Scalable Bayesian Optimization Accelerates Process Optimization of Penicillin Production

Anonymous Author(s)
Affiliation
Address
email

Abstract

While Bayesian Optimization (BO) has emerged as sample-efficient optimization method for accelerating drug discovery, it has rarely been applied to the process optimization of pharmaceutical manufacturing, which traditionally has relied on human-intuition, along with trial-and-error and slow cycles of learning. The combinatorial and hierarchical complexity of such process control also introduce challenges related to high-dimensional design spaces and requirements of larger scale observations, in which BO has typically scaled poorly. In this paper, we use penicillin prodcution as a case study to demonstrate the efficacy of BO in accelerating the optimization of typical pharmaceutical manufacturing processes. To overcome the challenges raised by high dimensionality, we apply a trust region BO approach (TuRBO) for global optimization of penicillin yield and empirically show that it outperforms other BO and random baselines. We also extend the study by leveraging BO in the context of multi-objective optimization, allowing us to further evaluate the trade-offs between penicillin yield, production time, and CO_2 emission as by-product. Through quantifying the performance of BO across high-dimensional and multi-objective optimization on drug production processes, we hope to popularize application of BO in this field, and encourage closer collaboration between machine learning and broader scientific communities.

1 Introduction

Process optimization of pharmaceutical manufacturing typically involves tuning many control and composition parameters to maximize yield. When new manufacturing processes are tested, practitioners often rely on human intuition for optimization and would have to go through many rounds of trial-and-error to gain a preliminary understanding. They would also optimize input features one-at-a-time, resulting in the loss of opportunity in reaching higher yields through global optimization. In addition, yield data from each set of input control parameter combinations usually takes days to obtain, resulting in large time and resource budgets that prohibit further optimization. Bayesian Optimization (BO) would be an ideal optimization algorithm for drug production processes because it excels in global optimization of systems with expensive cost functions using fewer experiments, and has applications ranging from hyperparameter tuning, molecule screening to materials optimization [1, 2, 3, 4, 5, 6]. However, success of BO has been limited to low-dimensional problems that require smaller sample budgets, whereas the pharmaceutical manufacturing often require optimization in high-dimensions to include the impact of temperature, time, and pH on final yield.

Global optimization in high-dimensional spaces is challenging for BO mostly due to the curse of dimensionality, where design space grows exponentially with n_{dim}. The number of local optima increases and global optimal becomes more elusive. As size of design space becomes intractable with respect to available budget, there will exist regions that have fewer observations, resulting in larger posterior uncertainty and thus over-exploration [7]. In addition, the commonly used surrogate model Gaussian Process (GP) [8] shows cubic time complexity $\mathcal{O}(n^3 + n^2 \cdot n_{\text{dim}})$ with data size n, making the task of fitting a heterogeneous black-box function with a global model computationally costly.

To overcome the high-dimensional optimization challenges and retain the sample-efficient trait of BO, we apply a trust region BO (TuRBO) for global optimization [7] and demonstrate its efficacy for pharmaceutical manufacturing process optimization. TuRBO simultaneously keeps a set of local optimizations campaigns through many local GP surrogate models, each fit independently to its

Submitted to NeurIPS 2021 AI for Science Workshop.

own local trust region, and utilizes a multi-armed bandit strategy at each learning cycle to allocate batch samples between each trust region. We have chosen the fed-batch penicillin fermentation process [9] as a representative drug production process to benchmark the performance of TuRBO and BO algorithms. We demonstrate the efficacy of TuRBO by quantitatively comparing it against state-of-the-art BO algorithms such as Ensemble Bayesian Optimization (EBO) [10], regular BO, and other commonly used sampling strategies such as random search and Latin Hypercube Sampling (LHS). Leveraging Expected Hypervolume Improvement (EHVI) [11] as acquisition function, we additionally perform multi-objective optimization via BO and demonstrate the evaluation of trade-offs between penicillin yield, time, and carbon dioxide (CO_2) by-product emission. We offer open source implementation of the penicillin model and benchmarking code to support future development [12].

