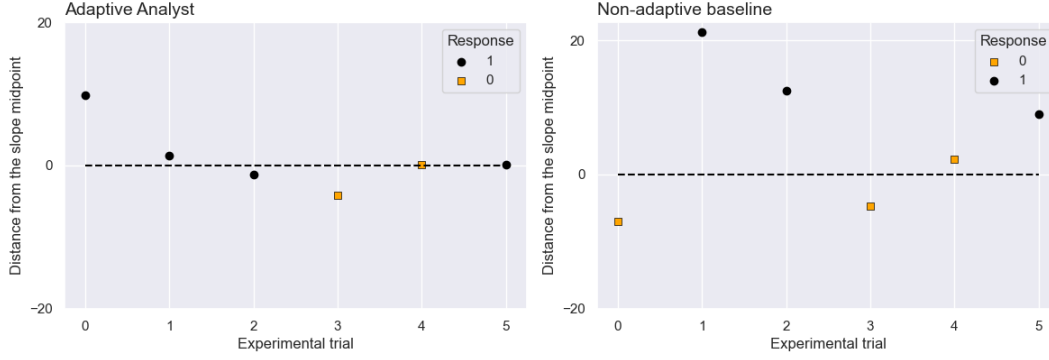


Figure 7: The Analyst learns to adapt to the data collected so far during the episode. Left panel: When the Analyst has information about the previous trials, it adapts to them and converges to the midpoint of the sigmoid function slope, which provides most information. Right panel: When information about the previous trial outcomes are not available, the Analyst learns a non-adaptive strategy in which the design values are sparsely distributed rather than focused where they are needed.

scatter plots ($a = 0.99, b = 0.91, R^2 = 0.80$ for panel (a), $a = 0.83, b = 0.01, R^2 = 0.77$ for panel (b), $a = 1.00, b = 0.03, R^2 = 0.54$ for panel (c)). The panels differ in the level of the perceptual noise, with low perceptual noise in panel (a), higher perceptual noise in panel (b) and highest perceptual noise in panel (c). Across all three panels it is clear that the Analyst is able to provide some level of estimate of the oculomotor noise, however, it is also clear that the estimates are much better when the perceptual noise is lower (panel (a)).

Figure 9 also illustrates the distribution of experimental design choices made for each level of perceptual noise. Panels (d) and (e), which correspond to panel (a), show that for low perceptual noise, experimental designs tend toward higher eccentricities and large targets (though there is some variance). However, as perceptual noise increases, the analyst is incentivised to select targets closer to the origin, in order to avoid corrupting the data with very noisy observations. This claim is supported by panels (f) and (h) which show more experiments with smaller eccentricities at higher perceptual noise values.

Finally, the quality of the design optimisation can be verified by training the analyst with random design values, and comparing the accuracy of the parameter estimations with an analyst trained to optimise design selections. The right panel of Figure 8 is a plot of the error in the estimate against the experiment number. With zero experiments there is no data and the parameter estimate is simply the learned prior of the parameter distribution. After the data from each experiment are incorporated into a new parameter estimate, the error decreases, and it decreases more quickly with Analyst designed experiments. It demonstrates a clear improvement in the accuracy of the parameter estimations with optimised designs.

6 STUDY 2: INFERENCE OF PERCEPTUAL AND MOTOR NOISE FROM GAZE MOVEMENTS

In Study 2, we applied the Analyst to a user model of eye-movements. Instead of using summary statistics as in Study 1, we allowed Analyst to observe the gaze fixations of each step during the episode. We also extended the model with a target detection requirement. As a consequence, whether or not the user model observes the target is

probabilistic. The probability of not observing the target increases when the target size is far away and target width is small.

In this study we report the results of the Analyst inferring three parameters: oculomotor noise, perceptual noise and movement time intercept. The observations include fixations at each time step, duration of each gaze, information about the target location and target width. A key difference between Study 1 and Study 2 is that, where there is a fixed amount of data per experiment in Study 1, in Study 2, experiments that with more distance and smaller targets can generate more data. As we will see, this fact impacts the selected designs.

6.1 Results

As with Study 1, we first measured the effect of changes in the user model parameters on the behaviour. When considering oculomotor and perceptual noise values, increasing one of these noise values while keeping another fixed should cause the user model to require more gaze fixations to reach the target. This is clearly visible in the panel (a) in Figure 10, where the perceptual noise value versus required gaze steps to reach the target is plotted. The near linear increase in number of steps with increasing noise suggests that the parameter should be readily recoverable. The oculomotor noise is equivalent to movement noise in study 1, the effect of which is shown in Figure 8 of Study 1.

