

ARTICLE OPEN



A physics informed bayesian optimization approach for material design: application to NiTi shape memory alloys

Danial Khatamsaz¹✉, Raymond Neuberger¹, Arunabha M. Roy¹, Sina Hossein Zadeh¹, Richard Otis¹ and Raymundo Arróyave^{1,3,4}

The design of materials and identification of optimal processing parameters constitute a complex and challenging task, necessitating efficient utilization of available data. Bayesian Optimization (BO) has gained popularity in materials design due to its ability to work with minimal data. However, many BO-based frameworks predominantly rely on statistical information, in the form of input-output data, and assume black-box objective functions. In practice, designers often possess knowledge of the underlying physical laws governing a material system, rendering the objective function not entirely black-box, as some information is partially observable. In this study, we propose a physics-informed BO approach that integrates physics-infused kernels to effectively leverage both statistical and physical information in the decision-making process. We demonstrate that this method significantly improves decision-making efficiency and enables more data-efficient BO. The applicability of this approach is showcased through the design of NiTi shape memory alloys, where the optimal processing parameters are identified to maximize the transformation temperature.

npj Computational Materials (2023)9:221; <https://doi.org/10.1038/s41524-023-01173-7>

INTRODUCTION

In many material design applications, complex computational models and/or experiments are employed to gain a better understanding of the material system under investigation or to improve its performance. High-fidelity models, however, often exhibit high non-linearity, effectively behaving as black-boxes that hinder intuitive understanding beyond input-output correlations. At the same time, experiments are inherently black-box in nature as intermediate linkages between inputs (e.g. chemistry, processing protocols) and outputs (i.e. properties or performance metrics) tend to be accounted for only in an implicit manner—there are some exceptions as shown in ref. ¹.

The ‘black-box’ nature of these interrogation tools is compounded by the significant cost associated with using these ‘information sources’² to query the materials space. An exhaustive exploration of the materials design space using sophisticated experimental or computational tools is thus infeasible. There is thus a growing need for novel data-efficient approaches that can effectively address these challenges while ensuring that the discovery and/or design process remains comprehensible and effective.

Recently, Bayesian optimization (BO) has emerged as a powerful optimization technique for handling expensive black-box functions, owing to its intrinsic capability to efficiently search the design space with minimal data, employing heuristic-based approaches to discover optimal design regions^{3–7}. BO frameworks typically consist of two main components: a surrogate model, commonly a Gaussian process (GP)⁸, and an acquisition function. The surrogate model represents our *uncertain* knowledge about the underlying objective function(s) and offers a cost-effective means of predicting outcomes at unobserved locations while capturing the uncertainty in these predictions. The acquisition function, which includes well-known examples such as Expected Improvement (EI), Upper Confidence Bound (UCB), and Probability

of Improvement (PI)⁵, leverages the uncertainty estimates from the surrogate model to guide the search for the next best experiment.

By striking a balance between the exploration of the design space (i.e., searching in uncertain regions) and the exploitation of the system’s current knowledge (i.e., refining the search around promising areas), BO utilizes GP’s probabilistic predictions to determine the next best experiment through a heuristic-based search or querying policy that maximizes the utility function. In many engineering design applications, resource constraints necessitate more data-efficient design approaches to reduce the resource requirements of these design tasks. The data efficiency of general BO approaches makes them appealing for solving challenging problems across various fields of science and engineering. In materials science, BO has emerged as a potent paradigm driving much of the recent progress in efficient materials discovery and design⁹.

Several approaches have been proposed to enhance the data efficiency of BO frameworks. One of these approaches exploits multi-fidelity BO techniques, which incorporate information fusion techniques^{10–24} to exploit information about a quantity of interest from multiple sources with varying fidelity and evaluation costs. By leveraging the correlation between different information sources, these techniques enable accurate inferences regarding a system’s true response (i.e., ground truth) with fewer expensive experiments by conducting cheaper ones. Applications of multi-fidelity BO in materials design include optimization of processing parameters, microstructure design, and materials selection^{1,25–28}.

Another challenge in BO is the deterioration of its performance when the design space’s dimensionality increases due to the relatively large learning and searching space. Dimensionality reduction techniques, such as subspace approximation methods^{29–32}, can be employed to ease the learning process by projecting the entire problem into a lower-dimensional space^{33–35}.

¹Materials Science and Engineering Department, Texas A&M University, College Station, TX 77843, USA. ²Engineering and Science Directorate, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA. ³Mechanical Engineering Department, Texas A&M University, College Station, TX 77843, USA. ⁴Department of Industrial & Systems Engineering, Texas A&M University, College Station, TX 77843, USA. ✉email: Khatamsaz@tamu.edu

For instance, active subspace-based BO has been introduced in refs. ^{27,36} to increase design efficiency by recognizing and exploring more informative regions, thereby eliminating non-informative or less-informative queries from the ground truth models. This approach allows BO to focus on a smaller set of dimensions that have the most impact on the optimization process.

Furthermore, BO frameworks generally struggle to handle high-dimensional design problems with very sparse supporting data due to the GP's inability to adequately represent underlying objective functions. Although GP's kernel parameters can be obtained using methods such as maximum likelihood or cross-validation, the minimal data typically available in many BO scenarios—especially at the beginning of a design campaign—means that these traditional methods do not yield reliable solutions, as the GP surrogate's hyperparameters (and their predictions) are highly sensitive to the data used for training. To address this issue, Batch Bayesian Optimization (BBO) has been introduced³⁷ and successfully implemented in materials design applications, such as alloy design and phase stability prediction^{26,38,39}. In BBO, rather than identifying optimal hyperparameter sets, a large number of GPs with varying hyperparameters are constructed to consider diverse possible function representations and smoothness levels. BBO remains agnostic regarding the region of the hyperparameter space most consistent with the limited available data, and consequently assumes that any region in the hyperparameter space could contain the hyperparameters that reflect the true behavior of the system being optimized. Once the BO problem is solved for each hyperparameter set considered, clustering is employed to reduce the number of selections to the available batch size in the experimental or computational framework, thereby enhancing the overall optimization process.

Despite the tremendous success of traditional BO algorithms in materials design, state-of-the-art BO-based algorithms typically treat the objective function as a black-box and rely solely on statistical information in the form of input and output data. As the number of design variables increases or the response surfaces exhibit less-smooth behaviors, these methods require more statistical information to accurately represent the attributes of a black-box objective function.

In some cases, even though the complete structure of a black-box objective function remains unknown, certain theoretical information can be gleaned from expert opinion or data-driven approaches. This allows for a partial understanding of the principles that a black-box objective function may follow, such as smoothness, the number of local and global extrema, periodic or exponential behavior, and/or sensitivity to different input variables. By inserting physically interpretable input-output relationships within the system, significant performance improvements can be achieved, which is not evident in purely data-driven approaches⁴⁰.