2 Penicillin Production

To demonstrate the efficacy of BO in increasing yield for a representative drug production process, we adopt a fed-batch process for penicillin production model proposed by Birol et al [9] and show the closed-loop optimization framework for the system in Figure 1. We realistically simulate the manufacturing process by taking the impact of temperature, pH into account due to the heat generation and acidity tendency during fermentation and microorganisms cell activity. In response to global movements in tracking greenhouse gas emissions, we also recorded evolution of CO_2, and would use it to evaluate additional trade-offs in such production processes.

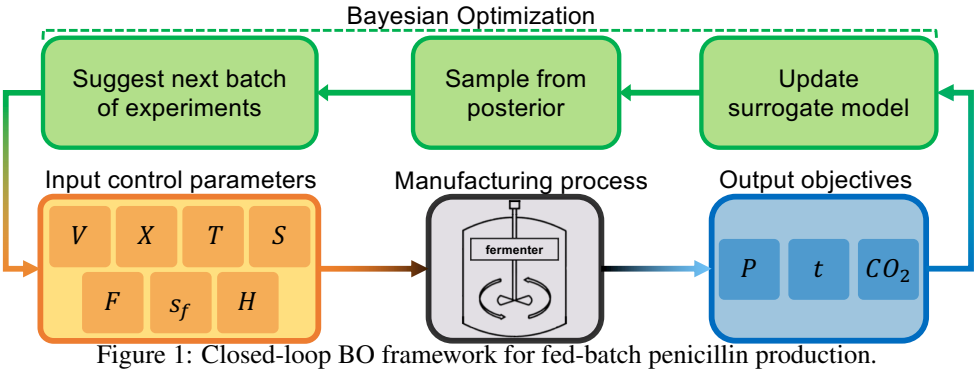

Figure 1: Closed-loop BO framework for fed-batch penicillin production.

We assume excess oxygen supply [13] and run simulations under closed-loop control of assigned temperature, pH and open-loop addition of given glucose substrate. At zero offset of P (penicillin concentration $\frac{g}{L}$), CO_2 (carbon dioxide concentration $\frac{mmol}{L}$) and t (reaction time hr)), the initial values of V (culture medium volume L), X (biomass concentration $\frac{g}{L}$), T (temperature K), S (glucose substrate concentration $\frac{g}{L}$), F (substrate feed rate $\frac{L}{hr}$), s_f (substrate feed concentration $\frac{g}{L}$) and H (H^+ concentration $\frac{mol}{L}$) form the design space as seen in Figure 1 and is governed by:

$$\frac{\mathrm{d}V}{\mathrm{d}t} = F - V\lambda[e^{5\cdot(\frac{T-T_0}{T_v-T_0})} - 1] \tag{1}$$

$$\mu = \frac{\mu_X}{1 + \frac{K_1}{H} + \frac{H}{K_2}} \cdot \frac{S}{K_X X + S} \cdot [k_g e^{(-\frac{E_g}{RT})} - k_d e^{(-\frac{E_d}{RT})}] \tag{2}$$

$$\frac{\mathrm{d}X}{\mathrm{d}t} = \mu X - \frac{X}{V} \cdot \frac{\mathrm{d}V}{\mathrm{d}t} \tag{3}$$

$$\mu_{pp} = \mu_p \cdot \frac{S}{K_p + S + \frac{S^2}{K_I}} \tag{4}$$

$$\frac{\mathrm{d}P}{\mathrm{d}t} = \mu_{pp} X - KP - \frac{P}{V} \cdot \frac{\mathrm{d}V}{\mathrm{d}t} \tag{5}$$