In Study 2, the goal of Analyst is to select the most informative experiments for inferring the oculomotor noise, perceptual noise and the intercept parameters for the movement time model. Analyst performance is illustrated in panels (e) to (g) of Figure 10. A linear regression was performed to assess the quality of the fits for each parameter ($a = 0.85, b = 0.01, R^2 = 0.76$ for panel (e), $a = 0.87, b = -0.01, R^2 = 0.84$ for panel (f), $a = 1.00, b = 0.04, R^2 = 0.99$ for panel (g)). The selected design values are illustrated in panels (b) and (c) in Figure 10. In this case, the analyst has learned to design experiments with small targets and large eccentricities.

Finally, the performance of the analyst is compared against an analyst trained by using random experimental designs. Panel (d)




Figure 8: Study 1 results. Panel (a), The number of steps required by the user model to reach the target against increasing movement noise values. Panel (b), The Analyst makes better parameter estimations the more it makes experiments. The analyst also performs better across stages compared to a baseline that samples designs uniformly. In panel (a), the shaded area represents ± 1 standard error over an evaluation batch with 1000 user episodes. In panel (b), the shaded area represents ± 1 standard error over an evaluation batch with 1000 models sampled from the prior.

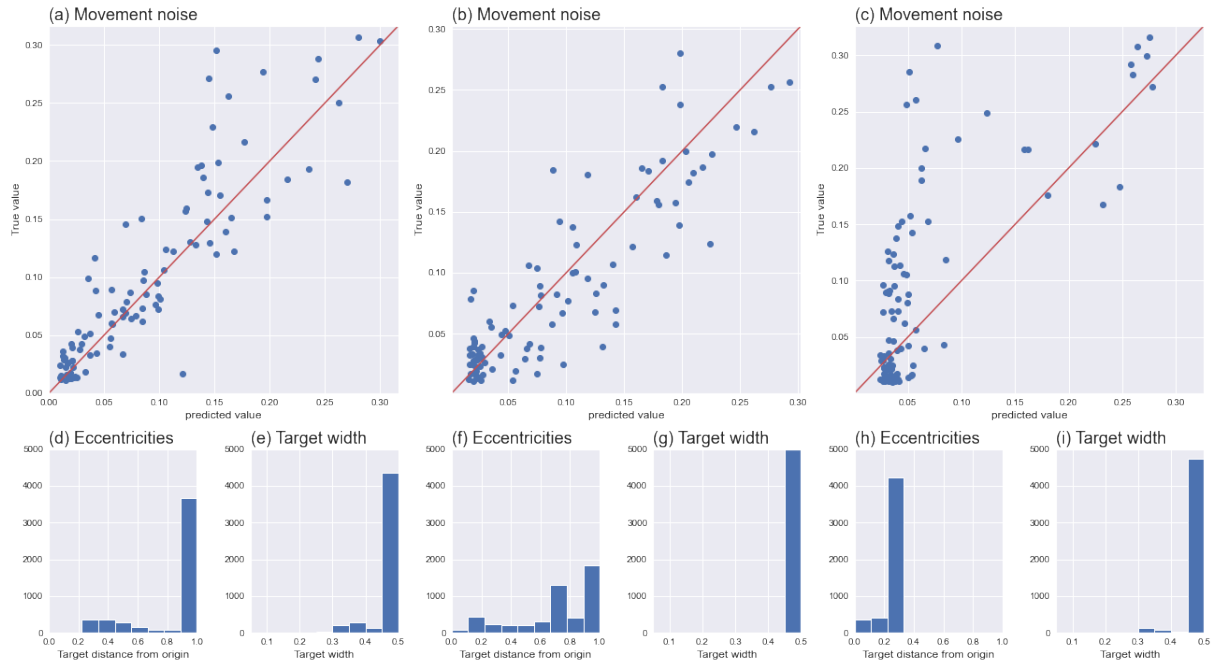


Figure 9: Study 1 results. Panels a-c illustrate the accuracy of the movement noise estimations when the user model has low (panel a), medium (panel b) or high (panel c) perceptual noise. Below the scatter plots, the corresponding design selections (target distance from origin and width) by the analyst are illustrated as histograms (panels d-i).

in Figure 10 shows a clear benefit of the optimised designs across experiments.

In summary, the results of Study 2 extend Study 1 by showing that the analyst can choose experimental designs and accurately infer multiple parameters at the same time. In addition, it can do so when each experiment returns a sequence of data (fixation locations and movement times) and not just point values. Finally, estimating parameters with selected designs outperforms doing so with random designs.

7 STUDY 3: INFERENCE OF PREFERENCES

In Study 3, we test to see whether Analyst can discover user model preferences. Preferences are important to HCI as they capture personal and sometimes discretionary aspects of how a person wants to interact with a computer. Preferences include preferences for music genre and/or movie directors, for example, but here we focus on speed-accuracy trade-offs. As we have said, the speed-accuracy trade-off is a significant determinant of how people choose to interact with computers [48] and is readily detectable in behaviour.

