

---

# Offline Contextual Bayesian Optimization for Nuclear Fusion

---

Youngseog Chung<sup>1\*</sup>, Ian Char<sup>1\*</sup>, Willie Neiswanger<sup>1</sup>, Kirthivasan Kandasamy<sup>2</sup>, Andrew Oakleigh Nelson<sup>3</sup>, Mark D Boyer<sup>3</sup>, Egemen Kolemen<sup>3</sup>, and Jeff Schneider<sup>1</sup>

<sup>1</sup>Department of Machine Learning, Carnegie Mellon University  
{*ichar, youngsec, willie, schneide*}@cs.cmu.edu

<sup>2</sup>Department of EECS, University of California Berkeley  
*kandasamy@eecs.berkeley.edu*

<sup>3</sup>Princeton Plasma Physics Laboratory  
{*anelson, mboyer, ekolemen*}@pppl.gov

\*These authors contributed equally to this paper.

## 1 Introduction

Nuclear fusion is regarded as the energy of the future since it presents the possibility of unlimited clean energy [Chen, 2011]. The most widespread method of realizing fusion reactions requires heating up isotopes of hydrogen to temperatures of hundreds of millions of degrees using a magnetic device called a *tokamak*. In this state, the nuclei of two nearby atoms may overcome electrostatic repulsion force between them to form a single nucleus, releasing energy. In contrast to nuclear fission, this process creates much less hazardous byproducts and has a plentiful fuel source [Chen, 2011, Clery, 2014]. One obstacle in utilizing fusion as a feasible energy source, however, is the stability of the reaction. At such high temperatures, the hydrogen atoms are in a plasma state, and any instability in the plasma can give rise to events called *disruptions*. If disruptions occur, the plasma is quickly lost and the nuclear reaction is halted. Ideally, one would have a controller for the tokamak that makes actions in response to the current state of the plasma in order to prolong the reaction as long as possible; however, such a controller is unknown as of now.

In this work, we make preliminary steps to learning such a controller. Since learning on a real world tokamak is infeasible, we tackle this problem by attempting to learn optimal controls offline via a simulator. In our experiments, we make the simplifying assumption that the plasma will only ever be in one of eight states. Finding the best action upon seeing one of the eight states (or contexts) bears similarities to the contextual bandit setting bandits [Krause and Ong, 2011, Agrawal and Goyal, 2013, Auer, 2002]; however, literature usually assumes that the context is decided by nature and cannot be set by the algorithm. Intuitively speaking, being able to decide how many queries are made in each context is advantageous because some contexts may be harder to learn the optimal action for or may occur more frequently in nature. Luckily, our tokamak simulator allows for contexts to be set, and since each simulation is expensive, it is crucial that we use an algorithm that prudently picks both actions and contexts to evaluate.

In this paper we introduce an algorithm that recommends a context and action pair to evaluate at every iteration. Our algorithm falls under the umbrella of *Bayesian optimization* (BO) since we introduce a prior over the reward structure of the context-action space and leverage this model for optimization. Unlike many other problems under the BO setting, our algorithm searches for an optimal action for each context rather than a single optimal action. This serves as a contrasting feature from other problems in multi-task BO [Swersky et al., 2013, Toscano-Palmerin and Frazier, 2018, Shah and Ghahramani, 2016], in which a single action that performs optimally across all objectives simultaneously is sought. Furthermore, our algorithm is theoretically grounded and can broadly be

applied to "offline" problems in which context can be set explicitly. In the remainder of the paper, we briefly describe the algorithm and apply it to the nuclear fusion problem.

