

Designing Composites with Target Effective Young's Modulus using Reinforcement Learning

Aldair E. Gongora*

Siddharth Mysore*

agongora@bu.edu

sidmys@bu.edu

Boston University

Boston, Massachusetts, USA

Elise F. Morgan

Keith A. Brown

Boston University

Boston, Massachusetts, USA

Beichen Li*

Wan Shou

Wojciech Matusik

Massachusetts Institute of Technology

Cambridge, Massachusetts, USA

Emily Whiting

Boston University

Boston, Massachusetts, USA

whiting@bu.edu

ABSTRACT

Advancements in additive manufacturing have enabled design and fabrication of materials and structures not previously realizable. In particular, the design space of composite materials and structures has vastly expanded, and the resulting size and complexity has challenged traditional design methodologies, such as brute force exploration and one factor at a time (OFAT) exploration, to find optimum or tailored designs. To address this challenge, supervised machine learning approaches have emerged to model the design space using curated training data; however, the selection of the training data is often determined by the user. In this work, we develop and utilize a Reinforcement learning (RL)-based framework for the design of composite structures which avoids the need for user-selected training data. For a 5×5 composite design space comprised of soft and compliant blocks of constituent material, we find that using this approach, the model can be trained using 2.78% of the total design space consists of 2^{25} design possibilities. Additionally, the developed RL-based framework is capable of finding designs at a success rate exceeding 90%. The success of this approach motivates future learning frameworks to utilize RL for the design of composites and other material systems.

CCS CONCEPTS

- Computing methodologies → Reinforcement learning; Modeling and simulation;
- Applied computing → Computer-aided design; Computer-aided manufacturing.

* Authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SCF '21, October 28–29, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9090-3/21/10...\$15.00

<https://doi.org/10.1145/3485114.3485123>

KEYWORDS

neural networks, composite design, finite element analysis, automated design, mechanical properties

ACM Reference Format:

Aldair E. Gongora, Siddharth Mysore, Beichen Li, Wan Shou, Wojciech Matusik, Elise F. Morgan, Keith A. Brown, and Emily Whiting. 2021. Designing Composites with Target Effective Young's Modulus using Reinforcement Learning. In *Symposium on Computational Fabrication (SCF '21), October 28–29, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3485114.3485123>

1 INTRODUCTION

Engineered composite materials and structures are able to achieve superior mechanical performance in comparison to their individual constituents [Dimas et al. 2013; Ghiasi et al. 2010, 2009; Lakes 1993; Wegst et al. 2015]. The ability to tailor their design to meet performance requirements has resulted in their widespread application in the aerospace, automotive, and maritime industries [Fratzl and Weinkamer 2007; Gu et al. 2016b; Guoxing and Tongxi 2004]. While the design process has previously relied on domain expertise, bio-mimicry, brute force exhaustive search, or iterative trial and error [Yeo et al. 2018], recent advancements in additive manufacturing (AM) have tremendously enhanced the realizable design space and challenged the conventional approaches to exploring design spaces [Gu et al. 2017a,b; Libonati et al. 2016]. The design freedom afforded by AM has significantly expanded the design space by enabling the fabrication of composites with arbitrary geometry and material distributions spanning various length scales. With this expansion comes challenges such as how to rapidly and efficiently exploring the vast and complex design space for optimal or targeted mechanical performance. While more traditional optimization techniques have been proposed, their robustness is often limited by the complexity of the design problems.

In order to overcome some of these design challenges, specifically those pertaining to exploring and modeling the design space, we propose using reinforcement learning (RL). RL algorithms learn to model the problem space through interaction and can be optimized to solve specific controls problems. We present a design framework that utilizes RL to design and find composite designs

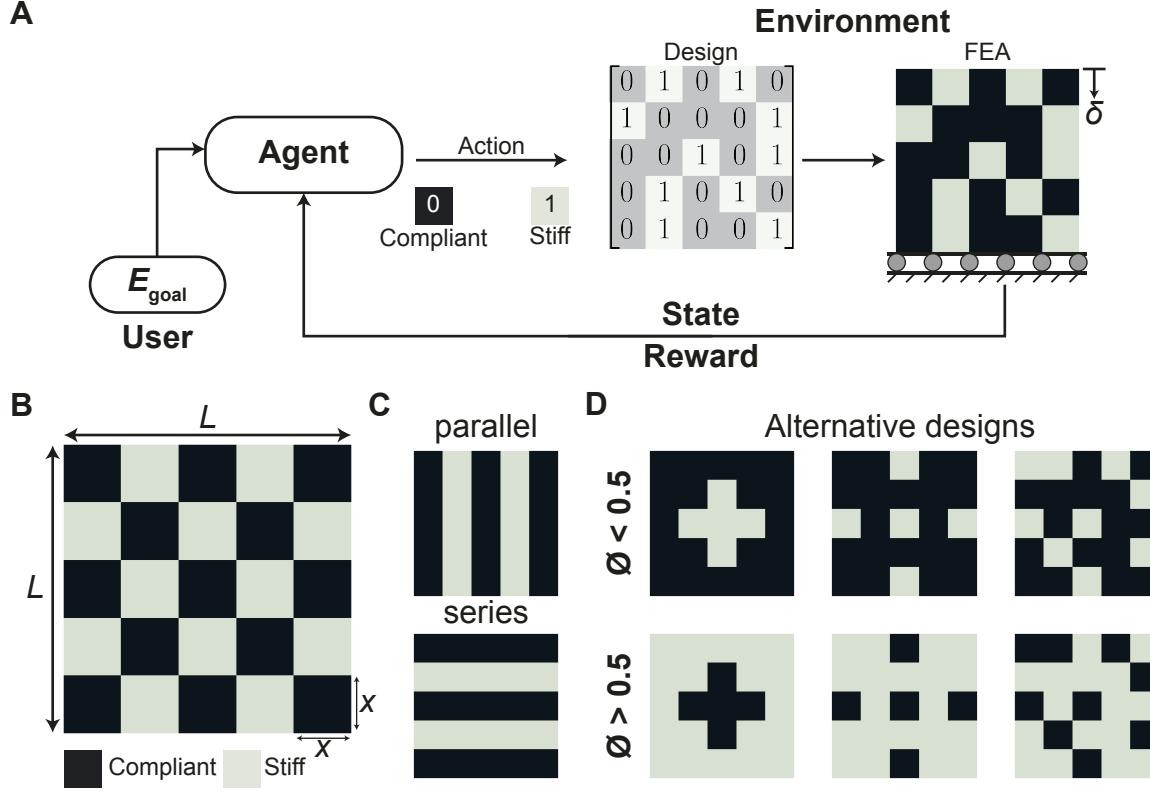


Figure 1: The Reinforcement Learning (RL) framework for the design of parametric composites consisting of stiff and compliant building blocks of constituent material to meet a specified target modulus E_{goal} (A). The design space considered in this work is built from a 5 by 5 arrangement of blocks of side length x where the composite sample has side lengths L (B). The design space consists of parallel and series designs (C) with a volume fraction of stiff constituent material $\phi = 0.40$. The design space also consists of alternative designs where the majority of the compliant or stiff constituent material is concentrated at the center, sparsely distributed along the edges, or randomly distributed (D).

that meet a specified performance target (Figure 1 A). While the models learned do not provide a general solution for arbitrary composite design problems, our work offers a guideline on successfully framing design problems as RL problems. While learning-based frameworks can be more data and computationally intensive than any one optimization run using more traditional techniques, they offer potential benefits at scale, where inference after training can be significantly less expensive.

In this work, we mainly consider a 2-material composite design problem where we attempt to optimize the composite design in order to achieve specific material properties. The RL-based design framework, takes an initial composite design and user-specified desired material properties, comprising of target elastic properties and material composition, and iteratively modifies the design until the desired properties are satisfied. The design is parameterized as a 2D grid of constituent materials and the learned policy adjusts the design one grid cell at a time. This naturally upper-bounds the number of required interactions by the number of cells in the design grid. We demonstrate that, despite only exploring less than 5% of the design space for our design problem during training, the learned RL policies are able to successfully solve over 95% of the design tasks in testing.

