Improving associative learning studies with Bayesian experimental design

Filip Melinscak[1−2] **(filip.melinscak@uzh.ch)** **Dominik R. Bach**[1−3] **(dominik.bach@uzh.ch)**

[1]Clinical Psychiatry Research, Department of Psychiatry, Psychotherapy, and Psychosomatics; Psychiatric Hospital; University of Zurich, 8032 Zurich, Switzerland

[2]Neuroscience Center Zurich; University of Zurich, 8057 Zurich, Switzerland

[3]Wellcome Trust Centre for Neuroimaging and Max Planck/UCL Centre for Computational Psychiatry and Ageing Research, University College London, London WC1 3BG, UK

Abstract

Associative learning theory has a rich tradition of formalized computational models. With increasingly complex and flexible models, it is becoming difficult to intuit experiments that can empirically distinguish between them. To address this issue, we propose a quantitative approach. Using the formalism of Bayesian experimental design, we tune experimental variables to maximize the utility of the experiment, i.e., to best discriminate computational models. We demonstrate the proposed method on two scenarios from the literature on models of classical conditioning. In both cases, optimized designs substantially outperform existing canonical designs in simulations: the odds of recovering the true model increase 15 times in one scenario, and 43 in the other. These results suggest that formally optimizing associative learning studies has potentially large benefits in terms of more accurate model selection.

Keywords: Bayesian experimental design, associative learning

Introduction

Recent concerns over reproducibility and statistical power in various scientific fields - including cognitive neuroscience (Szucs & Ioannidis, 2017) - have highlighted the importance of experimental design. To better address these issues, new tools for experimental design are being developed different domains, e.g., psychology (Cavagnaro, Myung, Pitt, & Kujala, 2010), neuroimaging (Daunizeau, Preuschoff, Friston, & Stephan, 2011), and neurophysiology (Lewi, Butera, & Paninski, 2009).

The study of associative learning is a discipline that has a rich tradition of theoretical and computational modeling but lacks formalized tools for experimental design. With today's increasingly complex and flexible models of associative learning, it is becoming difficult to manually devise experiments that would efficiently discriminate between them. Hence, to improve associative learning studies, methods of optimizing experimental designs for accurate model selection are necessary.

In this paper we present a formal computational approach to this problem, using simulation-based Bayesian experimental design. The simulation-based approach allows great flexibility in specifying the design problem: varied experiment structures, model spaces, study goals, and data analysis procedures can be accommodated. The Bayesian aspect of the method allows explicit and efficient use of prior knowledge to guide experimental design, and it facilitates updating of experimental designs upon observing new data.

To illustrate the proposed method, we take the example of classical conditioning as a special case of associative learning, and formalize the structure of experiments in a manner that facilitates optimization. In two sets of models drawn from the classical conditioning literature, we demonstrate that formally optimized experimental designs substantially outperform existing canonical designs in terms of model selection accuracy. These results suggest that our approach may greatly improve sensitivity and efficiency for model selection problems.

Methods

Simulation-based approach to design optimization

The simulation-based Bayesian experimental design framework (Wang & Gelfand, 2002) is outlined in fig. 1.A. In this framework, the user needs to specify following components of the optimization problem to make use of the proposed method. (a) Structure of the experiment, together with a set of tunable design variables and any constraints on these variables. (b) Set of candidate models: a set of generative models which will be used to simulate data and will subsequently be fitted to the data. (c) Sampling and fitting prior distributions over model parameters. The sampling prior is used to generate simulated data, and the fitting prior is used to fit models to simulated data (Brutti, De Santis, & Gubbiotti, 2014). The sampling prior reflects experimenters' prior knowledge, and can be as informative as desired, whereas the fitting prior should not be overly informative, so that it does not lead to foregone conclusions. (d) Utility function reflecting the goals of the experiment. (e) Analysis procedure which is to be applied to simulated datasets. (f) Optimization algorithm and its parameters.