$$\frac{\mathrm{d}S}{\mathrm{d}t} = -\frac{\mu}{Y_{X/S}} X - \frac{\mu_{pp}}{Y_{P/S}} X - m_X X + \frac{F \cdot s_f}{V} - \frac{S}{V} \cdot \frac{\mathrm{d}V}{\mathrm{d}t} \tag{6}$$

$$\frac{\mathrm{d}CO_2}{\mathrm{d}t} = \alpha_1 \frac{\mathrm{d}X}{\mathrm{d}t} + \alpha_2 X + \alpha_3 \tag{7}$$

where rest of the constants can be obtained from Table 2 of Birol et al [9]. To simulate real pharmaceutical manufacturing processes, the fermentation process is run open-loop for penicillin concentration P until P converges or the culture medium volume V exceed $V_{max} = 180L$. Each

process variable is updated at $\Delta t = 0.5hr$ intervals to account for the lag in off-line quantitative laboratory measurements often seen in bioprocesses. Our penicillin production model serves not only as a valuable toolbox for benchmarking optimization algorithms but also as a representative example for drug production processes including a variety of process variables on system dynamics.

3 Bayesian Optimization for Complex Process Optimization

Bayesian optimization (BO) [3, 8, 2] aims to solve the problem of finding a global optimum (min or max) of an unknown objective function g: $\vec{x}^* = \arg\max_x g(\vec{x})$ where $\vec{x} \in X$ and X is a domain of interest in $\mathcal{R}^{n_{dim}}$. There has been previous work in high-dimensional BO [14, 10, 15, 16, 17, 18, 19], yet many have relied on exploiting additive structure for the objective function [16, 17, 18, 10] such as EBO, and mapping to lower dimensional spaces [19, 20], which still require training a large number of GPs that scale poorly with the number of evaluations in high-dimensions. Evolutionary algorithms [21] are well designed for high-dimensional and multi-objective optimization, yet they require many more experimental budget than BO, making them not suitable for applications such as pharmaceutical manufacturing processes optimization that typically have costly experiments.

To scale BO for high-dimensional process optimization, we have chosen to apply TuRBO [7] to demonstrate the feasibility and efficacy of BO. TuRBO assembles the advantages of previous high-dimensional BO algorithms, notably replacing global model with multiple parallel and local models to avoid intractable time complexity and uneven posteriro uncertainty, and combines it with the use of trust regions (TR) [22, 23]. For TuRBO-m algorithm, m hyperrectangle TRs are held simultaneously, and the size of TRs can be dynamically adjusted based on experimental observation results to avoid over-exploration. The selection of experiments for subsequent batches are controlled by an implicit multi-armed bandit approach. b samples per batch are allocated unevenly into each TR, and at iteration t, the i^{th} new observation will be allocated to the most promising TR l if $x_i^{(t)} \in \arg\min_l(\arg\min_{x \in TR_l} f_l^{(i)})$ where $f_l^{(i)} \sim GP_l^t(\mu_l(x), k_l(x, x'))$.

In addition, we use BO in a multi-objective setting to perform multi-objective Bayesian Optimization (MOBO). Each objective will have an independent GP surrogate model, and hypervolume will be used as indicator for optimization success. We adapt EHVI from Yang et al. [11] as acquisition function that is not only fast but also naturally converge the experimental studies coverage of the true Pareto front [24, 25], which researchers can use to evaluate trade-offs between different objectives.