Incorporating physical information (or, more generally, prior scientific knowledge) into design frameworks reduces their dependency on statistical information, transforming *black-box optimization* into *gray-box optimization* as the information inside the box becomes partially observable^{41–44}. Examples of gray-box modeling for optimizing chemical processes and describing bioprocesses can be found in refs. ^{45–48}. The Physics-Informed Neural Network (PINN), for instance, is gaining popularity in numerous engineering fields due to its increased accuracy, faster learning^{49–61}, and the ability to employ smaller datasets for training⁵⁵. PINN combines traditional scientific computational modeling with a data-driven ML framework to embed physics into neural networks (NNs), enhancing the performance of learning algorithms by leveraging automatic and data-driven NN estimates⁵⁵. Applications of PINN in engineering studies are diverse, including predicting corrosion-fatigue crack propagation⁶², estimating creep-fatigue life of components at elevated temperatures⁶³,

and assessing bond quality and porosity in fused filament fabrication⁶⁴.

The intriguing question of whether domain knowledge can enhance data-only Bayesian Optimization (BO) merits consideration. In traditional BO, the function to optimize is treated as a black-box; our understanding of it only expands as we query it. Given that no initial information about the function's shape is available, learning its form is a sequential process based on successive queries. Nevertheless, recent insights reveal situations where partial information about the objective function can be gleaned.

Leveraging this partial information could potentially bolster the optimization process's efficiency and speed. This has been evidenced by the remarkable performance enhancement of 'gray-box' BO, as demonstrated in ref. ^{65,66}. Specifically, 'gray-box' BO has been applied to problems where the total objective function can be dissected into nested functions. The studies^{65,66} underscore how partial knowledge of one such function can dramatically expedite the convergence to the solution, outstripping the traditional 'black-box' BO approach.

The concept of a 'gray-box' holds profound potential for extension within the context of scientific problem-solving, such as in the field of materials discovery. Although the systems under study may be too intricately complex to be explicitly simulated, it is feasible to make robust assumptions about the underlying physical or chemical principles that bind isolated queries to a specific design space. Illustratively, these physical systems are constrained by fundamental laws and symmetries, which include but are not limited to invariance under symmetry operations, conservation of energy and momentum, and the requisite positivity in entropy production.

In a purely data-driven 'black-box' approach, a standard BO algorithm may struggle to fully grasp or learn these laws solely from data, particularly in a sparse data situation. Nevertheless, harnessing this knowledge could significantly accelerate the recovery of the underlying objective function. Consequently, the neglect of such knowledge could incur significant costs, highlighting the value of the 'gray-box' approach in the realm of complex scientific problem-solving.

In the application of Bayesian Optimization (BO), statistical data is supplied to Gaussian Processes (GPs) to construct models of the underlying objective function. The promising strategy here lies in adopting a gray-box approach, wherein GPs are fine-tuned using prior theoretical knowledge. This method seeks to augment their modeling capabilities, thereby enhancing the efficiency of the design process.

Gaussian Processes (GPs) are characterized by a mean function, which is typically set to a constant, often zero. The GP's response essentially converges to its mean at locations distant from any training data and outside the correlation zone established by length-scales. A basic manipulation of GPs involves shifting the mean function to match the average target value, thereby marginally enhancing the GP's performance. A more flexible strategy is proposed by Ziatdinov et al.⁶⁷, wherein they substitute the GP mean with a function of input variables, allowing it to vary across the space. These functions are constructed based on known physics of the target objective function, introducing a method termed "augmented BO." This approach enables designers to incorporate prior knowledge, such as the number of local extrema, gradient, and lower fidelity approximation of the true objective function. Consequently, a GP response converges to this prior knowledge in the absence of high-fidelity observations. Moreover, the GP is guided to capture the potential trend of objective function variability across the space using the augmented mean function. Some of the present authors have already demonstrated the utility of augmenting GP models by physics-informed prior means in the context of accelerated alloy modeling and discovery⁶⁸.

Kernel manipulation is another technique that modifies the GP's behavior by establishing the correlation between data points to predict outcomes at unknown locations. In Ladygin et al.'s work⁶⁹, GPs are used to reconstruct free-energy functions, and the kernel functions are modified based on the known physics of these free energy functions. The term "physics" here refers to the relationship between free energy and different variables. Integrating physical (theoretical) information with the statistical information of an objective function can significantly reduce the data dependency of GPs. In materials science, this physical information often exists in the form of prior knowledge or expert opinion. Incorporating such information into computational frameworks can significantly enhance their modeling capabilities, particularly in cases with limited data availability.

In light of the challenges and limitations inherent in existing Bayesian Optimization (BO)-based design frameworks, this study proposes a physics-informed BO framework. This framework introduces physics into the Gaussian Process (GP) kernel to explore potential efficiency enhancements in material system design and the discovery of optimal processing parameters. The proposed approach combines the advantages of traditional BO techniques—specifically, the extraction of complex relationships from data—with the benefits of employing known governing equations for physical modeling.

This work lays a foundation for the application of physics-infused kernel design within the BO framework, opening up new possibilities across various materials science applications. Initially, we illustrate how the modeling capability of GPs is enhanced when equipped with physics-infused kernels, using synthetic function examples. Subsequently, we apply this concept to construct a physics-informed BO, utilizing it to design NiTi shape memory alloys aimed at maximizing phase transformation temperature, while ensuring prescriptive characteristics of the distribution of second-phase (i.e. Ni₄Ti₃) strengthening particles.

RESULTS AND DISCUSSION

Incorporating physical information into Bayesian optimization

The performance of BO is largely dependent on the GP's ability to model the underlying objective function accurately. While expanding the statistical information, such as training data or observations, can enhance GP's probabilistic modeling performance, for many applications (such as materials discovery/optimization), generating sufficient data could pose a significant computational burden, potentially creating a bottleneck in the process.

As previously noted, adequately modeling the underlying function and its accompanying uncertainty is essential for a prudent selection of the next optimal experiment within a BO framework. GPs employ the concept of distance to draw correlations between existing observations and unobserved locations in the design space, using a kernel function. However, in black-box surrogate modeling, distance serves as the sole metric to establish the correlation between observed and unobserved locations, and to capture the input-output relationship. This approach does not inherently promote the incorporation of any accessible physical information, which could potentially enhance the model's efficacy.

Next, we'll illustrate how incorporating physical principles into kernels can enhance the probabilistic predictions made by GPs. In each scenario, kernels can be meticulously adjusted to incorporate the physical information intrinsic to the input-output relationship of the underlying objective function. This alteration can be perceived as adjusting the input space along various dimensions to govern the correlation between two points, taking into account not only their relative distance but also their specific locations.

To put this into practice, we've implemented our proposed physics-informed BO (PIBO) framework to efficiently identify processing parameters for NiTi shape-memory alloys. The objective is to maximize the temperature at which the shape-memory transformation occurs, thereby optimizing the performance of these materials.

Demonstration on synthetic functions

In this study, we have created two test scenarios to demonstrate how blending theoretical and statistical information can bolster the probabilistic modeling capabilities of GPs. For both black-box and physics-informed scenarios, we generate GP predictions using identical statistical information, that is, the same training data. In Gaussian Processes (GPs), the concept of distance plays a pivotal role in correlating an unobserved location with observed ones (i.e., training data) to estimate the response. This correlation is quantified using a kernel function, the behavior of which is governed by its hyperparameters. A deep comprehension of the response's characteristics can guide the selection of an appropriate kernel for modeling the objective function. Notably, the squared exponential kernel is frequently chosen as it adeptly captures the variability inherent in many physical systems; these systems typically do not exhibit jagged response surfaces. Furthermore, the infinitely differentiable nature of the squared exponential kernel renders it a valuable covariance function. Alternative kernels, such as the Matern and Ornstein-Uhlenbeck covariance functions^{70,71}, provide additional means to regulate the smoothness of the modeled objective function and to depict response surfaces with increased irregularities.