Once the user has fully specified the problem, design optimization proceeds by repeating the following steps, which constitute a single iteration. (1) Current values of design variables are used together with candidate models and the sampling prior to generate simulated datasets. (2) All the models are fitted to simulated datasets using the specified analysis procedure and the fitting prior. (3) The utility function is calculated based on the results of the fitting procedure (by averaging across simulations), and the estimated utility is passed to the optimizer. (4) Unless the termination criterion is satisfied, the optimizer proposes new values of design variables, based on the utilities of previously evaluated designs.

If we wish to provide users with a flexible method of experimental design, the simulation-based approach is instrumental, but it can lead to difficult optimization problems. Since the util-

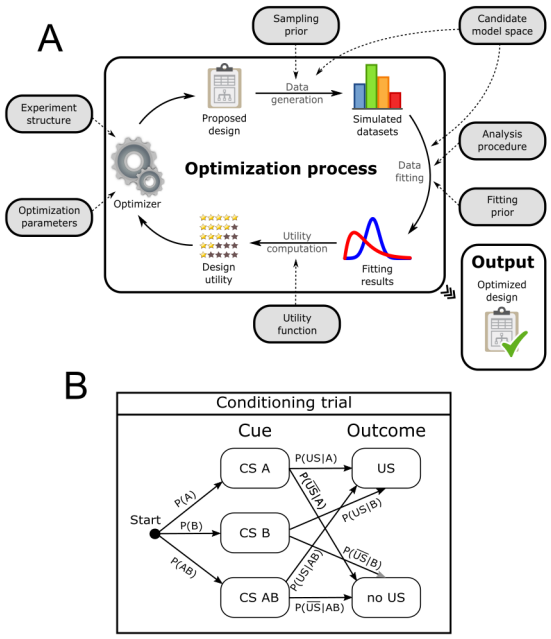

Figure 1: (A) Simulation-based Bayesian experimental design (see main text for details). Shaded elements denote user inputs. (B) Structure of a classical conditioning trial, with transition probabilities as design variables.

ity function is based on stochastic simulations, it is noisy and often computationally expensive to evaluate. Moreover, we usually cannot make strong assumptions about the utility function (like convexity) and we do not have access to its derivatives. Under these conditions, a state-of-the-art optimization algorithm is Bayesian optimization (Brochu, Cora, & de Freitas, 2010), and it was therefore our choice for optimizing designs. Bayesian optimization is a global optimization method which uses a surrogate model of the utility function to decide which values of optimized variables to evaluate next; this allows it to efficiently optimize over expensive utility functions, at the cost of continuously updating the surrogate model.

Optimizing the design of associative learning studies

To optimize associative learning studies we need to define the structure of such experiments. Here we considered classical conditioning, as an example of associative learning. One possible representation of conditioning experiments is shown in fig. 1.B. A conditioning trial consists of presenting a cue (conditioned stimulus, CS), followed by an outcome (unconditioned stimulus, US), which can be aversive or appetitive. The design of such an experiment can be specified in terms of cue probabilities ($P(\mathrm{CS})$), and conditional outcome probabilities ($P(\mathrm{US}|\mathrm{CS})$): these are the design variables which we seek to optimize. Since allowing the design variables of each trial to be optimized independently would result in an intractable, high-dimensional optimization problem, we divide the experiment into stages. Each stage is a block of trials in which the design variables are held constants. For our present purposes, we found that using two stages provided sufficient potential for design improvement.

Using a specified experimental design, we can stochastically generate a sequence of trials (cues and outcomes), which serve as an input to a computational model of associative learning. Models of associative learning are usually used to predict agents' responses to cues (i.e., conditioned responses, CRs). The simulated datasets are then obtained by running all the candidate models on the sequence of CSs and USs, and obtaining the predicted CRs (with added observation noise). All the candidate models are then fitted to each of the datasets. In results presented here, we used likelihood maximization (with the 'mfit' toolbox in Matlab) to fit the models, but fully Bayesian inference with proper priors could also be used. Based on the quality of fit, models are compared and the winning model is selected. Here we used the Bayesian Information Criterion (BIC) to compare models, as it takes into account both the model fit and model complexity. The model selection was performed on the basis of single simulations, corresponding to single-subject model comparison, which is generally more difficult than model selection with a group of subjects, due to the limited amount of data.