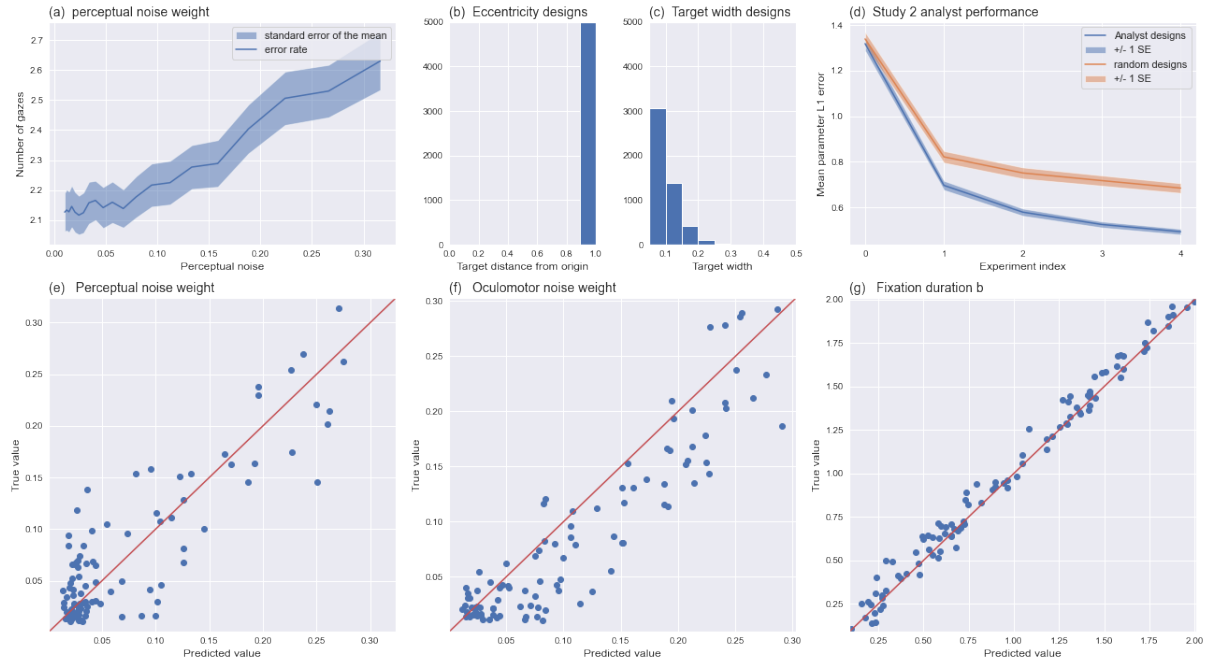


Figure 10: Results of Study 2. Panel (a) shows the effect of the user model’s perceptual noise parameter on number of steps taken. Panels (b) and (c) show the distribution of the experimental designs chosen by the Analyst. Panel (d) shows the improvement in parameter estimates with Analyst designed experiments versus with random experiments. The panels on the lower row show the accuracy of the parameter estimations for perceptual noise (e), oculomotor noise (f), movement time interception parameter (g). In panel (a), the shaded area represents ± 1 standard error over an evaluation batch with 1000 user episodes. In panel (d), the shaded area represents ± 1 standard error over an evaluation batch with 1000 models sampled from the prior.

The speed-accuracy trade-off determines the error rate which is a key property of interaction. The faster users attempt to perform a pointing task then the more errors that they make. In Study 3 we tested the extent to which Analyst was capable of estimating these preference parameters. To do so, we simulated a task in which the user model is able to end an episode by pressing a key on a keyboard. The observation space is extended to include information about whether the user has pressed a key or not.

7.1 Results

As with the previous studies, we first test whether the parameters, particularly the preference parameter, make an identifiable difference to the behaviour of the user model. In this case a useful metric for measuring the response of the parameters to the behaviour of the user agent is the error rate, which describes how often the user model ends the episode when the gaze is not in the true target, or the maximum number of gazes is reached without a keypress by the user. Figure 11, panels (a) to (d) show the error rate for four of the model parameters. Panel (c) shows how the error rate is larger when the preference is biased towards speed, and approaches zero when the preference is biased towards accuracy.

The lower row, panels (e)-(h) of Figure 11 shows the capability of the analyst to simultaneously infer all four parameters. As the Study 3 task is made more difficult by the increased number of

parameters, the error of the estimates is greater when compared to Study 2. However, there is a good correlation between ground truth and parameter estimate for all four parameters. Having said that the worst of the four estimates is for the preference parameter which has a noticeable skew.

Figure 12 shows the Analyst’s distribution of experimental design choices in two histograms. It always picks the smallest target (panel b) but shows a broader distribution of choice of distances. It is instructive to compare this distribution to that for Study 2.

In Study 2, a good strategy for the analyst was to select a small target, as far away from the fovea as possible. In contrast, less extreme distances are selected in Study 3. One reason for this difference may be that because of the speed-accuracy trade-off there is a risk that the smallest most distant targets are not selected at all and therefore there is little evidence gathered to inform the speed-accuracy trade off preference parameter. Therefore, in Study 2 it makes sense to select experimental designs in which the user model stands some chance of high accuracy. Lastly, the Analyst’s optimal experiments outperform an analysis conducted with data from random experimental designs (Figure 12 panel (c)).

