Acta Materialia 180 (2019) 260–271

Contents lists available at ScienceDirect

Acta Materialia

journal homepage: www.elsevier.com/locate/actamat

Full length article

Efficient use of multiple information sources in material design

Seyede Fatemeh Ghoreishi [a, d], Abhilash Molkeri [b], Raymundo Arróyave [a, b, c], Douglas Allaire [a], Ankit Srivastava [b, *]

[a] Department of Mechanical Engineering, Texas A&M University, College Station, TX, USA
[b] Department of Materials Science and Engineering, Texas A&M University, College Station, TX, USA
[c] Department of Industrial and Systems Engineering, Texas A&M University, College Station, TX, USA
[d] Institute for Systems Research, University of Maryland, College Park, MD, USA

ARTICLE INFO

Article history:
Received 19 May 2019
Received in revised form
3 September 2019
Accepted 8 September 2019
Available online 13 September 2019

Keywords:
Multiphase
Microstructure hull/Microstructure design
Modeling
Mechanical properties
Bayesian optimization

ABSTRACT

We present a general framework for the design/optimization of materials that is capable of accounting for multiple information sources available to the materials designer. We demonstrate the framework through the microstructure-based design of multi-phase microstructures. Specifically, we seek to maximize the strength normalized strain-hardening rate of a dual-phase ferritic/martensitic steel through a multi-information source Bayesian optimal design strategy. We assume that we have multiple sources of information with varying degrees of fidelity as well as cost. The available information from all sources is fused through a reification approach and then a sequential experimental design is carried out. The experimental design seeks not only to identify the most promising region in the materials design space relative to the objective at hand, but also to identify the source of information that should be used to query this point in the decision space. The selection criterion for the source used, accounts for the discrepancy between the source and the 'ground truth' as well as its cost. It is shown that when there is a hard constraint on the budget available to carry out the optimization, accounting for the cost of querying individual sources is essential.

1. Introduction

Integrated Computational Materials Engineering (ICME) [1,2], as currently understood, consists of the integration of multiple levels of computational tools, in combination with experiments, along the materials design/optimization chain to (i) establish quantitative Process-Structure-Property-Performance (PSPP) relationships; and (ii) exploit the so-established PSPP relationships for the acceleration of the materials design/optimization process. A challenge associated with ICME is the fact that the explicit integration of multiple computational materials tools remains an outstanding task [3].

While there have been some successes in terms of fully integrated ICME approaches to materials design [4–7], in most cases it is assumed that at each level/scale of simulation there is only one model that serves as a linkage along the PSPP relationship. This paradigm is somewhat restrictive as it is often the case that often times there are multiple models, with different levels of fidelity and associated (computational) costs, that could potentially be used as linkages to carry the modeling chain forward. Furthermore, existing frameworks do not explicitly account for the possibility of using other types of information, such as experiments, alongside computational models, within the same materials design/optimization framework.

Multifidelity methods have seen significant application in engineering design optimization. Most approaches build corrections to low-fidelity information from higher fidelity sources, such as adding global response surface corrections to low-fidelity models [8,9], using low-fidelity information for coarse-grained search while using high-fidelity function values for fine grained decisions [10,11], creating a response surface using both high- and low-fidelity results [12,13], and running higher-fidelity models when two or more lower-fidelity models disagree [14,15]. More formal multifidelity optimization frameworks use either a local approach, such as trust region model management [16–18], or a global approach constructed via interpolation of the high-fidelity objective function. For example, efficient global optimization, sequential Kriging optimization, and knowledge gradient optimization use a

* Corresponding author. .
E-mail address: ankit.sri@tamu.edu (A. Srivastava).

https://doi.org/10.1016/j.actamat.2019.09.009
1359-6454/

Gaussian process model to estimate the location of high-fidelity optima and guide multifidelity sampling [19–21]. Recent work by the authors has also considered the incorporation of correlated models and differing types of information sources [22–27]. Recent review articles (see e.g., Refs. [28,29]) also highlight the advances in multifidelity model fusion and the scope of the field.