The design utility is calculated based on the results of model selection. Here we simply used model selection accuracy, i.e., the proportion of simulations in which the selected winning model was the model which actually generated the data. To search for a design that maximizes the utility, we used the aforementioned Bayesian optimization (via the 'bayesopt' function in Matlab). The obtained design is then validated and compared to other reference designs by evaluating the utility function in a larger number of independent simulations.

Results

The seminal Rescorla-Wagner (RW) model (Rescorla & Wagner, 1972) has been successful in explaining many aspects of associative learning, but there are many phenomena which it cannot account for. To address some of these shortcomings, various modified versions of the RW model have been developed. One example is the Kalman Rescorla-Wagner (KRW) model (Dayan & Kakade, 2001), which is a probabilistic version of the RW model. Another example is the hybrid Rescorla-Wagner-Pearce-Hall (RWPH) model (Li, Schiller, Schoenbaum, Phelps, & Daw, 2011), which allows for flexible learning rates for each cue (termed "associabilities"). In light of these developments, we demonstrate the proposed design method in two scenarios in which the RW model is compared against its variants.

In Scenario 1, based on the simulation study of Kruschke (2008), RW and KRW models are compared using the backward blocking experimental design. Therefore, we took the backward blocking design as the reference design, against which we compared the formally optimized design. Scenario 2 is based on the empirical study of Li et al. (2011), which compared the RW model (labeled RW(V)) and three variants of

the RWPH model. The RWPH models differed by the quantity which was emitted as the conditioned response: associative weights (RWPH(V)), associabilities (RWPH(α)), or a mixture of weights and associabilities RWPH($V + \alpha$). The study used a reversal learning to obtain differential responses between the models, and we used this experimental design as the reference design for comparing the optimized design.

In both scenarios, we ran the Bayesian design optimization for 285 iterations. In each iteration the utility function was evaluated using 32 datasets simulated from each model. Using 32 CPU cores at 2.5 GHz, total optimization times for Scenario 1 and 2 were 4.1 h and 17 h, respectively. In Scenario 1, sampling priors for RW and KRW models were point priors placed at the parameter values used by Kruschke (2008). Similarly, in Scenario 2, we used a point prior placed at the estimated parameter values provided by Li et al. (2011) for the RWPH($V + \alpha$) model, and for the other models we used a subset of these parameters since the models are nested. Finally, we evaluated the optimized designs and reference designs in additional 256 simulations per each model-design combination. We also computed the responses of the models fitted to these simulated datasets. For visualization purposes, the responses were computed for all cues at each trial (although only one cue is observed by the model during data generation), and the responses were averaged across simulations.

For Scenario 1, model selection accuracies obtained with the reference backward blocking design and the optimized design are presented in fig. 2.A The results show that the probability of selecting the correct model increased from 56.64% (95% CI: [52.22%, 60.98%]) for the reference design, to 98.24% (95% CI: [96.69%, 99.19%]) for the optimized design. This improvement is mostly due to simulations in which the KRW model is the true model: in the simulations with the reference design, both models can adequately fit the data, but the RW model has lower complexity, hence it is often erroneously selected as the winning model.

For Scenario 2, comparison of model selection accuracies obtained with the reference reversal learning design and the optimized design is shown in fig. 2.D. Model selection accuracy improved from 59.38% (95% CI: [56.29%, 62.40%]) for the reference design, to 95.51% (95% CI: [94.05%, 96.69%]) for the optimized design. Largest contribution to the accuracy improvement stems from simulations in which RWPH($V + \alpha$) is the true model: using the reference design this model is often confused with the simpler RW(V) model, which is not the case with the optimized design (see inset of fig. 2.D).

Comparison of the reference and the optimized designs in both scenarios shows that the optimized design is less parsimonious (fig. 2.B and fig. 2.E), but it generally results in more distinct model responses (fig. 2.C and fig. 2.F). Exception to this are responses in Scenario 1, with RW model as the true model. In this case, even with the optimized design, the responses remain similar. However, the similarity of responses here is not problematic, because the RW model will be correctly selected due to its lower complexity, when compared with KRW. This shows that the optimization took into account not only the model space and the priors, but also the specifics of the applied analysis procedure.

