

COMMENT OPEN



MODEL, GUESS, CHECK: Wordle as a primer on active learning for materials research

Keith A. Brown

Research and games both require the participant to make a series of choices. Active learning is a process borrowed from machine learning for algorithmically making choices that has become increasingly used to accelerate materials research. While this process may seem opaque to researchers outside the field of machine learning, examining active learning in games provides an accessible way to showcase the process and its virtues. Here, we examine active learning through the lens of the game Wordle to both explain the active learning process and describe the types of research questions that arise when using active learning for materials research.

npj Computational Materials (2022)8:97; <https://doi.org/10.1038/s41524-022-00787-7>

A challenge common across research, whether computational or experimental, is deciding which experiment to perform next. While many of these choices are made ad hoc by a researcher in front of an instrument, there is a community, with roots tracing back centuries, that studies how to optimally select experiments¹. Widely used strategies include one factor at a time (OFAT) optimization and selecting experiments evenly across the available parameter space (e.g., grid-based searching). In contrast with strategies that select all experiments before performing an experimental campaign, active learning is a branch of machine learning in which experiments are selected sequentially using knowledge gained from all prior experiments^{2,3}. Encouragingly, recent studies have reported that active learning outperforms both design of experiments that define a set of experiments at the outset and experiments guided by human intuition^{4,5}. Indeed, examples of active learning in materials research have blossomed in recent years^{6–11}. One area that has especially shown a spotlight on this trend is autonomous experimentation in which a robotic system is used to carry out experiments that are chosen by active learning¹². Such systems have emerged in a broad range of fields including mechanics^{13,14}, biology^{15,16}, chemistry^{17,18}, nanotechnology¹⁹, and microscopy^{20,21}. In light of these advances in autonomous experimentation, and the demonstrated advantage of active learning, it is increasingly clear that there are tremendous opportunities for applying concepts of active learning in a wide range of research fields.

Despite the benefits of adopting active learning in research, there are still hurdles associated with learning how to implement active learning and communicating the advantages to researchers far removed from the fields of machine learning or data-driven science. Interestingly, the introduction and rapid popularization of the word game Wordle produced by Josh Wardle provides a fascinating platform for overcoming both of these challenges. In Wordle, the player has six guesses to identify a specific five-letter word. After each guess, feedback is provided about whether each letter is present in the target word and whether the letter is in the correct location. Part of the popularity of this game stems from the fact that it can only be played once a day, which further raises the stakes of each guess. At a glance, this game shares some extremely salient connections to experimental selection during a

research campaign. First, a Wordle player can guess any five-letter word, meaning that there is a large, but finite, parameter space that is known ahead of time. This is often true in a research campaign where the knobs that can be turned are known, even if their importance has yet to be determined. Second, the budget of available experiments is limited—to six by the programmer in Wordle—and by the natural constraints of time and other resources in research. Finally, the goal of Wordle, to identify the target word, mirrors the goal of finding a maximum or minimum property in some parameter space, which is a common task in many research studies. Despite these similarities, the goals of Wordle and research often have meaningful differences. For example, a researcher may not know when they have found the best experiment, the goal of a research campaign is often to learn the behavior of some property rather than simply find its extrema, and there can be more than one best experiment in a parameter space. Those caveats aside, the task of optimization is extremely common and an obvious first step in any more complex research goal.

Given the similarities between Wordle and research campaigns, it is our hope that one can gain insight into active learning and motivation for adopting it more generally by studying active learning in Wordle. Taking inspiration from a Bayesian optimization formalization of active learning, we can lay out an iterative cycle for selecting and interpreting guesses (Fig. 1A). The first task is to define a surrogate model that encompasses our knowledge about the system. As part of this, we must define the parameter space, which for Wordle, amounts to identifying all valid five-letter words, of which we identified 12,478. To illustrate why active learning is needed to play Wordle, given a parameter space this large, selecting guesses uniformly at random from the available words means that the player will win only 0.05% of the time. In a materials research campaign, defining the parameter space is akin to identifying all available materials and the valid ranges of experimental or computational parameters that can be tuned. Once we have defined the parameter space, we need to build a surrogate model of our belief about the space. In Wordle, this can be the belief of whether a given word is possibly correct, which is initially uniform for all available words. This surrogate model is implemented as a look up table in which each word is assigned a