4 Results and Discussions

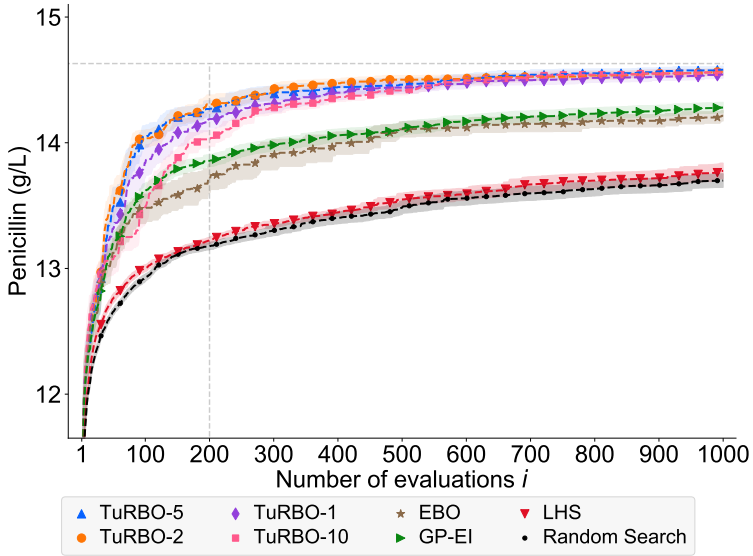

Figure 2: Performance of optimization algorithms on penicillin fermentation process. For each algorithm, initial sampling size and batch size are both 30. TuRBO-5 and TuRBO-2 stand out among all TuRBO algorithms, which all outperform EBO, regular BO methods, and random sampling baselines. Horizontal dash line shows the maximum penicillin concentration yield after sampling 10^6 evaluations from the penicillin model. Variation is visualized by plotting the median as well as shaded regions representing the 5^{th} to 95^{th} percentile of the aggregated 50 random run ensembles.

We show the quantitatively evaluated performance of TuRBO and other algorithmic baselines in Figure 2. Due to limited experimental budgets in pharmaceutical manufacturing, we would suggest practitioners to examine performance in the first 200 evaluations. We observe TuRBO algorithms considerably outperform EBO and regular BO, which all clearly show acceleration of process optimization over random sampling baselines. The improvement of holding multiple local models over a global surrogate model can be observed by comparing the performance of TurBO to that of regular BO. EBO uses an ensemble of additive GPs together with a batch acquisition function to scale BO to high-dimensions. The design space of penicillin production is only 7-dimensional, and thus does not clearly show the advantage of EBO over regular BO that utilizes GP as surrogate model and EI as acquisition function. We also observe that TuRBO with TR = 5 and TR = 2 are more effective than that of TR = 10. Depending on the batch size, having too many trust regions would allocate too few samples across each, resulting in inaccurate local surrogate model and subpar optimization performance. The number of TR should scale reasonably with available batch size per iteration.

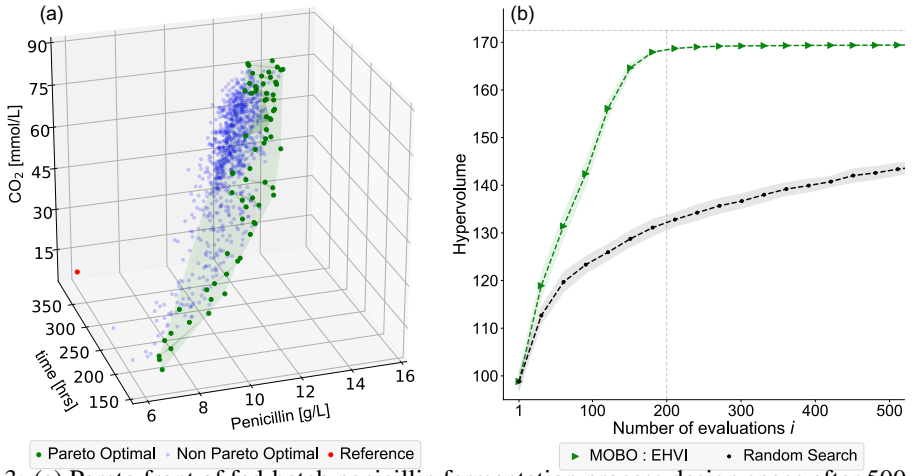

Figure 3: (a) Pareto front of fed-batch penicillin fermentation process design space after 500 evaluations. The three objectives of interest are final penicillin concentration yield, time, and CO_2 emission. Pareto optimal and pareto front are shown as green, non pareto optimal points as blue, and arbitrary reference point as red. (b) Performance of multi-objective optimization (w/ EHVI) on fed-batch penicillin fermentation process shown as normalized hypervolume vs. number of evaluations. Initial sampling size and batch size are both 30. Horizontal dash line shows the maximum normalized hypervolume after sampling 10^6 evaluations. Variation is visualized by plotting the median as well as shaded regions representing the 5^{th} to 95^{th} percentile of the aggregated 30 random run ensembles.