While the ability to handle differing information sources is critical to a fully integrated ICME framework, we also have to further consider that ICME poses the establishment of linkages along the PSPP chain as a necessary, but not sufficient, condition towards the acceleration of the materials design/optimization process. ICME-based PSPP model chains tend to be computationally costly and a key outstanding challenge is how to utilize these ICME tools to efficiently explore the materials design space. Currently, high-throughput (HT) experimental [30–32] and computational [33] exploration of the materials space constitute the dominant paradigm. These approaches, however, tend to be sub-optimal when there are constraints on the available resources.

To overcome the challenges associated with the open-loop exploration (computational or otherwise) of the materials design space, notions of optimal experimental design—based on Bayesian Optimization, for example—have been put forward [34,35] and have been shown to be quite effective in carrying out the targeted search of optimal materials solutions by balancing the exploration and exploitation of the materials design space. However, most Bayesian Optimization based frameworks, similar to ICME-based works, rely on a single source of information (i.e., model or experimental response surface) to query the materials design space. In the absence of sufficient data—as it is most often the case in materials design/optimization problems—relying on a single model from the start is highly risky. Moreover, such an approach is limited as it implicitly requires the use of the most accurate approximation to the 'ground truth' available to query the design space at every step of the sequential materials design/optimization task. However, a high degree of fidelity to the 'ground truth' often comes at a considerable cost in terms of resources and time.

A subset of the present authors [35,36] have recently sought to address this challenge by adaptively selecting competing (non-parametric) theories or models relating materials features to their performance [35,36]. The framework begins the exploration of the materials design space by assuming that each of the competing models has equal weight. The weight of individual models is then adjusted by computing the Bayesian evidence relative to the acquired data. Bayesian model averaging is then used to find the best next point to query in the design space as determined by all the competing models with their relative contribution to the utility function being adjusted by their updated weight. The framework, however, took a very conservative approach to information fusion and treated each model as *statistically independent*, ignoring the correlations between the models that could potentially be exploited in the materials design/optimization scheme. While this may not be a major issue when using non-parametric models (models that have no internal structure based on/constrained by physics), this may be a significant limitation when the models available are all connected by the common underlying physical phenomena they intend to represent.

State-of-the-art of approaches to the computer/data-enabled discovery and/or design of materials has moved forward significantly but there remain important issues to address. Specifically, model-based ICME approaches focus on integration of tools along the PSPP chain assuming there is only one model/tool relevant to each level (or linkage) of the problem. On the other hand, data-centric approaches tend to focus on the brute-force exploration of the materials design space, without much emphasis on being able to inform decisions on where to explore next based on the knowledge already acquired. Furthermore, both types of frameworks tend to discount the need for allocating resources in an efficient manner.

In this work, we present a framework that addresses the challenges outlined above. Specifically, we tackle the challenge of optimizing the features of a dual-phase microstructure for a specific (mechanical) property/performance metric. We further assume that we have at our disposal a set of models of varying complexity (or computational cost) and fidelity. We also consider that we have access to a 'ground truth', in the form of an expensive microstructure-based finite element model that can be queried at much more expense than the other sources of information and that, therefore, should be queried as infrequently as possible, provided the design space represented by such 'ground truth' is sampled efficiently.

The (inexpensive) sources used in this work are fused by accounting for their mutual correlation as well as their correlation with the 'ground truth'. The framework accounts for the value of individual information sources in relation to the property we wish to optimize. The queried source is used to construct a fuse model that represents our best estimate for the response of the 'ground truth'. We then incorporate the fused model within a Knowledge Gradient framework in order to carry out, in a principled manner, two decisions: (i) which source should be used to query the materials design space; and (ii) which location of the materials design space to query; taking into account a fixed budget (cost) for queries to carry out before assessing the 'ground truth'.

We note that we have previously presented preliminary work related to this problem [27] although a major difference between the previous approach and the one presented here is the fact that here we not only consider a significantly higher number of information sources (six versus three), but also explicitly account for the computational cost associated with individual sources and perform the optimization task under hard constraints in terms of computational resources (represented as total computing time) available to carry out the materials design/optimization task.

