Probing the Structure of String Theory Vacua with Genetic Algorithms and Reinforcement Learning

Alex Cole University of Amsterdam a.e.cole@uva.nl Sven Krippendorf Arnold Sommerfeld Center for Theoretical Physics LMU Munich sven.krippendorf@physik.uni-muenchen.de

Andreas Schachner Centre for Mathematical Sciences University of Cambridge as2673@cam.ac.uk Gary Shiu University of Wisconsin-Madison shiu@physics.wisc.edu

Abstract

Identifying string theory vacua with desired physical properties at low energies requires searching through high-dimensional solution spaces – collectively referred to as the string landscape. We highlight that this search problem is amenable to reinforcement learning and genetic algorithms. In the context of flux vacua, we are able to reveal novel features (suggesting previously unidentified symmetries) in the string theory solutions required for properties such as the string coupling. In order to identify these features robustly, we combine results from both search methods, which we argue is imperative for reducing sampling bias.

1 Motivation

While the set of solutions arising in string theory – known as the string landscape [1, 2] – encompasses a vast number of low energy effective field theories, only a small subset of them are phenomenologically viable. The landscape's enormity, with as many as 10^{500} [3, 4] to $10^{272,000}$ [5] states, makes exhaustive searches or random sampling impractical. While in simple string models *generic* solutions may be found easily, practitioners usually seek to construct solutions having special properties such as weak couplings, ensuring theoretical control, or similarity to our universe. These additional criteria impose an auxiliary landscape structure (i.e. optimization target) for which constrained systems of equations must be solved. The existence of a solution with precisely specified properties is not guaranteed – indeed, the discreteness of UV ingredients generally leads to a rich topological structure of voids and voxels in the landscape's low-energy properties [4, 6]. For this problem, the landscape's vastness is exacerbated by computational complexity; finding specific string vacua appears to be NP-hard [7–9], though for some toy examples heuristic algorithms may have success [10]. On this level, the situation of the string landscape is comparable to spin-glass systems [11–13] and protein folding [14–17].

It remains a top challenge to design efficient optimization methods to search for realistic vacua that may at the same time reveal hidden structure in the landscape. Regarding this latter point, string theorists are interested not only in the construction of realistic vacua but also the statistical distribution surrounding such vacua, which has implications for various dynamics-based proposals for the measure problem [8, 10, 18]. In recent years, stochastic optimization with Genetic Algorithms (GAs) [19–23] and Reinforcement Learning (RL) [24–27] have been utilized to search for viable string vacua, outperforming searches based on Metropolis-Hastings. The purpose of the present work is to further develop these approaches and compare their results. We consider a common search task

Fourth Workshop on Machine Learning and the Physical Sciences (NeurIPS 2021).

and apply dimensionality reduction via Principal Component Analysis (PCA) to the results of both algorithms. We identify two disconnected clusters in the neighborhood of the optimality condition, which are suggestive of a previously unknown symmetry in the distribution of vacua. We argue that comparing the results of qualitatively different search algorithms is crucial to ensuring the robustness of this structure, effectively reducing sampling bias.

2 Setup for string theory solutions

The data structure of the string landscape is characterized by a set of integer inputs, from which a set of continuous parameters are outputted by minimizing a potential energy function. Within the full landscape, we focus on flux vacua where the integer inputs are quantized fluxes $\vec{N} = (N_1, \ldots, N_k)$ which are subject to Gauss's law constraints (i.e., N_i , $i = 1, \ldots, k$ are bounded). The output is a vector of complex-valued fields $\vec{\phi} = (\phi_1, \ldots, \phi_m)$ obtained from minimizing a highly non-linear scalar potential $V(\vec{\phi}, \vec{N})$ whose functional form is discussed in [22, 26], see also e.g. [28, 29] for reviews. The number k of fluxes is related to the number m of scalar fields via k = 4m. In this note, we consider the 'conifold limit' in the geometry as defined in Sec. (5.4) of [30] with k = 8, which fixes the scalar potential V. Our methods generalize to other background geometries, which we plan to study in the future.