Discussion

We propose a novel method for improving studies of associative learning using simulations to evaluate experimental designs, and Bayesian optimization to tune them. In two scenarios drawn from the classical conditioning literature (Kruschke, 2008; Li et al., 2011), we demonstrate the effectiveness of the proposed method. In optimizing the designs, we explicitly take into account the prior knowledge about the problem, the goals of the study, and the analysis procedures, allowing us to obtain statistically more powerful designs. In the two presented scenarios, we chose single-subject model selection accuracy as the goal, and the optimized designs outperformed manual designs by a large margin: the odds of correctly identifying the true model increased 42.8 times in one scenario (RW vs. KRW) and 14.5 times in the other (RW vs. RWPH).

An important caveat to these results is that they indicate the potentially obtainable improvements. In present simulations the sampling prior is a point prior on model parameters, and the same values are used both in design optimization and evaluation. Therefore, the accuracy results obtained here should be understood as upper bounds, because in real use-cases the data-generating process is unlikely to have parameter values that precisely coincide with the ones used when planning the experiment. Nevertheless, the presented results suggest that there are large potential benefits in optimizing experimental designs. These benefits can be realized by iterating between experimentation, updating sampling priors, and optimizing designs, between subsamples of subjects, between individual subjects, or even within subjects.

There are also many future venues for further validating the proposed method and extending its scope. Since in practice we usually do not have point priors on the parameters of candidate models, it will be useful to investigate how does the performance of optimized designs change when prior knowledge is vague or deviates strongly from actual parameter values. It will also be important to extend the method to other potential experimental goals, such as model selection with hierarchical data (e.g., groups of subjects), or the goal of accurately estimating parameters in a single model, rather than selecting between multiple models. In this paper we have used a simple experimental structure with a small number of tunable parameters, but developing more complex experimental structures and still obtaining tractable optimization problems will be a crucial challenge. For such complex experiments it might be advantageous to use adaptive design optimization (Cavagnaro et al., 2010). Here the algorithm would not optimize the whole experiment at once, but would rather optimize the design variables of each trial on the fly, which results in a smaller number of design variables to be optimized.

Computational cognitive neuroscience has a rich modeling toolbox at its disposal, but a comparable set of methods for