¹Department of Mechanical Engineering, Physics Department, and Division of Materials Science & Engineering, Boston University, Boston, MA, USA. ✉email: brownka@bu.edu

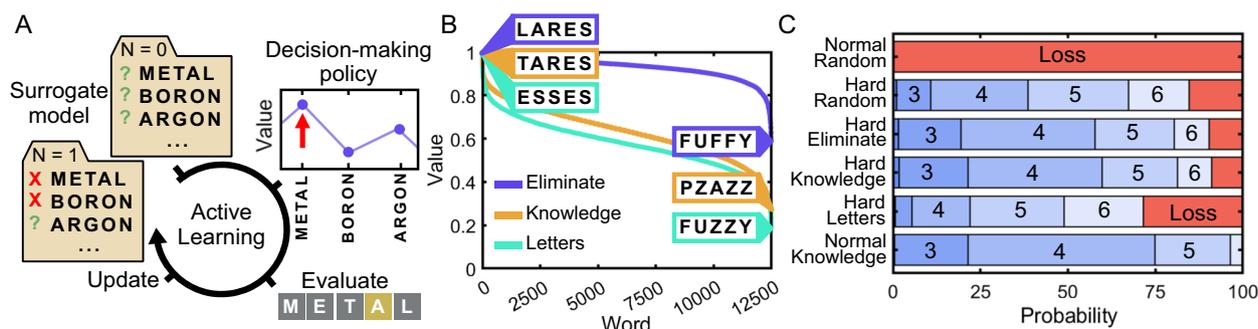


Fig. 1 **Playing Wordle using active learning.** **A** Diagram showing the anatomy of active learning in a game of Wordle. A surrogate model describes the present belief about the system. During the first turn ($N = 0$), all words are assumed to be candidates for the winning word. Next, a decision-making policy is employed to quantify the relative value of each word. Subsequently, the highest value word is guessed and the response is used to update the surrogate model. This process continues until the correct word is found or the player loses. **B** Ranked order evaluation of three decision-making policies for all words. Decision policies include 'Eliminate' in which each word's value is estimated by the number of words that could be eliminated if it was guessed, 'Knowledge' in which the value of each word is determined by the number of unique outcomes that could arise from guessing the word, and 'Letters' in which the value of each word is based on the abundance of its letters in the catalog of words. **C** Probability of different outcomes from playing Wordle guided by different strategies. Color indicates the number of guesses required to find the correct answer with red indicating that seven or more guesses were required. Strategies are either 'Hard' in that they require that words are selected from those that have not been ruled out or 'Normal' in that all words can be selected. 'Random' indicates that words were selected uniformly at random. The code used to produce these results can be found at <https://github.com/kabrownlab/wordle>.

probability that it is the correct word. If a given word is ruled out then its probability is set to 0 and the rest of the space is renormalized. More complex and versatile surrogate models, such as Gaussian process regressions, can be constructed that leverage the idea that points close to one another in parameter space should behave similarly, but a simple look up table is sufficient for illustrating the process. In either case, this surrogate model is updated as experiments are performed, which in Wordle amounts to ruling out words that are inconsistent with the responses to prior guesses. To show the value of iteratively incorporating this knowledge, if a player randomly selects words from this ever narrowing field of possibilities, they will win Wordle 85% of the time. This striking shift in outcomes illustrates the value of selecting each experiment using all available knowledge.

Randomly guessing from the words that remain in contention takes advantage of the information from the results of previous guesses, but not the information embodied by the parameter space itself. A major goal in active learning is determining the expected value of a given guess. In a materials research campaign, this could be related to a materials property, such as its Seebeck coefficient or fracture toughness. In Wordle, the value of a given word is related to how much it helps identify the correct word. Given this, there are many different ways to value a potential guess (Fig. 1B). Three examples include (1) assigning a score to each word based on how common the letters in the word are, (2) prioritizing words that allow the player to eliminate the most words, or (3) targeting words that lead to the most possible outcomes and therefore the most unique information. Each choice represents a decision-making policy that assigns a value to any given point in parameter space based on the current state of the surrogate model. Here, we find that selecting words based on how many possible words they eliminate or how much information they provide leads to victory >90% of the time, showing a ~5% improvement over randomly selecting from the available words. However, choosing based on the commonality of letters actually performs worse than randomly selecting from potential winners, highlighting the importance of choosing a decision-making policy wisely. As shown by this example, the benefits and risks associated with choosing different decision-making policies is both highly impactful on the pace of learning and a fascinating lens through which to view materials research.

