

[Re] Bandit Theory and Thompson Sampling-guided Directed Evolution for Sequence Optimization

Luka Žontar¹, 

¹University of Ljubljana, Ljubljana, Slovenia

Reproducibility Summary

Edited by
Koustuv Sinha,
Maurits Bleeker,
Samarth Bhargav

Received
04 February 2023

Published
20 July 2023

DOI
10.5281/zenodo.8173664

Scope of Reproducibility – The paper presents a novel DE approach using Thompson Sampling and Bandit theory, TS-DE. Our reproducibility study aims to confirm the 5 main claims of the original paper, including sublinear Bayesian regret, improved performance compared to basic DE, robustness to mutation rate changes, initial diversification, and the concentration of the population to the optimal value in later iterations, and iterative distribution shift towards optimal population fitness. Finally, we provide a reproducible environment to support the main claims of the original paper along with the source code of the proposed approach and all experiments, comprehensive documentation, and unit tests.

Methodology – No code was available beforehand for this article, thus we re-implemented the proposed approach by meticulously following the comprehensive explanations of the process in the original article. The experiments were run on a personal computer.

Results – We managed to reproduce all the experiments supporting the main claims of the original article. Additionally, we add uncertainty quantification to the results as we believe this is a crucial part to confirm any of the claims. Finally, we present the exploration-exploitation trade-off experiment in a more robust manner leveraging the nucleotide diversity metric to gain additional insight into how the proposed algorithm works.

What was easy – With comprehensive explanations in the original article, it was relatively easy to rewrite the main concepts from pseudo-code to Python code. Experiments were clearly explained with all the necessary hyperparameters.

What was difficult – Since no source code was available, every detail missing from the original article resulted in additional research and trial and error experimentation. To conclude, we recommend several improvements to the authors of the studied paper that could additionally improve the quality of their outstanding contribution. Some of the recommendations include better explanations of how the optimal solution is calculated, on which population the PCA is fitted, how to use θ^* and θ , and lastly improved documentation on how the basic DE was implemented.

Communication with original authors – We contacted the original authors on multiple occasions during the development of our reproducibility study but got no response.

Copyright © 2023 L. Žontar, released under a Creative Commons Attribution 4.0 International license.

Correspondence should be addressed to Luka Žontar (luka.zontar@outlook.com)

The authors have declared that no competing interests exist.

Code is available at <https://github.com/lukazontar/RC2022-Bandit-Theory-and-Thompson-Sampling-Guided-Directed-Evolution-for-Sequence-Optimization>. – SWH [swh:1.dir:7cffa7d64bb51bac56c43f657e2e3746a2539915](https://swh.cs.stanford.edu/swh/1.dir:7cffa7d64bb51bac56c43f657e2e3746a2539915).

Open peer review is available at https://openreview.net/forum?id=NE_x1dpz-Q.

1 Introduction

In this reproducibility challenge, we aim to independently verify the claims made in the studied paper regarding the effectiveness of a novel machine learning-assisted approach for Directed Evolution (DE), a wet-lab method used for discovering novel protein designs. The authors proposed a Thompson Sampling-guided Directed Evolution (TS-DE) framework for sequence optimization, where the sequence-to-function mapping is unknown and querying a single value is subject to costly and noisy measurements.

We plan to reproduce the experiments presented in the original paper, where the authors tested their approach on a simulated population of protein sequences represented as binary motifs, corresponding to favorable and unfavorable sites, respectively. The authors demonstrated that the TS-DE approach, leveraging bandit learning and Thompson sampling, reaches a nearly optimal regret bound and outperforms the basic DE approach by converging faster and being more robust to mutation rate scheduling.

In addition to reproducing the results, we aim to improve the original results by including uncertainty quantification, more accurately 95% confidence intervals, and showing how population diversity evolves over iterations of the TS-DE approach.

2 Scope of reproducibility

In this reproducibility study, we set out to rigorously validate the claims made by the authors in their original paper regarding the sublinear and nearly optimal regret bound of the proposed approach. Our primary objective is to replicate the key experiments presented in the original article and extend the evaluation to further test the efficacy of the proposed method.