While the framework is demonstrated in a computational context, the problem set up and its resolution closely mimics a typical materials design/optimization campaign: (i) there are more than one potential source of information (experimental and/or computational) about the system to be investigated; (ii) each of the information sources has different degree of fidelity with regards to the 'ground truth'; (iii) each source has different cost (monetary and/or otherwise); and (iv) there are hard constraints in time and resources available to complete the research campaign. The last point is often overlooked but constitutes a very important limitation to the effectiveness of any approach towards materials design/optimization. Indeed, at the end of the paper we provide ample (computational) evidence for the fundamental importance for accounting not only for the utility but for the cost of information sources when there are (as always) hard resource constraints.

2. Mechanical behavior of dual-phase microstructures

In this work, we demonstrate our materials design/optimization framework by focusing on the mechanical response of dual-phase advanced high strength steels. In these alloy systems, the microstructure consists of a relatively soft, ductile ferrite phase that is strengthened by the hard martensite phase [37]. The overall mechanical response of these dual-phase microstructures are controlled by the (non-linear) properties of the constituent phases, the microstructural features such as volume fraction, and the (non-linear) interaction among the constituent phases [38–41]. In principle, the overall response of these dual-phase microstructures can be tuned and optimized. Unfortunately, any approach to

microstructure design/optimization that does not rely on predictive models for the overall response of the microstructure will require extensive and expensive trial-and-error experimentation.

The overall mechanical response of composite dual-phase microstructures can be predicted with a high level of fidelity using single [38–41] or multi-scale [42,43] microstructure-based finite element calculations. Unfortunately, these models tend to be computationally costly, precluding their use to carry out direct optimization — of mechanical properties — over the entire microstructure space. The overall mechanical response of composite dual-phase microstructures can also be predicted using less computationally expensive reduced-order models with a lower level of fidelity. These reduced-order models include simple models that make strong assumptions on how strain [44], stress [45] or work of deformation [46] partitions between the constituent phases, and sophisticated micromechanical homogenization schemes that consider the dual-phase microstructure as a matrix with inclusions of another phase homogeneously dispersed within it [47]. The computational cost and predictive power of these reduced-order models also vary greatly, and the two are not necessarily correlated. The predictive power of a reduced-order model, here, refers to the range of microstructure space over which the model can accurately predict the microstructure-mechanical response correlation.

Recognizing the advantages and limitations of different potential cheap models used to predict the response of dual-phase microstructures, in this work we treat all of them—reduced-order models and micromechanical homogenization schemes—as sources of information with varying value and cost, with regards to the optimization problem at hand. In other words, in order to tackle the challenge of optimizing the dual-phase microstructures for enhanced mechanical performance, we assume that we have at our disposal a set of models of varying complexity (or computational cost) and fidelity. In order to demonstrate our framework, the high-fidelity microstructure-based finite element calculations are considered as 'ground truth'. Furthermore, the cost to query the 'ground truth' is taken to be much more expensive than the other sources of information.

Before we continue with the description of all the information sources and the finite element calculations, we note that the problem posed here is a drastic simplification of the real problem—even within a simulation-only setting—as we completely ignore the fact that microstructure is ultimately controlled via material chemistry and processing. A more realistic problem setup would consist of models incorporating the process/chemistry-microstructure connection as well as the effect of both the chemistry and the processing conditions on the properties of the constituent phases. For the sake of demonstration, in this work we ignore the process/chemistry-microstructure connection and limit our design space to simple microstructural descriptors.

2.1. Microstructure-based finite element modeling

The microstructure-based finite element modeling to predict the overall mechanical response of dual-phase microstructures is carried out using 3D representative volume elements (RVEs) as described in Refs. [27,41]. Fig. 1 shows two realizations of a 3D RVE of a dual-phase microstructure with about 50% volume fraction of the martensite (hard) phase. The RVEs are constructed using C3D8 brick elements of the ABAQUS/Standard element library [48], and have a dimension of $100\mu m \times 100\mu m \times 100\mu m$. The RVEs are subjected to monotonically increasing uniaxial tensile deformation under periodic boundary conditions.