The set of minima $\{\vec{\phi}\}\$ of V defines the flux landscape. Its elements (flux vacua) have varied physical properties (in the language of GAs, different phenotypes). We focus on the string coupling g_s (parametrizing the way strings interact) and the superpotential W_0 (setting the scale of masses). Our task is solving the inverse problem: given certain physical properties, i.e., values of (g_s, W_0) , what is the required input of quantized fluxes \vec{N} ? Due to the involved non-linearities, this problem cannot be addressed with closed form solutions and appropriate numerical search strategies have to be applied. In this work, we compare the efficiency of GAs and RL in solving this inverse problem, thereby aiming to address the following questions:

- 1. Can RL and stochastic optimization such as GAs efficiently search for desirable vacua?
- 2. Can these methods discover structure in the landscape?
- 3. How is it possible to understand such structure?

3 The search algorithms

We apply a generalized version of the GA described in [22] for generating flux vacua which consists of mainly three steps. We begin with a set of flux vectors $\{\vec{N}_1, \ldots, \vec{N}_p\}$, called a population of size p, with $\vec{N}_i \in [-30, 30]^8$. The GA selects individuals based on a fitness function, and individuals are in turn used to construct new flux vectors \vec{N}_l by performing crossover operations acting on pairs of vectors (\vec{N}_i, \vec{N}_j) . The third and arguably most important step concerns mutation, which alters the vectors \vec{N}_l by some randomized procedure. Lastly, one defines a novel population and repeats the above process. Each individual step involves a choice of operators, see e.g. [31, 32] for a comprehensive list. In our implementation, we utilize eight selection, nine crossover and nine mutation operators which, together with mutation, cloning and survival¹ rate, amounts to a total of 29 hyperparameters. They are determined for each task separately through sample runs for mini-batches and randomly chosen hyperparameters by extracting their correlations with the mean average distance of the final output to the optimal solution. This is a novel and powerful approach to optimizing hyperparameters for GAs that is easily generalized to other applications.

We compare the GA with the A3C RL implementation for flux vacua presented in [26] where it was found that A3C resulted in the best performance across several different RL algorithms. Starting from a random initial configuration $\vec{N_i} \in [-30, 30]^8$ the RL agent can increase or decrease in a single step one of the k entries by one, corresponding to a 16 dimensional action space. To identify models, the RL agent takes up to 500 steps or is reset. The policy network, consisting of four hidden dense layers,

¹Generically, it turns out to be useful to carry over a certain number of the fittest individuals (cloning) or replace the least fit individuals by fitter ones (survival of the fittest).



Figure 1: Left: Cluster structure in dimensionally reduced flux samples for RL and 25 GA runs (PCA on all samples of GA and RL). The colors indicate individual GA runs. Right: Dependence on flux (input) values (N_1 and N_5 respectively) in relation to principal components for a PCA fit of the individual output of GA and RL.

is optimized based on a reward function which balances the string theory consistency conditions (e.g. Gauss's law constraint) and the phenomenological requirements associated to the scalar potential, the string coupling, and the superpotential.

4 Results

For illustrative purposes we chose a target value of $|W_0| = 50,000$ which guarantees that many consistent flux vectors can be identified within reasonable effort. Both RL and GA approaches are found to be extremely efficient in identifying distinct string vacua with desirable properties, with RL producing for an arbitrary starting point an acceptable solution in 193 steps on average and respectively final GA populations display distinct and desired properties across the population after generating 20 new populations. For our GA samples it is most efficient to work with population sizes of 500 using 25 runs. The optimized choice of hyperparameters for this particular task led to a non-trivial combination of two selection, three crossover and four mutation operators.² Our A3C agent was trained for 15,000,000 episodes where the agent can take up to 500 steps before being reset to a random starting point. In both cases we utilize 10,000 samples for our analysis satisfying $|W_0| = 50,000 \pm 1,000.^3$

