Large Language Model-based Bayesian Optimization for Tokamak Stabilization

Tejus Gupta, Rohit Sonker, Efe Mert Karagözlü, Alexandre Capone,
Barnabás Póczos, Jeff Schneider
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Tokamak-based nuclear fusion offers a promising pathway toward sustainable, large-scale energy generation. Our work focuses on the task of predicting and suppressing tearing instabilities, which can cause the plasma to disrupt and potentially damage the tokamak, by optimizing Electron Cyclotron Heating (ECH) configurations. This task is challenging due to distribution shifts in the dynamics of the tokamak arising from hardware changes between experiments, actuator failures, and impurities in the plasma. We propose a large language model-informed Bayesian optimization scheme that aims to explore and find highly stable ECH configurations efficiently. The large language model allows us to leverage high-dimensional prior data and experiment logs written by scientists and operators, which would usually be impossible with conventional Bayesian optimization (BO) tools. We conduct preliminary offline evaluations on archived DIII-D tokamak data to compare our approach with conventional Bayesian optimization techniques.

1 Introduction

Nuclear fusion holds the promise of producing high quantities of energy with minimal to no environmental impact. In fusion, energy is released through a sustained, controlled reaction in which light atoms are combined to form heavier atoms. Among the various technologies used to achieve controlled fusion, the most promising are tokamaks [Kikuchi, 2010], which utilize magnetic fields to confine high-temperature and high-pressure plasma, thereby creating conditions for the fusion reaction to occur. However, maintaining high plasma pressures is challenging, one of the reasons being the occurrence of tearing instabilities in the plasma. These instabilities are magnetic islands that form in the plasma, which reduce the quality of confinement and can also grow, leading to plasma disruptions. Disruptions are harmful events that can severely damage the tokamak walls. Electron Cyclotron Heating (ECH) produced via gyrotrons has shown promise in suppressing these instabilities [Kolemen et al., 2014]. ECH is also used for other purposes such as core heating, and using ECH efficiently remains an open problem. This challenge is further compounded by fluctuations in the tokamak hardware, such as wall changes, sensor and actuator failures, and residual impurities resulting from prior experiments, which can affect the plasma dynamics and alter the effect of ECH. This makes it very difficult to plan experiments based on historical data. To address these challenges, online and adaptive methods such as Bayesian optimization have been used in this task [Sonker et al., 2025] and for other fusion tasks [Mehta et al., 2024]. However, these Bayesian optimization methods face difficulties due to high-dimensional search spaces, and are not able to readily leverage expert priors and experimental logs.

To address the aforementioned challenges, we present an LLM-based Bayesian optimization approach for tokamak control. Our approach leverages observations collected during past experiments to

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: Machine Learning and the Physical Sciences.

explore and find heating profiles that maximize the time-to-instability. Moreover, unlike classical Bayesian optimization, our LLM-based approach enables the processing of logs written by scientists between experiments. We present preliminary results obtained from a simulated Bayesian optimization setup informed by past experiments and logs.

2 Related Work

2.1 Tokamak control and tearing instabilities

A key challenge in tokamaks is controlling neoclassical tearing modes (also referred to as tearing instabilities), which are magnetic islands that form within the plasma. These instabilities degrade plasma confinement [Sauter et al., 1997] and can grow unless suppressed, leading to plasma disruptions [Westerhof, 1990] and potentially damaging the reactor wall. Plasma disruptions are one of the major concerns for stable operations in future large tokamaks, e.g., ITER [Lehnen et al., 2015]

To counteract tearing instabilities, Electron Cyclotron Heating (ECH) is often used to create localized current drive and heating at the location of instabilities. To create these, gyrotrons are aimed at specified locations in the plasma [Kolemen et al., 2014] [Nelson et al., 2020]. Gyrotrons can also be used pre-emptively to stabilize the plasma [Bardóczi et al., 2023]. Since tokamak devices regularly undergo system upgrades, past data is potentially unreliable, thus, Bayesian optimization becomes a strong candidate for pre-emptive ECH suppression, where the tokamak is treated as a black box [Sonker et al., 2025].