In summary, Study 3 demonstrates that Analyst can simultaneously estimate multiple parameters, and importantly, it can estimate both capacity parameters (e.g. oculomotor and perceptual noise) and preference parameters (e.g. speed-accuracy trade off) – from

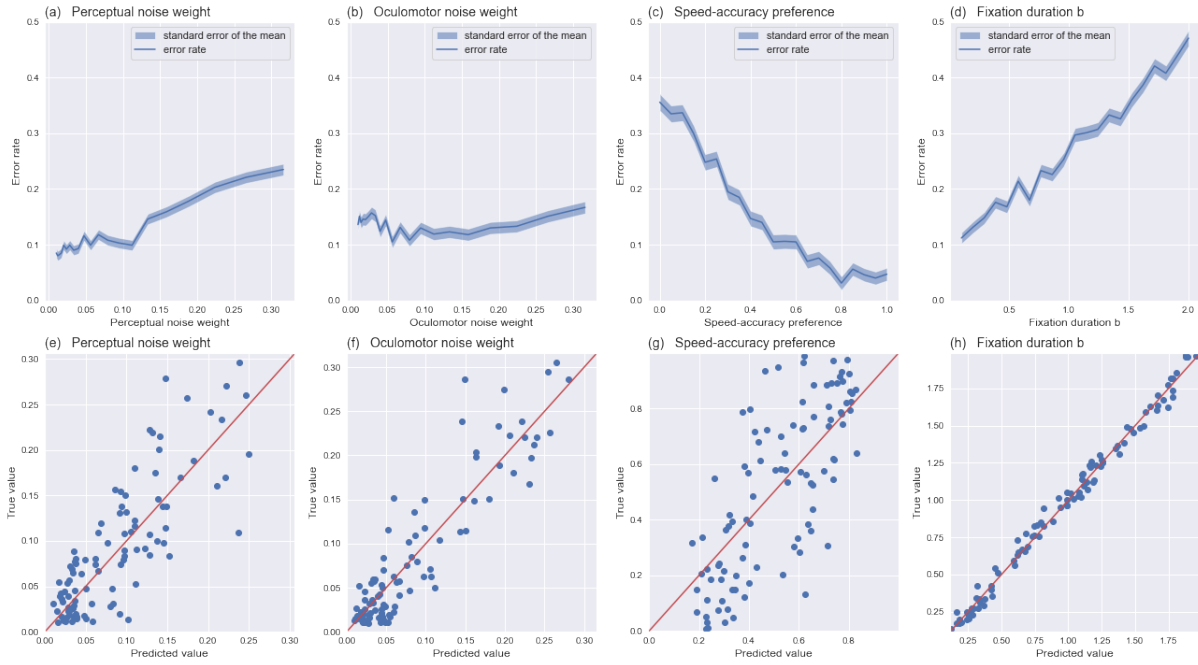


Figure 11: Results of Study 3. The upper row illustrates the adaptation of the user model to parameters for the perceptual noise (a), oculomotor noise (b), speed-accuracy preference (c) and movement time intercept (d). The lower row illustrates the corresponding accuracy of the parameter estimates. In panels (a) - (d), the shaded area represents ± 1 standard error over an evaluation batch with 1000 user episodes.

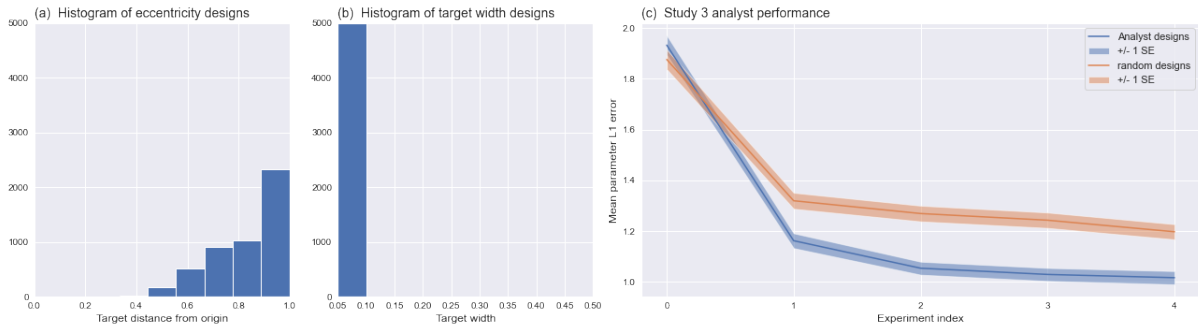


Figure 12: Results for Study 3. Panels (a) and (b) show the histograms of Analyst selected experimental designs. Panel (c) shows the performance gain that follows from inferring parameters on the basis of the designs selected by the analyst compared to random designs. In panel (c), the shaded area represents ± 1 standard error over an evaluation batch with 1000 models sampled from the prior.

the same experiments. While there is a noticeable decrease in the quality of the parameter estimates with the increase in the number of parameters under consideration, compelling correlations are still generated. As with the other two studies the analyst selected experimental designs outperformed the random designs.