In the black-box modeling example, we employ the squared exponential kernel function, and for the physics-informed case, we use a modified version of this function. This kernel modification transforms the space along various dimensions, ensuring that the correlation between observations depends not just on their relative (Euclidean) distance but also on the location of each point, offering a more comprehensive perspective. The schematic of both scenarios is illustrated in Fig. 1.

Our first example is given by Eq. (1). In contrast with black-box modeling—where the only information available is a limited set of observations linking some x values to their corresponding $f(x)$ values—, in the physics-informed scenario we can introduce additional knowledge about the relationship between $\sqrt{x}$ and $f(x)$. On the other hand, in black-box modeling, the correlation between different observations solely depends on their relative distance, disregarding their actual location. However, given the same distance, it's reasonable to assume a stronger correlation between observations with larger x values as the function's rate of change diminishes with increasing x .

$$f(x) = 1 + 2 \times \sqrt{x} \quad (1)$$

As a result, we can adjust the kernel function to use $\sqrt{x}$ instead of x for correlation calculations, as shown in Eq. (2). In essence, we're compressing the input space so that points within the same distance in this compressed space will have the same correlations. The kernel function does more than just control correlation; it also transforms the response into a linear form. This linear representation is advantageous because linear functions are among the simplest to model with minimal observations. Therefore, the process of learning a linear function and subsequently mapping it back to the original space proves to be substantially more efficient than attempting to learn the form of the original objective function directly. As such, the kernel function is tailored so that the transformed objective function aligns closely with a linear representation when the exact equation is known. In cases where the objective function's variability is only partially understood, the kernel ensures the function manifests as a smoother curve,




Fig. 1 Schematic representation comparing black-box and gray-box Bayesian optimization. In the gray-box approach, the Bayesian optimization framework is enriched with underlying physical information derived from within the system under study. By integrating statistical data with physical insights, a more data-efficient design framework is achieved. Consequently, the optimum design is uncovered with fewer experimental iterations, and the overall modeling uncertainty is significantly reduced.

mirroring a quasi-linear behavior.

$$k(x, x') = \sigma_s^2 \exp\left(-\frac{(\sqrt{x} - \sqrt{x'})^2}{2l^2}\right) \quad (2)$$

Figure 2 showcases a comparison between black-box and physics-informed modeling across various length scales. Evidently, with the same set of training data, physics-informed modeling excels in tracking the function's variability trend. This superior performance is attributed to the refined correlation computations achieved by kernel modification. Our experiments with different length scales reveal that while enhancing the correlation strength uniformly might improve smoothness, it doesn't guarantee a more precise capture of function variability. Though such an approach may seem effective in this specific instance, it could compromise the GP's probabilistic modeling, especially on irregular response surfaces, thereby risking issues of underfitting.

The second example delves into a more complex mathematical representation, as depicted in Eq. (3). In this scenario, we assume a partial awareness of the underlying theoretical (or physics-based) information, particularly regarding the relationship among $g(x)$, $\frac{1}{x}$, and $\log(x)$. This test problem mirrors scenarios frequently encountered in real-world design tasks, where the exact impact of design variables might remain elusive. Nonetheless, a designer might possess preliminary understanding or insights related to the physical principles at play, such as the association between the outcome and the inverse of a design variable.

$$g(x) = \exp\left(\frac{2}{x}\right) \times \left(1 - \exp\left(\frac{x^{1.4} - \log(x)}{x}\right)\right) \quad (3)$$

Accordingly, the kernel is modified as in Eq. (4) to import partially-known physical information to the kernel. This includes the impact of $\frac{1}{x}$ and $\log(x)$ on $g(x)$.

$$k(x, x') = \sigma_s^2 \exp\left(-\frac{\left(\frac{1}{x} - \frac{1}{x'}\right)^2 + (\log(x) - \log(x'))^2}{2l^2}\right) \quad (4)$$

Figure 3 showcases the results obtained using various length scales for both black-box and physics-informed scenarios. While

increasing smoothness aids in achieving a better fit in this instance, it doesn't necessarily capture the correct form of the response surface. In contrast, physics-informed modeling excels, adeptly capturing the pronounced decrease near the input space's lower boundary and precisely representing the extremum of $g(x)$.

Demonstration on precipitation modeling

Tailoring the transformation temperature of NiTi shape-memory alloys is a materials design problem that has received significant attention in recent years due to the many potential applications of shape-memory actuators with finely controlled transformation temperatures^{72–74}. The transformation temperature of near-equiatomic NiTi is known to be highly sensitive to the amount of free nickel in the matrix, with an increase of 1% Ni corresponding to as much as a 90 K decrease in the transformation temperature⁷⁵. Furthermore, the Ni-content of the matrix — and thus, the transformation temperature of the SMA — can be controlled through the use of heat treatments to form Ni-rich Ni_4Ti_3 precipitates⁷⁶. It's noteworthy that beyond simply regulating the transformation temperature, the attributes of the strengthening precipitate population, such as interparticle distance, offer additional means to adjust the transformation's characteristics. For instance, shorter interparticle distances can mitigate the hysteresis of the transformation, owing to the elastic strain fields surrounding coherent Ni_4Ti_3 precipitates^{77,78}. A testament to this approach's efficacy can be seen in recent work that leveraged this strategy to achieve stable transformation in additively manufactured NiTi alloys⁷⁹.

The precipitation behavior during heat treatment—which influences the final matrix composition and transformation temperature—is intricately governed by initial composition, time, and temperature. Historically, finding the best heat treatments for precipitation-strengthened SMAs has relied on extensive experimentation due to a lack of predictive physics-based models linking chemistry, processing history, and microstructure⁸⁰. Even when such models exist, utilizing them optimally can be computationally burdensome. Previously, one of our authors, along with collaborators, applied non-physics-aware Bayesian




Fig. 2 Comparison of physics-informed and black-box modeling of Eq. (1). All Gaussian processes are initialized with identical training data at different length scales to examine the impact of smoothness as well as utilization of physics-infused kernel. These results show that increasing smoothness does not have the same impact as importing physics in modeling a function and may result in underfitting issues in more complex examples.

Optimization (BO) to pinpoint the best microstructures in SMAs⁸¹. However, this method primarily addressed the property-microstructure relationship without delving into the crucial linkage between microstructure and its origin in chemistry/processing.

Given the complexity of Ni_4Ti_3 precipitate evolution within off-stoichiometric NiTi-based SMAs, optimizing heat treatment is daunting. This makes NiTi-based SMAs an exemplary candidate to illustrate the potential of our proposed Physics-Informed Bayesian Optimization (PIBO) method. We emphasize that while our framework targets computational ‘black-boxes’, it’s adaptable to purely experimental scenarios: For PIBO, the origin of the data isn’t vital, since the connections between inputs and outputs are established through stochastic models—e.g. Gaussian Processes (GPs)—that act as surrogates of the real (or virtual) system to optimize.