Being able to efficiently generate samples, we now report on how we can determine and analyze the structure of the set of 8-dimensional flux vectors. This is imperative to develop an understanding of the string landscape and to extract constraints for UV parameters, here our flux vectors, associated with phenomenological properties, i.e., $|W_0| = 50,000$ for our search. To this end, we perform a PCA on the combined GA and RL samples which is shown on the left of Fig. 1. Uniformly across GA and RL samples we observe two clusters which provides a non-trivial cross-check of this structure. This is remarkable as the sampling dynamics associated to GA and RL are different and we hence consider these features as less likely to be artifact of our sampling procedure. We checked that equivalent structures emerge when performing t-sne [33]. We note that the samples generated from outputs of different GA runs show clear variations. They can be traced back to the way GA operations produce new flux vectors.

The presence of two clusters with uniformly distributed phenomenological properties (superpotential $|W_0|$ and string coupling g_s) suggests the emergence of a discrete symmetry whose exact understanding provides an exciting future direction. Analytically the structure within the clusters can be characterized with scalings of flux components as shown on the right of Fig. 1 where a PCA was performed on the GA and RL samples individually. While the results are largely equal for N_5 , the scaling

²To be more precise, we used tournament and rank selection, random k-point, uniform and whole arithmetic recombination crossover as well as random, random k-point, insertion and displacement mutation.

³Our samples are gauge inequivalent under $SL(2, \mathbb{Z})$ which is enforced through explicit gauge fixing in the GA [6, 22] or through a suitable reward structure in RL [26].



Figure 2: Correlation of flux parameters for vacua with $|W_0| = 50,000 \pm 1,000$. We use both GA and RL samples.

with respect to N_1 highlights distinguishing features between the GA and RL samples. We find that this structure on the "consistent" flux vectors highly depends on the imposed phenomenological properties, i.e., when sampling different superpotential values the structure changes.

For the PCA analysis we observe the variance ratios for the respective PCA components PC_i listed in Tab. 1. As this low-variance of the higher PCA components suggests, the sample is effectively lower dimensional. This can be understood further by checking the correlations among the flux vacua. Similar to [26], we compare them for both search strategies which is capable of highlighting universal structures specific to a given class of vacua.

Table 1:	Variance	of com	ponents	in	PCA
----------	----------	--------	---------	----	-----

Data	PC_1	$ PC_2 $	PC ₃	PC_4	$PC_5 - PC_8$
GA+RL	0.52	0.22	0.13	0.1	<0.03
GA	0.60	0.19	0.13	0.05	<0.03
RL	0.68	0.13	0.09	0.08	<0.02

Combining samples from both search algorithms before computing correlations allows us to reduce sampling bias. In Fig. 2, we show the combined correlation maps for the two methods. While certain aspects agree in the individual GA and RL correlations, variations suggest that we ought to consider the combined data set. In particular, the combined correlation maps has strong correlations which are also shared with the individual correlation maps from GA and RL. This includes the correlations of (N_2, N_8) , (N_3, N_7) ,

 (N_5, N_7) and (N_3, N_5) , but, in both GA and RL, we also observe different correlations not present in the respective other sample. We note that the flux pairs (N_2, N_8) and (N_3, N_5) appear as products in the D3-charge contribution to the tadpole. In addition, the dominant PC directions tell us how to move in flux space to find flux vectors with similar physical properties, providing a new direction on studying the statistical properties of this landscape in a dimensionally reduced subspace.

5 Perspective

We have demonstrated that by utilizing both RL and GA, it is possible to identify novel structures in the string landscape. To reduce sampling bias and identify these structures robustly, we found it more beneficial to use multiple search strategies. At this stage we cannot single out one method as being preferred. Overall, it should be stressed that the performance is subject to hyperparameters. Here we added and optimized hyperparameter searches in the GA analysis extending previous studies in [22] and used RL hyperparameters which emerged from a hyperparameter search in [26].