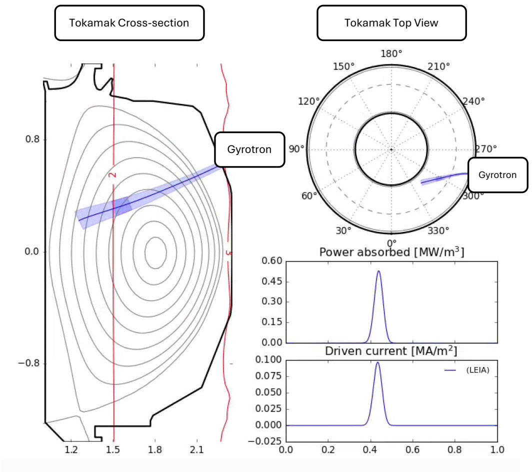

Figure 1: Gyrotron action on the Plasma inside the Tokamak. The bottom 2 curves indicate the heating power absorbed (ECH profile) and current driven in the plasma (ECCD profile) from the core to outer region of the plasma.

In this work, we adopt a similar approach to that in [Sonker et al., 2025] to mitigate tearing instabilities in plasma experiments by selecting the optimal shape of the feedforward ECH heating profile, which we approximate using a Gaussian curve. An example of the angle of the gyrotron and the resulting ECH deposition profile is shown in Fig. 1. Gyrotrons at DIII-D are used to produce Electron Cyclotron Heating (ECH), and additionally, depending on the angle of incidence, they also lead to Electron Cyclotron Current Drive (ECCD). Here, we only optimize for the Electron Cyclotron Heating (ECH) profile produced by the combined effect of all gyrotrons. We focus on $m/n = 2/1$ tearing instabilities, which are very common and highly prone to cause disruptions. m, n refer to the poloidal and toroidal mode numbers respectively. We also consider experiments exclusively in a high-q$_{min}$ plasma scenario, a type of high-pressure scenario that supports long-pulse steady-state plasma operation.

2.2 LLM-Augmented Bayesian Optimization

Classical Bayesian Optimization (BO) methods have been deployed for adaptive experimental design across a range of settings such as adaptive clinical trials [Berry, 2006], ecological monitoring [Diggle and Lophaven, 2006]. Despite these successes, these BO methods often fail to leverage the potentially rich task-specific prior information available to us. This has motivated a recent surge in interest using language models as a tool for more informed Bayesian optimization [Liu et al., 2024, Chang et al., 2025]. Our methodology primarily follows previous work on directly prompting the language models with prior experiment history and the observed values of the variable of interest, and prompting it to generate the next experiment [Zhang et al., 2023]. Our approach contrasts with these in that we incorporate detailed experimental logs into the experimental design observations.

The use of language in the fusion community offers a rich source of data [Mehta et al., 2023]. After every experiment, scientists create logs that contain information about the experiment. We investigate if using this additional information can improve decision making under a BO setup.

3 Problem Setup

3.1 Contextual Bayesian Optimization (CBO)

We now formulate our problem in a contextual Bayesian optimization setup. Our goal is to find an ECH profile a_{ech} that maximizes the time-to-(tearing)instability given the context. The ECH profile is approximated by a Gaussian, parameterized by a 3-dimensional vector. In this work, we use the target normalized plasma pressure β_N as the context, which strongly correlates with plasma stability and is frequently adjusted between experiments by physicists. Our goal is to perform well on any given target β_N.

The contextual decision problem is then as follows: every round t, a context $\beta_{N,t}$ is determined prior to action selection. The agent then chooses a heating profile configuration x_t and measures the (noisy) time-to-instability

$$y_t \;=\; f(x_t, \beta_{N,t}) + \varepsilon_t,$$

where f is the unknown objective, i.e., the expected time-to-instability, and ε_t is the observation noise. The goal is to learn a policy π that maps contexts to actions with high utility, i.e., $\pi(\beta_N) \approx \arg\max_x f(x, \beta_N)$. This is achieved by choosing actions that trade off exploration and exploitation over multiple rounds.

3.2 Auxiliary Feedback from experiments

In addition to the observations y_t, we also include experiment logs written down by scientists after the experiment is performed. The information presented here encompasses a wide range of specific details about the experiment and tokamak, including neutral beam failures, pellet injection, and gyrotron lag. Importantly, these logs can contain information that is impossible to express in the low-dimensional search space of traditional BO tools. Logs written during and after experiments are often written in shorthand; hence, we first pass these logs through a Llama-3-8B model to eliminate short forms and turn them into well-formed, complete sentences. Some examples are shown in table 1. These logs contain all information observed during the shot and are not restricted to our goal of tearing mode avoidance and ECH control. This information is present in some logs, however, not in all. Inaccuracies may arise as a result of LLM hallucinations. For example in the first log in 1, the LLM incorrectly assumes that numbers 5000 and 3599-3612 correspond to power, when they refer to timings in milliseconds. Thus, the auxiliary information provided is inherently noisy.