To model the impact of heat treatments on transformation times, Kawin⁸², an open-source implementation of the Kampmann-Wagner Numerical (KWN) model for phase precipitation, was used to model the volume fraction and distribution of Ni_4Ti_3 precipitates resulting from a given heat treatment. KWN is a mean-field precipitation approach that traces the evolution of precipitate populations in a bulk volume, rather than individual particles. In this model, a particle size distribution is represented as a continuous function, with particles binned based on their size. During each simulation step, processes like nucleation, growth, and dissolution are accounted for to update the particle size distribution, while ensuring adherence to mass balance and continuity equations. Nucleation is framed according to the Classical Nucleation Theory (CNT) where nucleation events are driven by heterophase fluctuations in a metastable solid solution, determining the critical nucleus size and driving force for

nucleation. Precipitate growth is modeled by coupling the matrix composition to the growth of the precipitate, whose coupling is given by the growth rate equation:

$$\frac{dR}{dt} = \sum_j \frac{D_{ij} x_j^\infty - x_j^a}{R x_i^\beta - x_i^a} \quad (5)$$

where R is precipitate radius, D_{ij} is the diffusivity, x_j^∞ is far-field composition, x_i^a is the interfacial composition in the matrix phase, and x_i^β is the interfacial composition in the precipitate phase. The model naturally encapsulates coarsening as smaller particles dissolve, shifting rightward the overall size distribution. Numerically, particles below a specific size are deemed fully dissolved and excluded from further computations. In this work, the thermodynamic and kinetic model for Ni-Ti from⁸³ is used. Further details on the general numerical modeling approach can be found in the Kawin paper⁸². For each simulated heat treatment condition, a final matrix composition was calculated and used with the model for NiTi transformation temperature developed by Frenzel et. al. (Eq. (6))⁷⁵:

$$T_{ms}(x_{\text{Ni}}) = \left(A + \frac{(B * x_{\text{Ni}})}{\text{at}} \% + C * D^{x_{\text{Ni}}/\text{at}\% - 50} \right) \quad (6)$$

where $A = 4511.2373$, $B = -83.42425$, $C = -0.04753$, and $D = 204.86781$. $T_{ms}(x_{\text{Ni}})$ is the start temperature for the martensitic transformation of NiTi and was used as the observable to optimize in the following Bayesian Optimization. The design space in question was defined by initial compositions ranging from 0.50 to 0.52 Ni, heat treatment temperatures ranging from 650K to 1050K, and heat treatment times ranging from 0 to 18,000 s (five hours).

Given that the precipitation of second-phase particles is the phenomenon that ultimately affects the transformation




Fig. 3 Comparison of physics-informed and black-box modeling of Eq. (3). All Gaussian processes are initialized with identical training data at different length scales to examine the impact of smoothness as well as utilization of physics-infused kernel. As shown, using a physics-infused kernel helps to catch the correct form of function variability and extremum approximations. When the function displays irregular variability between observations, black-box modeling is insufficient to accurately represent the function's true form.

temperature, it is crucial to understand how the volume fraction of precipitates affects the transformation temperature of the matrix. In this context, we can assume that the transformation temperature depends on the composition of the matrix, which can in turn be modulated through the precipitate volume fraction—in short, the higher the volume of precipitates, the greater the depletion of Ni in the matrix, which subsequently results in a decrease in transformation temperature. The precipitate distribution is a function of time (as well as temperature and initial solution-annealed matrix composition)^{84,85}. The precipitation process, in turn, is due to nucleation and growth where the growth is a time-varying diffusion-controlled process.

To this end, if we assume that atoms in a solid are transported through long-range diffusion, a random walk can be a simplistic, yet effective, first approximation to capture atomic diffusion in a lattice^{86,87}. It is noteworthy to mention that the phenomenological Langevin equation that represents a minimalist description of the stochastic motion of a random walk reveals that diffusion-controlled processes must be proportional to $\sqrt{x_{time}}$ ⁸⁸. We can thus exploit this scientific prior knowledge to design the kernel used to describe the dependence of observations on time.

We note that there are several other monotonically increasing forms (i.e. power laws) describing diffusional processes that could be potential candidates to design our kernel. For example, cubic law dependence (i.e., $\sqrt[3]{x_{time}}$) is typically observed when examining the diffusional growth of (oxide) layers in which the dominant transport mechanism is through grain boundary diffusion^{89,90}. In cases in which grain boundary diffusion is accompanied by grain coarsening—such as concurrent intermetallic layer growth⁹¹—the power law describing diffusional processes can take the form $\sqrt[4]{x_{time}}$. However, these other types of power laws depend on

processes that are not directly mapped to random bulk diffusion and do not constitute the most straightforward or fitting assumptions for a plausible kernel structure for *this specific problem*.

The relationship between heat treatment temperature and precipitate volume fraction can also be exploited to further refine our kernel function. Every set of (composition, temperature) inputs corresponds to an equilibrium precipitate phase fraction. While the rate at which the equilibrium phase fraction is approached (and the final volume fraction attained within the proscribed treatment time) is highly non-monotonic and dependent on the kinetics, the equilibrium phase fraction itself is always inversely proportional to the heat treatment temperature. At higher temperatures the solubility limit of the matrix increases, decreasing the driving force for precipitation and the equilibrium volume fraction. As such, simply scaling our kernel function by $1/T$ can provide a decent first approximation of the overall trend between temperature and phase fraction even though the true nature of this relationship is more complex and hard to quantify.

Composition, on the other hand, is best left directly proportional to the objective in the kernel function. Increasing solute concentration will also increase the driving force for precipitation, equilibrium phase fraction, and growth rate of the precipitates, directly increasing the final volume fraction. While this is again only a simplistic approximation of the true relationship, it describes the overall trend between composition and precipitation behavior across the input space.

From these three relationships a modified squared exponential kernel function can be created, shown in Eq. (7), that adjusts the design space along the dimensions of heat treatment time and temperature to better account for the impact of each variable on




Fig. 4 Maximum transformation temperature found using physics-informed and black-box BO. The results are averaged over 50 replication of simulations, showing the mean and 95% confidence intervals. Each simulation ran for 30 iterations. About 15–20 iterations are sufficient for physics-informed BO to reach the maximum.

precipitation behavior.

$$k(\mathbf{x}'', \mathbf{x}') = \sigma_s^2 \exp\left(-\frac{(x''_{\text{comp}} - x'_{\text{comp}})^2}{2l_{\text{comp}}^2}\right) \times \exp\left(-\frac{(\sqrt{x''_{\text{time}}} - \sqrt{x'_{\text{time}}})^2}{2l_{\text{time}}^2}\right) \times \exp\left(-\left(\frac{1}{x''_{\text{temp}}} - \frac{1}{x'_{\text{temp}}}\right)^2 \times \frac{l_{\text{temp}}^2}{2}\right) \quad (7)$$

To ensure that the BO produces non-trivial solutions, constraints were imposed on the design problem. Trivial solutions were defined as input sets that result in a negligible amount of Ni_4Ti_3 phase precipitation, with the transformation temperature being dominated only by the initial composition rather than the heat treatment parameters. Two classifiers were constructed to filter the design space, using labeled data from previously evaluated precipitation simulations across 2000 sets of compositions and processing parameters. The feasible space was defined as regions where the Ni_4Ti_3 volume fraction was greater than 0.01 and the mean inter-particle distance was lower than 5×10^{-8} m. Only potential inputs within the space identified by the classifiers and satisfying these conditions were considered for the BO.