Optimization becomes very difficult when the reward structure is sparse. We have considered an example where the optimization task is feasible. For model building purposes, it is interesting to consider questions of fine-tuning and rare solutions. For instance, certain string theory models (e.g.

[34]) assume the existence of rare vacua with $|W_0| \ll 1$. It will be interesting to compare GA and RL strategies to those developed by human string phenomenologists in the search for these vacua (e.g. [35]). It is important to stress that our sampling techniques in this example clearly demonstrate that lower dimensional subsets of the landscape are no obstacle, even though they generally populate a measure zero subset of the state space. The true challenge is presented by the local structure near the global optimum which is generally unknown a priori. We plan to utilize our current methods to explore the local structure near rare string theory solutions obtained in the literature.

Revealing the structure behind and constraints on string theory solutions has a significant potential for our understanding of string theory. Potentially, such constraints could rule out phenomenologically viable combinations of low-energy parameters when no vacua can be constructed in such a region. Hints of such constraints are clearly visible in the current analysis. However, our analysis also shows that – at this stage – numerical samples have to be considered carefully, as we observe variations across different methods (cf. Fig. 1). When correlations can be identified robustly, they can be used to efficiently initialize sampling strategies (e.g. proposal densities) for vacua with shared properties. Explicitly constructing the symmetries identified by search algorithms presents another interesting avenue of future work.

Finally, we look forward to investigating how these early results can be transferred to other more complex backgrounds and different phenomenological requirements.

Acknowledgments and Disclosure of Funding

We thank the authors and maintainers of Jupyter [36], Matplotlib [37], NumPy [38], Seaborn [39], Pandas [40], Openai gym [41], Scikit-learn [42] and Chainer-rl [43]. SK would like to thank R. Kroepsch and M. Syvaeri for helpful discussions. AS acknowledges support by the German Academic Scholarship Foundation and by DAMTP through an STFC studentship. GS is supported in part by the DOE grant DE-SC0017647.

References

- Raphael Bousso and Joseph Polchinski. Quantization of four form fluxes and dynamical neutralization of the cosmological constant. *JHEP*, 06:006, 2000. doi: 10.1088/1126-6708/ 2000/06/006.
- [2] Leonard Susskind. The Anthropic landscape of string theory. pages 247–266, 2 2003.
- [3] Sujay Ashok and Michael R. Douglas. Counting flux vacua. *JHEP*, 01:060, 2004. doi: 10.1088/1126-6708/2004/01/060.
- [4] Frederik Denef and Michael R. Douglas. Distributions of flux vacua. JHEP, 05:072, 2004. doi: 10.1088/1126-6708/2004/05/072.
- [5] Washington Taylor and Yi-Nan Wang. The F-theory geometry with most flux vacua. JHEP, 12: 164, 2015. doi: 10.1007/JHEP12(2015)164.
- [6] Alex Cole and Gary Shiu. Topological Data Analysis for the String Landscape. JHEP, 03:054, 2019. doi: 10.1007/JHEP03(2019)054.
- [7] Frederik Denef and Michael R. Douglas. Computational complexity of the landscape. I. *Annals Phys.*, 322:1096–1142, 2007. doi: 10.1016/j.aop.2006.07.013.
- [8] Frederik Denef, Michael R. Douglas, Brian Greene, and Claire Zukowski. Computational complexity of the landscape II—Cosmological considerations. *Annals Phys.*, 392:93–127, 2018. doi: 10.1016/j.aop.2018.03.013.
- [9] James Halverson and Fabian Ruehle. Computational Complexity of Vacua and Near-Vacua in Field and String Theory. *Phys. Rev. D*, 99(4):046015, 2019. doi: 10.1103/PhysRevD.99.046015.
- [10] Ning Bao, Raphael Bousso, Stephen Jordan, and Brad Lackey. Fast optimization algorithms and the cosmological constant. *Phys. Rev. D*, 96(10):103512, 2017. doi: 10.1103/PhysRevD.96. 103512.
- [11] David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Physical review letters*, 35(26):1792, 1975.