4 Experiment Setup

In this section, we describe our experiments.

We employ an end-to-end LLM to perform Bayesian optimization, with the difference that we include written logs in the observations. We utilize historical data from 281 past experiments carried out at the DIII-D tokamak facility between 2012 to 2023. These experiments correspond to high q_{min} plasma scenario with β_N values ranging from 2.6 to 3.9.

This data is used to simulate live experiments. At each step we sample a target plasma pressure β_N^t from the range of historical values. We then select a subset of historical experiments whose β_N lies between $[\beta_N^t - \epsilon, \beta_N^t + \epsilon]$, where $\epsilon = 0.05$. Now, the Bayesian optimization task corresponds to selecting the ECH value corresponding to any experiment in this subset. After selecting the ECH value, we treat the corresponding experiment and its time to tearing mode as an evaluated observation, which is then used to perform the posterior update. To evaluate each method, we employ contextual cumulative regret

$$\mathcal{R}_T \;=\; \sum_{t=1}^{T} \left[f^{\max} - f(x_t, \beta_{N,t}) \right] ,$$

where $f^{\max}$ denotes the maximal time-to-instability and T denotes the current round.

Example # **Experiment Logs**

1 **Raw Log:**

```
Shotnumber = 186497
PCS CHANGES:
ECH: change parameter data:  test parameters/ShotStart/ECH Shutdown
periodic action:  ECH Control Params/ShotStart/Gyrotrons
Shot Type = Plasma
periodic action:  Beam Settings/ShotStart/System
release release/3.0.30----------------------------------------
Gyro  DELAY  REQUEST  ACTUAL  COMMENTS
Leia  1500  5000  3611  DIII-D Permit
R2D2  1500  5000  3599  DIII-D Permit
Luke  1500  5000  3612  DIII-D PermitRepeat.
Result:  OK, good shot.  Plasma shot Ok
BEAMS:
- Repeat
RESULTS:
- Beams all work
- Feedback seems to be doing well until mode takes bite out of betan
```

Summarized Log: The experiment logs indicate that a plasma shot was conducted at the DIII-D National Fusion Facility, utilizing three gyrotrons, Leia, R2D2, and Luke, to control the plasma. The gyrotrons were set to operate with a delay of 1500 microseconds, a requested power of 5000 megawatts, and an actual power of approximately 3599–3612 megawatts. The plasma was successfully initiated, and all beams were functioning as expected until a mode began to affect the beta value. The plasma shot was deemed successful, with the result being classified as “OK, good shot.”

2 **Raw Log:**

```
Shotnumber = 155543
Last shot of the hybrid part of todays experiment.  Increase density
to the same as 155541 and increase co-beam power to maximum.
A very good shot until 4.5 s when a 2/1 mode grew up (although not a
disaster as beta_N remained above 3.2 until the beam started turning
off at 4.9 s).
Between 3.0 s and 4.5 s the average values were beta_N = 3.64, H_89P
= 2.4 and surface loop voltage of -0.003 V.
The totals for this morning are as follows:
11 plasma shots, that consisted of
1 reference shot
6 clean up shots at two different dRsep values
4 good shots with beta_N>3.2, three of which had negative loop
voltage @ 1.0 MA
...  (showing only one fourth of the full log due to space
limitations)
```

Summarized Log: The last shot of the hybrid part of today’s experiment was completed by increasing the plasma density to the same level as shot 155541 and maximizing the co-beam power. This shot was very successful until 4.5 seconds, when a 2/1 mode began to grow, although the plasma’s betaN remained above 3.2 until the beam was turned off at 4.9 seconds. Between 3.0 and 4.5 seconds, the average plasma properties were betaN = 3.64, H89P = 2.4, and a surface loop voltage of -0.003 volts. Today’s experiment consisted of 11 plasma shots, including 1 reference shot, 6 clean-up shots at two different dRsep values, and 4 successful shots with a betaN of over 3.2, three of which had a negative loop voltage at 1.0 MA.

Table 1: Examples of experiment logs which are usually written post experiment. Note that these logs have been processed by an LLM to convert them into full sentences and eliminate short-forms. This can lead to some hallucinations which are marked in red.