8 DISCUSSION

We have explored the properties of a new method of user model parameter estimation and shown that, for three progressively complex pointing tasks, it can make rapid estimates of parameter values for individual simulated users. It does so by learning a near-optimal policy for choosing the experiments that are most likely to generate informative data.

All three of the studies showed that the learned policies lead to more accurate parameter estimates than random experiments.

Parameters could be estimated both with data that summarised a sequence of user submovements (Study 1), as well as with data that included multiple steps in a single experimental observation (Studies 2 and 3) – two important types of user data found in HCI. Further, in Study 1, 2 and 3, each of four successive designed experiment led to an improvement in the parameter estimates and a concomitant reduction in the prediction error. This improvement was more rapid for designed experiments than for random experiments. The estimated parameter values were highly correlated with the true value after only four experiments.

While we have demonstrated the viability of the amortised approach for pointing tasks, more work is needed to verify that it generalises to other HCI tasks. While we believe that the approach is structured so as to provide design of experiments and inference of parameters for any complex simulation-based user models, further empirical work is needed. If generalisation is possible then, while very high computational costs are paid during training (between 3 and 9 hours of wall time on a laptop for the simulations reported above), the result is a very fast (milliseconds), data-lean (four observations) deployment. In other words, the approach trades the high cost of training an ensemble model and training the analyst, for a subsequent reduction in the time required to infer user model parameters. Post-training, not only can the best next experiment be determined in milliseconds but in addition the analyst can minimise the total number of experiments required to determine best fitting parameter values. However, further empirical work is required to see whether this promise is delivered for a broad range of HCI tasks.

Additional properties of our approach include that it provides adaptive, non-myopic designs (See the experiments in Section 4) and that it can be trained with arbitrary non-differentiable simulators with intractable likelihoods. As we see in the next Section, these properties compare well with other work in this area.

8.0.1 Comparison to related approaches. The work reported above was inspired and informed by recent work in both HCI and machine learning. Of particular importance was work on ensemble user models [27, 34] and work on reinforcement learning based experimental design and inference [5]. Also of importance is the work on Bayesian approaches to optimal experimental design [7, 17, 25, 26, 36] and work on relation nets [44] which was crucial to tractable reinforcement learning.

Bayesian experimental design has the clear potential advantage of mathematical rigour, as well as estimates of the posterior distribution of parameter values, rather than point estimates. The Bayesian framework, reviewed briefly above, tackles the problem of experimental design by optimising the expected information gain e.g. how much more certain we will become about the values of the parameters we are fitting. EIG is equivalent to maximising the Mutual Information (MI) between the parameters and data when performing the experiment design. However, EIG does not account for inaccuracies resulting from the amortisation in the parameter estimation. Thus, it is not clear that optimising EIG leads to good designs in our setting. Instead of optimising EIG, in the RL approach it is possible to directly optimise the designs for amortised parameter estimation through a joint objective. Also, estimating the MI in the conventional BOED framework is doubly intractable [21, 41]. Due to this computational complexity, estimating Bayes optimal designs

is not feasible when doing experiments. This has led to approaches that amortise the cost to a pre-trained deep network by using a tractable approximation to the MI [5, 17, 21]. Tractable computation of the MI objective for implicit models is further complicated by the fact that the likelihood function of the parameter is not known [14, 29]. Ivanova et al. [21] tackle this problem by introducing a separate critic network. Their approach however requires the use of a differentiable simulator. Using RL alleviates this need, as the score function gradient estimator directly calculates gradients from the specified reward.

8.1 Future work

The results reported above represent a preliminary investigation of the potential of RL-based experimental design and inference in HCI. They suggest a number of future studies. Perhaps most significantly, the effectiveness of the method must be tested with human participants. While the simulated participants used to test the approach above are sampled from the distribution of real users and previous studies have demonstrated the human-like behaviour of these simulations [11], further work is needed. While human studies are beyond the scope of the current article, the software that we have built makes it very easy to deploy the Analyst learned policies in interactive software with an eye tracker. This software would – without lag – choose the best experimental design (target distance and width), observe a user's saccades and fixations, update model parameters, update the observation history and repeat.

The approach must also be tested on a broader range of tasks so as to empirically establish its generality. While we have formalised the method in terms that we believe to be fully general, in this paper, we have only tested it on abstract problems and pointing tasks. A broader range of HCI-related tasks would include menu-search tasks [12], decision making tasks [13], and biomechanical control tasks [20], to name but three.

