

A new method to dynamically choose questions and minimize costs in behavioral experiments

Abstract

Adaptive design optimization algorithms use Bayesian statistics and machine learning to identify the experimental designs (e.g., stimulus properties and combinations, testing schedule, rewards, etc.) that maximize the expected information gain, leading to rapid accumulation of information with the fewest number of trials. To the best of our knowledge, however, none of the existing algorithms are cost-aware, in the sense that none of them consider the possibility that different designs and results may have different associated costs. Here we introduce a new cost-aware algorithm to dynamically choose questions and minimize costs in behavioral experiments. We also present ground true simulations showing the drastic reduction in cost that can be achieved when studying hackers' behavior. Importantly, in all cases the decision about what design to use took milliseconds or a few seconds in a conventional desktop computer, which makes the algorithm suitable for on the fly design optimization even when running online massive experiments.

Introduction

Adaptive Design Optimization (ADO) is a Bayesian optimal design procedure that ensures optimal efficiency in data collection and in the integrity of inference about the phenomenon of interest [Cavagnaro, 2009; Wang, 2010; Myung, 2013]. It takes into account previous results to update prior beliefs and estimate the design parameters (e.g., stimulus properties and combinations, testing schedule, rewards, etc.) that maximize the expected information gain in the next trials in order to achieve the experimental objective with the fewest numbers of observations possible.

Different versions have been developed and successfully applied since the early 1990s, for instance, in visual psychophysics [Gu, 2016; Hou, 2016], or in a classical behavioral economics experiment, the delayed discount task [Ahn, 2020] (see [Ryan, 2016] for a review on Bayesian design optimization), but its adoption has been slow, in part because of the nontrivial technical skills required to implement them [Yang, 2020] and, in part, because of its limitations since the existing algorithms are time-consuming and are only practical if trying to infer among a limited set of parametric models [Myung, 2013; Chang, 2021].

With the explosion of massive online experiments [Hartshorne, 2019; Almaatouq, 2021; Zimmerman, 2016; Rieznik, 2017] it is foreseeable that the implementation of ADO for on the fly design optimization will exponentially increase. Since many of these online experiments measure complex behaviors, they lack a proper expected parametrized model and an ADO approach would only be practical if it is data-driven and nonparametric, i.e., if it is able to optimally select designs independently of the functional form of the data-

generating model. To the best of our knowledge there is only one article proposing a nonparametric ADO algorithm: a recent article by *Chang et al* [Chang, 2021]. Also, at the BitTrap Attacker Behavior Analysis (ABA) Labs we already developed and implemented our own new nonparametric ADO algorithm, and wrote an article introducing it that is currently under revision in a scientific journal.

With this advances, ADO implementations will soon be ubiquitous. And one question will naturally arise, a question that, to the best of our knowledge, has not yet been answered in the scientific literature: how can we take into account that not all measures have the same cost? When different measures have different costs, what should be optimized at each trial is no longer the information gain but the information gain *per unit cost*. In this article we present an algorithm to do exactly that, and that can be easily applied using current ADO implementations with few modifications to the algorithms.

A cost-aware ADO framework

The goal of any ADO implementation is to discovery, with the fewest possible trials, which of a set of M models is the one that best represent the reality, i.e., the one that is the true data-generating model. Suppose $m = \{1, 2, \dots, M\}$ is one of a set of M models being considered, d is a design, and y is the outcome of an experiment with design d under model m . Traditional ADO formulations select optimal designs by maximizing the expected information gain [Cavagnaro, 2009; Wang, 2010; Myung, 2013]. If a measure is done at d with result y , the information gain is given according to the Kullback–Leibler divergence by

$$\sum_{m=1}^N p(m|y, d) \log_2 \left(\frac{p(m|y, d)}{p(m)} \right) \quad (1)$$

Note that $p(m|y, d)$ is the posterior probability of model m , while $p(m)$ is its prior. So, the *expected* information gain is given by the sum of the information gains that would be obtained if the result of a measure y is weighed by the marginal probability for y :

$$I(d) = \sum_{y=0}^1 p(y|d) \sum_{m=1}^M p(m|y, d) \log_2 \left(\frac{p(m|y, d)}{p(m)} \right) \quad (2)$$