To evaluate the efficacy of both black-box and physics-informed BO, we carried out 50 design simulations, each encompassing 30 iterations. Figure 4 presents the results, illustrating the average transformation temperature achieved at each iteration, complete with 95% confidence intervals. As evident from the figure, the physics-informed BO consistently surpasses the black-box BO. Specifically, it identifies combinations of compositions and processing parameters that lead to higher transformation temperatures in fewer iterations. On average, the physics-informed approach reaches the peak theoretical transformation temperature (around 340K) in just 15 iterations. In contrast, the black-box BO doesn't achieve this benchmark even within the span of 30 iterations. This disparity arises because the black-box method lacks insights into the fundamental physics of the problem, necessitating more iterations to pinpoint optimal design regions.

The solution sets corresponding to each run are shown in Fig. 5. It can be seen that physics-informed BO is able to explore the design space more effectively by producing a wider range of solutions within the same number of iterations. This also indicates




Fig. 5 The solutions corresponding to maximum transformation temperature in all 50 replications of simulations. Panels (a) and (b) show the initial composition and time as the solutions to 50 replications of each scenario. In the case of physics-informed BO, the design space has been searched more effectively to discover sets of processing parameters and compositions that result in the maximum transformation temperature. Black-box BO hesitates to explore as it is solely relying on statistical information, and our investigation shows that 30 iterations are not sufficient to discover optimum solutions in the design space under a physics-agnostic scheme.

that numerous potential combinations of time, temperature, and composition can be used to produce the same final transformation temperature, a conclusion that might be missed using only black-box BO.

For each of these solutions, Fig. 6 illustrates the volume fraction and mean inter-particle distance of the precipitates, verifying the classifiers used by showing that all values are within the desired ranges defined by this design problem. Figure 7 displays the optimal regions pinpointed by both physics-informed and black-box BO approaches. Leveraging the underlying physics, the PIBO framework excels in navigating the design space, consistently identifying superior solutions compared to the physics-agnostic BO, as illustrated in Fig. 4. PIBO's proposed solutions showcase greater variance, especially evident along the 'treatment time' axis. In contrast, the black-box BO tends to confine its explorations to a limited window, yielding less optimal solutions. The informed assumptions that the time and temperature axes should scale as $\sqrt{t}$ and $1/T$ respectively nudges PIBO to consider designs with extended treatment durations and more varied temperatures. Meanwhile, black-box BO seems hesitant to venture too far along either axis.




Fig. 6 Volume fraction and mean inter-particle distance of discovered sets of solutions shown in Fig. 5. Panels (a) and (b) show the final volume fraction and mean inter-particle distance values as the design constraints associated with solutions of 50 replications of each scenario. The results confirm all solutions satisfy the design constraints as both volume fraction and mean inter-particle distances are within the desired range that results in non-trivial solutions.

The final concentration of Ni left in the matrix after heat treatment, which is directly linked to the transformation temperature via Eq. (6), is ultimately dependent on the kinetics of the precipitation reaction and distribution of precipitates throughout the matrix as the reaction passes successively through the nucleation, growth, and coarsening-dominated regimes. While some sets of conditions initially rapidly approach the equilibrium (maximum) phase fraction, slower reactions dominated by growth instead of coarsening may reach a higher phase fraction (and therefore lower matrix Ni% and higher transformation temperature) within the same heat treatment period.

This phenomenon can be clearly seen in Fig. 5, where the lack of scaling led the black-box BO to consistently pick solutions with a similarly low initial composition, temperature, and heat-treatment time. These solutions correspond to a scenario where low temperature creates a high nucleation driving force, leading to the rapid creation of a fine dispersion of small particles throughout the matrix (seen in Fig. 6 as a small inter-particle distance). Once enough solute has been depleted to halt the nucleation of new particles, the reaction becomes largely driven by sluggish low-temperature coarsening kinetics, and the rate at which transformation temperature increases slows considerably.

As seen in Fig. 7 the physics-informed BO was instead able to generate a wide variety of potential solutions ranging from high-temperature, high-duration solutions with a smaller number of




Fig. 7 Optimal solutions discovered by physics-informed and black-box BO scenarios. While physics-informed scenario can explore the design space efficiently to discover any optimal solutions, black-box BO is not capable of getting close to those regions within the 30-iteration limit. Larger marker size indicates a higher heat treatment temperature which is also emphasized by the color.

particles growing consistently in a controlled manner to low-temperature, low-duration solutions that nucleated rapidly but efficiently, leaving little excess solute behind after the initial nucleation phase. From a materials design standpoint, many of these solutions are more desirable than those found via black-box BO because precipitation occurs in a more controlled manner, helping to identify regions of the input space where the final property (transformation temperature) can potentially be finely tuned.

Adjusting the kernel function in BO frameworks slightly affects computational cost. Yet, the physics-informed approach may extend simulation times, exploring a broader design space. On average, using Intel Xeon 6248R, 3.0GHz processors, a BO run takes about 24 min for 30 iterations, while a PIBO run takes roughly 26 min, and such a difference is shown in Fig. 5b. It's worth noting that this doesn't directly translate to efficiency. In our tests, PIBO rapidly achieved the maximum transformation temperature, whereas black-box BO failed to do so within 30 iterations.

Our study highlights the importance of incorporating physical knowledge into BO frameworks to improve their performance in discovering optimal design regions. By infusing statistical information with theoretical insights, we strengthened the GP's probabilistic modeling capability, resulting in reduced data dependency and faster convergence to the optimal design. This is particularly important when experimental data collection is required instead of computational simulations. The incorporation of physical knowledge not only improves the performance of BO frameworks, but also allows for a deeper understanding of the underlying physics governing the system, which can lead to more informed and efficient design decisions. A future research direction could focus on developing data-driven approaches that can autonomously comprehend and leverage physical laws, without the need for human intervention, to further improve the performance of BO frameworks in materials design applications.

METHODS

In this study, we proposed to inject partially known physics of precipitation to BO by manipulating GP's kernel function. As a result, the precipitation model is assumed as a gray-box rather

than a black-box due to partially known information inside the box. In the following, the different components of the BO framework and precipitation model are discussed.

Gaussian process regression

Gaussian processes are widely used for modeling objective functions in various fields due to their ability to provide probabilistic predictions with low computational costs⁸. The underlying concept of GPs relies on the correlation between data points in an input space, which enables the prediction of model uncertainty. This property is essential in Bayesian approaches that require probabilistic predictions in unobserved regions of the input space. Additionally, GPs are flexible and can be easily manipulated to suit the modeling requirements of different applications.