## 2 Algorithm

Generally, one can think of the optimization problem for each context as a task; we are then faced with a multi-task optimization problem. Let  $\mathcal{X}$  be the finite collection of tasks and let  $\mathcal{A}$  be the compact set of possible actions. For our application, we assume that the set of actions is the same for each task. Let  $f : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$  be the reward function, where  $f(x, a)$  is the reward for performing action  $a$  in task  $x$ . It is assumed that this reward function is always bounded. Let  $\hat{h} : \mathcal{X} \rightarrow \mathcal{A}$  be our estimated mapping from task to action. Our goal is then to find an  $\hat{h}$  which maximizes  $\sum_{x \in \mathcal{X}} f(x, \hat{h}(x)) \omega(x)$ , where  $\omega(x) \geq 0$  is some weighting on  $x$  (e.g. probability of seeing  $x$  at evaluation time or the importance of  $x$ ). At round  $t$  of optimization, we pick a task  $x_t$  and an action  $a_t$  to perform a query  $(x_t, a_t)$  and observe a noisy estimate of the function  $y_t = f(x_t, a_t) + \epsilon_t$ , where  $\epsilon_t \sim N(0, \sigma_\epsilon^2)$  and is iid. Let  $D_t$  be the sequence of queried tasks, actions, and rewards up to time  $t$ , i.e.  $D_t = \{(x_1, a_1, y_1), \dots, (x_t, a_t, y_t)\}$ . Additionally, define  $\hat{y}_t(x)$  to be the best reward observed for task  $x$  up to time  $t$ ,  $\hat{a}_t(x)$  to be the action made to see this corresponding reward, and  $\mathcal{A}_t(x)$  to be the set of all actions made for task  $x$  up to time  $t$ . In this work, we use Gaussian Processes (GP) to model the reward function. When tasks are correlated, one can use a single GP to jointly model tasks and actions; however, for this paper we only consider a fixed finite set of tasks and opt to model each task,  $x$ , with an independent GP with mean function  $\mu_x$  and covariance function  $\sigma_x$ . For more information about GPs, we refer the reader to Rasmussen and Williams [2005].

Our proposed algorithm (shown in Algorithm 1), named Multi-Task Thompson Sampling (MTS), extends the classic Thompson sampling strategy [Thompson, 1933] to the multi-task setting. The algorithm, simply put, acts optimally with respect to samples drawn from the posterior. That is, at every round a GP sample of the reward function is drawn for each task, and these samples are used as if they were the ground truth reward function to identify the task in which the most improvement can be made. After repeating this for  $T$  iterations, we return the estimated mapping  $\hat{h}$  such that  $\hat{h}(x) = \hat{a}_T(x)$  if an evaluation was made for task  $x$ ; if no evaluations have been made for the task,  $\hat{h}(x)$  maps to an  $a \in \mathcal{A}$  drawn uniformly at random.

---

### Algorithm 1 Multi-Task Thompson Sampling (MTS)

---

**Input:** capital  $T$ , initial capital  $t_{init}$ , mean functions  $\{\mu_x\}_{x \in \mathcal{X}}$ , kernel functions  $\{\sigma_x\}_{x \in \mathcal{X}}$ .  
Do random search on tasks in round-robin fashion until  $t_{init}$  evaluations are expended.  
**for**  $t = t_{init} + 1$  **to**  $T$  **do**  
    Draw  $f(x, \cdot) \sim GP(\mu_x, \sigma_x) | D_{t-1} \quad \forall x \in \mathcal{X}$ .  
    Set  $x_t = \operatorname{argmax}_{x \in \mathcal{X}} \left[ \left( \max_{a \in \mathcal{A}} \tilde{f}(x, a) - \max_{a \in \mathcal{A}_t(x)} \tilde{f}(x, a) \right) \omega(x) \right]$ .  
    Set  $a_t = \operatorname{argmax}_{a \in \mathcal{A}} \tilde{f}(x_t, a)$ .  
    Observe  $y_t = f(x_t, a_t)$ .  
    Update  $D_t = D_{t-1} \cup \{(x_t, a_t, y_t)\}$ .  
**end for**  
**Output:**  $\hat{h}$

---

One benefit of this algorithm is that it comes with theoretic guarantees. For the following, define  $a_t^*(x)$  to be the past action played for task  $x$  up to time  $t$  that yields the largest expected reward. That is,

$$a_t^*(x) := \begin{cases} \operatorname{argmax}_{a \in \mathcal{A}_t(x)} f(x, a) & \mathcal{A}_t(x) \neq \emptyset \\ \operatorname{argmin}_{a \in \mathcal{A}} f(x, a) & \text{else} \end{cases}$$

Note that  $a_t^*(x)$  has an implicit dependence on  $f$ .