In the finite element calculations, it is assumed that both the ferrite and the martensite phase follow an isotropic elastic-plastic constitutive relation, with identical Young's modulus—$E = 200$GPa—and Poisson's ratio—$\nu = 0.3$, and a Ludwik type strain-hardening response,

$$\tau^p = \tau_o^p + K^p \left(\varepsilon_{pl}^p\right)^{n^p}, \tag{1}$$

where τ^p is the flow stress, ε_{pl}^p is the plastic strain, τ_o^p is the yield strength, K^p is the strengthening coefficient, and n^p is the strain-hardening exponent of phase p. The values for these parameters are given in Table 1. The models are qualitative in nature and the parameters used are chosen in order to have a microstructure that represents a soft phase (ferrite) with a lower initial yield strength and a higher strain-hardenability than the hard phase (martensite) [38–41].

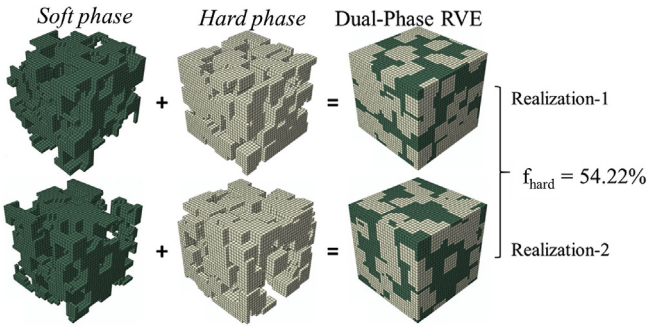

Fig. 1. Two realizations of the representative volume element (RVE) of a dual-phase microstructure with 54.22% (by volume) of a hard (martensite) phase. The two realizations refer to different distributions of the hard phase particles in the RVE with a fixed phase volume fraction.

2.2. Reduced-order models

The overall mechanical response of a dual-phase microstructure can also be predicted using reduced-order models constructed under different assumptions regarding the partitioning of strain, stress or work of deformation among the constituent phases. Three such reduced-order models considered here are: (i) the Voigt/Taylor model that assumes isostrain partitioning, (ii) the Reuss/Sachs model that assumes equipartitioning of stress and (iii) the isowork model that assumes both phases undergo identical works of deformation as the composite microstructure is deformed [49]. For all three reduced-order models the constitutive relation, $\tau^p = f(\varepsilon_{pl}^p)$, is assumed to follow Eq. (1), with the values of the parameters given in Table 1.

2.3. Micromechanical models

In addition to the simple reduced-order models presented above, we also exploit three more sophisticated micromechanical homogenization schemes as sources of information that provide estimates of the overall mechanical response of dual-phase microstructures. The first two homogenization schemes employed

Table 1
Parameterization of the Ludwik power law for the constituent phases of the dual-phase microstructure.

Constituent Phase, p	τ_o^p [MPa]	K^p [MPa]	n^p
Soft (ferrite)	300	2200	0.5
Hard (martensite)	1500	450	0.06

here are referred to as 'secant method - 1' and 'secant method - 2', whereas the third one is referred to as the 'elastic constraint' method. The secant method proposed by Weng [47], predicts the mechanical response of a two-phase composite microstructure based on Hill's weakening constraint power in a plastically-deforming matrix. In a dual-phase microstructure, where both phases are capable of undergoing plastic deformation, and the phase constitution covers the entire range, i.e., zero-to-one (0% phase 1 & 100% phase 2) to one-to-zero (100% phase 1 & 0% phase 2), it is difficult to decide which phase should be considered as an inclusion and which phase should be considered to be the matrix. Thus, the 'secant method - 1' is the secant method proposed by Weng [47], where the ferrite phase is considered the matrix, while the scenario when martensite is considered the matrix is referred to as 'secant method - 2'. It is expected that these two variants of the secant method will be valid in opposite regions in the microstructure design space. The third homogenization scheme, referred to as 'elastic constraint', is based on Kröner's treatment of the matrix-inclusion system under elastic constraints [47]. For this method, the final prediction does not depend on which phase, ferrite or martensite, is assumed as inclusion or the matrix. For all three homogenization schemes, 'secant method - 1', 'secant method - 2' and 'elastic constraint', the inclusion-inclusion interaction at finite concentration are accounted for by the Mori-Tanaka method. For all three homogenization schemes, the constitutive relation, $\tau^p = f(\varepsilon^p_{pl})$, is assumed to follow Eq. (1), with the values of the parameters given in Table 1.

