



AFRL-AFOSR-VA-TR-2022-0044

Robust Autonomous Adaptive Experimentation

Pitt, Mark
OHIO STATE UNIVERSITY
1960 KENNY RD
COUMBUS, OH,
US

09/13/2021
Final Technical Report

DISTRIBUTION A: Distribution approved for public release.
--

Air Force Research Laboratory
Air Force Office of Scientific Research
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 13-09-2021		2. REPORT TYPE Final		3. DATES COVERED (From - To) 15 Apr 2016 - 14 Apr 2021	
4. TITLE AND SUBTITLE Robust Autonomous Adaptive Experimentation				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER FA9550-16-1-0053	
				5c. PROGRAM ELEMENT NUMBER 61102F	
6. AUTHOR(S) Mark Pitt				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) OHIO STATE UNIVERSITY 1960 KENNY RD COUMBUS, OH US				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AF Office of Scientific Research 875 N. Randolph St. Room 3112 Arlington, VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR RTA2	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-AFOSR-VA-TR-2022-0044	
12. DISTRIBUTION/AVAILABILITY STATEMENT A Distribution Unlimited: PB Public Release					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Four Specific Aims were proposed to develop, implement, and empirically validate Bayesian learning algorithms for autonomous adaptive experimentation in cognitive science and materials science. Aims 1 and 2, which focused on algorithm development, were achieved by implementing a robust autonomous adaptive system (RAAS, i.e., an experimentation framework) using three algorithms: (1) adaptive design optimization (ADO, a model-based algorithm, Aim 1), (2) Bayesian optimization (BO, model-free algorithm, Aim 2), and (3) Gaussian Process Active Learning (GPAL), a second model-free approach to optimal experimental design that is solely data-driven. Aim 3 tested ADO and GPAL in the fields of decision making and numerical estimation. Aim 4 tested the use of BO to improve the growth of carbon nanotubes and improve the precision of 3D printing (Aim 4 was carried out in collaboration with Dr. Benji Maruyama of the Materials and Manufacturing Directorate at AFRL). In all application domains, we have successfully demonstrated the robustness and efficiency of these algorithms in achieving the research objectives. This work advances the current state of the art in autonomous research in the cognitive and materials sciences.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			HAL GREENWALD
U	U	U	UU	8	19b. TELEPHONE NUMBER (Include area code) 703-588-8441

AFOSR Final Performance Report

Project Title:

Robust Autonomous Adaptive Experimentation

Award Number:

FA9550-16-1-0053

Project Period:

15 April 2016 – 14 April 2021

Program Officer:

Dr. Hal S. Greenwald
Computational Cognition & Machine Intelligence Program
Air Force Office of Scientific Research (AFOSR/RTA)
875 N. Randolph St.
Arlington, VA 22203-1768
(703) 588-8441

Principal Investigator:

Prof. Mark A. Pitt
Department of Psychology
The Ohio State University
1835 Neil Avenue
Columbus, OH 43210
Pitt.2@osu.edu

Report Date:

14 July 2021

Robust Autonomous Adaptive Experimentation

Abstract

Four Specific Aims were proposed to develop, implement, and empirically validate Bayesian learning algorithms for autonomous adaptive experimentation in cognitive science and materials science. Aims 1 and 2, which focused on algorithm development, were achieved by implementing a robust autonomous adaptive system (RAAS, i.e., an experimentation framework) using three algorithms: (1) adaptive design optimization (ADO, a model-based algorithm, Aim 1), (2) Bayesian optimization (BO, model-free algorithm, Aim 2), and (3) Gaussian Process Active Learning (GPAL), a second model-free approach to optimal experimental design that is solely data-driven. Aim 3 tested ADO and GPAL in the fields of decision making and numerical estimation. Aim 4 tested the use of BO to improve the growth of carbon nanotubes and improve the precision of 3D printing (Aim 4 was carried out in collaboration with Dr. Benji Maruyama of the Materials and Manufacturing Directorate at AFRL). In all application domains, we have successfully demonstrated the robustness and efficiency of these algorithms in achieving the research objectives. This work advances the current state of the art in autonomous research in the cognitive and materials sciences.