In the future, user models with rapid parameter estimation could help enhance interaction for each individual user. Mouse gain functions, text completion, icon sizes, colour pallet, etc. are almost always never tuned to an individual's preferences and capacities. Instead, interfaces provide settings by which the interaction can be 'adapted' manually requiring the user to actively choose configurations in accordance with their beliefs about what is good for them. We believe that this process could be complemented with automatic personalisation methods based upon the RL-based Analyst reported above.

In addition, further work is needed on picking good hyperparameters for the Analyst. As the reported studies became more complex, training of the neural networks became more challenging and some amount of hyperparameter tuning was involved in generating the results reported above. Finding the very best possible performance with extensive hyperparameter search was not within the scope of the current article. Therefore, the performance for both optimised Analyst and Analyst using random designs could be improved.

Lastly, further work is needed to explore the implications of RL-based amortised parameter estimation for a range of HCI-related problems. A/B testing, for example, can be enhanced by first fitting a user model and then selecting an interaction design accordingly. With amortised methods, it may be possible to do this in

real time. Similarly, recommender systems and decision support systems might be enhanced by fitting simulation-based user models to user preferences and selecting recommendations with the aid of the fitted user model.

In conclusion, we have demonstrated, in-silico, the potential utility of amortised experimental design and parameter estimation based on an RL algorithm that learns how to choose experiments for estimating user model parameters. By doing so, it can reduce the subsequent time cost of interaction with humans and the amount of human data.

ACKNOWLEDGMENTS

This work was supported by the Academy of Finland (Flagship programme: Finnish Center for Artificial Intelligence FCAI) and ELISE Networks of Excellence Centres (EU Horizon:2020 grant agreement 951847) and Bitville Oy. The authors want to thank the Probabilistic Machine Learning (PML) and the User Interfaces research groups at Aalto University for fruitful discussions and feedback.