2.4. Comparison of the predictions of 'information sources' and the 'ground truth'

In this work, we treat all the reduced-order models, Section 2.2, and micromechanical homogenization schemes, Section 2.3, as sources of information, while the microstructure-based finite element calculations, Section 2.1, are treated as 'ground truth'. The predicted flow curve of a dual-phase microstructure with 25% volume fraction of the hard (martensite) phase, using the reduced-order models and micromechanical homogenization schemes are compared against the finite element results in Fig. 2. As can be seen from the figure, compared to the finite element predictions, the isostress model significantly under-predicts the stress values at nearly all plastic strain levels. The isostrain and isowork models, on the other hand, over-predict the stress at low strain levels. However, at large strain levels their predictions are comparable with the finite element predictions. In addition, none of the reduced-order models are able to correctly predict the overall strain-hardening response of the dual-phase microstructure. Of all the micromechanical homogenization schemes, predictions of the Weng [47] secant model with the soft phase considered as the matrix (secant method - 1) are comparable with the finite element calculations at low strain levels. The predictions of the Weng [47] secant model with hard phase considered as the as matrix (secant method - 2) and the elastic constraint model are roughly the same, and at low strain levels both over predict the stress values.

The flow strength and the strain-hardenability of a material are two very important mechanical properties with practical implications for both its performance and manufacturability. These two fundamental mechanical properties can be reduced to a single metric by introducing a strength normalized strain-hardening rate, given as $(1/\tau)(d\tau/d\varepsilon_{pl})$. This quantity provides an indication of the ductility and formability of the material, with higher values corresponding to higher ductility and formability. Fig. 3 shows how $(1/\tau)(d\tau/d\varepsilon_{pl})$ varies with the volume fraction of the hard phase, f_{hard}, estimated at a plastic strain level of $\varepsilon_{pl} = 0.9\%$ from the microstructure-based finite element calculations. The figure shows that $(1/\tau)(d\tau/d\varepsilon_{pl})$ exhibits a maxima at a finite volume fraction of the hard phase. The variation of $(1/\tau)(d\tau/d\varepsilon_{pl})$ in the figure also exhibits small perturbations, indicated as error bars, arising from seven different realizations of the RVE representation of the dual-phase microstructure with identical phase constitutions. The response of $(1/\tau)(d\tau/d\varepsilon_{pl})$ predicted using the three reduced-order models and the three micromechanical homogenization schemes are also shown in Fig. 3. As can be seen, none of the approximate models or information sources are capable of reproducing the response predicted using the microstructure-based finite element calculations over the entire phase constitution space. The 'Secant method-1' approximation is capable of reproducing the response at small volume fractions of the hard phase but the discrepancy with