First, we list the main claims that the authors make in the original paper:

- **Claim 1:** TS-DE reaches sublinear regret bound that decreases with the number of sequences in the population.

In the original paper, the claim is formally proven in Section 5.2. The experiment is shown in the left-hand side of Figure 6.1.

- **Claim 2:** TS-DE outperforms the basic DE when evaluating average population fitness.

The experiment that supports this claim is shown on the right-hand side of Figure 6.1 in the original article.

- **Claim 3:** TS-DE is less sensitive to mutation scheduling than the basic DE.

The experiment that supports this claim is shown on the right-hand side of Figure 6.1 in the original article.

- **Claim 4:** In TS-DE the population fitness distribution shifts towards optimal during evolution.

The experiment that supports this claim is shown on the right-hand side of Figure 6.2 in the original article.

- **Claim 5:** TS-DE first diversifies and then concentrates around the optimal solution.

The experiment that supports this claim is shown in the left-hand side of Figure 6.2 in the original article.

Note that all claims were tested on a simulated dataset of zeros and ones, corresponding to unfavorable and favorable sites, respectively.

3 Methodology

The work we attempt to reproduce does not provide any source code and the authors were unresponsive to our attempts to further discuss the article. Hence we decided to reproduce the studied paper leveraging the thorough documentation of the process provided in the article. Although we did not have access to the original source code, we made every effort to match the implementation as described in the paper. No additional documentation and no GPUs were used in the process of reproducing the studied paper. The main focus of our reproducibility work was to review the validity of the claims made in the original study. Our primary focus was to test whether we can confirm those claims by carefully following the procedure proposed in the main article.

In addition to testing the original claims, we deliver the source code of the original article to the broader research community. To ensure the code is accessible and usable by others, we include comprehensive documentation and unit tests.

3.1 Model descriptions

Our reproducibility work includes the implementation and evaluation of two Directed Evolution (DE) algorithms, TS-DE and the basic DE. Both aim to optimize utility of the population. The proposed TS-DE algorithm enhances the performance by incorporating the posterior update of the utility function parametrization, effectively incorporating learning from historical feedback data to drive the posterior toward the optimal θ^* .

Parameters: Basic DE –

- d - Number of protein motifs in a sequence.
- S_0 - Initial population consisting of M candidate sequences.
- θ^* - optimal θ - parametrization of the linear Bayesian utility model for which we aim to optimize the protein design.
- μ - Mutation rate, $\in (0, 1)$ - a function of T .
- σ - Standard deviation used in noisy feedback evaluation in Assumption 3.5 in the original article.
- T - Number of iterations we run in the TS-DE (or basic DE) algorithm.
- f - Protein utility function that we are trying to maximize. The function includes a parameter theta that we are optimizing.

Parameters: TS-DE – All the parameters as in the Basic DE. Additionally, TS-DE requires a few parameters more:

- M - Population size.
- λ - Scalar that controls the trade-off between exploitation and exploration in the optimization process.

3.2 Datasets

All experiments use binary datasets, populations of sequences, where the initial population consists of all zeros. The only case where the initial population is randomly initialized and consists of both zeros and ones is the case where we are testing the basic DE approach with mutation rate $\mu = 0$ to ensure the test case is non-trivial. Datasets differ from experiment to experiment due to different parameters d (sequence length) and M (population size).

3.3 Hyperparameters

Hyperparameters were carefully documented in the original article. Since our goal was to test the validity of the claims in the original article, we focused on testing the given hyperparameter values.

When testing the claim whether TS-DE outperforms the basic DE, authors did not provide the value of the population size. $M = 40$ was manually chosen. Figure 5 uses the same setup as Figures 3 and 4.

3.4 Experimental setup and code

We constructed our experiments by meticulously following the comprehensive explanations of the process in the original article. Moreover, we provide uncertainty quantification in all experiments to be more convinced in the validity of the claims we test. For this reason, we run all experiments 100 times to be able to get 95% confidence intervals in all the experiments. Confidence intervals are calculated using bootstrapping.