2 RELATED WORK

To overcome the challenges of design optimization in additive manufacture, techniques combining computational methods and optimization algorithms, such as topology optimization, have enabled the design of tailored composites with targeted structural and material design requirements such as elastic properties, Poisson's ratio, and tunable stress-strain curves [Chen et al. 2018; Gu et al. 2016a; Sigmund and Maute 2013; Sigmund and Petersson 1998; Zhu et al. 2017]. Although these approaches have achieved success in specific classes of problems, they are often limited by the complexity of the design space and associated computational cost.

Recently, machine learning (ML) based design frameworks have achieved success in the design and discovery of materials and structures with optimal or targeted properties [Abueidda et al. 2019; Chen and Gu 2019; Gongora et al. 2021, 2020; Gu et al. 2018; Mao et al. 2020; Nikolaev et al. 2016]. Specifically in the field of composite design, ML techniques utilizing artificial neural networks and deep learning have been used in classification applications [Chen and Gu 2020] and in inverse design applications [Yang et al. 2020]. These applications have demonstrated the ability to train ML-based models with a fraction of observations from the design space to successfully assess or predict mechanical performance. A non-trivial

challenge that faces the development of accurate ML models is the selection of appropriate training data, as inferior models can emerge from training on insufficient or poorly selected data. In practice, the training data is often selected using uniform random sampling or a design of experiments approach, such as Latin hypercube sampling; however, the applicability of these approaches are limited by the size of the design space and the computational or experimental cost [Jin et al. 2020; Silver et al. 2017].

The dependence of the success of a ML model on appropriately selected training data motivates the development and application of an approach that automatically selects the training data. An approach that can address this challenge is reinforcement learning (RL) [Andrychowicz et al. 2017; Botvinick et al. 2019; Leinen et al. 2020; Mnih et al. 2015; Popova et al. 2018]. RL models, also called RL agents, iteratively interact with a system in order to affect the state of the system. In lieu of curating a dataset of behaviors, a reward function can be defined to reflect the quality of actions taken by agents. The behavior of RL agents is conditioned to maximize the received reward. This shifts some of the burden from dataset curation to reward function engineering, but the latter can be significantly more intuitive and tractable when dealing with large data-spaces.

3 DESIGN SPACE AND FINITE-ELEMENT ANALYSIS

In order to frame composite design as a reinforcement learning problem, we needed to establish a suitable state-space representation and build an environment for the RL algorithm with which to interact with. Training in simulation allows for faster feedback than fabrication and physical analysis, so we use finite element analysis (FEA) as our primary analytical tool for estimating design properties. Before employing the RL algorithm, it is imperative to validate the finite element predictions of Young's modulus E . This section discusses the composite design space considered and the development and validation of FEA pipeline.

3.1 The Composite Design Space

Given that the set of all possible designs is large and intractable, we consider here a more constrained problem amenable to experimental validation so as to establish a viable framework for RL-automated design. We consider a composite design space built from a 5×5 binary arrangement of material blocks where each block can be composed of one of two materials that differ in their elastic moduli. The stiffer of these two materials we term "stiff" and the more compliant of these materials we term "compliant". The material blocks of the composite have a side length $x = 5$ mm while the total composite has a side length $L = 25$ mm and depth of 5 mm (Figure 1 B). Without considering any geometric symmetries, a total of 2^{25} distinct designs exist in the composite design space. Classical parallel and series composite designs can be found in the design space by varying the volume fraction ϕ of stiff material (Figure 1 C). Alternative designs also exist; for instance, designs where the majority of the compliant or stiff constituent material is concentrated at the center, sparsely distributed along the edges, or randomly distributed (Figure 1 D). With the ability to fabricate alternative designs, a range of Young's modulus values can be achieved for a given volume fraction.

3.2 Development and Validation of a Predictive Model using Finite Element Analysis

The model we developed uses 2-D explicit FEA, which was implemented in C++ using the Taichi Graphics Library [Hu 2018], to predict the effective Young's modulus of the composite in the 5×5 design space. Here, we used a Neo-Hookean model for both the compliant and stiff constituent materials and a static compressive strain of $1e-4$. The stiff and compliant constituent materials considered in this work were VeroWhitePlus (VW+) and a volume percentage mixture of 50% VW+ and 50% TangoBlackPlus (TB+), respectively. The effective Young's modulus of the composite was estimated from the predicted stress at a prescribed compressive strain of $1e-4$, which was kept constant for all simulations in this study. The effective Young's modulus was calculated by dividing the predicted stress by the compressive strain. To simulate the uniaxial compression of the composite design, displacement boundary conditions were imposed on the nodes of the top surface of the composite. The coefficients of the Neo-Hookean model were computed from the Young's modulus and Poisson's ratio of the constituent materials. The Young's modulus and Poisson's ratio of the stiff constituent material was 1818 MPa and 0.33, respectively, while the Young's modulus and Poisson's ratio of the compliant constituent material was 364 MPa and 0.49, respectively. Additionally, an explicit solver was selected over an implicit solver in the FEA pipeline to avoid potential computational bottlenecks arising from the need to invert the stiffness matrix of the course of the applied deformation in implicit analysis. In the explicit solver, the time step was set to $2.3e-7$, which was sufficiently small to prevent numerical instability in forward Euler integration. A fixed 2000 iterations was used in the solver which was determined to be sufficiently large to guarantee convergence without compromising computational time based on preliminary tests. Furthermore, the parameters used in the solver were determined to be appropriate based on the reasonable agreement observed between Young modulus predictions and experimental measurements.

The selected composite designs in this study were fabricated from the aforementioned stiff and compliant constituent materials and experimental measurements were taken in triplicate for each composite design where E is the mean effective Young's modulus of the measurements. To assess the model using FEA, the predicted effective Young's modulus \tilde{E} of selected composite designs were compared to experimental measurements of the effective Young's modulus E from quasi-static uniaxial compression tests conducted on a universal testing machine (Instron 5984). The first set of selected composites designs were designs with only compliant constituent material $\phi = 0$, only stiff constituent material $\phi = 1$, and classical parallel and series composite designs with $\phi = 0.2, 0.4, 0.6$, and 0.8 (Figure 2 A). The absolute relative error between \tilde{E} and E for composite designs with $\phi = 0$ and $\phi = 1$ were 6.97% and 3.35%, respectively, while the average absolute relative error between \tilde{E} and E for the classical parallel and series composite designs with $\phi = 0.2, 0.4, 0.6$, and 0.8 was 12.10% (Figure 2 B). Evaluating the entire dataset with $\phi = 0, 0.2, 0.4, 0.6, 0.8$, and 1.0 , the root mean squared error (RMSE) was 101.79 MPa with an $R^2 = 0.96$ where R is the Pearson correlation coefficient. For these composite designs, E can also be analytically estimated via the Voigt and Reuss

