

## Monte Carlo tree search for materials design and discovery

**Thaer M. Dieb**, National Institute for Materials Science, Tsukuba, Japan; Graduate School of Frontier Sciences, the University of Tokyo, Kashiwa, Japan; RIKEN, AIP, Tokyo, Japan

**Shenghong Ju**, National Institute for Materials Science, Tsukuba, Japan; Department of Mechanical Engineering, the University of Tokyo, Tokyo, Japan  
**Junichiro Shiomi**, National Institute for Materials Science, Tsukuba, Japan; Department of Mechanical Engineering, the University of Tokyo, Tokyo, Japan; CREST, JST, Tokyo, Japan

**Koji Tsuda**, National Institute for Materials Science, Tsukuba, Japan; Graduate School of Frontier Sciences, the University of Tokyo, Kashiwa, Japan; RIKEN, AIP, Tokyo, Japan

Address all correspondence to Koji Tsuda at [tsuda@k.u-tokyo.ac.jp](mailto:tsuda@k.u-tokyo.ac.jp)

(Received 15 January 2019; accepted 18 March 2019)

### Abstract

Materials design and discovery can be represented as selecting the optimal structure from a space of candidates that optimizes a target property. Since the number of candidates can be exponentially proportional to the structure determination variables, the optimal structure must be obtained efficiently. Recently, inspired by its success in the Go computer game, several approaches have applied Monte Carlo tree search (MCTS) to solve optimization problems in natural sciences including materials science. In this paper, we briefly reviewed applications of MCTS in materials design and discovery, and analyzed its future potential.

### Introduction

The ability to design a material with desired properties a priori using computational methods has been promised by computational materials science for many years.<sup>[1]</sup> This problem can be framed as selecting the optimal composite material structure that meets certain quality metrics from a space of candidates.<sup>[2,3]</sup> One example is the structural determination of a substitutional alloy problem in solid-state materials design,<sup>[4,5]</sup> where optimal atoms (or vacancies) assignment in a crystal structure is determined to maximize or minimize a target property.

To accelerate this process, researchers have emphasized data-driven and machine learning approaches as the fourth paradigm of science.<sup>[6,7]</sup> Data-driven materials design approaches are iterative design algorithms. Given a space of candidates  $S$ , the algorithm aims to find the optimal candidate  $p_{\text{bst}}$  that optimizes a black-box function  $f(p)$  (usually the target property). Starting with a random selection, and within a predefined number of iterations, the algorithm evaluates a set of selected candidates and obtains feedback for a more informed selection on the next iteration. The function  $f(p)$  is evaluated by experiment or simulation and it is computationally expensive to query. It is necessary to reach the optimal candidate with as few queries as possible<sup>[8]</sup> (Fig. 1).

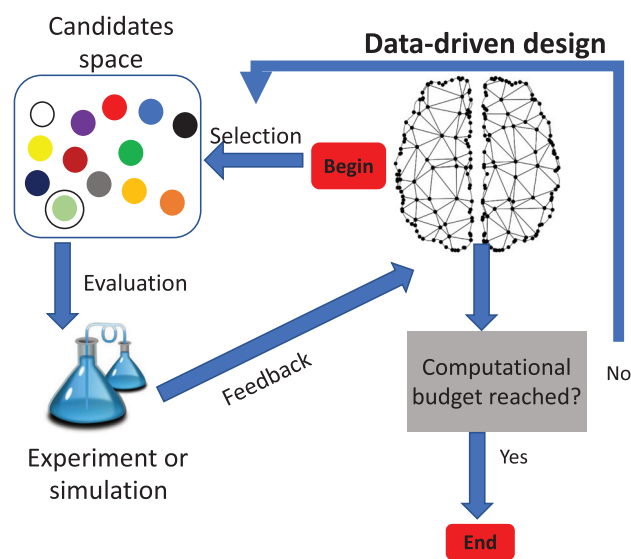
Several methods have been applied for data-driven materials design. Evolutionary algorithms, such as genetic algorithms<sup>[9,10]</sup> that use human evolution mechanisms (such as crossover and mutation), have been used to solve optimization problems at a large scale. Genetic algorithms need adequate

data available a priori to tune many parameters for optimal performance. However, a priori data are often limited in materials design and discovery.