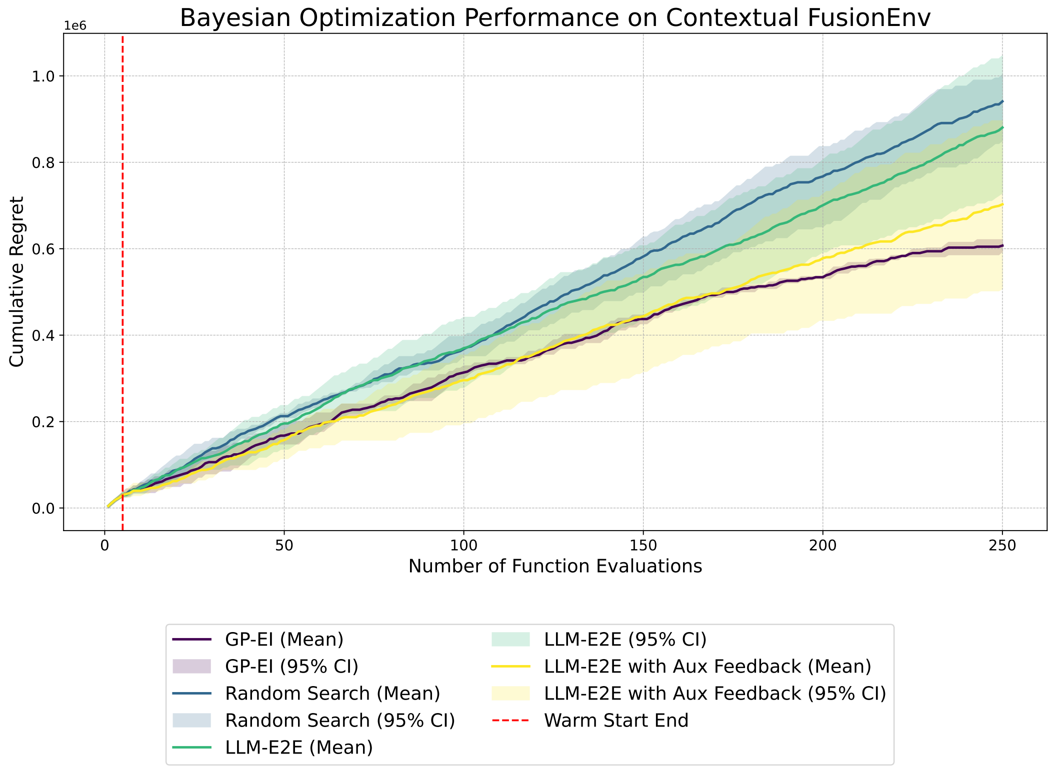

Figure 2: Contextual Cumulative Regret for LLM-guided BO and other baselines with an optimization budget of 250 steps. All methods were run for 3 random seeds.

Methods Compared We compare the following approaches -

- **GP-EI :** Traditional Bayesian optimization with a Gaussian Process surrogate model and Expected Improvement as the Acquisition Function
- **LLM-E2E :** LLM as an end-to-end Bayesian Optimizer. Here, we provide the history of past observations, the current context and subset of ECH configurations to select from and ask the LLM to pick the next ECH configuration. We do this iteratively while expanding the past observation set.
- **LLM-E2E with Aux Feedback :** Similar to LLM-E2E, with the LLM being provided experiment logs as auxiliary feedback for each evaluation point.

The LLM used for experiments was Llama-3.3-70B instruction-tuned model [Grattafiori et al., 2024].

Fig. 2 shows the performance of our LLM-guided BO method and baselines with an optimization budget of 250 timesteps. As expected, random search shows linear regret, and GP-EI shows sublinear regret. LLM-E2E performs similarly to random search, which suggests vanilla usage of LLMs do not improve the optimization performance for this task. However, we observe that auxiliary feedback improves the LLMs performance significantly, making it track the mean performance of the GP-EI closely until at least 150 steps. We hypothesize that the degradation in performance after 150 steps is due to the lack of an explicit Bayesian update mechanism with LLMs, and the long context windows.

5 Conclusion

We have presented an LLM-based BO approach for tokamak plasma stabilization. Our method incorporates freeform text into the optimization loop, contrasting with vanilla Bayesian optimization methods. Preliminary results using a simulated setup show promise compared to an LLM-based approach without experimental logs. However, vanilla Bayesian optimization outperforms our strategy in the long run, indicating that improvements can be achieved. This motivates our future work, where we will combine both methods in a way that simultaneously leverages their advantages.

6 Acknowledgements

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award(s) DE-FC02-04ER54698 and DE-SC0024544.