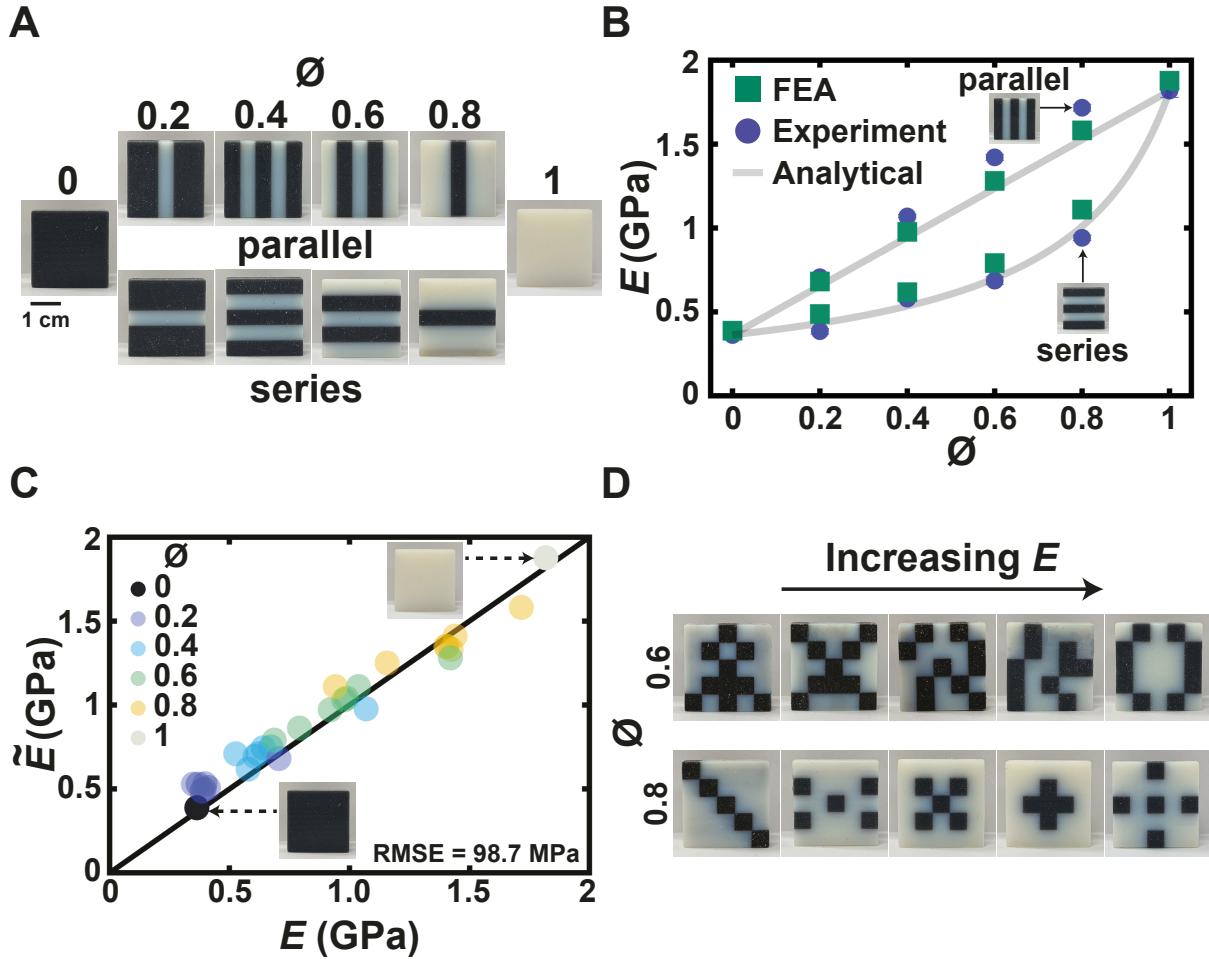


Figure 2: Fabricated composite designs with volume fraction between $\phi = 0$ and $\phi = 1$ with classical parallel and series composite designs with $\phi = 0.2, 0.4, 0.6$, and 0.8 (A). Comparison of experimentally measured effective Young's modulus E , FEA predicted effective Young's modulus \tilde{E} , and analytical predictions of effective Young's modulus for parallel and series composite designs with varying ϕ (B). Comparison of E and \tilde{E} for previously tested designs and five alternative designs for each $\phi = 0.2, 0.4, 0.6$, and 0.8 (C). Example designs with $\phi = 0.6$ and $\phi = 0.8$ in ascending order based on E (D).

approximations for parallel and series designs, respectively. The approximations also yielded reasonable predictions of E with an average absolute relative error of 9.07% when compared to experimental measurements. While the Voigt and Reuss approximations enable rapid predictive approximations of E , the models are only expected to be accurate for parallel and series composite designs.

To further assess the predictions using FEA, the \tilde{E} of five alternative designs for each $\phi = 0.2, 0.4, 0.6$, and 0.8 were compared to experimental measurements from uniaxial compression testing. Adding these results to the classical parallel and series dataset, the agreement between \tilde{E} and E was evaluated. The RMSE was computed to be 98.69 MPa with $R^2 = 0.97$ and the average absolute relative error between \tilde{E} and E of all the selected designs was calculated to be 13.8% (Figure 2 C). Based on these results, we concluded that the predictions of Young's modulus from FEA were in good agreement with experimental measurements. Interestingly, from the E of alternative designs with varying volume fraction, a broad

range of achievable modulus values with non-trivial designs were observed. Specifically, for $\phi = 0.6$ the minimum E was 792.3 MPa and the maximum E was approximately 31% larger at 1.04 GPa. Additionally, for $\phi = 0.8$, the minimum E was 1.16 GPa and the maximum E was approximately 24% larger at 1.44 GPa (Figure 2 D). While this range was observed for the experimental measurements of alternative designs, brute-force evaluations of all the designs suggest that for a given ϕ there is a range of achievable E via alternative designs. The overlap in ranges for varying ϕ demonstrates the ability to vary performance for a given ϕ , which is a key factor in designing composites to meet performance requirements. Note that \tilde{E} for a composite design only corresponds to quasi-static uniaxial compression conditions. While designs such as the alternative composite designs may possess E that depend on the loading direction, only quasi-static compression conditions as shown in Figure 1 A are considered in this study.

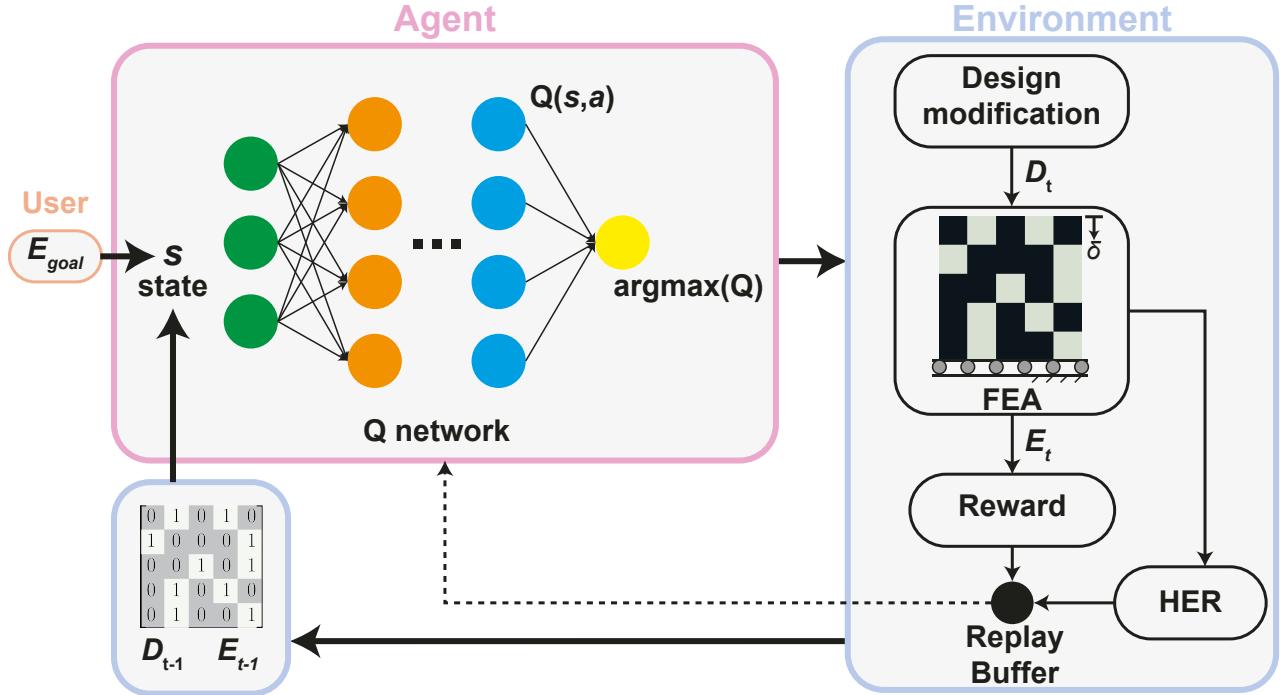


Figure 3: Overview of the reinforcement learning (RL) for composite design framework

4 DEVELOPING THE REINFORCEMENT LEARNING (RL) FRAMEWORK FOR COMPOSITE DESIGN

Reinforcement learning (RL) algorithms seek to solve a sequential decision-making problem, which are typically formulated as a Markov decision process (MDP). There are two major components of any RL setup – (i) the training environments, and (ii) the agent(s), i.e., the learned behavior model. Agents take actions in response to the current state, s_t , of the environment and the state of the environment changes in response to the action taken. The agent decides its actions by attempting to maximize the cumulative reward, received from the environment. In order to successfully train an RL agent on a task, it is important to balance the complexities of the agent policy as well as the task environment. While more complex policy functions can have an increased modeling capacity which may allow learning of more complex tasks, this comes with increased computational costs and training data requirements. We designed our task environment with this in mind.