Another approach is Bayesian optimization (BO)<sup>[11,12]</sup> that iteratively selects an optimal candidate from a search space that optimizes an expensive black-box function  $f(p)$ . A deceive advantage of BO is that, it uses the predicated merit of the candidate as well as the uncertainty of the prediction when selecting the next promising candidate. This feature has been successful in several materials design problems.<sup>[2,5,13–15]</sup> However, as the uncertainty of the prediction needs to be assessed for all candidates in the search space, excessive computation in BO faces serious challenges in large-scale problems, a common case in materials design.

The recent success of Monte Carlo tree search (MCTS)<sup>[16]</sup> in computer games<sup>[17,18]</sup> has encouraged researchers to apply it to natural science domains. For example, Yang et al. used MCTS for *de novo* molecular generation by selecting an optimal design with a predefined desired criterion from a vast chemical space.<sup>[19]</sup> Segler et al. applied MCTS to discover retrosynthetic routes to plan the syntheses of small organic molecules.<sup>[20]</sup>

MCTS has also been applied in materials science and engineering. In this study, we reviewed the utilization of MCTS in materials design and discovery, particularly, to analyze its ability to solve large-scale optimization problems. To the best of our knowledge, no such review has been published yet. This paper is organized into four sections. The section “Monte Carlo tree search” presents the MCTS algorithm. In the section



**Figure 1.** A data-driven materials design approach. Starting randomly, an algorithm selects a candidate set for experimentation within a computational budget. The experimental feedback is then used for a more informed selection in the next iteration.

“Applications in materials design”, we discuss the utilization of several variations of MCTS for finding optimal materials design. Finally, “Concluding remarks” section summarizes the paper with forward-looking remarks.

## Monte Carlo tree search

MCTS<sup>[16]</sup> is an iterative, guided, random best-first tree search algorithm that systemically searches a space of candidates to obtain an optimal solution that maximizes the black-box function  $f(p)$ . A distinguishing feature of MCTS is that it encodes the search space into a shallow tree that iteratively expands in the direction of the promising solutions, eliminating the need to construct a full tree. MCTS then utilizes guided randomization to obtain full solutions from this shallow tree.

Within a computational budget, MCTS explores the search space over multiple iterations. Initially, only the root node exists. Then, each iteration consists of four steps: selection, expansion, simulation, and back-propagation (Fig. 2). In the selection step, the tree is traversed from the root to the most promising leaf using a comparative score. The chosen leaf is then expanded by adding a child node to it in the expansion step. Only partial solutions can be obtained from this shallow tree. In the simulation step, a full solution is generated by random rollout<sup>[16]</sup> starting from a node; the remaining variables are randomly determined. The merit of the obtained solution is evaluated. In the back-propagation step, the nodes’ information is updated along the path back to the root. MCTS balances exploration against exploitation using a hyper parameter that depends on the range of maximum and minimum merits values of the candidates.

## Applications in materials design A Python package for an optimal atom assignment problem

An open source implementation of MCTS was developed by Dieb et al. using the Python programming language.<sup>[21]</sup> The package MDTS (materials design using tree search) is available at <https://github.com/tsudalab/MDTS>. MDTS solves structure determination problems by optimal atom (or vacancy) assignment with composition constraints. This method assumes a material structure  $P$  with  $n$  positions  $\{p_1, p_2, \dots, p_n\}$  that need to be assigned with several types of atoms (or vacancies). It searches for the optimal configuration  $p_{bst}$  (subject to composition constraints) that maximizes a black-box function  $f(p)$  (usually the target property). MDTS uses the MCTS model to encode the candidate space into a tree where a node at level  $l$  represents a possible atom assignment in the position  $l$  in the structure. MDTS uses the upper confidence bound (UCB) score<sup>[16]</sup> to compare child nodes when traversing the tree. The UCB score is computed following Eq. (1):