Part of the value of active learning is that it opens the door for the algorithm to make choices based on relationships inside the data that might be difficult for human users to intuit. Conceptually, this can amount to selecting experiments that themselves are not likely to lead to high performance but improve chances for future success. To apply this concept in Wordle, consider choosing between two words as potential guesses: word 1 has some probability of being a winning word but would provide little actionable information if it is not the target word. In contrast, word 2 has already been ruled out from being a winning word, but the response from this word would on average greatly increase the player's chance of success by providing actionable information such as ruling out other words. This latter approach, namely selecting the words with the most possible outcomes even among those that have been ruled out produces success a remarkable 99.7% of the time. This interesting result can be conceptualized as coming from the dichotomy between exploration and exploitation in acknowledging that if your budget is known, which of course in Wordle it is, it is beneficial to spend earlier guesses exploring parameter space before focusing on regions in which success is expected. In active learning, there are decision-making policies that naturally balance exploration and exploitation such as the expected improvement policy that selects points in parameter space that are judged to be most likely to increase, in the case of maximization, the current value of the maximum.

While analyzing the algorithmic process of sequentially selecting guesses for Wordle provides insight into the process for materials development, it also raises a key shortcoming of purely algorithmic active learning. The discussed algorithms assume that all words initially have an equal chance of being correct, which discounts the bias against esoteric words in the target word selection by the game's creator. There are two ways one could envision trying to take advantage of this information. First, one could imagine including information about word popularity as prior knowledge in the active learning loop. This could amount to, for example, weighting the initial probability of each word in proportion to how often it is used in literature. It is also worth emphasizing that prior knowledge has already been introduced by only considering combinations of five-letter words that are English words. Second, this could also be dynamically addressed using human-in-the-loop (HITL) active learning in which the human-machine partnership is leveraged to further accelerate the

learning process. Generally, HITL entails finding a productive way to combine the best attributes of each member of the partnership. This has been productively used to outperform algorithms or humans alone in fields including radiology²² and robotics²³. In the present example, HITL could entail using the algorithm to identify a short list of high value words and allow the human to select from these words as a way of considering the relative popularity of each word. Advances in this area are particularly interesting in materials research where the insight of researchers may be difficult to quantify but could be productively employed through such a partnership.

Formalizing the selection of choices, whether in games or materials research, forces one to define and consider the goals and important information present in a system. Because much materials research still relies on relatively simple heuristics for experimental selection, it is our hope that this comment, and the broader efforts to introduce active learning into materials research that it reflects, will spark the curiosity of new researchers and engage them in the process of looking more deeply at how experiments are selected. Such interest can be complemented by easy-to-implement active learning packages in a variety of programming languages and high quality tutorial articles^{24,25}. Along these lines, the code used to generate the results presented in this work are posted at <https://github.com/kabrownlab/wordle>. Perhaps most importantly, the incorporation of active learning elevates the conversation and thoughts in the research enterprise. For example, rather than thinking about which word to select in Wordle, instead the player thinks about what defines a word's value. This level of discourse centered around selecting and refining decision-making policies and surrogate models is an exciting prospect for the materials community.

DATA AVAILABILITY

The code used to produce the data in this paper can be found at <https://github.com/kabrownlab/wordle>.