Figure 3 provides an overview of the RL framework developed to design composites with a target Young's modulus E_{goal} and volume fraction ϕ_{goal} . We began by considering similar problems already solvable by RL and found close parallels with the problem of bit-flipping – i.e., modifying one binary string to match another given one, one element at a time. Like in bit-flipping, our design problem includes a current state and a desired future state, where our 2-material design, D_t , can be represented as a binary string (where an $n \times m$ grid can be represented by a nm-dimensional binary vector). Unlike the bit-flipping problem; however, we do not

know the desired design ahead of time, only the desired properties. By including the current (at iteration t) and desired material properties, E_t and E_{goal} , respectively, in the state representation we equivalently indicate necessary information about a goal state to allow the agent to model and achieve desired material properties. Formulating the design problem as a sequential task in this way enables two key advantages: (i) it maintains the Markovian properties of states being independent of their histories, thus allowing RL to solve this problem as an MDP, and (ii) it allows for a natural representation of termination state and computational limit and consequently lends itself to an intuitive reward signal.

Defining the design problem as a sequential RL problem allows us to side-step difficulties that would normally arise with procuring a training set for supervised learning. In a typical supervised learning approach for a design problem such as this, or indeed any search problem, one of the key issues would be in determining the 'quality' of arbitrary states with respect to a desired goal state. Computing the scalar difference between the desired and current material properties does not directly translate to a strategy on traversing the design space towards a viable design that satisfies the design goals, prohibiting direct optimization. Without a clear notion of distance, or first knowing what an ideal search trajectory is, it would be difficult to quantify the quality of any given state, thus making it difficult to define loss signals to facilitate training. In contrast, RL allows the use of a significantly more intuitive binary reward signal to guide training: if the properties of the current design are within a tolerance of the desired properties, the agent gets a zero-reward, and it gets a negative reward (-1) in all other cases. It is then left to the RL algorithm to learn to evaluate how to

best traverse the manifold of the design space in order to satisfy design requirements.

Our agents are represented by Q-networks, which gained prominence when [Mnih et al. 2015] demonstrated the viability of deep RL in learning to play a diverse set of video-games. Video-games require agents to achieve and maintain desirable states, and this is typically achieved through sequential interaction, much like the design problem we consider here. The Q-network learning framework is developed as an extension of the tabular Q-learning problem [Sutton and Barto 2018; Watkins and Dayan 1992], where the RL model learns a Q-value function mapping states and actions to expected rewards. If an agent is able to accurately model the Q-value, it can optimize its actions at every state to maximize rewards. Deep Q-networks extend tabular Q-learning by representing the Q-function not as a table enumerating all states and actions but rather a neural network that predicts the Q-value for any possible action given an input state. The agent is also given a strict limit of nm iterations to solve the problem for an $n \times m$ grid – as the maximum design distance between a desired design and the current one (or any arbitrary two designs) cannot exceed the number of material cells. It naturally follows from the Q-learning optimization that the best way for the agents to learn to maximize rewards is to solve the problem as quickly as possible.

While suitable training environment and RL agent design enable effective learning, they still may not be efficient enough to be practical as deep RL can be highly data intensive [Leinen et al. 2020]. Given the vast number of variations in the configuration space and the not insignificant time that it takes to analyze a composite with FEA, it was important to improve data efficiency during training. To that end, we utilize Hindsight Experience Replay (HER) [Andrychowicz et al. 2017], a technique developed for goal based deep RL, where training data is synthetically augmented based on data gathered while agents explore their training environments. Briefly, HER treats the final and/or intermediate states reached in a sequence of environment interactions as synthetic target goals. While the true goal may not be reached, the sequence of interactions still provides agents information on how to reach the states that were visited, should they become relevant in future. In effect, the virtual re-contextualization of prior experience allows every interaction with the environment to be several times more useful as a data-point, thus improving data efficiency.

For our 5×5 grid designs, agents are trained in cycles of experience gathering and optimization over 750 cycles with 50 episodes per cycle, for a total of 37500 episodes. Episodes are capped by the total number of cells in the material design (in our case 25) as this is the maximum number of changes that should be required for any design. In all, agents make fewer than 1 million calls to the finite element analysis (FEA) solver, which, while a significant number, corresponds to sampling at most 2.8% of the 2^{25} design options during training. While the true unique design space is likely smaller when accounting for design symmetries, our method still offers relatively low data complexity for learning. Further details on setup and training are provided in Appendix B.

5 RESULTS AND DISCUSSION

To evaluate the performance of RL framework for composite design, we consider the previously described composite design space built from a 5×5 binary arrangement of materials blocks where each block can be either compliant or stiff material. The user specifies a target Young's modulus E_{goal} and volume fraction ϕ_{goal} and the trained agent has a maximum budget of 25 iterations to return a final design – corresponding to the total number of cells in the grid.

For example, given $E_{goal} = 500$ MPa, $\phi_{goal} = 0.20$, and a randomly selected initial design configuration with an initial $E_0 = 808$ MPa and $\phi_0 = 0.52$, the agent returns a design with final $E = 512$ MPa and $\phi = 0.24$ in seven sequential design changes (Figure 4 A). The absolute relative error in E is 2.4%. Due to the discrete nature of the 5×5 design space, the error in ϕ was the addition of one stiff material block which resulted in the final ϕ being larger than the target ϕ by 0.04. Due to the number of possible design variations for volume fractions $\phi \in [0.24, 0.48]$, we observed that agents were typically least performant in this region of the design space. Additionally, two examples are shown in Figure 4 B where one example, Figure 4 B1, has a $\phi_{goal} = 0.4$ and $E_{goal} = 600$ MPa and another example, Figure 4 B2, where $\phi_{goal} = 0.80$ and $E_{goal} = 1300$ MPa. For the example in Figure 4 B1, the final returned configuration exactly satisfies $\phi_{goal} = 0.4$ and slightly exceeds $E_{goal} = 600$ MPa by 5%. In the second case, Figure 4 B2, the returned configuration exactly satisfies $\phi_{goal} = 0.8$ and slightly exceeds $E_{goal} = 1300$ MPa by 0.08%. The agent begins with a randomly selected initial design and it then finds designs that meet both specified goals with an absolute relative error less than or equal to 5%. From these examples, it is observed that agents have preserved aspects of the original design while making modifications to achieve the user specified goal. The observed behavior is indicative that the agents have not just learned to memorize solutions but to modify designs to achieve desired properties while under the prescribed budget, and ideally with the fewest changes. This follows as a natural result of the RL optimization criteria however, as the optimal strategy that maximizes learning rewards is the one that is able to satisfy the design problem in the fewest iterations.

While the previous results are promising, they represent a single trained agent. To further assess the performance of the RL framework for the design of composites, we tested our learning framework with 20 independently trained agents over which we achieved a consistent and median success rate P of 95% (Figure 5 A). The 95% median success rate reflects the average fraction of 2600 of design problems successfully solved by each agent (a total of 52,000 randomly conducted tests) within desired tolerances. This is achieved while requiring exploration of only 3% of the design space – aided significantly by using HER. Notably, there is a low variance in the performance among agents reflecting the stability of the RL framework in designing composite.