- [12] Francisco Barahona. On the computational complexity of ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10):3241, 1982.
- [13] Marc Mézard, Giorgio Parisi, and Miguel Angel Virasoro. Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications, volume 9. World Scientific Publishing Company, 1987.
- [14] Cyrus Levinthal. How to fold graciously. *Mossbauer spectroscopy in biological systems*, 67: 22–24, 1969.
- [15] Ron Unger and John Moult. Finding the lowest free energy conformation of a protein is an np-hard problem: proof and implications. *Bulletin of mathematical biology*, 55(6):1183–1198, 1993.
- [16] Joseph D Bryngelson, Jose Nelson Onuchic, Nicholas D Socci, and Peter G Wolynes. Funnels, pathways, and the energy landscape of protein folding: a synthesis. *Proteins: Structure, Function, and Bioinformatics*, 21(3):167–195, 1995.
- [17] Michael Socolich, Steve W Lockless, William P Russ, Heather Lee, Kevin H Gardner, and Rama Ranganathan. Evolutionary information for specifying a protein fold. *Nature*, 437(7058): 512–518, 2005.
- [18] Justin Khoury and Onkar Parrikar. Search Optimization, Funnel Topography, and Dynamical Criticality on the String Landscape. JCAP, 12:014, 2019. doi: 10.1088/1475-7516/2019/12/014.
- [19] Johan Blåbäck, Ulf Danielsson, and Giuseppe Dibitetto. Fully stable dS vacua from generalised fluxes. *JHEP*, 08:054, 2013. doi: 10.1007/JHEP08(2013)054.
- [20] Johan Blåbäck, Ulf Danielsson, and Giuseppe Dibitetto. Accelerated Universes from type IIA Compactifications. JCAP, 03:003, 2014. doi: 10.1088/1475-7516/2014/03/003.
- [21] Steven Abel and John Rizos. Genetic Algorithms and the Search for Viable String Vacua. JHEP, 08:010, 2014. doi: 10.1007/JHEP08(2014)010.
- [22] Alex Cole, Andreas Schachner, and Gary Shiu. Searching the Landscape of Flux Vacua with Genetic Algorithms. *JHEP*, 11:045, 2019. doi: 10.1007/JHEP11(2019)045.
- [23] S. AbdusSalam, S. Abel, M. Cicoli, F. Quevedo, and P. Shukla. A systematic approach to Kähler moduli stabilisation. *JHEP*, 08(08):047, 2020. doi: 10.1007/JHEP08(2020)047.
- [24] James Halverson, Brent Nelson, and Fabian Ruehle. Branes with Brains: Exploring String Vacua with Deep Reinforcement Learning. JHEP, 06:003, 2019. doi: 10.1007/JHEP06(2019)003.
- [25] Magdalena Larfors and Robin Schneider. Explore and Exploit with Heterotic Line Bundle Models. *Fortsch. Phys.*, 68(5):2000034, 2020. doi: 10.1002/prop.202000034.
- [26] Sven Krippendorf, Rene Kroepsch, and Marc Syvaeri. Revealing systematics in phenomenologically viable flux vacua with reinforcement learning. 7 2021.
- [27] Andrei Constantin, Thomas R. Harvey, and Andre Lukas. Heterotic String Model Building with Monad Bundles and Reinforcement Learning. 8 2021.
- [28] Michael R. Douglas and Shamit Kachru. Flux compactification. *Rev. Mod. Phys.*, 79:733–796, 2007. doi: 10.1103/RevModPhys.79.733.
- [29] Mariana Grana. Flux compactifications in string theory: A Comprehensive review. *Phys. Rept.*, 423:91–158, 2006. doi: 10.1016/j.physrep.2005.10.008.
- [30] Oliver DeWolfe, Alexander Giryavets, Shamit Kachru, and Washington Taylor. Enumerating flux vacua with enhanced symmetries. JHEP, 02:037, 2005. doi: 10.1088/1126-6708/2005/02/037.
- [31] Melanie Mitchell. L.d. davis, handbook of genetic algorithms. *Artificial Intelligence*, 100(1): 325–330, 1998. ISSN 0004-3702.
- [32] Fabian Ruehle. Data science applications to string theory. *Phys. Rept.*, 839:1–117, 2020. doi: 10.1016/j.physrep.2019.09.005.
- [33] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [34] Shamit Kachru, Renata Kallosh, Andrei D. Linde, and Sandip P. Trivedi. De Sitter vacua in string theory. *Phys. Rev. D*, 68:046005, 2003. doi: 10.1103/PhysRevD.68.046005.