$$ucb_i = \frac{z_i}{v_i} + C \sqrt{\frac{2 \ln v_{parent}}{v_i}} \quad (1)$$

where  $z_i$  is the accumulated merit of the node (i.e., the sum of the immediate merits of all downstream nodes),  $v_i$  is the visit count of the node,  $v_{parent}$  is the visit count of the parent node, and  $C$  is the constant to balance exploration and exploitation.


MDTS adaptively sets  $C$  at each node following similar idea to<sup>[22]</sup>:

$$C = \frac{\sqrt{2}J}{4} (f_{max} - f_{min}) \quad (2)$$

where  $J$  is a meta-parameter initially set to one, which increases whenever the algorithm encounters a dead-end leaf, to allow for more exploration. The parameters  $f_{max}$  and  $f_{min}$  are the maximum and minimum immediate merits in the downstream nodes, respectively. The automatic setting of  $C$  allows MDTS to work parameter-free with optimal performance on several applications. For example, MDTS successfully designed large silicon–germanium (Si–Ge) alloy structures when BO could not be used due to its high computational cost.

## Grain boundary segregation

Analysis of the grain boundary segregation behavior of impurities, dopants, and vacancies is very important for broad materials development due to its effect on material properties. Kiyohara et al.<sup>[23]</sup> recently applied the MCTS method to determine a stable segregation configuration of copper  $\Sigma 5[001]/(210)$  and  $\Sigma 37[001]/(750)$  with silver impurities. A binary Monte Carlo tree was constructed where a node represented either a copper or silver atom assigned to a segregation site; the process searched for an optimum candidate with minimal segregation energy. A stable copper  $\Sigma 5[001]/(210)$  configuration was reached by searching only 1% of all candidate



**Figure 2.** MCTS encodes the search space as a shallow decision tree. MCTS repeats four steps. In the selection step, a promising leaf node is chosen by following the child node with the best score. The expansion step adds a child node to the selected node. During simulation, a full solution is created by random rollout and its merit is evaluated. The back-propagation step updates information for the nodes along the path back to the root for better selection in the next iteration.

configurations (Fig. 3). This result was identical to a previous study using theoretical calculations, where empirical potentials for copper–silver alloys were generated to systematically investigate the segregation.<sup>[24]</sup> The optimal solution of copper  $\Sigma 37$  [001]/(750) (which had a significantly larger search space) was achieved after exploring only 0.34% of the search space.

It is reported that the search efficiency was significantly affected by the order of the search (a search by order of segregation energy was more efficient than a random search). Additionally, the analysis of the search path and the number

of rollouts were important for understanding the background of the search.


### Optimizing the energy gap for graphene nanoflakes

The study of the band gap in graphene is important for nanoelectronic applications. Cao et al. investigated the optimization of graphene nanoflakes' (GNFs) energy gaps by selecting the optimal structure.<sup>[25]</sup> This study was challenged by the large number of candidate structures increasing as the number of carbon atoms ( $N_c$ ) in the graphene increased and the expensive computation of the energy gap ( $E_g$ ) for a specific structure. The researchers proposed an effective tight-binding model for the electronic properties of the hydrogen-terminated GNFs focusing on the effect of carbon atom local bonding. The parameters of this model were fit using first-principles calculation results on structures with  $N_c \leq 34$ . For larger search spaces ( $N_c > 34$ ), they used the MCTS method in combination with the congruence check to search for larger GNFs structures with large  $E_g$  based on the properties of the smaller GNFs.