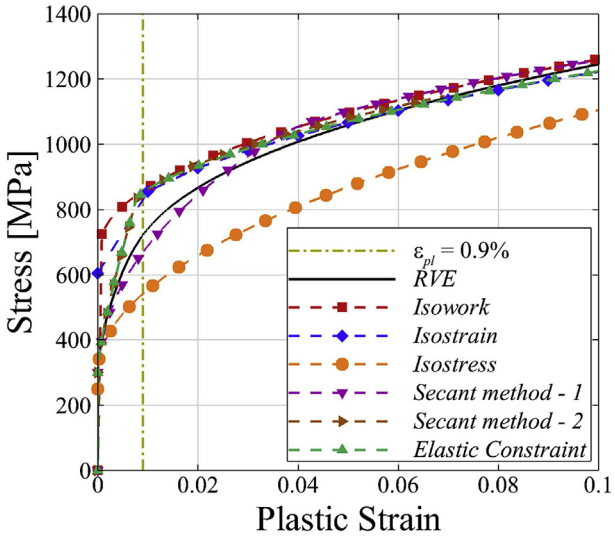

Fig. 2. Comparison of the predicted stress-plastic strain curves by the three reduced-order models and the three micromechanical homogenization schemes with the microstructure-based finite element calculations (RVE) of a dual-phase microstructure with 25% volume fraction of the hard phase.

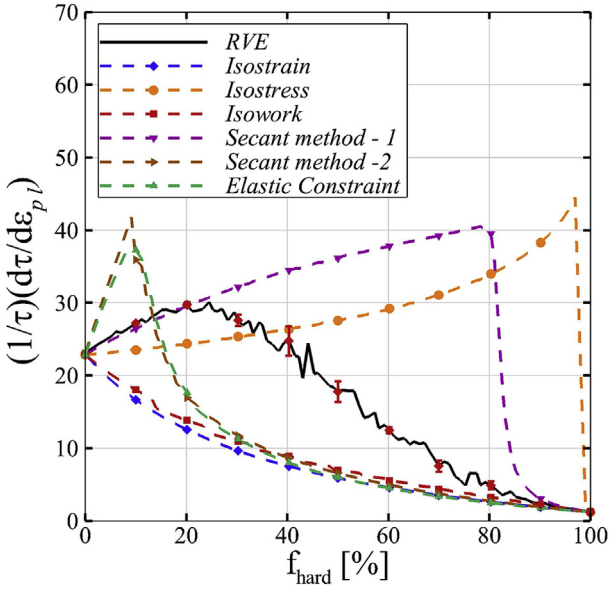

Fig. 3. Strength normalized strain-hardening rate $(1/\tau)(d\tau/d\varepsilon_{pl})$ at $\varepsilon_{pl} = 0.9\%$ as a function of the volume fraction of the hard phase, f_{hard}. Predictions of the three reduced-order models and three micromechanical homogenization schemes are compared to the microstructure-based finite element calculations (RVE).

the 'ground truth' becomes significant at volume fractions above 20%. Most models tend to converge towards the finite element predictions when the majority of the microstructure consists of the hard phase. In this regime, the microstructure exhibits no strain-hardenability.

The results presented in Fig. 3 clearly show that all the information sources—reduced-order models and micromechanical homogenization schemes—are incapable of reproducing the 'ground truth' with an acceptable level of fidelity. On the other hand, these information sources are cheap and their predictions differ from the 'ground truth' in a systematic manner. Thus, it may be possible to learn this discrepancy and use this knowledge to arrive at a more robust estimation of the 'ground truth' at an extremely low cost. The computational cost of each information source and the 'ground truth' is calculated based on their run time on a PC equipped with an Intel® Xeon® E5-2670 v2 (Ivy Bridge-EP) processor. The run time of a given information source was determined by tracking the wall-clock time using the 'tic' and 'toc' functions in MATLAB. To avoid the effect of external load due to memory constraints, each information source was queried 1,000 times and the average run time over 1,000 queries was considered as the run time or cost of each information source. The final assigned computational cost to each information source are given in Table 2. While all the information sources considered here are computational models and the costs for each information-source is simply the computational time, our approach is generic enough to incorporate any other type of information source or cost, as long as they can be represented in a consistent fashion.