1 Introduction

Advances in science depend on researchers being able to make inferences about their experimental results with confidence. A constant challenge in achieving this objective (i.e., good inference) is that experiments are difficult to design because the consequences of design decisions are not known in advance of data collection. Computational methods such as optimal experimental design in statistics and active learning in machine learning can assist in improving scientific inference. Algorithms integrate information about the experimental design with the data collected thus far (e.g., measurements, such as sensor readings) to forecast combinations of design parameters (e.g., physical inputs) that are likely to be most informative about the experiment's objective (e.g., maximize output or classification accuracy). When applied repeatedly across data collection episodes, their adaptive nature makes it possible to obtain highly informative data in a short amount of time.

The goal during this funding period was to develop, implement, and empirically validate methods of optimal experimental design (OED) in two scientific fields: materials science and cognitive science. As part of the grant effort, we have collaborated with Dr. Benji Maruyama's team at the Air Force Research Laboratory (AFRL) at Wright-Patterson Air Force Base (Dayton, OH), to apply OED algorithms to materials science experiments in two domains: carbon nanotubes synthesis and additive manufacturing (3D print). In addition, in our own lab at Ohio State, we applied essentially the algorithms in behavioral experiments in two domains of cognitive science: decision making and numerical cognition.




Figure 1: Schematic illustration of the robust autonomous adaptive system (RAAS).

A primary achievement during the grant period was the development and application of a general-purpose, data-driven approach to adaptive experimentation, dubbed the robust autonomous adaptive system (RAAS). RAAS is in essence a nonparametric Bayesian method that simultaneously learns the underlying model that generated the data (model learning) while optimizing the experimental design (design optimization) to learn the model efficiently, as illustrated in Figure 1. It is a nonparametric (thus model-free) extension of its predecessor, adaptive design optimization (ADO) that we developed previously in our lab. In contrast to RAAS, ADO a model-based parametric adaptive experimentation method. Both RAAS and ADO are closed-loop experimentation systems that we applied and to research questions in materials science and cognitive science.

In what follows, we summarize major accomplishments of the research effort.

2 Accomplishments

2.1 Materials Science Applications

We repurposed RAAS for the problem of optimizing materials science experiments by implementing Bayesian optimization (BO) as its core algorithm, so dubbed the Bayesian-optimized robust autonomous adaptive system (BORAAS). Bayesian optimization is a machine-learning framework for globally optimizing a black-box objective function that is expensive to evaluate. We applied and empirically validated the BORAAS framework in two areas of materials science research, the main results of which are summarized in sections 2.1.1 and 2.1.2.

2.1.1 Optimizing Autonomous Carbon Nanotubes Synthesis

We applied BORAAS to the problem of maximizing the growth rate of carbon nanotubes (CNTs). The closed-loop experiment setup is shown in Figure 2. BORAAS explored the space of four input (design) parameters (Ethylene, Hydrogen, water, temperature) to identify the combination that yielded the highest CNT growth rate, as measured by the G band (wavelength) in the Raman spectrum.




Figure 2: Schematic illustration of the carbon nanotube synthesis autonomous experiment system (CNT-ARES). On the left is the experiment chamber, which includes a laser whose output is recorded by a Raman spectrometer. CNT growth rates, as measured by the G band, are fed into BO on the right. The algorithm then generates a new set of input parameters to run in the next CNT-ARES experiment.