Assuming there are N previously observed data denoted by $\{\mathbf{X}_N, \mathbf{y}_N\}$, where $\mathbf{X}_N = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ and $\mathbf{y}_N = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_N))$, then the GP prediction at unobserved location $\mathbf{x}$ is given by the normal distribution:

$$f_{\text{GP}}(\mathbf{x})|\mathbf{X}_N, \mathbf{y}_N \sim \mathcal{N}(\mu(\mathbf{x}), \sigma_{\text{GP}}^2(\mathbf{x})) \quad (8)$$

where

$$\begin{aligned} \mu(\mathbf{x}) &= K(\mathbf{X}_N, \mathbf{x})^T [K(\mathbf{X}_N, \mathbf{X}_N) + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{y}_N \\ \sigma_{\text{GP}}^2(\mathbf{x}) &= k(\mathbf{x}, \mathbf{x}) - K(\mathbf{X}_N, \mathbf{x})^T [K(\mathbf{X}_N, \mathbf{X}_N) + \sigma_n^2 \mathbf{I}]^{-1} K(\mathbf{X}_N, \mathbf{x}) \end{aligned} \quad (9)$$

with k as a real-valued kernel function, $K(\mathbf{X}_N, \mathbf{X}_N)$ as a $N \times N$ matrix with m, n entry as $k(\mathbf{x}_m, \mathbf{x}_n)$, and $K(\mathbf{X}_N, \mathbf{x})$ is a $N \times 1$ vector with m^{th} entry as $k(\mathbf{x}_m, \mathbf{x})$. The term σ_n^2 is used to model observation error. A popular choice for kernel function is squared exponential:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_s^2 \exp \left(- \sum_{h=1}^d \frac{(x_h - x'_h)^2}{2l_h^2} \right) \quad (10)$$

where d is the dimensionality of the input space, σ_s^2 is the signal variance, and l_h , where $h = 1, 2, \dots, d$, is the characteristic length-scale to determine correlation strength between observations within dimension h .

Precipitation modeling

Kawin, the software utilized in this framework to model precipitation behavior, is based on the Kampmann-Wagner Numerical (KWN) model⁸³ for phase precipitation. This model employs discretized size "bins" to simulate the evolution of a particle size distribution during nucleation, growth, and coarsening. It strikes a balance between the computational efficiency of mean-field modeling and the data accuracy of phase-field modeling.

The KWN model is rooted in classical nucleation theory, which calculates the size and nucleation rate of particles at each time step based on the thermodynamic and physical properties of the matrix and precipitate phases. The model discretizes the particle size distribution into size "bins" and adds newly nucleated particles to the corresponding bin. All particles within a bin are assumed to have the same radius. The model then simulates growth and coarsening as the flux between size bins caused by diffusion-controlled growth and the Gibbs-Thomson effect. The matrix composition and size distribution are then updated using mass-balance and continuity equations. Overall, the KWN model provides a balance between the computational efficiency of mean-field modeling and the data fidelity of phase-field modeling.

Ni_4Ti_3 is a metastable phase that exists in a state of high lattice distortion, forming semi-coherent plate-shaped precipitates oriented only along the $\langle 111 \rangle$ planes of the B2 matrix phase⁹². Kawin was chosen to model this behavior because of its native elastic energy calculations and ease with which it could be modified to suit such an atypical precipitation case. The

thermodynamic characterization of the NiTi system was obtained from Povoden-Karadeniz et al.⁹³, crystallographic information from Naji et al.⁹⁴, elastic properties from Wagner and Windl⁹⁵, and diffusivity from Bernardini et al.⁹⁶.

DATA AVAILABILITY

Codes and data generated in this work are available at the following Github repository: <https://github.com/Danialkh26/PIBO>.

Received: 17 May 2023; Accepted: 21 November 2023;

Published online: 13 December 2023

REFERENCES

- Molkeri, A. et al. On the importance of microstructure information in materials design: Psp vs pp. *Acta Mater.* **223**, 117471 (2022).
- Ghoreishi, S. F., Molkeri, A., Arróyave, R., Allaire, D. & Srivastava, A. Efficient use of multiple information sources in material design. *Acta Mater.* **180**, 260–271 (2019).
- Moćkus, J. On bayesian methods for seeking the extremum. In *Optimization techniques IFIP technical conference*, 400–404 (1975).
- Jones, D. R., Schonlau, M. & Welch, W. J. Efficient global optimization of expensive black-box functions. *J. Global Optim.* **13**, 455–492 (1998).
- Frazier, P. I. A tutorial on bayesian optimization. Preprint at <https://arxiv.org/abs/1807.02811> (2018).
- Snoek, J., Larochelle, H. & Adams, R. P. Practical bayesian optimization of machine learning algorithms. In Pereira, F., Burges, C., Bottou, L. & Weinberger, K. (eds.) *Adv. in Neural Inf. Process. Syst.*, vol. 25, 2951–2959 (Curran Associates, Inc., 2012).
- Frazier, P. I. Bayesian optimization. In *Recent advances in optimization and modeling of contemporary problems*, 255–278 (Informa, 2018).
- Rasmussen, C.E., Williams, C.K.I. Gaussian processes for machine learning (AdaptiveComputation and Machine Learning), 8–29. (The MIT Press, Cambridge, MA, USA, 2005).
- Arróyave, R. et al. A perspective on bayesian methods applied to materials discovery and design. *MRS Commun.* **12**, 1037–1049 (2022).
- Clyde, M.A. Model Averaging 2nd edn, 320–335, Ch. 13 (Wiley-Interscience, Hoboken, NJ, USA, 2003).
- Hoeting, J. A., Madigan, D., Raftery, A. E. & Volinsky, C. T. Bayesian model averaging: a tutorial (with comments by m. clyde, david draper and ei george, and a rejoinder by the authors. *Stat. Sci.* **14**, 382–417 (1999).
- Leamer, E. Specification Searches: Ad Hoc Inference with Nonexperimental Data. John Wiley & Sons, New York, NY. (1978).
- Madigan, D. & Raftery, A. E. Model selection and accounting for model uncertainty in graphical models using occam's window. *J. Am. Stat. Assoc.* **89**, 1535–1546 (1994).
- Mosleh, A. & Apostolakis, G. The assessment of probability distributions from expert opinions with an application to seismic fragility curves. *Risk Anal.* **6**, 447–461 (1986).
- Reinert, J. M. & Apostolakis, G. E. Including model uncertainty in risk-informed decision making. *Ann. Nuclear Energy* **33**, 354–369 (2006).
- Riley, M. E., Grandhi, R. V. & Kolonay, R. Quantification of modeling uncertainty in aeroelastic analyses. *J. Aircraft* **48**, 866–873 (2011).
- Zio, E. & Apostolakis, G. Two methods for the structured assessment of model uncertainty by experts in performance assessments of radioactive waste repositories. *Reliab. Eng. Syst. Safety* **54**, 225–241 (1996).
- Julier, S. J. & Uhlmann, J. K. A non-divergent estimation algorithm in the presence of unknown correlations. In *Proceedings of the 1997 American Control Conference (Cat. No. 97CH36041)*, vol. 4, 2369–2373 (IEEE, 1997).
- Geisser, S. A bayes approach for combining correlated estimates. *J. Am. Stat. Assoc.* **60**, 602–607 (1965).
- Morris, P. A. Combining expert judgments: A bayesian approach. *Manag. Sci.* **23**, 679–693 (1977).
- Winkler, R. L. Combining probability distributions from dependent information sources. *Manag. Sci.* **27**, 479–488 (1981).
- Forrester, A. I., Sobester, A. & Keane, A. J. Multi-fidelity optimization via surrogate modelling. *Proc. Royal Soc. Math. Phys. Eng. Sci.* **463**, 3251–3269 (2007).
- Tran, A., Wildey, T. & McCann, S. smf-bo-2cogp: A sequential multi-fidelity constrained bayesian optimization framework for design applications. *J. Comput. Inf. Sci. Eng.* **20** (2020).
- Kennedy, M. C. & O'Hagan, A. Predicting the output from a complex computer code when fast approximations are available. *Biometrika* **87**, 1–13 (2000).