References

L Bardóczi, RJ La Haye, EJ Strait, NC Logan, SP Smith, NJ Richner, and JD Callen. Direct preemptive stabilization of neoclassical tearing modes by electron cyclotron current drive in the diii-d low-torque iter baseline scenario. *Nuclear Fusion*, 63(9):096021, 2023.

Donald A Berry. Bayesian clinical trials. *Nature reviews Drug discovery*, 5(1):27–36, 2006.

Chih-Yu Chang, Milad Azvar, Chinedum Okwudire, and Raed Al Kontar. `LLINBO`: Trustworthy llm-in-the-loop bayesian optimization. *arXiv preprint arXiv:2505.14756*, 2025.

Peter Diggle and Søren Lophaven. Bayesian geostatistical design. *Scandinavian Journal of Statistics*, 33(1):53–64, 2006.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Mitsuru Kikuchi. A review of fusion and tokamak research towards steady-state operation: A jaea contribution. *Energies*, 3(11):1741–1789, 2010.

Egemen Kolemen, AS Welander, RJ La Haye, NW Eidietis, DA Humphreys, J Lohr, V Noraky, BG Penaflor, R Prater, and F Turco. State-of-the-art neoclassical tearing mode control in diii-d using real-time steerable electron cyclotron current drive launchers. *Nuclear Fusion*, 54(7):073020, 2014.

M. Lehnen, K. Aleynikova, P.B. Aleynikov, D.J. Campbell, P. Drewelow, N.W. Eidietis, Yu. Gasparyan, R.S. Granetz, Y. Gribov, N. Hartmann, E.M. Hollmann, V.A. Izzo, S. Jachmich, S.-H. Kim, M. Kočan, H.R. Koslowski, D. Kovalenko, U. Kruezi, A. Loarte, S. Maruyama, G.F. Matthews, P.B. Parks, G. Pautasso, R.A. Pitts, C. Reux, V. Riccardo, R. Roccella, J.A. Snipes, A.J. Thornton, and P.C. de Vries. Disruptions in iter and strategies for their control and mitigation. *Journal of Nuclear Materials*, 463:39–48, 2015. ISSN 0022-3115. doi: https://doi.org/10.1016/j.jnucmat.2014.10.075. URL `https://www.sciencedirect.com/science/article/pii/S0022311514007594`. PLASMA-SURFACE INTERACTIONS 21.

Tennison Liu, Nicolás Astorga, Nabeel Seedat, and Mihaela van der Schaar. Large language models to enhance bayesian optimization. *arXiv preprint arXiv:2402.03921*, 2024.

Viraj Mehta, Joseph Abbate, Allen Wang, Andrew Rothstein, Ian Char, Jeff Schneider, Egemen Kolemen, Cristina Rea, and Darren Garnier. Towards llms as operational copilots for fusion reactors. In *NeurIPS 2023 AI for Science Workshop*, 2023.

Viraj Mehta, Jayson Barr, Joseph Abbate, Mark D Boyer, Ian Char, Willie Neiswanger, Egemen Kolemen, and Jeff Schneider. Automated experimental design of safe rampdowns via probabilistic machine learning. *Nuclear Fusion*, 64(4):046014, 2024.

A. O. Nelson, N. C. Logan, W. Choi, E. J. Strait, and E. Kolemen. Experimental evidence of electron-cyclotron current drive-based neoclassical tearing mode suppression threshold reduction during mode locking on diii-d. *Plasma Physics and Controlled Fusion*, 62(9), 07 2020. ISSN ISSN 0741-3335. doi: 10.1088/1361-6587/ab9b3b. URL `https://www.osti.gov/biblio/1642432`.

Olivier Sauter, Robert J La Haye, Zuoyang Chang, David A Gates, Yutaka Kamada, Hartmut Zohm, Anders Bondeson, Dominique Boucher, James D Callen, Ming S Chu, et al. Beta limits in long-pulse tokamak discharges. *Physics of Plasmas*, 4(5):1654–1664, 1997.

Rohit Sonker, Alexandre Capone, Andrew Rothstein, Hiro Josep Farre Kaga, Egemen Kolemen, and Jeff Schneider. Multi-timescale dynamics model bayesian optimization for plasma stabilization in tokamaks. In *Forty-second International Conference on Machine Learning*, 2025.

E Westerhof. Tearing mode stabilization by local current density perturbations. *Nuclear Fusion*, 30 (6):1143, 1990.

Michael R Zhang, Nishkrit Desai, Juhan Bae, Jonathan Lorraine, and Jimmy Ba. Using large language models for hyperparameter optimization. *arXiv preprint arXiv:2312.04528*, 2023.