**Theorem 1.** Define the maximum information gain to be  $\gamma_T := \max_{D_T} I(D_T; f)$ , where  $I(\cdot; \cdot)$  is the Shannon mutual information. Assume that  $\mathcal{X}$  and  $\mathcal{A}$  are finite. Then if Algorithm 1 is played for

$T$  rounds where  $t_{init} = 0$

$$\mathbb{E}[R_{T,f}] \leq |\mathcal{X}| \left( \frac{1}{T} + \sqrt{\frac{|\mathcal{X}||\mathcal{A}|\gamma_T}{2T}} \right)$$

where the expectation is with respect to the data sequence collected and  $f$ , and where  $R_{T,f}$  is defined to be

$$R_{T,f} := \frac{\sum_{x \in \mathcal{X}} \omega(x) (\max_{a \in \mathcal{A}} f(x, a) - f(x, a_T^*(x)))}{\sum_{x \in \mathcal{X}} \omega(x) (\max_{a \in \mathcal{A}} f(x, a) - \min_{a \in \mathcal{A}} f(x, a))}$$

when the denominator is not 0. Otherwise,  $R_{T,f}$  takes the value of 0.

The proof relies on ideas from Kandasamy et al. [2019], and the details can be found in Appendix A of Char et al. [2019]. The result gives a bound on the normalized simple regret summed across tasks where the  $\sqrt{|\mathcal{X}||\mathcal{A}|}$  factor in the theorem accounts for the number of actions that can be taken at every step, and the  $\sqrt{\gamma_T}$  factor characterizes the complexity of the prior over the tasks. Intuitively, this result shows that there is no task in which we will have especially bad results, and when  $\gamma_T = o(T)$ , the normalized simple regret converges to 0 in expectation for every task. Finally, we note that these types of results can usually be generalized to infinite action spaces via known techniques [Russo and Van Roy, 2016, Bubeck et al., 2011]

### 3 Application to Nuclear Fusion

#### 3.1 Tokamak Simulator (TRANSP) and Overview of Problem

We use the TRANSP program [Grierson et al., 2018] to simulate fusion reactions on the DIII-D tokamak, a tokamak in San Diego that is operated by General Atomics. TRANSP is a time-dependent transport code used for interpretive analysis and predictive simulations of tokamaks. Access to TRANSP and running TRANSP experiments were possible thanks to our collaborators at Princeton Plasma Physics Lab. TRANSP operates by simulating real-world experiments (referred to as "shots") that were conducted on DIII-D. By running the predictive module of TRANSP, we are able to predict how changes in controls would affect the plasma. When simulating a given shot (a simulation on TRANSP is referred to as a "run"), we can identify variables at each time step that correspond to the state of the plasma. One such variable that we focus on is  $\beta_n$ , which is a ratio of the pressure of the plasma to the magnetic energy density.  $\beta_n$  serves as a proxy for the economic output of the reaction. Besides this quantity, we also consider the *total energy eigenvalues*, which represent the amount of change in energy within and outside the plasma due to certain perturbations. In particular, we focus on the minimum value of the total energy eigenvalues, which we will refer to as  $\Delta\omega$ .  $\Delta\omega$  serves as a proxy for the stability of the plasma. When conducting a simulation, we apply controls that specify parameters of the neutral beams, which include power, energy, full energy fraction and half energy fraction. The DIII-D tokamak has a total of 8 neutral beams, 6 of which are co-current beams (inject in the same direction as the plasma current) and 2 of which are counter-current beams (inject in the opposite direction of the plasma current). In our experiments, we confine the action space to 2 dimensions: power coefficient of co-current beams and counter-current beams, each with domain [0.001, 1.0]. These power coefficients are applied by multiplying the maximum power of the set of beams by the coefficient. By ranging the power coefficient from 0.001 to 1.0, we essentially scale the beam powers from the minimum to the maximum power level possible.

In our experiments, we consider 8 distinct states of the plasma, which are represented by 8 shots. In all 8 shots, a common instability called *tearing* occurred. Ideally, we would like to perform preventative measures once we sense a tear is about to occur. Therefore, we start the simulation 150 ms before time of tearing and run the simulator until 150 ms after the tearing. After the run completes, we extract the  $\Delta\omega$  and  $\beta_n$  values at 5ms increments throughout the duration of the run (total 300 ms) and average them to produce  $\overline{\Delta\omega}$  and  $\overline{\beta_n}$ . In order to balance between stability and the pressure in the tokamak, we set our reward to be  $10\overline{\beta_n} + 100\overline{\Delta\omega}$ , where we chose coefficients based on the scales of each value to make the two objective components roughly equal. In summary, we optimize a combination of pressure and stability of the plasma, for each of the 8 different plasma states (8 tasks or contexts) simultaneously, by changing the power level of the co-current and counter-current beams (2D controls).