While approximately 5% of cases are categorized as “failed”, they did not fail catastrophically but rather failed to achieve a design within the desired tolerances within 25 iterations (Figure 5 B). Specifically, the average error of E of the agents was 61 ± 8 MPa while the average error of ϕ was $8 \pm 2\%$. In all cases, the tolerances for success were set at $E = 50$ MPa and $\phi = 0.04$ (corresponding to a 4% error in the volume fraction). The tolerance enabled the agents

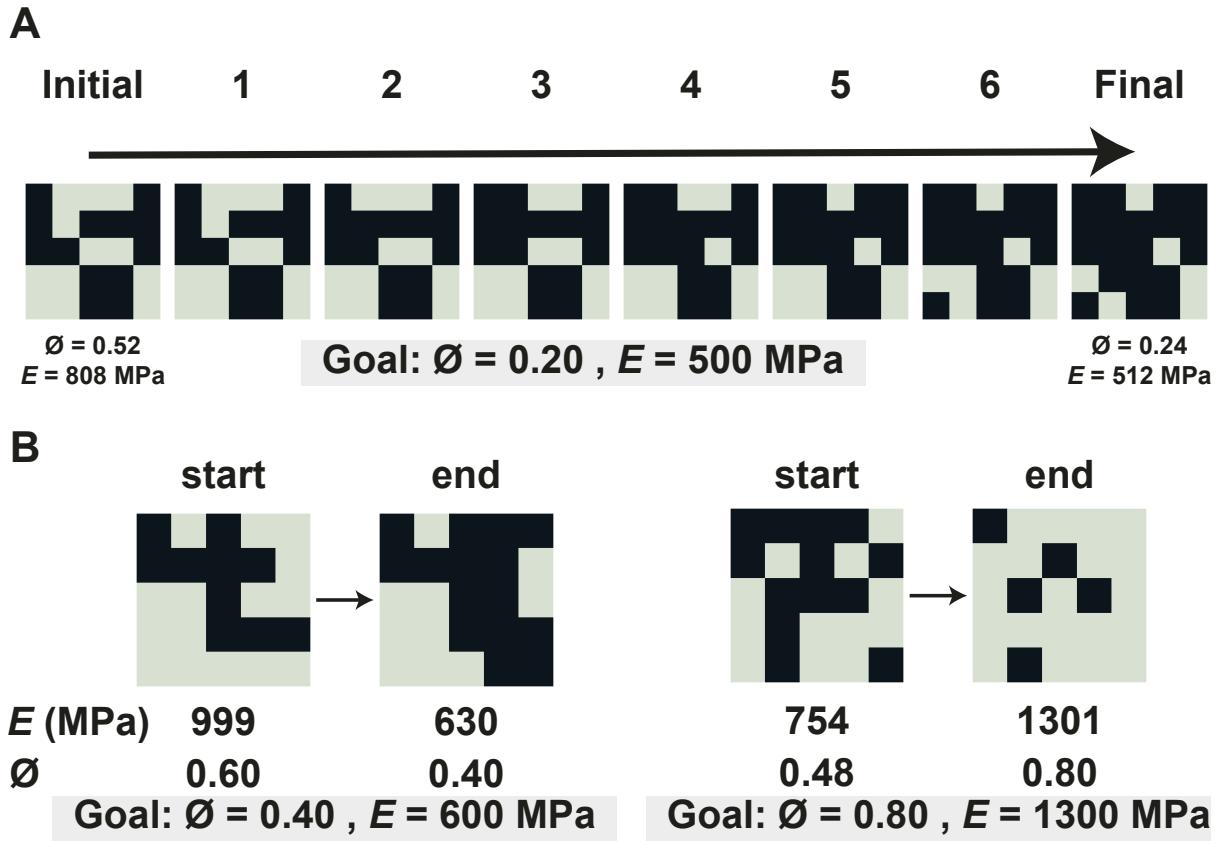


Figure 4: Sequential design changes from a randomly selected initial to final design made by a trained agent to meet a goal $E = 500 \text{ MPa}$ and goal $\phi = 0.2$ (A). Two examples with randomly selected initializations and the final designs returned by the trained agent (B).

to rapidly locate designs that met specified user requirements in the vast design space. Additionally, considering the prescribed tolerances, even in cases of ‘failure’ the error in the final model’s properties marginally outside the tolerance thresholds. While average P , is an important metric, the variation of P as a function of ϕ can also yield insight into the performance of the RL framework (Figure 5 C). We also compare the RL approach to nearest neighbor based solutions based on sampling the design space using as many samples as the upper bound of unique configurations visited by the RL approach (i.e. 2.78%) (details in Section B.4). For ϕ in the range of 0.24 and 0.6, P is observed to slightly fall below 0.90. Specifically, $\phi = 0.48$ has the lowest median $P = 0.89$. The decrease in P reflects that this area in the design space possesses the largest number of possible configurations. When comparing this against the nearest neighbor approaches, we see that the RL agents do perform marginally worse, on average, in the regions of the design space where there are many possible design configurations. This was to be expected, since sampling from the design space, either uniformly at random, or randomly within the possible designs for each possible ϕ , would result in a large proportion of the total samples being in the range of $\phi \in [0.4, 0.6]$, as this ϕ -range represents over 77% of the design space. Crucially though, we see that the RL

method is able to maintain relatively consistent levels of P even in the regions of design space that are represented by fewer possible design configurations. Since the RL-based solver effectively builds its own internalized model of the system, it can compensate for the relative lack of representation in the training data. In contrast, highly data-dependent techniques, such as nearest neighbor search, would fail to generalize outside of the available data, allowing the RL approach to to achieve a higher average $P = 0.95$ as compared to $P = 0.84$ and $P = 0.71$ for the ϕ -biased and naive nearest neighbor searches, respectively. Furthermore, the RL approach tries to maintain design similarity to the original input, which would be difficult to achieve through sampling-based approaches, given their data reliance.

6 CONCLUSION

In this work, we developed and utilized a RL-based framework for the design of composite structures to meet user-specified target Young’s modulus values E . The effectiveness of this approach was evaluated using a 5×5 composite design space that was comprised of stiff and compliant constituent materials. Using the RL-based approach, the models were trained using approximately 2.78% of the