3. Design framework

An optimization problem can be simply written as,

$$\mathbf{x}^* = \underset{\mathbf{x}\in\chi}{\operatorname{argmax}}\, f(\mathbf{x}), \tag{2}$$

where f is the objective function, and $\mathbf{x}$ is a set of design variables in the vector space χ, and $\mathbf{x}^*$ is the final design point. In real-world applications, an analytical form of the objective function is often unknown and/or expensive to evaluate, and there are oftentimes constraints placed on decision-making. Additionally, in most real-world applications, there are potentially multiple ways in which one can query the design space. One can use, for example, a combination of numerical simulations, experiments and expert opinions, to approximate the objective function with varying fidelity or accuracy, and with varying cost (monetary or otherwise). Here, we present a framework that exploits multiple information sources based on a trade off between the cost and fidelity of a specific information source by augmenting Eq. (2). Our framework seeks the value of $\mathbf{x}^*$ under an overall resource (allocated monetary budget or time) constraint.

Table 2
Computational cost of the different information sources used.

Information Source	Cost (seconds)	Normalized Cost
Reduced-order Models		
iso-strain	2.3×10^{-4}	1
iso-stress	1.0×10^{-3}	4.4
iso-work	4.7×10^{-1}	2.0×10^{3}
Micromechanical models		
secant method 1	3.8×10^{1}	1.7×10^{5}
secant method 2	8.4×10^{1}	3.7×10^{5}
elastic constraint	3.6×10^{1}	1.6×10^{5}
Finite element (RVE) or 'ground truth'		
RVE	7.2×10^{3}	3.1×10^{7}

A schematic of our framework is shown in Fig. 4. As shown in the figure, we first model the response of each information source through Gaussian process (GP)-based surrogates. Next, we proceed to fuse them using standard approaches for the fusion of normally distributed data. The fused means and variances in the input design space, χ, are then used to construct a fused GP model. The fused GP model is now used to determine the next design point and the information source to query while balancing the cost of the query and the value of such query relative to the objective function. Once the selected design point has been identified and a query has been made, the corresponding fused GP is updated. This loop is continued until the objective of the optimization problem has been met or the resource allocated to this optimization task has been exhausted. The framework is further described in detail in the subsequent paragraphs.

We assume that we have S information sources, $f_i(\mathbf{x})$, where $i\in\{1,2,\ldots,S\}$, available that can be used to approximate the objective function, $f(\mathbf{x})$, at $\mathbf{x}$. The response of each of the S information sources are first modeled through GP-based surrogates, Fig. 4. GPs for regression purposes consist of a nonparametric Bayesian approach to conditioning a probability distribution to training data [50]. GP regression models are widely used as surrogate models in engineering due to their flexibility and ability to be updated as more information is gained. In our work, the GPs are fit using data from previous queries, $\{\mathbf{X}_{N_i},\mathbf{y}_{N_i}\}$, where $\mathbf{X}_{N_i}=(\mathbf{x}_{1,i},\ldots,\mathbf{x}_{N_i,i})$ corresponds to the N_i input samples used to query the response of source i and $\mathbf{y}_{N_i}$ corresponds to the output—the domain of the information source correspond to the microstructural degrees of freedom (i.e., phase fractions) available. The posterior GPs distributions of each f_i, $f_{GP,i}(\mathbf{x})$, at any point $\mathbf{x}$ are

$$f_{GP,i}(\mathbf{x})\Big|\mathbf{X}_{N_i},\mathbf{y}_{N_i}\sim\mathcal{N}\Big(\mu_i(\mathbf{x}),\sigma^2_{GP,i}(\mathbf{x})\Big), \tag{3}$$

where

$$\mu_i(\mathbf{x})=K_i(\mathbf{X}_{N_i},\mathbf{x})^T\Big[K_i(\mathbf{X}_{N_i},\mathbf{X}_{N_i})+\sigma^2_{n,i}I\Big]^{-1}\mathbf{y}_{N_i}, \tag{4}$$

and

$$\sigma^2_{GP,i}(\mathbf{x})=k_i(\mathbf{x},\mathbf{x})-K_i(\mathbf{X}_{N_i},\mathbf{x})^T\Big[K_i(\mathbf{X}_{N_i},\mathbf{X}_{N_i})+\sigma^2_{n,i}I\Big]^{-1}K_i(\mathbf{X}_{N_i},\mathbf{x}). \tag{5}$$