25. Khatamsaz, D. et al. Efficiently exploiting process-structure-property relationships in material design by multi-information source fusion. *Acta Mater.* **206**, 116619 (2021).
26. Couperthwaite, R. et al. Materials design through batch bayesian optimization with multisource information fusion. *Jom* **72**, 4431–4443 (2020).
27. Khatamsaz, D. et al. Adaptive active subspace-based efficient multifidelity materials design. *Mater. Design* **209**, 110001 (2021).
28. Khatamsaz, D. et al. Multi-objective materials bayesian optimization with active learning of design constraints: design of ductile refractory multi-principal-element alloys. *Acta Mater.* **236**, 118133 (2022).
29. Gill, P. E., Murray, W. & Wright, M. H. *Practical optimization* (SIAM, Philadelphia, PA, 2019).
30. Russi, T. M. *Uncertainty quantification with experimental data and complex system models*. Ph.D. thesis, UC Berkeley (2010).
31. Constantine, P. G., Dow, E. & Wang, Q. Active subspace methods in theory and practice: applications to kriging surfaces. *SIAM J. Sci. Comput.* **36**, A1500–A1524 (2014).
32. Constantine, P. G. *Active subspaces: Emerging ideas for dimension reduction in parameter studies* (SIAM, Philadelphia, PA, 2015).
33. Shan, S. & Wang, G. G. Survey of modeling and optimization strategies to solve high-dimensional design problems with computationally-expensive black-box functions. *Struct. Multidiscip. Optim.* **41**, 219–241 (2010).
34. Chinesta, F., Huerta, A., Rozza, G. & Willcox, K. *Model Reduction Methods*, 1–36 (John Wiley & Sons, Ltd, 2017).
35. Benner, P., Gugercin, S. & Willcox, K. A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Rev.* **57**, 483–531 (2015).
36. Ghoreishi, S. & Allaire, D. Adaptive uncertainty propagation for coupled multi-disciplinary systems. *AIAA J.* 3940–3950 (2017).
37. Joy, T. T., Rana, S., Gupta, S. & Venkatesh, S. Batch bayesian optimization using multi-scale search. *Knowledge-Based Syst.* **187**, 104818 (2020).
38. Khatamsaz, D., Arroyave, R. & Allaire, D. L. Materials design using an active subspace-based batch bayesian optimization approach. In *AIAA SCITECH 2022 Forum*, 0075 (2022).
39. Khatamsaz, D. et al. Bayesian optimization with active learning of design constraints using an entropy-based approach. *NPJ Comput. Mater.* **9**, 1–14 (2023).
40. Lakshminarayanan, M. et al. Comparing data driven and physics inspired models for hopping transport in organic field effect transistors. *Sci. Rep.* **11**, 23621 (2021).
41. Bohlin, T. P. *Practical grey-box process identification: theory and applications* (Springer Science & Business Media, 2006).
42. Kroll, A. Grey-box models: concepts and application. *N. Front. Comput. Intel. Appl.* **57**, 42–51 (2000).
43. Sohlberg, B. & Jacobsen, E. W. Grey box modelling—branches and experiences. *IFAC Proc. Vol.* **41**, 11415–11420 (2008).
44. Hauth, J. *Grey-box modelling for nonlinear systems*. Ph.D. thesis, Technische Universität Kaiserslautern (2008).
45. Asprion, N. et al. Gray-box modeling for the optimization of chemical processes. *Chem. Ing. Tech.* **91**, 305–313 (2019).
46. Molga, E. Neural network approach to support modelling of chemical reactors: problems, resolutions, criteria of application. *Chem. Eng. Process. Process Intensif.* **42**, 675–695 (2003).
47. Psychogios, D. C. & Ungar, L. H. A hybrid neural network-first principles approach to process modeling. *AIChE J.* **38**, 1499–1511 (1992).
48. Thompson, M. L. & Kramer, M. A. Modeling chemical processes using prior knowledge and neural networks. *AIChE J.* **40**, 1328–1340 (1994).
49. Cuomo, S. et al. Scientific machine learning through physics-informed neural networks: where we are and what's next. *J. Sci. Comput.* **92**, 88 (2022).
50. Wang, S., Teng, Y. & Perdikaris, P. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM J. Sci. Comput.* **43**, A3055–A3081 (2021).
51. Sahli Costabal, F., Yang, Y., Perdikaris, P., Hurtado, D. E. & Kuhl, E. Physics-informed neural networks for cardiac activation mapping. *Front. Phys.* **8**, 42 (2020).
52. Kashinath, K. et al. Physics-informed machine learning: case studies for weather and climate modelling. *Philos. Trans. R. Soc. A* **379**, 20200093 (2021).
53. Haghighat, E., Raissi, M., Moure, A., Gomez, H. & Juanes, R. A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics. *Comput. Methods Appl. Mech. Eng.* **379**, 113741 (2021).
54. Lu, L. et al. Physics-informed neural networks with hard constraints for inverse design. *SIAM J. Sci. Comput.* **43**, B1105–B1132 (2021).
55. Karniadakis, G. E. et al. Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440 (2021).
56. Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
57. Cai, S., Wang, Z., Wang, S., Perdikaris, P. & Karniadakis, G. E. Physics-Informed Neural Networks for Heat Transfer Problems. *J. Heat Transfer* **143**, 060801 (2021).
58. Mao, Z., Jagtap, A. D. & Karniadakis, G. E. Physics-informed neural networks for high-speed flows. *Comput. Methods Appl. Mech. Eng.* **360**, 112789 (2020).
59. Zhang, E., Dao, M., Karniadakis, G. E. & Suresh, S. Analyses of internal structures and defects in materials using physics-informed neural networks. *Sci. Adv.* **8**, eabk0644 (2022).
60. Roy, A. M., Bose, R., Sundararaghavan, V. & Arróyave, R. Deep learning-accelerated computational framework based on physics informed neural network for the solution of linear elasticity. *Neural Networks* **162**, 472–489 (2023).
61. Roy, A. M. & Guha, S. A data-driven physics-constrained deep learning computational framework for solving von mises plasticity. *Eng. Appl. Artif. Intell.* **122**, 106049 (2023).
62. Dourado, A. & Viana, F. A. Physics-informed neural networks for missing physics estimation in cumulative damage models: a case study in corrosion fatigue. *J. Comput. Inf. Sci. Eng.* **20**, 061007 (2020).
63. Zhang, X.-C., Gong, J.-G. & Xuan, F.-Z. A physics-informed neural network for creep-fatigue life prediction of components at elevated temperatures. *Eng. Fract. Mech.* **258**, 108130 (2021).
64. Kapsuzoglu, B. & Mahadevan, S. Physics-informed and hybrid machine learning in additive manufacturing: application to fused filament fabrication. *Jom* **72**, 4695–4705 (2020).
65. Astudillo, R. & Frazier, P. I. Thinking inside the box: a tutorial on grey-box bayesian optimization. In *2021 Winter Simulation Conference (WSC)*, 1–15 (IEEE, 2021).
66. Palmerin, S. T. *Grey-Box Bayesian optimization: improving performance by looking inside the Black-Box*. Ph.D. thesis (Cornell University, 2020).
67. Ziatdinov, M. A., Ghosh, A. & Kalinin, S. V. Physics makes the difference: Bayesian optimization and active learning via augmented gaussian process. *Mach. Learn.: Sci. Technol.* **3**, 015022 (2022).
68. Vela, B., Khatamsaz, D., Acemi, C., Karaman, I. & Arróyave, R. Data-augmented modeling for yield strength of refractory high entropy alloys: A bayesian approach. *Acta Mater.* **261**, 119351 (2023).
69. Ladygin, V., Beniya, I., Makarov, E. & Shapeev, A. Bayesian learning of thermodynamic integration and numerical convergence for accurate phase diagrams. *Phys. Rev. B* **104**, 104102 (2021).
70. Genton, M. G. Classes of kernels for machine learning: a statistics perspective. *J. Mach. Learn. Res.* **2**, 299–312 (2001).
71. Bibbona, E., Panfilio, G. & Tavella, P. The ornstein–uhlenbeck process as a model of a low pass filtered white noise. *Metrologia* **45**, S117 (2008).
72. McCue, I. D. et al. Controlled shape-morphing metallic components for deployable structures. *Mater. Design* **208**, 109935 (2021).
73. Cao, Y. et al. Large tunable elastocaloric effect in additively manufactured ni–ti shape memory alloys. *Acta Mater.* **194**, 178–189 (2020).
74. Zadeh, S. H. et al. An interpretable boosting-based predictive model for transformation temperatures of shape memory alloys. *Comput. Mater. Sci.* **226**, 112225 (2023).
75. Frenzel, J. et al. Influence of ni on martensitic phase transformations in niti shape memory alloys. *Acta Mater.* **58**, 3444–3458 (2010).
76. Khalil-Allafi, J., Dlouhy, A. & Eggeler, G. Ni4ti3-precipitation during aging of niti shape memory alloys and its influence on martensitic phase transformations. *Acta Mater.* **50**, 4255–4274 (2002).
77. Zhu, J., Gao, Y., Wang, D., Zhang, T.-Y. & Wang, Y. Taming martensitic transformation via concentration modulation at nanoscale. *Acta Mater.* **130**, 196–207 (2017).
78. Zhu, J. et al. Making metals linear super-elastic with ultralow modulus and nearly zero hysteresis. *Mater. Horiz.* **6**, 515–523 (2019).
79. Lu, H. et al. Stable tensile recovery strain induced by a ni4ti3 nanoprecipitate in a ni50. 4ti49. 6 shape memory alloy fabricated via selective laser melting. *Acta Mater.* **219**, 117261 (2021).
80. Takabayashi, S., Tanino, K. & Kitagawa, K. Heat treatment effect on transformation properties of tini shape memory alloy film. *J. Soc. Mater. Sci. Japan* **46**, 220–224 (1997).
81. Solomou, A. et al. Multi-objective bayesian materials discovery: Application on the discovery of precipitation strengthened niti shape memory alloys through micromechanical modeling. *Mater. Design* **160**, 810–827 (2018).
82. Ury, N. et al. Kawin: An open source kammann–wagner numerical (kwn) phase precipitation and coarsening model. *Acta Mater.* **255**, 118988 (2023).
83. Kampmann, R., Eckerlebe, H. & Wagner, R. Precipitation kinetics in metastable solid solutions—theoretical considerations and application to cu–ti alloys. *MRS Online Proceedings Library (OPL)* **57**, 525 (1985).
84. Deschamps, A. & Hutchinson, C. Precipitation kinetics in metallic alloys: Experiments and modeling. *Acta Mater.* **220**, 117338 (2021).
85. Laplanche, G. et al. Phase stability and kinetics of σ -phase precipitation in crmnfecon high-entropy alloys. *Acta Mater.* **161**, 338–351 (2018).
86. Balluffi, R. W., Allen, S. M. & Carter, W. C. *Kinetics of materials* (John Wiley & Sons, 2005).