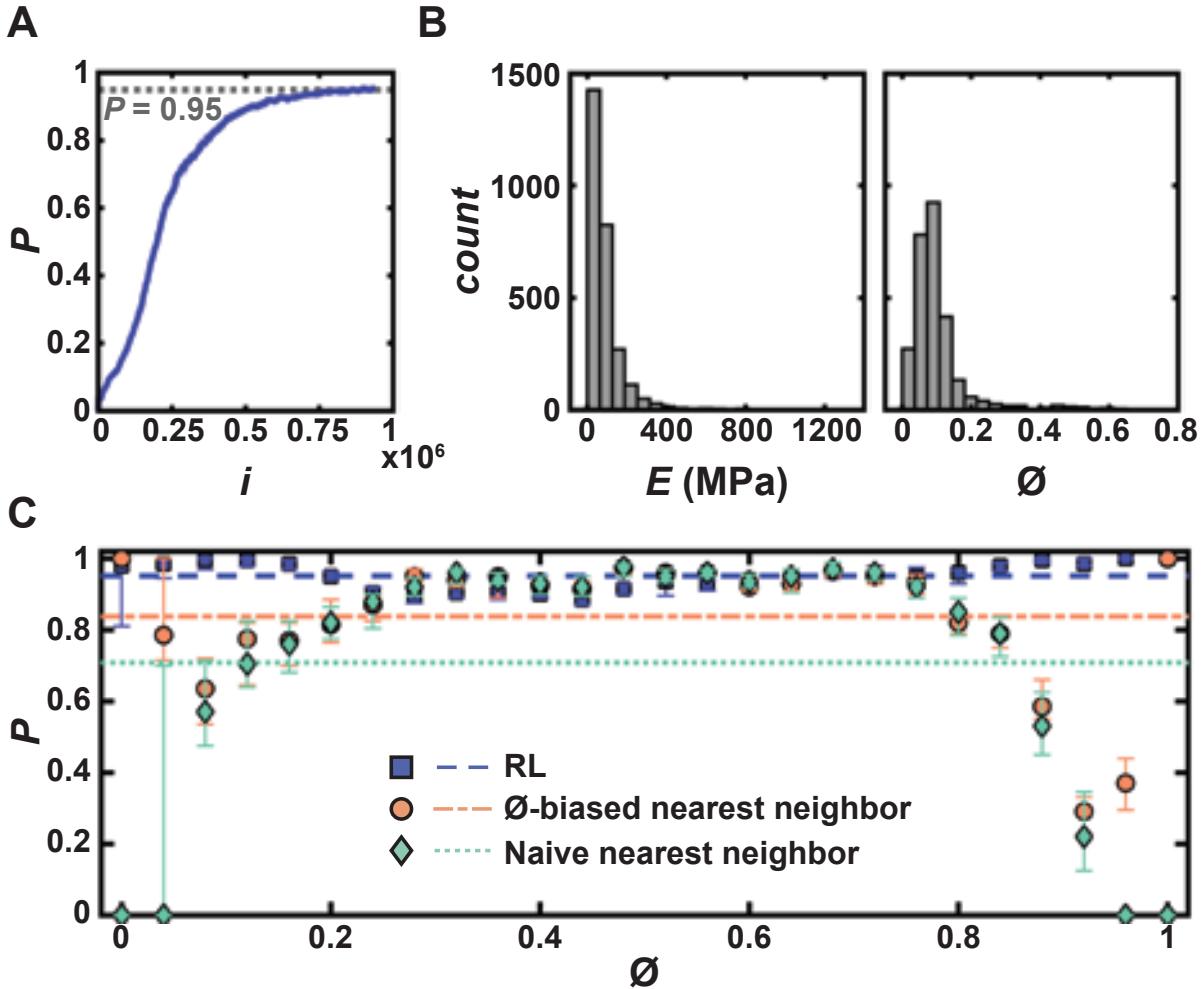


Figure 5: Median success rate P as function of the number of environment interactions i during training evaluated over 20 independent agents with the shaded region represents the 25th and 75th percentiles (A). Distribution of deviation from desired properties in failure cases during testing (B). P as a function of ϕ for the 52,000 random tests (C).

2^{25} design space. Tested on a total of 52000 test cases, 20 independently trained RL agents successfully solved the design optimization problem in 95% of the tests. In totality, this work demonstrated the promise of RL in materials design since traditional design of experiment approaches are limited by the size of the design space and supervised machine learning approaches are limited by the quality of the training data. While we recognize that there are still open questions on the extensibility of such an approach to more complex designs and design requirements, this work is meant to establish a framework and the viability of RL applied to automated design. Additionally, the results described in this work demonstrate the viability of RL for composite design motivating future work and applications considering the distribution and arrangement of multiple configurations of composite designs composed of material blocks. Notably, in these cases, size and boundary effects become imperative to consider in tandem with designing the RL framework for a larger design space. An RL-based design framework circumvents the challenges faced by traditional design of experiment approach

and supervised ML approaches by being able to automatically generate its own training set and solve subsequent design problems. This further motivates design approaches in materials design and discovery to focus on exploring larger design spaces.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1813319, the Alfred P. Sloan Foundation: Sloan Research Fellowship, Google LLC, the Boston University Rafik B. Hariri Institute for Computing and Computational Science and Engineering (2017-10-005), and the U.S. Army CCDC Soldier Center (contract W911QY2020002).

REFERENCES

- Diab W. Abueidda, Mohammad Almasri, Rami Ammourah, Umberto Ravaioli, Iwona M. Jasik, and Nahil A. Sobh. 2019. Prediction and optimization of mechanical properties of composites using convolutional neural networks. *Composite Structures* 227, June (2019), 111264. <https://doi.org/10.1016/j.compstruct.2019.111264>