Figure 1: **Fusion Simulation Experiments.** Each of the above show average values and standard error from 10 trials. (a) shows the log total regret summed across all states and (b) log regret achieved in each state. Note that curves differ in length for (b) since different amounts of resources were allocated for each state. Note that regret was approximated by treating the optimum as the maximum value observed for each state, plus a small  $\epsilon$  term.

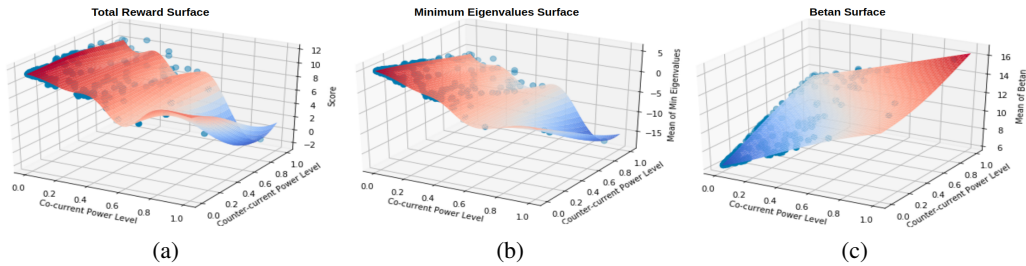


Figure 2: **Reward and Reward Component Surfaces.** The surfaces have been estimated by fitting a GP to the queried points from shot 149205, which corresponds to plasma state 1. Each point in the space shows a value returned by the simulator.

### 3.2 Tokamak Control Optimization

We optimize the reward produced from TRANSP simulations using both MTS and a method that chooses shots (or plasma states) uniformly at random then performs the standard Thompson sampling procedure. The optimization experiment results presented in Figure 1 are averaged over 10 trials, each with 125 query capital. In each trial and for each shot, 5 initial points are drawn uniformly at random for evaluation to form an initial GP. Each task is modeled by an independent GP with an RBF kernel and hyperparameters are tuned for each GP every time an observation is seen for its corresponding shot by maximizing the marginal likelihood. We parallelized the experiments by having a batch of 20 workers, each making evaluations according to the respective algorithms using a shared pool of collected data. This process is suboptimal since workers operate asynchronously (i.e. they do not wait to see the data other workers will collect); however, Kandasamy et al. [2018] showed that this approach is not unreasonable for the standard Thompson sampling setting. The results demonstrate how MTS is able to achieve better performance by focusing its resources intelligently. Looking at Figure 1 (b), in contexts where reward (and hence regret) levels off quickly (e.g. plasma states 4 and 5), MTS is able to recognize that resources should be allocated in other contexts where higher improvement is expected. This is reflected with more queries and better optimization in plasma states 2 and 3.

### 3.3 Discussion of Results and Future Directions

Figure 2 shows the total reward and reward component surfaces for one of the 8 shots from the experiments conducted in Section 3.2. From the total reward surface (Figure 2 (a)), we can see that

beam power should be scaled down for maximum total reward. However, the surfaces of each reward component indicate the negative correlation between plasma stability (Figure 2 (b)) and pressure (Figure 2 (c)): as beam power is scaled up, stability increases but pressure decreases. Hence, there is a fine balance in optimizing a combination of these components, which is dependent on the weighting between the two objectives. This raises further questions about exactly what weighting would be optimal for actual plasma behavior over a longer period of time. In addition, these results provide interesting insights to the fusion community. Although there has been some previous work applying machine learning to fusion in the past [Cannas et al., 2013, Tang et al., 2016, Montes et al., 2019, Kates-Harbeck et al., 2019, Baltz et al., 2017], to the best of our knowledge this is the first work towards learning a tokamak controller offline with no human intervention.

In the future, we hope to discover more interesting results by increasing our action space and forming different reward functions. We are also working to expand this work to finding a policy over a continuous set of plasma states, rather than just a subset of eight. From an algorithmic standpoint, this problem has been explored by Ginsbourger et al. [2014] and Pearce and Branke [2018]. We have preliminary evidence showing that a variation of our algorithm performs competitively with theirs. Lastly, readers may remark that the problem of tokamak control is actually a reinforcement learning problem, since we should be searching for an optimal policy that makes a sequence of actions rather than a single action. Because the simulator is expensive (it takes approximately 2 hours to simulate one control evaluation lasting 300 ms), we chose to limit the scope of the problem in this work; however, we wish to revisit this in the future.