The design maximizing this equation is selected in traditional ADO formulations. Now, suppose each result y of a measure has an associated cost. Let's call this cost $C_{y,d}$. The subscripts y and d indicate that the cost may depend on both the trial result and design. Now, we can add this information in the first sum of equation (2) in order to maximize the expected information gain per unit cost:

$$I_{CostAware}(d) = \sum_{y=0}^1 \frac{p(y|d)}{C_{y,d}} \sum_{m=1}^M p(m|y,d) \log_2 \left(\frac{p(m|y,d)}{p(m)} \right) \quad (3)$$

The design maximizing this new equation will then maximize the expected information gain per unit cost. In the next section we present ground true simulations showing the advantage of using this algorithm when compared against a classical ADO algorithm or a random strategy: the same amount of information can be obtained but at a much lower cost.

Results

Suppose we want to know the probability that successful a hacker [succesfully] compromising a computer will accept a given amount of money in order to end the attack and reveal their position. In other words, we want to know the shape of the curve P (acceptance probability) vs. U\$\$ (offered money). In order to measure this curve we can offer different money amounts when attacks occur, and measure the proportion of times the offer is accepted. Suppose that, given our knowledge about the word, ours believes about the possible shapes of this curve are represented by the head map shown in Figure 1a. Suppose now that the shape of the true curve is given by the dotted blue curve in Figure 1b, which shows that, for instance, over 50% of the hackers will accept USD250 to reveal their position and nearly 100% will do it for USD2000. The question we ask is how much money we need to spend in order to modified our believes in such a way that after some measures (i.e., after some money offering and results of acceptance or rejection) our posterior believes are represented by the head map in Figure 1b, and so are much closer to the true generating model (blue points) than our priors (Figure 1a). The heat map in Figure 1b is obtained after an information gain of 50% relative to the initial uncertainty (measured as entropy through the well-known Shannon formula).




Figure 1: (a) Heat map representing our prior believe about the probability of acceptance as a function of the offered money. (b) Heat map and (background) and true generating data curve (dotted blue) when the uncertainty is halved, i.e., when the information gained equals 50% of the initial uncertainty.

In order to study this question we perform ground true simulations: we simulate a population of hacker that behave according to the ground true curve shown in blue in Figure 1b and observe how many trials would be needed to achieve convergence towards the heat map shown in Figure 1b if our prior beliefs are given by the heat map shown in

Figure 1a. When running the experiment we study three strategies: choosing randomly among different money amounts; choosing the amounts that optimizes the expected information gain using equation (2), where d , the experimental design to be chosen, is the amount of offered money; and choosing the amounts that optimizes the expected information gain *per unit cost* using equation (3). After each trial, we update our beliefs (represented as heat maps as done in Figure 1) using the well-known Bayes formula, and chose a new design (amount) for the next trial. How much it cost in each case to obtain an information gain of 50% relative to the initial uncertainty? In Figure 2 we plot the results of the simulations.

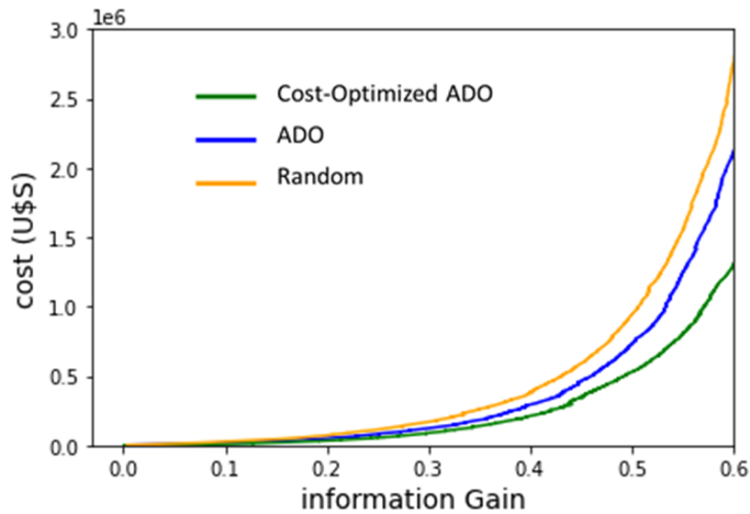


Figure 2 Information gain vs. cost using, using the prior shown in Figure 1a and the ground true shown in blue in Figure 1b, for three different strategies: a conventional ADO strategy, a cost-aware ADO, and choosing randomly among the different offered amounts.