The algorithm was embedded in the CNT-ARES system developed by Dr. Maruyama's team at AFRL, and two experimental campaigns were run after piloting, the results of which are in Figure 3 (Chang, Nikolaev, et al., 2020). In the figure, green circles denote seed trial that were used to initialize BORAAS. In the first campaign (bottom graphs) they were generated by the experimenter, and in the second (top graphs) they were generated randomly. The results are similar in both cases. Across approximately 100 closed-loop iterations (known as a “campaign”), the algorithm improved growth rate by up to a factor of 8. The middle graphs show a 13-trial moving average (red line) along with the growth rate predicted by BORAAS (black line). The convergence of these functions well before the end of the experiment shows the algorithm quickly learned the underlying relationship between the four inputs.




Figure 3. Results of two BORAAS-based campaigns using four print parameters carried out on CNT-ARES. (a) The raw growth rate of seed and planned experiments for the two campaigns, BORAAS-1 (bottom panel) and BORAAS-2 (top panel), increased as BORAAS optimized the objective function v . The inset in (a) shows example growth curves obtained by ARES from a seed and planned experiment. (b) Central moving average of v ($v_{c.m.a.}$, calculated using the experimental data in panel (a) with a sample window size of 13 data points) and predicted growth rates ($v_{pred.}$, provided by BO). (c) Normalized difference (Δ) between the central moving average and predicted growth rate for the two campaigns. BORAAS improved the growth rate after about ~ 100 experiments regardless of how the seed was generated

2.1.2 Optimizing Autonomous Additive Manufacturing (3D Printing)

The CNT-ARES set-up was retooled and extended to an additive manufacturing (AM; 3D printing) project, also directed by Dr. Maruyama at AFRL. This AM-ARES system (Figure 4) provided another opportunity in which to apply BORAAS. BORAAS was integrated into the AM-ARES set-up in which the objective was to optimize printing a cylinder whose shape filled the green-outlined area in one stroke (right side of Figure 5).




Figure 4. A simplified flowchart is provided as an overview to the prototype Additive Manufacturing Autonomous Research System (AM-ARES) closed-loop autonomous printing process.

The 3D printer itself consists of a vertical pen whose movement and extrusion of caulk is machine-controlled by 25+ parameters. Four key design parameters were selected and varied in the experiment: x coordinate, y coordinate, print speed, and prime delay. The printing challenge is to avoid underfilling or overfilling the target area, which is especially tricky when printing starts. After each print iteration, imaging software outputs an estimate of the proportion of the target area filled by caulk (objective score), which was then fed as input to BORAAS for the next iteration of the experiment campaign.

The results of a representative campaign are shown on the left in Figure 5 (Deneault, Chang, et al., 2021). Blue dots show BORAAS performance. A near optimal score of 0.94 was obtained in 100 iterations, with the region containing optimal parameter values being reached in 90. Optimal performance can also be seen in the parameter values (other colored dots) plateauing simultaneously toward the end of the campaign. The results were replicated in two additional 100-iteration campaigns (not shown). In contrast, campaigns using gradient descent achieved an objective score of only 0.5.




Figure 5: Results of a BORAAS-based, AM-ARES campaign using four print parameters. After running a campaign of 100 experiments, BORAAS converged to an objective score of 0.94.

A unique aspect of this project was that BORAAS did not reside on a computer adjacent to AM-ARES, but instead was on a cloud server. AM-ARES is housed in Dr. Maruyama's satellite lab

outside AFRL, making remote access possible. We developed a protocol for remote closed-loop experiments.

SUMMARY of Materials Science Applications: During the funding period, we established a productive collaboration between OSU and AFRL that advanced the priorities of the Air Force in materials science. We designed machine-learning algorithms (BORAAS) to optimize autonomous experimentation in ARES, and then validated their performance in two research domains. In both CNT synthesis and additive manufacturing, the BORAAS-guided ARES experiments achieved the twin goals of achieving the desired objective (maximizing growth rate, precise printing) while being efficient (short campaigns).