## 4 Acknowledgements

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE1252522 and DGE1745016. Willie Neiswanger is also supported by NSF grants CCF1629559 and IIS1563887. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Youngseog Chung is supported by the Kwanjeong Educational Foundation.

The authors would also like to thank the reviewers for their helpful feedback.

## References

- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135, 2013.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- EA Baltz, E Trask, M Binderbauer, M Dikovskiy, H Gota, R Mendoza, JC Platt, and PF Riley. Achievement of sustained net plasma heating in a fusion experiment with the optometrist algorithm. *Scientific reports*, 7(1):6425, 2017.
- Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. X-armed bandits. *Journal of Machine Learning Research*, 12(May):1655–1695, 2011.
- Barbara Cannas, Alessandra Fanni, A Murari, Alessandro Pau, Giuliana Sias, and JET EFDA Contributors. Automatic disruption classification based on manifold learning for real-time applications on jet. *Nuclear Fusion*, 53(9):093023, 2013.
- Ian Char, Youngseog Chung, Willie Neiswanger, Kirthevasan Kandasamy, Andrew Oakleigh Nelson, Mark D Boyer, Egemen Kolemen, and Jeff Schneider. Offline contextual bayesian optimization. In *Advances in Neural Information Processing Systems*, 2019.
- Francis Chen. *An indispensable truth: how fusion power can save the planet*. Springer Science & Business Media, 2011.
- Daniel Clery. *A Piece of the sun: the quest for fusion energy*. Abrams, 2014.

- David Ginsbourger, Jean Baccou, Clément Chevalier, Frédéric Perales, Nicolas Garland, and Yann Monerie. Bayesian adaptive reconstruction of profile optima and optimizers. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):490–510, 2014.
- BA Grierson, X Yuan, M Gorelenkova, S Kaye, NC Logan, O Meneghini, SR Haskey, J Buchanan, M Fitzgerald, SP Smith, et al. Orchestrating transp simulations for interpretative and predictive tokamak modeling with omfit. *Fusion Science and Technology*, 74(1-2):101–115, 2018.
- Kirthevasan Kandasamy, Akshay Krishnamurthy, Jeff Schneider, and Barnabás Póczos. Parallelised bayesian optimisation via thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 133–142, 2018.
- Kirthevasan Kandasamy, Willie Neiswanger, Reed Zhang, Akshay Krishnamurthy, Jeff Schneider, and Barnabas Poczso. Myopic posterior sampling for adaptive goal oriented design of experiments. In *Proceedings of the 36th International Conference on Machine Learning*. JMLR. org, 2019.
- Julian Kates-Harbeck, Alexey Svyatkovskiy, and William Tang. Predicting disruptive instabilities in controlled fusion plasmas through deep learning. *Nature*, page 1, 2019.
- Andreas Krause and Cheng S Ong. Contextual gaussian process bandit optimization. In *Advances in Neural Information Processing Systems*, pages 2447–2455, 2011.
- Kevin Joseph Montes, Cristina Rea, Robert Granetz, Roy Alexander Tinguely, Nicholas W Eidietis, O Meneghini, Dalong Chen, Biao Shen, Bingjia Xiao, Keith Erickson, et al. Machine learning for disruption warning on alcator c-mod, diii-d, and east. *Nuclear Fusion*, 2019.
- Michael Pearce and Juergen Branke. Continuous multi-task bayesian optimisation with correlation. *European Journal of Operational Research*, 270(3):1074–1085, 2018.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning series)*. The MIT Press, 2005. ISBN 9780262182539.
- Daniel Russo and Benjamin Van Roy. An information-theoretic analysis of thompson sampling. *The Journal of Machine Learning Research*, 17(1):2442–2471, 2016.
- Amar Shah and Zoubin Ghahramani. Pareto frontier learning with expensive correlated objectives. In *International Conference on Machine Learning*, pages 1919–1927, 2016.
- Kevin Swersky, Jasper Snoek, and Ryan P Adams. Multi-task bayesian optimization. In *Advances in neural information processing systems*, pages 2004–2012, 2013.
- William Tang, Matthew Parsons, Eliot Feibush, A Murari, J Vega, A Pereira, and J Choi. Big data machine learning for disruption predictions. In *26th IAEA Fusion Energy Conference-IAEA CN-234, Paper Number EX/P6-47*, 2016.
- William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Saul Toscano-Palmerin and Peter I Frazier. Bayesian optimization with expensive integrands. *arXiv preprint arXiv:1803.08661*, 2018.