REFERENCES

- [1] John R Anderson, Daniel Bothell, Michael D Byrne, Scott Douglass, Christian Lebiere, and Yulin Qin. 2004. An integrated theory of the mind. *Psychological Review* 111, 4 (2004), 1036–60. <https://doi.org/10.1037/0033-295X.111.4.1036>
- [2] Anthony Atkinson, Alexander Donev, and Randall Tobias. 2007. *Optimum experimental designs, with SAS*. Vol. 34. OUP Oxford.
- [3] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. 2018. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261* (2018).
- [4] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and Devon Hjelm. 2018. Mutual information neural estimation. In *International conference on machine learning*. PMLR, 531–540.
- [5] Tom Blau, Edwin V Bonilla, Iadine Chades, and Amir Dezfouli. 2022. Optimizing Sequential Experimental Design with Deep Reinforcement Learning. In *International Conference on Machine Learning*. PMLR, 2107–2128.
- [6] S K Card, T P Moran, and A Newell. 1983. *The Psychology of Human-Computer Interaction*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- [7] Daniel R Cavagnaro, Richard Gonzalez, Jay I Myung, and Mark A Pitt. 2013. Optimal decision stimuli for risky choice experiments: An adaptive approach. *Management science* 59, 2 (2013), 358–375.
- [8] Kathryn Chaloner and Isabella Verdinelli. 1995. Bayesian experimental design: A review. *Statist. Sci.* (1995), 273–304.
- [9] Haiyang Chen, Hyung Jin Chang, and Andrew Howes. 2021. Apparently Irrational Choice as Optimal Sequential Decision Making. In *AAAI'21 Proceedings of the Thirty-fifth AAAI Conference on Artificial Intelligence*. AAAI Press.
- [10] Jialei Chen, Simon Mak, V Roshan Joseph, and Chuck Zhang. 2019. Adaptive design for Gaussian process regression under censoring. *arXiv preprint arXiv:1910.05452* (2019).
- [11] Xiuli Chen, Aditya Acharya, Antti Oulasvirta, and Andrew Howes. 2021. An Adaptive Model of Gaze-based Selection. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [12] Xiuli Chen, Gilles Bailly, Duncan P Brumby, Antti Oulasvirta, and Andrew Howes. 2015. The Emergence of Interactive Behaviour: A Model of Rational Menu Search. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 4217–4226.
- [13] Xiuli Chen, Sandra Dorothee Starke, Chris Baber, and Andrew Howes. 2017. A cognitive model of how people make decisions through interaction with visual displays. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 1205–1216.
- [14] Kyle Crammer, Johann Brehmer, and Gilles Louppe. 2020. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences* 117, 48 (2020), 30055–30062.
- [15] Peter Dayan and Nathaniel D Daw. 2008. Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience* 8, 4 (2008), 429–453.
- [16] Seungwon Do, Minsuk Chang, and Byungjoo Lee. 2021. A Simulation Model of Intermittently Controlled Point-and-Click Behaviour. In *CHI Conference on Human Factors in Computing Systems*. 1–21.
- [17] Adam Foster, Desi R Ivanova, Ilyas Malik, and Tom Rainforth. 2021. Deep adaptive design: Amortizing sequential bayesian experimental design. In *International Conference on Machine Learning*. PMLR, 3384–3395.
- [18] Christoph Gebhardt, Antti Oulasvirta, and Otmar Hilliges. 2020. Hierarchical Reinforcement Learning Explains Task Interleaving Behavior. *Computational Brain & Behavior* (2020), 1–21.
- [19] Andrew Howes, Richard L Lewis, and Alonso Vera. 2009. Rational adaptation under task and processing constraints: Implications for testing theories of cognition and action. *Psychological review* 116, 4 (2009), 717.
- [20] Aleksi Ikkala, Florian Fischer, Markus Klar, Miroslav Bachinski, Arthur Fleig, Andrew Howes, Perttu Hämäläinen, Jörg Müller, Roderick Murray-Smith, and Antti Oulasvirta. 2022. Breathing Life Into Biomechanical User Models. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–14.
- [21] Desislava Ivanova, Adam Foster, Steven Kleinegesse, Michael U Gutmann, and Thomas Rainforth. 2021. Implicit Deep Adaptive Design: Policy-Based Experimental Design without Likelihoods. *Advances in Neural Information Processing Systems* 34 (2021).
- [22] Jussi Jokinen, Aditya Acharya, Mohammad Uzair, Xinhui Jiang, and Antti Oulasvirta. 2021. Touchscreen Typing As Optimal Supervisory Control. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [23] Antti Kangasrääsiö and Samuel Kaski. 2018. Inverse reinforcement learning from summary data. *Machine Learning* 107, 8 (2018), 1517–1535.
- [24] David E Kieras and Anthony J Hornof. 2014. Towards accurate and practical predictive models of active-vision-based visual search. In *ACM CHI'14*. 3875–3884.
- [25] Steven Kleinegesse and Michael U Gutmann. 2020. Bayesian experimental design for implicit models by mutual information neural estimation. In *International Conference on Machine Learning*. PMLR, 5316–5326.
- [26] Steven Kleinegesse and Michael U Gutmann. 2021. Gradient-based Bayesian experimental design for implicit models using mutual information lower bounds. *arXiv preprint arXiv:2105.04379* (2021).
- [27] Minhae Kwon, Saurabh Daptardar, Paul R Schrater, and Xaq Pitkow. 2020. Inverse rational control with partially observable continuous nonlinear dynamics. *Advances in neural information processing systems* 33 (2020), 7898–7909.
- [28] Richard L Lewis, Andrew Howes, and Satinder Singh. 2014. Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in cognitive science* 6, 2 (2014), 279–311.
- [29] Jarno Lintusaari, Michael U Gutmann, Ritabrata Dutta, Samuel Kaski, and Jukka Corander. 2017. Fundamentals and recent developments in approximate Bayesian computation. *Systematic biology* 66, 1 (2017), e66–e82.
- [30] I S MacKenzie. 1992. Fitts' law as a research and design tool in human-computer interaction. *Human-computer interaction* 7, 1 (1992), 91–139.
- [31] D.E. Meyer, D.E. Irwin, A.M. Osman, and J. Kounios. 1988. The dynamics of cognition and action: mental processes inferred from speed-accuracy decomposition. *Psychological Review* 95, 2 (1988), 183–237.
- [32] D.E. Meyer, R.A. Abrams, S. Kornblum, C.E. Wright, J.E.K. Smith, and J E Keith Smith. 1988. Optimality in human motor performance: ideal control of rapid aimed movements. *Psychological Review* 95, 3 (1988), 340.
- [33] Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih. 2020. Monte Carlo Gradient Estimation in Machine Learning. *J. Mach. Learn. Res.* 21, 132 (2020), 1–62.
- [34] Hee-Seung Moon, Seungwon Do, Wonjae Kim, Jiwon Seo, Minsuk Chang, and Byungjoo Lee. 2022. Speeding up Inference with User Simulators through Policy Modulation. In *CHI Conference on Human Factors in Computing Systems*. 1–21.
- [35] Jay I Myung, Daniel R Cavagnaro, and Mark A Pitt. 2013. A tutorial on adaptive design optimization. *Journal of mathematical psychology* 57, 3-4 (2013), 53–67.
- [36] Jay I Myung and Mark A Pitt. 2016. Model comparison in psychology. *The Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience* 5 (2016).
- [37] Antti Oulasvirta, Jussi PP Jokinen, and Andrew Howes. 2022. Computational Rationality as a Theory of Interaction. In *CHI Conference on Human Factors in Computing Systems*. ACM, 1–14.
- [38] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019).
- [39] Ben Poole, Sherjil Ozair, Aaron Van Den Oord, Alex Alemi, and George Tucker. 2019. On variational bounds of mutual information. In *International Conference on Machine Learning*. PMLR, 5171–5180.
- [40] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. 2021. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research* 22, 268 (2021), 1–8. <http://jmlr.org/papers/v22/20-1364.html>
- [41] Tom Rainforth, Rob Cornish, Hongseok Yang, Andrew Warrington, and Frank Wood. 2018. On nesting monte carlo estimators. In *International Conference on Machine Learning*. PMLR, 4267–4276.
- [42] Rajesh PN Rao. 2010. Decision making under uncertainty: a neural model based on partially observable markov decision processes. *Frontiers in computational neuroscience* 4 (2010), 146.