In Figure 3, we show the pareto front and hypervolume improvement at each evaluation when performing multi-objective optimization using BO equipped with GP surrogate model and EHVI acquisition function. Compared to the performance of random search, we clearly observe the efficacy of multi-objective BO in converging towards pareto front and expanding hypervolume in the first 200 evaluations. We see that experiments with highest penicillin yield are not necessarily the most optimal when other experiments can show high penicillin yield but take significant less time to ferment and result in lower CO_2 emissions. While the exact trade-off criteria and chosen optimal process control parameters will be decided by practitioners, we hope to demonstrate how BO can accelerate process optimization and offer comparison of trade-offs in a multi-objective optimization setting.

5 Conclusions

We demonstrated the efficacy of BO algorithms on optimizing yield of a fed-batch penicillin production process, which serves as a representative case for many more pharmaceutical manufacturing processes. We have applied trust region BO (TuRBO) to address the challenges of high-dimensional optimization for BO and observe its advantage over other BO algorithms and random baselines. We observe that a suitable number of trusted regions corresponding to designated batch size most be chosen for best optimization efficacy. In addition, we show the efficacy of BO in combination with multi-objective optimization, which allowed us to evaluate trade-offs between additional objectives such as reaction time, and CO_2 besides penicillin yield with fewer experiments. Our observations indicate that variations of BO are suitable tools for optimization drug production processes in high-dimensions. We hope to share our insights with the field, and popularize the use of BO and machine learning in future drug production processes.

6 Broader Impact

Successful benchmarks of scalable BO algorithms for process optimization of Penicillin production show us an example of how machine learning algorithms can help pharmaceutical and potentially broader industrial areas. Break-through in multi-objective optimization algorithms in high dimensions will have direct impact on the efficiency of those processes. Considering usually the high cost to validate proposed conditions, when combined with efficient experimental technologies (automated labs for example), the optimization algorithms will bring huge economical values to those fields. In addition, successful optimization algorithms can also be extended to broader areas, which share similar characteristics with the Penicillin production we discussed here, such as multiple constraints (e.g. budget, EHS (Environment, Health, and Safety), etc.), and large number of control parameters. We hope our work will not only motivate algorithm development and adoption across broader scientific communities, but also promote the integrated applications of algorithms and cutting edge experimental technologies in broader industrial scenarios.

References

[1] Peter I Frazier and Jialei Wang. Bayesian optimization for materials design. In *Information Science for Materials Discovery and Design*, pages 45–75. Springer, 2016.

[2] Peter I Frazier. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

[3] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.

[4] Jost Tobias Springenberg, Aaron Klein, Stefan Falkner, and Frank Hutter. Bayesian optimization with robust bayesian neural networks. *Advances in neural information processing systems*, 29:4134–4142, 2016.

[5] Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.

[6] David E Graff, Eugene I Shakhnovich, and Connor W Coley. Accelerating high-throughput virtual screening through molecular pool-based active learning. *Chemical Science*, 2021.

[7] David Eriksson, Michael Pearce, Jacob Gardner, Ryan D Turner, and Matthias Poloczek. Scalable global optimization via local Bayesian optimization. In *Advances in Neural Information Processing Systems*, pages 5496–5507, 2019.

[8] Carl Edward Rasmussen and Hannes Nickisch. Gaussian processes for machine learning (gpml) toolbox. *The Journal of Machine Learning Research*, 11:3011–3015, 2010.