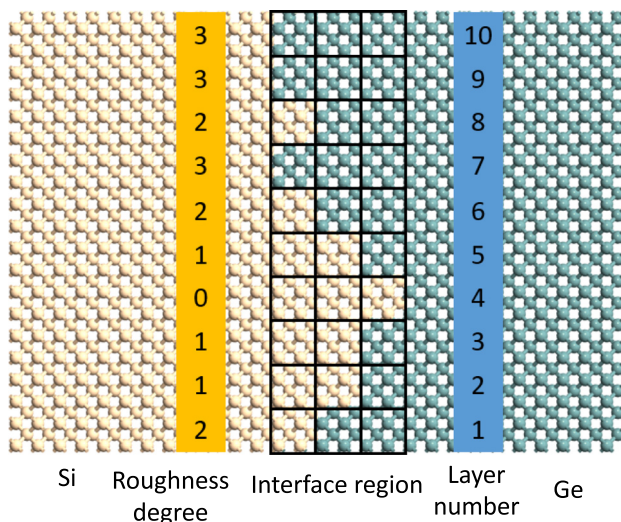
Using this approach, they efficiently obtained optimal structures with the largest  $E_g$  for each  $N_c$ . The reported results were confirmed with first-principles calculations.

### Surface/interface roughness optimization

More recently, Ju et al.<sup>[26]</sup> employed MCTS to study the effect of surface/interface design in nanostructures for optimal thermal transport. They investigated the limits of inhibition and enhancement of phonon transport in the interfacial roughness between a silicon–germanium (Si–Ge) bi-layer nanofilm (Fig. 4). In this setting, the number of possible roughness



**Figure 3.** (a) The promising configuration of the grain boundary structure of copper  $\Sigma 5$ [001]/(210) up to the 16th site. Yellow and gray circles represent silver and copper atoms, respectively. (b) Strain map at each site at the grain boundary. The sites with a positive strain are larger spatially (red); the reverse holds for the negative strain (blue). Reprinted from Kiyohara and Mizoguchi<sup>[23]</sup>, with the permission of AIP Publishing.



**Figure 4.** Interfacial roughness in a bi-layer nanofilm of Si-Ge. The number of possible configurations increased exponentially with the number of layers and the degree of roughness.


configurations increased exponentially with the number of layers in the nanofilm and the degree of roughness (e.g., there were 1,048,576 candidates for 10 layers and four degrees).

Researchers used the MDTS package<sup>[21]</sup> with the atomistic Green's function<sup>[27,28]</sup> to search for candidates with maximum and minimum interfacial thermal conductance. The reported optimal surface/interface roughness configurations were non-intuitive, and they were in the middle range between flat and very rough. This study confirmed the scalability and efficiency of MCTS because it applied MCTS to a large search space.

### Determination of boron-doping in a graphene structure

Tuning material properties by incorporating additional elements is a common practice in materials science and engineering. Both optimal composition and spatial distribution of incorporated atoms affect the target property. Dieb et al. investigated the most stable structures of doped boron atoms in graphene (for a boron concentration up to 31.25%) using a MCTS-based method in conjunction with atomistic simulations.<sup>[29]</sup> They considered a supercell constructed of a hexagonal unit cell of graphene and carbon substituted by boron.

To increase the efficiency of MCTS, researchers engineered the rollout in the expansion step with a more sophisticated solution using BO<sup>[11,12]</sup> (i.e., a Bayesian rollout). In this mechanism, a random pool of full candidates is enumerated under the expanded node with a predefined size,  $Z$ . BO is then used iteratively to select the optimal candidate from the pool. BO maintains a Gaussian process (GP)<sup>[30]</sup> as the surrogate model of the objective function. As the MCTS progresses, more data points are observed and included in the GP for training, which optimizes the model for the target function (Fig. 5).



**Figure 5.** Bayesian rollout. A random pool of full candidates is generated under the expanded node. Initially, GP uses a random selection of data points for evaluation. As the search progresses down the tree, more observations are accumulated in the GP for a more informed future selection. To determine the next selection, the Bayesian rollout uses an acquisition function that considers the predicted value and prediction uncertainty.