In the experimental setup, we use three different measures to support the proposed claims:

1. Bayesian regret: defined in the original article. In the article it is rather unclear which θ is used where, which is why we rewrite the definition here:

$$BayesRGT(T, M) = E \left[\sum_{t=1}^T \sum_{i=1}^M (f_{\bar{\theta}_t}(x^*) - f_{\theta^*}(x_{t,i})) \right]$$

2. Average population fitness, where fitness is defined as $f(x) = \theta \cdot x$. θ^* is used as θ in all our experiments.
3. Nucleotide diversity [1]: measures polymorphic population diversity. The metric is adapted to our definition of sequences using binary protein motifs. Note that x_i and x_j are frequencies of i-th and j-th sequences and π_{ij} is the number of sites where those sequences differ. Definition:

$$\pi = \frac{M}{M-1} \sum_{i,j} x_i x_j \pi_{ij}$$

Source code is available on GitHub.

3.5 Computational requirements

All the experimentation was done on a Lenovo Thinkpad X1 with the following specifications:

- Windows 11 Pro
- Processor: 11th Gen Intel(R) Core(TM) i7-1165G7, 2.80GHz
- RAM: 32GB

Additionally, we provide required CPU time for each of the experiments in Section 4.

4 Results

Our reproducibility study successfully reproduced the experiments supporting the main claims of the original paper. We observe significant differences in our results, likely due to missing specifics from the original source code. Despite that our conclusions align with the original article. We are unable to talk about how precise we were in reproducing the original article due to missing source code.




Figure 1. Population accumulated Bayesian regret evolving over iterations of the TS-DE process. We test multiple population sizes M .

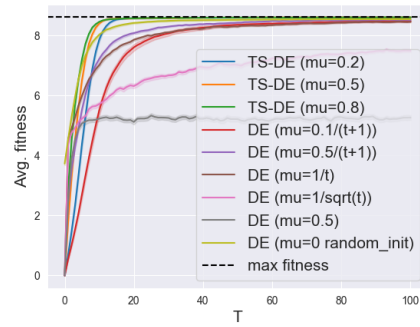


Figure 2. Average fitness evolution comparing TS-DE and the basic DE approach. With each approach we test different parameters μ

4.1 Results reproducing original paper

Bayesian regret evolution with varying population sizes – (*CPU time: $\approx 15min$*) supports **Claim 1**. As shown in Figure 1, we see that Bayesian regret curves converge to sublinear regret bound which is the main claim of the original article. We also note that as M increases the Bayesian regret decreases. This supports the claim made with the left-hand side in Figure 6.1 in the original article. To show how certain we can be in confirming the claim, we include uncertainty quantification in our reproduced experiment. Results show that as M decreases so does the size of the confidence intervals and our certainty in the validity of the claim increases.

Average fitness evolution: TS-DE vs. basic DE – (*CPU time: $\approx 17min$*) supports **Claims 2 and 3**. In this task, we verify the claim that TS-DE outperforms the basic DE in convergence speed. Our results, shown in Figure 2, confirm the claim. To reduce randomness effects and increase certainty in our results, we averaged 100 runs of the experiment and found that the results supported the original study. We also observed that the basic DE is much more sensitive to mutation rate scheduling than TS-DE, which is consistent with the original article’s findings.

Fitness distribution shift – (*CPU time: $\approx 3min$*) supports **Claim 4**. This experiment supports the claim that fitness distribution shifts towards the optimal value over iterations. While the authors tested the claim on just a single evaluation, our experiment was run on 100 evaluations to make the claim validity more robust. Additionally, different time steps were used in the process due to differences between the original and reproduced results.

Population evolution: diversification and convergence – (*CPU time: $< 1min$*) supports **Claim 5**. In this task, we focus on reproducing the evolution trajectory of a population. Our end goal is to visualize multiple snapshots of high-dimensional populations and map them to 2D space using KDE density contour plot. With this plot, we support the claim from the original article and show how TS-DE balances the exploration-exploitation trade-off using initial diversification followed by approaching and concentrating on the optimal solution. Figure 3 shows how the population first diversifies in steps $t = 5$, $t = 10$, and then converges to the optimal solution in steps $t = 20$, $t = 35$.