[9] Gülnur Birol, Cenk Ündey, and Ali Cinar. A modular simulation package for fed-batch fermentation: penicillin production. *Computers & chemical engineering*, 26(11):1553–1565, 2002.

[10] Zi Wang, Clement Gehring, Pushmeet Kohli, and Stefanie Jegelka. Batched large-scale bayesian optimization in high-dimensional spaces. In *International Conference on Artificial Intelligence and Statistics*, pages 745–754. PMLR, 2018.

[11] Kaifeng Yang, Michael Emmerich, André Deutz, and Thomas Bäck. Efficient computation of expected hypervolume improvement using box decomposition algorithms. *Journal of Global Optimization*, 75(1):3–34, 2019.

[12] Qiaohao Liang. Scalable bayesian optimization accelerates process optimization of penicillin production. `https://github.com/HarryQL/TuRBO-Penicillin`, 2021.

[13] Jasdeep Mandur and Hector Budman. Robust optimization of chemical processes using bayesian description of parametric uncertainty. *Journal of Process Control*, 24(2):422–430, 2014.

[14] Ziyu Wang, Frank Hutter, Masrour Zoghi, David Matheson, and Nando de Feitas. Bayesian optimization in a billion dimensions via random embeddings. *Journal of Artificial Intelligence Research*, 55:361–387, 2016.

[15] Bo Chen, Rui Castro, and Andreas Krause. Joint optimization and variable selection of high-dimensional gaussian processes. *arXiv preprint arXiv:1206.6396*, 2012.

[16] Kirthevasan Kandasamy, Jeff Schneider, and Barnabás Póczos. High dimensional bayesian optimisation and bandits via additive models. In *International conference on machine learning*, pages 295–304. PMLR, 2015.

[17] Paul Rolland, Jonathan Scarlett, Ilija Bogunovic, and Volkan Cevher. High-dimensional bayesian optimization via additive models with overlapping groups. In *International conference on artificial intelligence and statistics*, pages 298–307. PMLR, 2018.

[18] Mojmír Mutnỳ and Andreas Krause. Efficient high dimensional bayesian optimization with additivity and quadrature fourier features. *Advances in Neural Information Processing Systems 31*, pages 9005–9016, 2019.

[19] Amin Nayebi, Alexander Munteanu, and Matthias Poloczek. A framework for bayesian optimization in embedded subspaces. In *International Conference on Machine Learning*, pages 4752–4761. PMLR, 2019.

[20] Roman Garnett, Michael A Osborne, and Philipp Hennig. Active learning of linear embeddings for gaussian processes. *arXiv preprint arXiv:1310.6740*, 2013.

[21] Yaochu Jin and Jürgen Branke. Evolutionary optimization in uncertain environments-a survey. *IEEE Transactions on evolutionary computation*, 9(3):303–317, 2005.

[22] Naijia Anna Dong, David J Eckman, Xueqi Zhao, Shane G Henderson, and Matthias Poloczek. Empirically comparing the finite-time performance of simulation-optimization algorithms. In *2017 Winter Simulation Conference (WSC)*, pages 2206–2217. IEEE, 2017.

[23] Ya-xiang Yuan. A review of trust region algorithms for optimization. In *Iciam*, volume 99, pages 271–282, 2000.

[24] Kaifeng Yang, Andre Deutz, Zhiwei Yang, Thomas Back, and Michael Emmerich. Truncated expected hypervolume improvement: Exact computation and application. In *2016 IEEE Congress on Evolutionary Computation (CEC)*, pages 4350–4357. IEEE, 2016.

[25] Ivo Couckuyt, Dirk Deschrijver, and Tom Dhaene. Fast calculation of multiobjective probability of improvement and expected improvement criteria for pareto optimization. *Journal of Global Optimization*, 60(3):575–594, 2014.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [No]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] https://github.com/HarryQL/TuRBO-Penicillin
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]