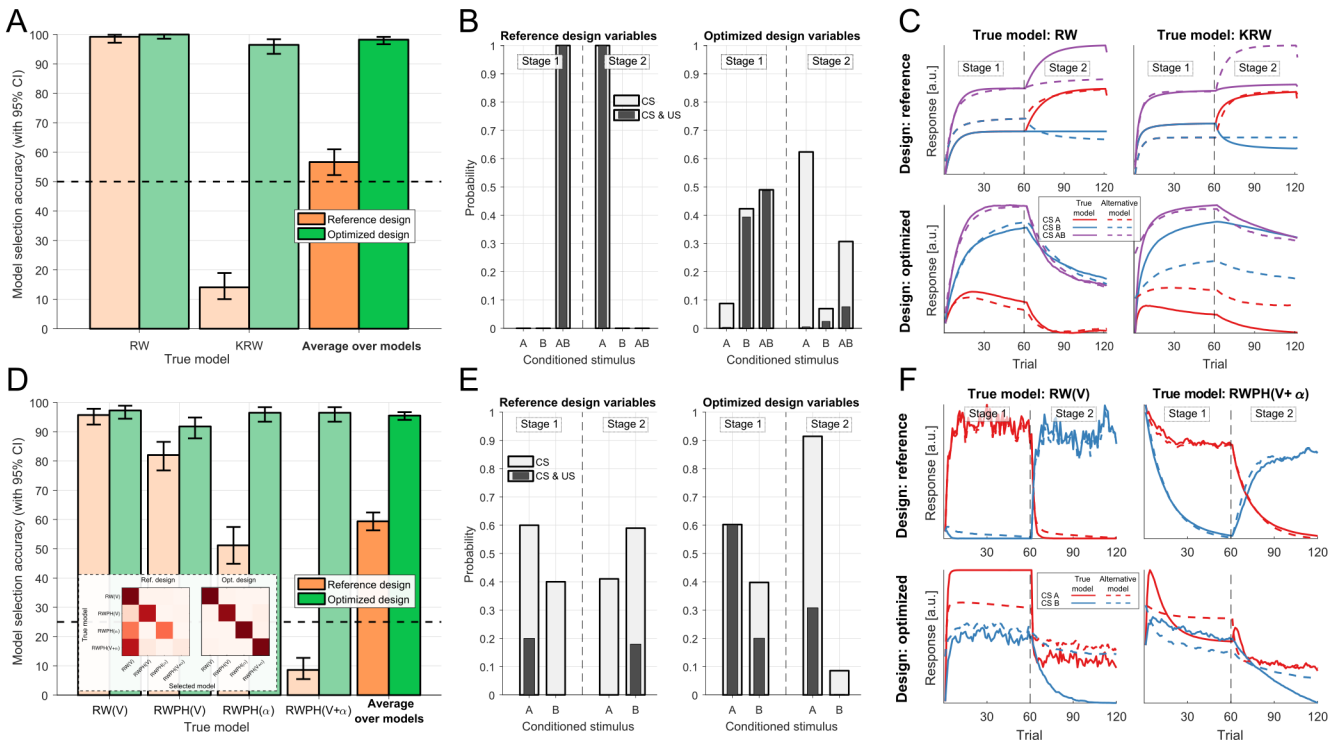

Figure 2: Comparison of reference and optimized designs. (A,D) Model selection accuracy, (B,E) design variables' values, and (C,F) model responses for Scenarios 1 and 2, respectively. Inset in (D) shows the confusion matrix between models.

experimental design is lacking. We believe that the formal design optimization approach presented here is a valuable addition to this toolbox. The outlined method has potential to improve not only basic research on associative learning, but a similar approach could have substantial translational value in the nascent field of computational psychiatry where increasing sample size to increase statistical power is often not an option. In this context, formal optimization could be used to develop accurate and efficient behavioral assays which would link computational models and their parameters to clinically relevant variables.

References

Brochu, E., Cora, V. M., & de Freitas, N. (2010). A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. *arXiv preprint arXiv:1012.2599*.

Brutti, P., De Santis, F., & Gubbiotti, S. (2014). Bayesian-frequentist sample size determination: a game of two priors. *METRON*, *72*(2), 133–151.

Cavagnaro, D. R., Myung, J. I., Pitt, M. A., & Kujala, J. V. (2010). Adaptive Design Optimization: A Mutual Information-Based Approach to Model Discrimination in Cognitive Science. *Neural Computation*, *22*(4), 887–905.

Daunizeau, J., Preuschoff, K., Friston, K., & Stephan, K. (2011). Optimizing Experimental Design for Comparing Models of Brain Function. *PLoS Computational Biology*, *7*(11), e1002280.

Dayan, P., & Kakade, S. (2001). Explaining away in weight space. In *Advances in neural information processing systems 14* (pp. 451–457).

Kruschke, J. K. (2008). Bayesian approaches to associative learning: From passive to active learning. *Learning and Behavior*, *36*(3), 210–226.

Lewi, J., Butera, R., & Paninski, L. (2009). Sequential Optimal Design of Neurophysiology Experiments. *Neural Computation*, *21*(3), 619–687.

Li, J., Schiller, D., Schoenbaum, G., Phelps, E. A., & Daw, N. D. (2011). Differential roles of human striatum and amygdala in associative learning. *Nature Neuroscience*, *14*(10), 1250–1252.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical Conditioning II Current Research and Theory*, *21*(6), 64–99.

Szucs, D., & Ioannidis, J. P. A. (2017). Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLOS Biology*, *15*(3), e2000797.

Wang, F., & Gelfand, A. E. (2002). A simulation-based approach to Bayesian sample size determination for performance under a given model and for separating models. *Statistical Science*, *17*(2), 193–208.