Figure 3. Population evolution shown as KDE density contour plot of population mapped to 2D space using PCA method.

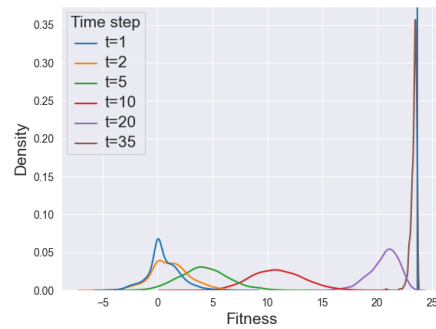


Figure 4. TS-DE shifting population fitness distribution towards optimal value over iterations of the TS-DE process.

4.2 Results beyond original paper

Population evolution: diversification and convergence – (*CPU time: $\approx 4min$*) supports **Claim 5**. While the left-hand side in Figure 6.2 in the original article well represents how TS-DE balances the exploration-exploitation trade-off, we decided to test this claim in a more robust environment by running 100 evaluations instead of just one. We decided that it might be interesting to show population evolution in terms of diversity as can be seen in Figure 5. In the plot, we see how population diversity first increases gradually and then starts to decrease as the population converges toward the optimal solution.




Figure 5. Population diversity evolution over iterations of the TS-DE process. We show the exploration-exploitation trade-off progress.

5 Discussion

To conclude, we were able to support all the claims of the paper and provide the source code in addition with comprehensive documentation and test cases. Additionally, we

managed to include uncertainty quantification in all the results. We also included a metric for nucleotide diversity (adapted for binary sequences of protein motifs) with which we attempt to clarify the exploration-exploitation trade-off in TS-DE.

5.1 What was easy

The authors' pseudo-code was clearly written and the general concepts were easily translatable into Python code. Pseudo-code is written for multiple substeps and functions defined in the article. The notation in the article is clearly understood with some pre-knowledge of Bayesian statistics. The authors also did a very good job of documenting the pseudo-code with explanations and listing the hyperparameters used in experiments.

5.2 What was difficult

Even though pseudo-code was clearly written and explained, we want to stress the importance of sharing the source code and thoroughly documenting the experiments. Even a simple parameter, such as the seed of the random number generation process, can make it impossible to perfectly reproduce the results. Not sharing the source code results in the need to include every single detail in the article. We list some recommendations to authors to make the article more easily reproducible:

1. In most experiments, authors show the maximal value or the optimal solution, however, it is unclear both how this is calculated. To reproduce the results, we would want to state the optimal solution or explain the process of how the optimal solution was retrieved. In our experiments, the optimal θ^* is sampled from the definition in Assumption 3.2 in the studied paper.
2. As mentioned before, we would recommend adding the complete hyperparameter setup for Figure 6.2. The M definition is missing in the original article. We used $M = 40$.
3. To reproduce the right-hand side of Figure 6.2 in the original article, it would be beneficial to learn what data was used to fit PCA to the population in the experiment. In our experiment, we used all the sequences from all the populations that were generated in the optimization process.
4. Additionally, we noted that the Bayesian regret definition does not include where to use θ^* and where to use $\tilde{\theta}$. Even though this is derived later in the article in the proof of the regret bound in 5.2, we would recommend making this part clearer for the reader.
5. While authors write down the main differences between TS-DE and the basic DE, it would help to include the full definition with pseudo-code. Alternatively, the authors could publish the source code for the basic DE as well.

5.3 Communication with original authors

We contacted the original authors on multiple occasions during our attempts to reproduce the studied article, however, there was no response.

References

1. M. Nei and W. Li. "Mathematical model for studying genetic variation in terms of restriction endonucleases." In: **Proceedings of the National Academy of Sciences of the United States of America** 76 10 (1979), pp. 5269–73.