87. Lawler, G. F. & Limic, V. Introduction. In *Random walk: a modern introduction*, 1–20 (Cambridge University Press, 2010).
88. Gao, Y., Yu, T. & Wang, Y. Phase transformation graph and transformation pathway engineering for shape memory alloys. *Shape Mem. Superelasticity* **6**, 115–130 (2020).
89. Suvaci, E., Oh, K.-S. & Messing, G. Kinetics of template growth in alumina during the process of templated grain growth (tgg). *Acta Materialia* **49**, 2075–2081 (2001).
90. Xu, H. et al. Behavior of aluminum oxide, intermetallics and voids in cu–al wire bonds. *Acta Mater.* **59**, 5661–5673 (2011).
91. Park, M. S. & Arroyave, R. Formation and growth of intermetallic compound cu₆sn₅ at early stages in lead-free soldering. *J. Electron. Mater.* **39**, 2574–2582 (2010).
92. Ke, C., Cao, S. & Zhang, X. Three-dimensional phase field simulation of the morphology and growth kinetics of ni₄ti₃ precipitates in a niti alloy. *Model. Simul. Mater. Sci. Eng.* **22**, 055018 (2014).
93. Povoden-Karadeniz, E., Cirstea, D., Lang, P., Wojcik, T. & Kozeschnik, E. Thermodynamics of ti–ni shape memory alloys. *Calphad* **41**, 128–139 (2013).
94. Naji, H., Khalil-Allafi, J. & Khalili, V. Microstructural characterization and quantitative phase analysis of ni-rich niti after stress assisted aging for long times using the rietveld method. *Mater. Chem. Phys.* **241**, 122317 (2020).
95. Wagner, M. F.-X. & Windl, W. Elastic anisotropy of ni₄ti₃ from first principles. *Scr. Mater.* **60**, 207–210 (2009).
96. Bernardini, J., Lexcelent, C., Daróczy, L. & Beke, D. Ni diffusion in near-equiatomic ni–ti and ni–ti (–cu) alloys. *Philos. Mag.* **83**, 329–338 (2003).

ACKNOWLEDGEMENTS

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-22-2-0106. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government Purposes notwithstanding any copyright notation herein. Department of Energy (DOE) ARPA-E ULTIMATE Program through Project DE-AR0001427. DK acknowledges the support of NSF through Grant No. CDSE-2001333. RA acknowledges the support from Grants No. NSF-CISE-1835690 and NSF-DMREF-2119103. Part of this research was carried out at the Jet Propulsion Laboratory (JPL), California Institute of Technology, under a contract with the National Aeronautics and Space Administration

(80NM0018D0004). This research was supported by the JPL Strategic University Research Partnership (SURP) program.

AUTHOR CONTRIBUTIONS

R.A. and D.K. designed the problem. D.K. implemented the Bayesian optimization framework. R.N. and R.O. developed the precipitation model. D.K. and R.N. wrote the first version of the manuscript. Finally, all the authors edited and reviewed the final version of the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Correspondence and requests for materials should be addressed to Danial Khatamsaz.

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) 2023