- [35] Mehmet Demirtas, Manki Kim, Liam Mcallister, and Jakob Moritz. Vacua with Small Flux Superpotential. *Phys. Rev. Lett.*, 124(21):211603, 2020. doi: 10.1103/PhysRevLett.124.211603.
- [36] Thomas Kluyver, Benjamin Ragan-Kelley, Fernando Pérez, Brian E Granger, Matthias Bussonnier, Jonathan Frederic, Kyle Kelley, Jessica B Hamrick, Jason Grout, Sylvain Corlay, et al. *Jupyter Notebooks-a publishing format for reproducible computational workflows.*, volume 2016. 2016.
- [37] John D Hunter. Matplotlib: A 2d graphics environment. *Computing in science & engineering*, 9(03):90–95, 2007.
- [38] Stefan Van Der Walt, S Chris Colbert, and Gael Varoquaux. The numpy array: a structure for efficient numerical computation. *Computing in science & engineering*, 13(2):22–30, 2011.
- [39] Michael L Waskom. Seaborn: statistical data visualization. *Journal of Open Source Software*, 6 (60):3021, 2021.
- [40] Wes McKinney et al. pandas: a foundational python library for data analysis and statistics. *Python for high performance and scientific computing*, 14(9):1–9, 2011.
- [41] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint: 1606.01540*, 2016.
- [42] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikitlearn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- [43] Seiya Tokui, Kenta Oono, Shohei Hido, and Justin Clayton. Chainer: a next-generation open source framework for deep learning. In *Proceedings of workshop on machine learning systems* (*LearningSys*) in the twenty-ninth annual conference on neural information processing systems (NIPS), volume 5, pages 1–6, 2015.

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] We comment on the limitations in our final perspectives section.
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] This work has no negative societal impacts as it is concerned with questions in theoretical high energy physics and cosmology.
 - (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? [No] Our results apply machine learning techniques on data arising in string theory. We refer to the relevant string theory literature where detailed derivations and underlying assumptions to these string theory constructions can be found.
 - (b) Did you include complete proofs of all theoretical results? [No] We have included references for all relevant theoretical results which are of relevance for our string theory models.
- 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)? [No] (Post-processed) results reproducing our analysis will be available on github at a later stage.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [No] We included the relevant details and references on our hyperparameters but we did not include ALL the training details in this short paper.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] There are several sources of errors and we comment in particular on the selection and sampling bias produced by our different methods (cf. Section on Results)
 - (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] Our code has been run on standard desktops with no extensive (cluster/GPU) computational resources.
- 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] We used the following packages: Jupyter [36], Matplotlib [37], NumPy [38], Seaborn [39], Pandas [40], Openai gym [41], Scikit-learn [42] and Chainer-r1 [43]. These references are provided in the acknowledgments of the final version.
 - (b) Did you mention the license of the assets? [No] We added the references to these packages where the detailed licence information can be found. If required, we can add this information in a revised version.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [No] We plan to make these assets available in the future when they can be run on more background geometries, i.e. are more widely applicable.
 - (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? [Yes] There is no such information or offensive content.
- 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? $[\rm N/A]$