Received: 19 February 2022; Accepted: 13 April 2022;

Published online: 02 May 2022

REFERENCES

- Hicks, C. R. Fundamental concepts in the design of experiments. (Holt, New York, 1964).
- Cohn, D. A., Ghahramani, Z. & Jordan, M. I. Active learning with statistical models. *J. Artif. Intell. Res.* **4**, 129–145 (1996).
- Pereira-Santos, D., Prudêncio, R. B. C. & de Carvalho, A. C. Empirical investigation of active learning strategies. *Neurocomputing* **326**, 15–27 (2019).
- Liang, Q. et al. Benchmarking the performance of Bayesian optimization across multiple experimental materials science domains. *npj Comput. Mater.* **7**, 188 (2021).
- Chen, D., Skouras, M., Zhu, B. & Matusik, W. Computational discovery of extremal microstructure families. *Sci. Adv.* **4**, eaao7005 (2018).
- Lookman, T., Balachandran, P. V., Xue, D. & Yuan, R. Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. *npj Comput. Mater.* **5**, 1–17 (2019).
- Bassman, L. et al. Active learning for accelerated design of layered materials. *npj Comput. Mater.* **4**, 74 (2018).
- Tian, Y. et al. Role of uncertainty estimation in accelerating materials development via active learning. *J. Appl. Phys.* **128**, 014103 (2020).
- Tran, K. & Ulissi, Z. W. Active learning across intermetallics to guide discovery of electrocatalysts for CO₂ reduction and H₂ evolution. *Nat. Catal.* **1**, 696 (2018).
- Kusne, A. G. et al. On-the-fly closed-loop materials discovery via Bayesian active learning. *Nat. Commun.* **11**, 1–11 (2020).
- Sun, S. et al. Accelerated development of perovskite-inspired materials via high-throughput synthesis and machine-learning diagnosis. *Joule* **3**, 1437–1451 (2019).

- Stach, E. et al. Autonomous experimentation systems for materials development: a community perspective. *Matter* **4**, 2702–2726 (2021).
- Gongora, A. E. et al. Using simulation to accelerate autonomous experimentation: a case study using mechanics. *iScience* **24**, 102262 (2021).
- Gongora, A. E. et al. A Bayesian experimental autonomous researcher for mechanical design. *Sci. Adv.* **6**, eaaz1708 (2020).
- King, R. D. et al. Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature* **427**, 247 (2004).
- King, R. D. et al. The automation of science. *Science* **324**, 85–89 (2009).
- Coley, C. W. et al. A robotic platform for flow synthesis of organic compounds informed by AI planning. *Science* **365**, eaax1566 (2019).
- Häse, F., Roch, L. M. & Aspuru-Guzik, A. Next-generation experimentation with self-driving laboratories. *Trends Chem.* **1**, 282–291 (2019).
- Nikolaev, P. et al. Autonomy in materials research: a case study in carbon nanotube growth. *npj Comput. Mater.* **2**, 16031 (2016).
- Ziatdinov, M.A. et al. Hypothesis learning in automated experiment: application to combinatorial materials libraries. *Adv. Mater. Early View*, 2201345 (2022).
- Roccapriore, K. M., Dyck, O., Oxley, M. P., Ziatdinov, M. & Kalinin, S. V. Automated experiment in 4D-STEM: exploring emergent physics and structural behaviors. Preprint at <https://doi.org/10.48550/arXiv.2112.04479> (2021).
- Patel, B. N. et al. Human-machine partnership with artificial intelligence for chest radiograph diagnosis. *npj Digit. Med.* **2**, 1–10 (2019).
- Zhang, J. et al. Human-in-the-loop optimization of exoskeleton assistance during walking. *Science* **356**, 1280–1284 (2017).
- Frazier, P. I. A tutorial on Bayesian optimization. Preprint at <https://doi.org/10.48550/arXiv.1807.02811> (2018).
- Astudillo, R. & Frazier, P. I. Thinking inside the box: a tutorial on grey-box Bayesian optimization. In *2021 Winter Simulation Conference (WSC) 1-15 (2021)* Preprint at <https://doi.org/10.48550/arXiv.2201.00272> (2022).

ACKNOWLEDGEMENTS

We acknowledge fruitful discussions with Sam Ribnick, Alan Gardner, and Kelsey Snapp.

AUTHOR CONTRIBUTIONS

K.A.B. conceptualized this work, performed the simulations, and wrote the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Correspondence and requests for materials should be addressed to Keith A. Brown.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022