It is clear that the cost-aware strategy converges much faster towards the true generating model, at a much faster rate per unit cost. We stress that in all cases the decision about what design to use, i.e., how much money to offer in the next trial, took milliseconds or a few seconds in a conventional desktop computer, which makes this algorithm suitable for on the fly design optimization even when running online massive experiments.

Conclusion

We introduced a new algorithm to dynamically choose questions and minimize costs in behavioral experiments. We performed ground true simulations studying a specific case in which different designs and results have different costs, showing that the implementation of our algorithm can drastically reduce the cost of the experiment. With the explosion of massive online experiments, we forecast a rapid growing adoption of cost-aware ADO strategies in behavioral science.

References

- Ahn, W. Y., Gu, H., Shen, Y., Haines, N., Hahn, H. A., Teater, J. E., ... & Pitt, M. A. (2020). Rapid, precise, and reliable measurement of delay discounting using a Bayesian learning algorithm. *Scientific reports*, 10(1), 1-10.
- Almaatouq, A., Becker, J., Houghton, J. P., Paton, N., Watts, D. J., & Whiting, M. E. (2021). Empirica: a virtual lab for high-throughput macro-level experiments. *Behavior Research Methods*, 1-14.
- Cavagnaro, D. R., Myung, J. I., Pitt, M. A., & Kujala, J. V. (2010). Adaptive design optimization: Amutual information based approach to model discrimination in cognitive science. *Neural Computation*, 22(4), 887–905.
- Chang, J., Kim, J., Zhang, B. T., Pitt, M. A., & Myung, J. I. (2021). Data-driven experimental design and model development using Gaussian process with active learning. *Cognitive Psychology*, 125, 101360.
- Gu, H., Kim, W., Hou, F., Lesmes, L. A., Pitt, M. A., Lu, Z. L., & Myung, J. I. (2016). A hierarchical Bayesian approach to adaptive vision testing: A case study with the contrast sensitivity function. *Journal of Vision*, 16(6), 15-15.
- Hartshorne, J. K., de Leeuw, J. R., Goodman, N. D., Jennings, M., & O'Donnell, T. J. (2019). A thousand studies for the price of one: Accelerating psychological science with Pushkin. *Behavior research methods*, 51(4), 1782-1803.
- Hou, F., Lesmes, L. A., Kim, W., Gu, H., Pitt, M. A., Myung, J. I., & Lu, Z. L. (2016). Evaluating the performance of the quick CSF method in detecting contrast sensitivity function changes. *Journal of vision*, 16(6), 18-18.
- Myung, J. I., Cavagnaro, D. R., & Pitt, M. A. (2013). A tutorial on adaptive design optimization. *Journal of Mathematical Psychology*, 57, 53–67.
- Rieznik, A., Moscovich, L., Frieiro, A., Figini, J., Catalano, R., Garrido, J. M., ... & Gonzalez, P. A. (2017). A massive experiment on choice blindness in political decisions: Confidence, confabulation, and unconscious detection of self-deception. *PloS one*, 12(2), e0171108.
- Ryan, E. G., Drovandi, C. C., McGree, J. M., & Pettitt, A. N. (2016). A review of modern computational algorithms for Bayesian optimal design. *International Statistical Review*, 84(1), 128-154.

Wang, S. W., Filiba, M., & Camerer, C. F. (2010). Dynamically optimized sequential experimentation (dose) for estimating economic preference parameters. Manuscript submitted for publication.

Yang, J., Pitt, M. A., Ahn, W. Y., & Myung, J. I. (2021). ADOpy: a python package for adaptive design optimization. *Behavior Research Methods*, 53(2), 874-897.

Zimmerman, F., Shalom, D., Gonzalez, P. A., Garrido, J. M., Alvarez Heduan, F., Dehaene, S., ... & Rieznik, A. (2016). Arithmetic on your phone: A large scale investigation of simple additions and multiplications. *Plos one*, 11(12), e0168431.