Using this method, researchers have reported atomic structures of the most stable configurations of B-graphene at different B concentrations. The stability of these structures has also been verified by the density functional theory calculations with the use of the Vienna Ab initio Simulation Package (VASP).<sup>[31]</sup>

### Concluding remarks

Our review of MCTS in materials design and discovery confirmed that MCTS can successfully solve large-scale optimization problems in materials design. An interesting feature of MCTS is that it does not require a complicated descriptor or a large dataset to tune many parameters. Instead, a single hyper parameter can be tuned automatically and adaptively based on the target application. On the other hand, MCTS can be most useful when experimenting (simulation) time is short. A significantly long experiment time will wipe out the advantage of quick design time of MCTS. Additionally, current implementations of MCTS are only available for discrete search spaces. Enhancing MCTS for continuous spaces can support a wider range of materials design applications.

MCTS may have potential not yet fully exploited in materials science. A recent trend of combining MCTS with other machine learning methods, such as Bayesian learning and neural networks, has emerged to increase MCTS's efficiency, and this has shown promising results.

### Acknowledgment

This work was supported by the "Materials research by Information Integration" Initiative (MI2I) project, the CREST Grant No. JPMJCR16Q5 from the Japan Science and Technology Agency (JST), and KAKENHI Grants No. 16H04274 from the Japan Society for the Promotion of Science (JSPS).

### References

1. S.B. Sinnott: Material design and discovery with computational materials science. *J. Vac. Sci. Technol. A* **31**, 050812 (2013).