2.2 Cognitive Science Applications

RAAS was adapted to problems in cognitive science that would benefit from the application of optimal experimental design. The core algorithm used was a Gaussian Process (GP), which is a nonparametric Bayesian method for inferring the functional form of an unknown system in a model-free manner (i.e., without making parametric modeling assumption about the functional form of the underlying cognitive process). We combined GP with active learning in machine learning into a novel optimal experimental design (OED) framework, which we dubbed Gaussian Process Active Learning (GPAL). We applied and empirically validated GPAL in two subdomains of cognitive science, as summarized in sections 2.2.1 and 2.2.2.

2.2.1 Optimally Designing Decision Making Experiments

Delay discounting (DD), a common cognitive task in decision making and cognitive neuroscience research, is a preferential choice task that is often employed to measure an individual's ability to delay gratification by quantifying the preference of a sooner-smaller reward against a later-larger reward (e.g., "Do you prefer \$10 today or \$40 dollars in two weeks?"). Task performance has been linked to various clinical disorders such as impulsivity, gambling, addiction, and ADHD.




Figure 6: Two-dimensional plots of the probability of choosing the later-larger option over the smaller-sooner option. The plot on the left are data from a rational decision maker. Those on the right are from a person who is inconsistent when the options are extreme.

We applied GPAL to infer the latent discounting function by optimizing the task trials over two design parameters (i.e., reward amount and time delay). The results indicated that GPAL was able to quickly map an individual's 2-dim discounting function, as shown in Figure 6 (Chang, Kim, et al., 2019 & 2021). The figure includes two observed examples of the discounting function, one normal pattern on the left and the other abnormal pattern on the right. Each is a plot of the probability of choosing the later-larger reward over the smaller-sooner reward. The plot on the left is an example of a monotonic discounting function that is indicative of a rational decision maker. The plot on the right is a nonmonotonic function from an individual that violates one of the fundamental axioms of rational choice theory. In short, GPAL's flexibility and sensitivity to unique response behavior makes it well suited for measuring individual differences in decision making and behavioral tasks more generally. As such, we believe that GPAL is a promising tool for developing unbiased models of cognition.

2.2.2 Optimally Designing Numerical Cognition Experiments

We applied the GPAL framework to map the mental representation of number using a number-line estimation task conducted with children and adults. This work was done in collaboration with Prof. John Opfer's developmental psychology lab at Ohio State.

Here, we are interested in the form of the psychophysical function that links numbers to the subjective estimation of numerical magnitudes, in particular, among children. It is well established in the field that children tend to have logarithmically compressed numerical estimates. For example, they think that \$1000, \$100, and \$10 are more or less equally spaced in their mental world. They therefore tend to believe that receiving a \$90 raise from a \$10 weekly allowance is equally attractive as getting a \$900 raise from a \$100 weekly allowance. This tendency persists until children reach the 3rd and 4th grade.

A thorough comparison of the mental representation of number requires participants to make numerical judgments over a wide range of number (e.g., 0-50, 0-400). Such experiments require many trials and are impractical with young children, whose attention span is short (20 minutes). GPAL's efficiency and sensitivity to individual differences made it possible to carry out such an experiment. More concretely, all our experiments with kids lasted just 15 mins or less, instead of 1 or 2 hours that would be needed without GPAL.




Figure 7: The logarithmicity measure of GP estimates for the three age groups against upper bounds.

The results obtained indicated that GPAL quickly characterized an individual’s mental representation of numbers, specifically, the extent to which the representation is logarithmic or linear as shown in Figure 7 (Lee et al., 2021). The graph shows that logarithmicity decreases with age, although even adults are not completely linear when large numbers (> 100) are involved.

SUMMARY of Cognitive Science Applications: We have demonstrated the promise of GPAL, a data-driven cognitive modeling framework, and have empirically validated its robustness and efficiency in two domains of cognitive science, decision making and numerical cognition. In both fields, GPAL efficiently selects stimuli (e.g., pairs of monetary choices) to ensure that over a comparatively short number of trials, a participant’s response profile is precisely captured. GPAL’s data-driven orientation brings out differences among individuals, which can lead to greater understanding of behavior and more precise (unbiased) modeling of behavior.