- arXiv:1906.00094
- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. 2017. Hindsight experience replay. *Advances in Neural Information Processing Systems* 2017-Decem, Nips (2017), 5049–5059. arXiv:1707.01495
- Matthew Botvinick, Sam Ritter, Jane X. Wang, Zeb Kurth-Nelson, Charles Blundell, and Demis Hassabis. 2019. Reinforcement Learning, Fast and Slow. *Trends in Cognitive Sciences* 23, 5 (2019), 408–422. <https://doi.org/10.1016/j.tics.2019.02.006>
- Chun-Teh Chen and Grace X. Gu. 2019. Machine learning for composite materials. *MRS Communications* (2019), 1–11. <https://doi.org/10.1557/mrc.2019.32>
- Chun Teh Chen and Grace X. Gu. 2020. Generative Deep Neural Networks for Inverse Materials Design Using Backpropagation and Active Learning. *Advanced Science* 1902607 (2020). <https://doi.org/10.1002/advs.201902607>
- Desai Chen, Mélina Skouras, Bo Zhu, and Wojciech Matusik. 2018. Computational discovery of extremal microstructure families. *Science Advances* 4, 1 (jan 2018). <https://doi.org/10.1126/sciadv.aa07005>
- Leon S. Dimas, Graham H. Bratzel, Ido Eylon, and Markus J. Buehler. 2013. Tough composites inspired by mineralized natural materials: Computation, 3D printing, and testing. *Advanced Functional Materials* 23, 36 (2013), 4629–4638. <https://doi.org/10.1002/adfm.201300215> arXiv:9907039 [gr-qc]
- Peter Fratzl and Richard Weinkamer. 2007. Nature's hierarchical materials. *Progress in Materials Science* 52, 8 (2007), 1263–1334. <https://doi.org/10.1016/j.pmatsci.2007.06.001> arXiv:34548501731
- Kunihiro Fukushima. 1980. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics* 36, 4 (1980), 193–202. <https://doi.org/10.1007/BF00344251>
- Hossein Ghasi, Kazem Fayazbakhsh, Damiano Pasini, and Larry Lessard. 2010. Optimum stacking sequence design of composite materials Part II: Variable stiffness design. *Composite Structures* 93, 1 (2010), 1–13. <https://doi.org/10.1016/j.compstruct.2010.06.001>
- Hossein Ghasi, Damiano Pasini, and Larry Lessard. 2009. Optimum stacking sequence design of composite materials Part I: Constant stiffness design. *Composite Structures* 90, 1 (2009), 1–11. <https://doi.org/10.1016/j.compstruct.2009.01.006>
- Aldair E Gongora, Kelsey L Snapp, Emily Whiting, Patrick Riley, Kristofer G Reyes, Elise F Morgan, and Keith A Brown. 2021. Using simulation to accelerate autonomous experimentation: A case study using mechanics. *Iscience* 24, 4 (2021), 102262.
- Aldair E Gongora, Bowen Xu, Wyatt Perry, Chika Okoye, Patrick Riley, Kristofer G Reyes, Elise F Morgan, and Keith A Brown. 2020. A Bayesian experimental autonomous researcher for mechanical design. *Science Advances* 6, 15 (2020). <https://doi.org/10.1126/sciadv.aaz1708>
- Grace X. Gu, Chun Teh Chen, and Markus J. Buehler. 2018. De novo composite design based on machine learning algorithm. *Extreme Mechanics Letters* 18 (2018), 19–28. <https://doi.org/10.1016/j.eml.2017.10.001>
- Grace X. Gu, Leon Dimas, Zhao Qin, and Markus J. Buehler. 2016a. Optimization of Composite Fracture Properties: Method, Validation, and Applications. *Journal of Applied Mechanics* 83, 7 (2016), 071006. <https://doi.org/10.1115/1.4033381>
- Grace X. Gu, Flavia Libonati, Susan D. Wettermark, and Markus J. Buehler. 2017a. Printing nature: Unraveling the role of nacre's mineral bridges. *Journal of the Mechanical Behavior of Biomedical Materials* 76, May (2017), 135–144. <https://doi.org/10.1016/j.jmbbm.2017.05.007>
- Grace X. Gu, Mahdi Takaffoli, and Markus J. Buehler. 2017b. Hierarchically Enhanced Impact Resistance of Bioinspired Composites. *Advanced Materials* 29, 28 (2017), 1–7. <https://doi.org/10.1002/adma.201700060>
- Grace X. Gu, Mahdi Takaffoli, Alex J. Hsieh, and Markus J. Buehler. 2016b. Biomimetic additive manufactured polymer composites for improved impact resistance. *Extreme Mechanics Letters* 9 (2016), 317–323. <https://doi.org/10.1016/j.eml.2016.09.006>
- Lu Guoxing and Yu Tongxi. 2004. Energy absorption of structures and materials. *International Journal of Impact Engineering* 30, 7 (2004), 881–882. <https://doi.org/10.1016/j.ijimpeng.2003.12.004>
- Hado Hasselt. 2010. Double Q-learning. In *Advances in Neural Information Processing Systems*, J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta (Eds.), Vol. 23. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2010/file/091d584fcfed301b442654dd8c23b3fc9-Paper.pdf>
- Yuanming Hu. 2018. Taichi: An Open-Source Computer Graphics Library. arXiv:1804.09293 [cs.GR]
- Zeqing Jin, Zhizhou Zhang, Kahraman Demir, and Grace X. Gu. 2020. Machine Learning for Advanced Additive Manufacturing. *Matter* 3, 5 (2020), 1541–1556. <https://doi.org/10.1016/j.matt.2020.08.023>
- Piaras Kelly. 2013. Solid mechanics part I: An introduction to solid mechanics. *Solid mechanics lecture notes* (2013), 241–324.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. BT - 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings. <http://arxiv.org/abs/1412.6980>
- Roderic Lakes. 1993. Materials with structural hierarchy. *Nature* 361, 6412 (1993), 511–515. <https://doi.org/10.1038/361511a0>
- Philipp Leinen, Malte Esders, Kristof T. Schütt, Christian Wagner, Klaus Robert Müller, and F. Stefan Tautz. 2020. Autonomous robotic nanofabrication with reinforcement learning. *Science Advances* 6, 36 (2020). <https://doi.org/10.1126/sciadv.abb6987>
- Flavia Libonati, Grace X. Gu, Zhao Qin, Laura Vergani, and Markus J. Buehler. 2016. Bone-Inspired Materials by Design: Toughness Amplification Observed Using 3D Printing and Testing. *Advanced Engineering Materials* 18, 8 (2016), 1354–1363. <https://doi.org/10.1002/adem.201600143>
- Yunwei Mao, Qi He, and Xuanhe Zhao. 2020. Designing complex architecture materials with generative adversarial networks. *Science Advances* 6, 17 (2020). <https://doi.org/10.1126/sciadv.aaz4169>
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533. <https://doi.org/10.1038/nature14236>
- Vinod Nair and Geoffrey E Hinton. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning (ICML '10)*. Omnipress, Madison, WI, USA, 807–814.
- Pavel Nikolaev, Daylond Hooper, Frederick Webber, Rahul Rao, Kevin Decker, Michael Krein, Jason Poleski, Rick Barto, and Benji Maruyama. 2016. Autonomy in materials research: A case study in carbon nanotube growth. *npj Computational Materials* 2, August (2016). <https://doi.org/10.1038/npjcompumats.2016.31>
- Mariya Popova, Olexandr Isayev, and Alexander Tropsha. 2018. Deep reinforcement learning for de novo drug design. *Science Advances* 4, 7 (2018), 1–15. <https://doi.org/10.1126/sciadv.aap7885> arXiv:1711.10907
- Eftychios Sifakis and Jernej Barbic. 2012. FEM simulation of 3D deformable solids: a practitioner's guide to theory, discretization and model reduction. In *ACM SIGGRAPH 2012 Courses*, 1–50.
- Ole Sigmund and Kurt Maute. 2013. Topology optimization approaches: A comparative review. *Structural and Multidisciplinary Optimization* 48, 6 (2013), 1031–1055. <https://doi.org/10.1007/s00158-013-0978-6>
- O. Sigmund and J. Petersson. 1998. Numerical instabilities in topology optimization: A survey on procedures dealing with checkerboards, mesh-dependencies and local minima. *Structural Optimization* 16, 1 (1998), 68–75. <https://doi.org/10.1007/BF01214002>
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George Van Den Driessche, Thore Graepel, and Demis Hassabis. 2017. Mastering the game of Go without human knowledge. *Nature* 550, 7676 (2017), 354–359. <https://doi.org/10.1038/nature24270>
- Richard S Sutton and Andrew G Barto. 2018. *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA.
- Hado van Hasselt, Arthur Guez, and David Silver. 2015. Deep Reinforcement Learning with Double Q-learning. arXiv:1509.06461 [cs.LG]
- Christopher J C H Watkins and Peter Dayan. 1992. Q-learning. *Machine Learning* 8, 3 (1992), 279–292. <https://doi.org/10.1007/BF00992698>
- Ulrike G.K. Wegst, Hao Bai, Eduardo Saiz, Antoni P. Tomsia, and Robert O. Ritchie. 2015. Bioinspired structural materials. *Nature Materials* 14, 1 (2015), 23–36. <https://doi.org/10.1038/nmat4089>
- Charles Yang, Youngsoo Kim, Seunghwa Ryu, and Grace X. Gu. 2020. Prediction of composite microstructure stress-strain curves using convolutional neural networks. *Materials and Design* 189 (2020), 108509. <https://doi.org/10.1016/j.matdes.2020.108509>
- Jingjie Yeo, Gang Seob Jung, Francisco J. Martin-Martinez, Shengjie Ling, Grace X. Gu, Zhao Qin, and Markus J. Buehler. 2018. Materials-by-design: Computation, synthesis, and characterization from atoms to structures. *Physica Scripta* 93, 5 (2018). <https://doi.org/10.1088/1402-4896/aaab4e2>
- Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. 2017. Two-Scale Topology Optimization with Microstructures. 36, 5 (2017). arXiv:1706.03189 <http://arxiv.org/abs/1706.03189>

A MATERIALS MODELING, FABRICATION AND ANALYSIS METHODS

A.1 Fabrication and Mechanical Testing

All composite specimens used in this study were printed using a multi-material 3D printer (Objet260 Connex). The specimens were printed using two constituent materials. The first constituent material was VeroWhitePlus (VW+) and the second constituent material was a volume percentage mixture of 50% VeroWhitePlus and 50% TangoBlackPlus (TB+). Compression tests were conducted using a universal testing system (Instron 5984) at a loading rate of 3 mm/min. Three samples for each selected composite design were fabricated and tested to obtain a mean measurement of Young's modulus, which was measured as the slope of the linear portion of the tested specimen's stress-strain curve.

A.2 Analytical analysis

The achievable range of Young's moduli E of a composite design comprised of two constituent materials is bounded by the Voigt and Reuss analytical models. The Voigt model assumes that composite designs in parallel can be modelled as springs acting in parallel and the Reuss model assumes that the composite designs in series can be modelled as springs in series. The models are a function of the stiff constituent's volume fraction ϕ_s , the modulus of the stiff constituent material E_s and the modulus of the compliant constituent material E_c . The Voigt model can be expressed as $E = \phi_s E_s + (1 - \phi_s)E_c$. The Reuss model can be expressed as $E = (E_s E_c) / (\phi_s E_c + (1 - \phi_s)E_s)$. The modulus of the constituent materials E_s and E_c are 1,818 MPa and 364 MPa, respectively.