2. A. Seko, A. Togo, H. Hayashi, K. Tsuda, L. Chaput, and I. Tanaka: Prediction of low-thermal-conductivity compounds with first-principles anharmonic lattice-dynamics calculations and Bayesian optimization. *Phys. Rev. Lett.* **115**, 205901 (2015).
3. P.V. Balachandran, D. Xue, J. Theiler, J. Hogden, and T. Lookman: Adaptive strategies for materials design using uncertainties. *Sci. Rep.* **6**, 19660 (2016).
4. K. Okhotnikov, T. Charpentier, and S. Cadars: Supercell program: a combinatorial structure-generation approach for the local-level modeling of atomic substitutions and partial occupancies in crystals. *J. Cheminf.* **8**, 17 (2016).
5. S. Ju, T. Shiga, L. Feng, Z. Hou, K. Tsuda, and J. Shiomi: Designing nanostructures for phonon transport via Bayesian optimization. *Phys. Rev. X* **7**, 021024 (2017).
6. A. Agrawal and A. Choudhary: Perspective: Materials informatics and big data: Realization of the “fourth paradigm” of science in materials science. *APL Mater.* **4**, 053208 (2016).
7. M. Drosback, Materials Genome Initiative: Advances and Initiatives, *JOM*, **66**, 334–335, (2014).
8. T.M. Dieb and K. Tsuda: Machine learning-based experimental design in materials science. In *Nanoinformatics*, edited by I. Tanaka (Springer, Singapore, 2018). pp. 65–74.
9. T.K. Patra, V. Meenakshisundaram, J. Hung, and D. Simmons: Neural-network-biased genetic algorithms for materials design: Evolutionary algorithms that learn. *ACS Comb. Sci.* **19**, 96 (2017).
10. W. Paszkowicz, K.D. Harris, and R.L. Johnston: Genetic algorithms: A universal tool for solving computational tasks in Materials Science. *Comput. Mater. Sci.* **45**, ix (2009).
11. J. Snoek, H. Larochelle, and R. Adams: Practical Bayesian optimization of machine learning algorithms. *Adv. Neural Inf. Process. Syst.* **25**, 2951–2959 (2012).
12. D.R. Jones, M. Schonlau, and W.J. Welch: Efficient global optimization of expensive black-box functions. *J. Global Optim.* **13**, 455 (1998).
13. T. Ueno, T. Rhone, Z. Hou, T. Mizoguchi, and K. Tsuda: COMBO: an efficient Bayesian optimization library for materials science. *Mater. Discov.* **4**, 18–21 (2016).
14. S. Kiyohara, H. Oda, K. Tsuda, and T. Mizoguchi: Acceleration of stable interface structure searching using a kriging approach. *Jpn. J. Appl. Phys.* **55**, 045502 (2016).
15. R. Aggarwal, M.J. Demkowicz, and Y.M. Marzouk: Bayesian inference of substrate properties from film behavior. *Modell. Simul. Mater. Sci. Eng.* **23**, 015009 (2015).
16. C. Browne, E. Powley, D. Whitehouse, S.M. Lucas, P.I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton: A survey of Monte Carlo tree search methods. *IEEE Trans. Comput. Intell. AI Games* **4**, 1–43 (2012).
17. D. Silver, A. Huang, C. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis: Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484 (2016).
18. J. Mehat, and T. Cazenave: Combining UCT and nested Monte Carlo search for single-player general game playing. *IEEE Trans. Comp. Intell. AI Games* **2**, 271 (2010).
19. X. Yang, J. Zhang, K. Yoshizoe, K. Terayama, and K. Tsuda: ChemTS: an efficient python library for de novo molecular generation. *Sci. Technol. Adv. Mater.* **18**, 972 (2017).
20. M.H.S. Segler, M. Preuss, and M. P. Waller: Planning chemical syntheses with deep neural networks and symbolic AI. *Nature* **555**(7698), 604–610 (2018).
21. T.M. Dieb, S. Ju, K. Yoshizoe, Z. Hou, J. Shiomi, and K. Tsuda: MDTs: automatic complex materials design using Monte Carlo tree search. *Sci. Technol. Adv. Mater.* **18**, 498 (2017).
22. L. Kocsis and C. Szepesvári: *Bandit based Monte-Carlo Planning in Machine Learning: ECML 2006* (Springer, Berlin, Heidelberg, 2006) pp. 282–293.
23. S. Kiyohara and T. Mizoguchi: Searching the stable segregation configuration at the grain boundary by a Monte Carlo tree search. *J. Chem. Phys.* **148**, 241741 (2018). <https://doi.org/10.1063/1.5023139>.
24. S. Kiyohara and T. Mizoguchi: Investigation of segregation of silver at copper grain boundaries by first principles and empirical potential calculations. *AIP Conf. Proc.* **1763**, 040001 (2016). <https://doi.org/10.1063/1.4961349>.
25. Z. Cao, Y. Zhao, J. Liao, and X. Yang: Gap maximum of graphene nano-flakes: a first principles study combined with the Monte Carlo tree search method. *RSC Adv.* **7**, 37881 (2017).
26. S. Ju, T.M. Dieb, K. Tsuda, and J. Shiomi: *Optimizing Interface/Surface Roughness for Thermal Transport*. Machine Learning for Molecules and Materials NIPS 2018 Workshop (2018).
27. W. Zhang, T. S. Fisher, and N. Mingo: Simulation of interfacial phonon transport in Si–Ge heterostructures using an atomistic Green’s function method. *J. Heat Transfer* **129**, 483–491, (2006).
28. J. Wang, J. Wang, and N. Zeng: Nonequilibrium Green’s function approach to mesoscopic thermal transport. *Phys. Rev. B* **74**, 033408, (2006).
29. T.M. Dieb, Z. Hou, and K. Tsuda: Structure prediction of boron-doped graphene by machine learning. *J. Chem. Phys.* **148**, 241716 (2018). <https://doi.org/10.1063/1.5018065>.
30. C.E. Rasmussen and C.K.I. Williams, eds.: *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006).
31. G. Kresse, and J. Furthmüller: Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. *Comput. Mater. Sci.* **6**, 15 (1996).