- [43] Elizabeth G Ryan, Christopher C Drovandi, James M McGree, and Anthony N Pettitt. 2016. A review of modern computational algorithms for Bayesian optimal design. *International Statistical Review* 84, 1 (2016), 128–154.
- [44] Adam Santoro, David Raposo, David GT Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. 2017. A simple neural network module for relational reasoning. In *Advances in Neural Information Processing Systems*.
- [45] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [46] Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems* 30 (2017).
- [47] Simon Valentin, Steven Kleinegesse, Neil R Bramley, Michael U Gutmann, and Christopher G Lucas. 2021. Bayesian Optimal Experimental Design for Simulator Models of Cognition. *arXiv preprint arXiv:2110.15632* (2021).
- [48] Shumin Zhai, Jing Kong, and Xiangshi Ren. 2004. Speed–accuracy tradeoff in Fitts’ law tasks—on the equivalency of actual and nominal pointing precision. *International journal of human-computer studies* 61, 6 (2004), 823–856.
- [49] Yunfeng Zhang and Anthony J Hornof. 2014. Understanding multitasking through parallelized strategy exploration and individualized cognitive modeling. In *ACM CHI’14*. 3885–3894.

A APPENDIX: ARCHITECTURE AND HYPERPARAMETERS DETAILS

All simulations were run using PyTorch deep learning [38] and Stable Baselines3 (SB3) Reinforcement Learning [40] libraries. As the focus on this paper is to show the capability of our method, we mainly use network architectures and hyperparameters that are typical for neural network training with only minor hyperparameter search.

A.1 User model

We use the same user model in Studies 1 and 2. The user model in Study 3 includes additional functionality for a key press. The policy functions use a feature extractor network.

Layer	description	dimension	activation
H1	fully connected	32	ReLU
H2	fully connected	64	ReLU
H3	fully connected	128	ReLU
H4	fully connected	128	ReLU

The feature extractor is followed by the policy and value heads, which have identical architectures:

Layer	description	dimension	activation
H1	fully connected	128	ReLU
H2	fully connected	64	ReLU

SB3 produces Gaussian distributions for actions and value predictions, from where mean values are used as point estimates for actions and value predictions. We used the Proximal Policy Optimization (PPO) algorithm [45] to train the user model with Reinforcement Learning. The most important PPO parameters are given in the table below.

Parameter	value
Training iterations	3 000 000
learning rate	2e-4
gamma	0.99
clip range	0.18
entropy coefficient	0.001
maximum gradient norm	0.55

A.2 Analyst model

As described in the body text, the Analyst uses a relation network with a mean pooling MLP network. The Analyst has the following layers:

Layer	description	dimension	activation
H1	fully connected	32	ReLU
H2	fully connected	64	ReLU
H3	fully connected	128	ReLU
H4	fully connected	256	ReLU
H5	-	pooling layer	no activation
H6	fully connected	256	ReLU
H7	fully connected	64	ReLU

The feature extraction network is followed by identical policy and value networks with the following dimensions:

Layer	description	dimension	activation
H1	fully connected	64	ReLU
H2	fully connected	64	ReLU

As with the user model, SB3 produces Gaussian distributions for actions and value predictions, from where mean values are used as point estimates for actions and value predictions. Similarly, the Analyst is trained with the PPO algorithm. The most important PPO parameters are given in the table below.

Parameter	value
Training iterations	4 000 000
exponentially decaying learning rate	start 5e-5
gamma	0.99
clip range	0.10
entropy coefficient	0.01
maximum gradient norm	0.55

A.3 Non-myopic demonstration

The Non-myopic demonstration was trained using the Soft Actor Critic (SAC) algorithm with SB3 default hyperparameters. The Analyst network is simplified so that the relation network is replaced by a multilayer perceptron, therefore data point pairwise calculations (defined in equations (2) and (3)) are not used. The policy and Q networks have 2 layers, each with 128 hidden units and ReLU activations.

A.4 Adaptation demonstration

The Adaptation demonstration uses PPO with SB3 default settings. The EMA network (described in Section 3.3) is not used. The feature extractor network consists of 4 Layers followed by ReLU activations. The dimensions of the layers are 256 for layer H1, 256 for layer H2, 128 for layer H3 and 256 for layer H4. This is followed by a mean pooling layer H5 and a layer H6 of width 256 followed by ReLU activations. The policy and value heads are identical with two layers with a width of 256 units followed by ReLU activations.