A.3 Finite-element analysis

In order to predict the E of the composite designs, we performed 2-D finite-element analysis (FEA). The 5×5 composite design was represented using a 40×40 grid of quad-finite elements where each cell of the composite was represented by an 8×8 grid. We adopted a Neo-Hookean material model to compute the hyper-elastic response of each element, where the strain energy density function is given by [Sifakis and Barbic 2012].

$$W = \frac{\mu}{2} (I_1 - 2 - \ln J) + \frac{\lambda}{2} (\ln J)^2. \quad (1)$$

Here, μ and λ are the Lamé parameters; I_1 is the first invariant of the right Cauchy-Green deformation tensor; J is the determinant of the deformation gradient. The material coefficients of the stiff and compliant base materials ($\mu_s, \mu_c, \lambda_s, \lambda_c$) were computed from Young's moduli (E_s, E_c) and Poisson's ratios (v_s, v_c) as follows [Kelly 2013].

$$\begin{aligned} \mu &= \frac{E}{2(1 + v')}, \\ \lambda &= \frac{E'v'}{(1 + v')(1 - 2v')}, \\ v' &= \frac{v}{1 + v}, \\ E' &= \frac{E(1 + 2v)}{(1 + v)^2}. \end{aligned} \quad (2)$$

We set E_s, E_c the same as experimental measurements with $v_s = 0.33$ and $v_c = 0.49$. Additionally, we applied negative exponential

damping to the nodal velocities ($v_{i,j}$) to address material viscosity in the FEM.

$$v'_{i,j} = v_{i,j} e^{-\gamma \Delta t_{\text{sim}}}. \quad (3)$$

The damping constant (γ) was set to 1e5 which leads to desirable damping effects in constituent material simulations. Based on these settings, the deformation was solved using an explicit solver under the Dirichlet boundary constraint of a static compressive strain equal to 1e-4 on the displacement in the loading direction. The time step (Δt_{sim}) was set to 2.3e-7 which is sufficiently small to prevent numerical instability in forward Euler integration. Above all, we ran a fixed 2,000 iterations to guarantee convergence. The estimated Young's modulus was then derived from the gauge stress measured at boundary nodes. The simulator was implemented in C++ with Taichi Graphics Library [Hu 2018]. Furthermore, we note that the simulator is not the bottleneck of our computational pipeline, since each simulation takes well less than a second on a single CPU core.

B REINFORCEMENT LEARNING FRAMEWORK

B.1 Task Environment Design

At design-step t , the task environment state is represented numerically by a 29-dimensional $s_t \in S$ where $s_t^i \in [0, 1]$ and $S \subset R^{29}$. The first 25 components capture the 5×5 material design grid, D_t . The 26th and 27th components capture the current Young's modulus and volume fraction of material composition, while the 28th and 29th components capture the desired modulus and volume fractions respectively. Neural networks are known to be sensitive to the scale of the inputs, so we scale the Young's moduli by dividing it by the maximum value of the two materials' moduli, i.e. $f(E) = \frac{E}{\max(E_1, E_2)}$. The action taken by the agent, a_t , in effect, selects a material cell in the design grid to 'flip' between the two materials. The resultant design matrix is evaluated for its material properties and this, along with the modified design is returned as the next state, s_{t+1} , to the agent. During training, s_{t+1} is also used to compute the binary reward of -1 or 0 depending on whether the state is within tolerance of the desired properties. We use a Young's modulus tolerance of 50 MPa and a volume fraction tolerance of 4%, i.e. being one material cell off from the desired composition.

For each episode of the training run, which is limited to 25 iterations as that is the maximum number of changes that could theoretically be required to achieve a goal for a 5×5 grid, a goal volume fraction and desired Young's modulus is selected randomly from the range of possible volume fractions and Young's moduli as defined by the Voigt and Reuss approximations (as shown in Figure 2). A design, chosen uniformly at random, is also instantiated and evaluated. If the design's properties start too close to the goal properties, the goals are reset until the desired Young's modulus is at least 100 MPa away from the current design's modulus.

B.2 Network Design and Training

In order to approximate the Q-function, we use a neural network with an input layer, 3 fully connected hidden layers and an output layer. The input layer has 29 input nodes to read in the 29-dimensional state provided by the training environment. The first and second hidden layers have 128 and 64 neurons respectively with

Rectified Linear Unit (ReLU) activations [Fukushima 1980; Nair and Hinton 2010]. These are passed to a 26-dimensional layer which estimates the Q-value of 26 possible actions (one for each material cell ‘flip’ and including a null action where the agent does nothing). The output action is determined by an argmax over the predicted Q-values to select action with the highest expected value. We use a Q-value discount factor of $\gamma = 0.99$ and the networks are trained with an Adam optimizer [Kingma and Ba 2015] with a learning rate of 10^{-3} , a batch size of 320 and a training buffer size of 2×10^6 . Training is run over 750 cycles with 50 episodes per cycle, with each episode consisting of 25 environment interactions (episodes are not terminated immediately upon success so that agents learn to not deviate from a good design). 500 optimization steps are performed at the end of every cycle.

To mitigate over-fitting and over-estimation in the Q network, we employ Double Deep Q-learning. Double Q learning [Hasselt 2010; van Hasselt et al. 2015] uses two networks in training instead of just one. The networks start with equal weights but only one is updated consistently during training (the ‘main’ or ‘live’ network) while the other (‘target’) network is held fixed for set intervals. When estimating Q-values for the training loss, the target network is used to estimate the temporal difference error [Sutton and Barto 2018] and is updated with weights from the live network only periodically. During training, we update the target network once per cycle with a 95% interpolation factor on the weights between the target and live networks.

B.3 Developing an Exhaustive Dataset

In order to enable quick testing and iteration on experimental design, we developed an exhaustive dataset of all with the 2^{25} possible designs and their associated FEA values. While any individual agent requires under 1 million data samples (less than 3% of the design space), each agent may experience a different set of data, as is the nature of RL training. Each FEA call takes time however and, in an academic setting, this was found to limit early development of our RL framework. To mitigate this, we parallelized the FEA computation of every possible design to build a library of material properties for all the designs which could be looked up by the RL agents during training. This ensured that we were not re-computing the FEA for any previously computed model and it also allowed us to massively parallelize the training of multiple agents when verifying the efficacy of our method, to demonstrate that the performance was consistent and repeatable. We stress again that this was only done to facilitate the academic pursuit of a performance analysis of the method – any individual agent would only need to be connected directly to the FEA solver during training and would only need to be trained once.

B.4 Testing against brute-force methods

The development of an exhaustive dataset also allowed us to compare our RL method against brute-force methods to solve a similar design problem. A naive brute-force algorithm which keeps with the bit-flipping inspired problem formulation would present with too much uncertainty in the run-time and could require a lot of sampling from the dataset - which would, in a real application, correspond to significant time spent on the FEM solver. Instead, we

compare our method against a nearest-neighbor approach, where a subset of designs are sampled from the full set of possible designs and given design requirements, the nearest neighbor from the sampled subset is used to present a possible solution. The nearest-neighbor solver is easily bounded to use a similar amount of total data as the RL algorithm by sampling 3% of the dataset.

We employed two sampling strategies for building the nearest-neighbor search subset: (i) Naive Nearest Neighbor: using 3% of the data, sampled uniformly from the full dataset, and (ii) ϕ -biased Nearest Neighbor: using 3% of the possible samples for each of the 26 possible volume fractions, totaling to 3% of the full design space; both cases are rounded up to their nearest integer. The former is an admittedly naive approach but is likely more reflective of the sampling the RL agents might be expected to encounter, since we make no explicit attempt to condition how the RL algorithms explore the design space. The latter however is a minor change which uses a little bit of our prior knowledge of the design space to better shape the nearest-neighbor solution. Since we do not track the number of unique states visited by the RL agents during training, a comparison of sampling strategies on strategy alone is more difficult to judge, but we are able to compare the performances on the design tasks. Given that the nearest-neighbor approaches do not attempt to model the design space, we do not compare the similarity of the final solutions to any initial state, as that is not something this technique can be reasonably expected to account for in this particular design problem.