Here, k_i is a real-valued kernel function associated with information source i over the input space, $K_i(\mathbf{X}_{N_i},\mathbf{X}_{N_i})$ is the $N_i\times N_i$ matrix whose m,n entry is $k_i(\mathbf{x}_{m,i},\mathbf{x}_{n,i})$, $K_i(\mathbf{X}_{N_i},\mathbf{x})$ is the $N_i\times 1$ vector whose m^{th} entry is $k_i(\mathbf{x}_{m,i},\mathbf{x})$ for information source i, and the term $\sigma^2_{n,i}$ can be used to model observation error of information source i or to guard against numerical ill-conditioning. A major ingredient of GPs is the prior information about the degree of correlation between different points in the input space. This information is encoded in a kernel function. In this work we assume a relatively smooth response surface based on the behavior of the 'ground truth' shown in Fig. 3 and thus we use the squared exponential kernel,

$$k_i(\mathbf{x},\mathbf{x}')=\sigma_s^2\exp\left(-\sum_{h=1}^{d}\frac{(x_h-x'_h)^2}{2l_h^2}\right), \tag{6}$$

where d is the dimension of the input space, σ_s^2 is the variance, and l_h, where $h=1,2,\ldots,d$, is the characteristic length-scale that measures the degree of correlation in the input space. We assume that each dimension of the input space, h, has its own stationary characteristic length-scale—the formulation is general but in this case we attempt to solve an optimization problem with a one

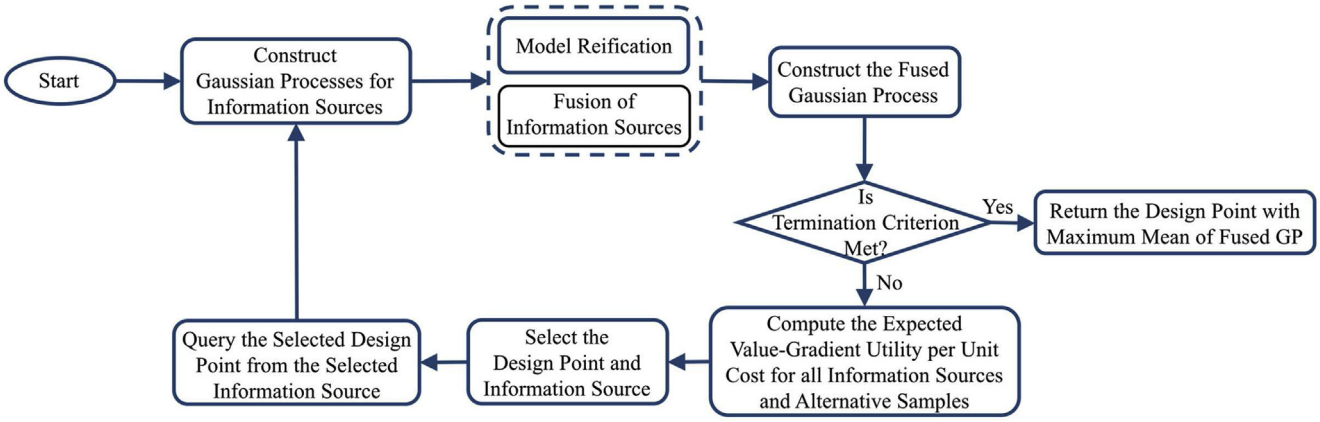

Fig. 4. Flow chart of the proposed efficient multi-information source optimization framework.

dimensional design/input space. The GP models for each information source are fit through maximum likelihood or Bayesian techniques [50].

The uncertainty of each information source with respect to the 'ground truth' can be estimated from the intrinsic variance of the GP as well as the discrepancy between the information source and the 'ground truth':

$$\sigma_i^2(\mathbf{x}) = \sigma_{GP,i}^2(\mathbf{x}) + \sigma_{f,i}^2(\mathbf{x}), \tag{7}$$

where $\sigma_{f,i}^2(\mathbf{x})$ is the variance of the discrepancy of information source i. We note that this variance can in principle vary over the input space.

While it is to be expected that low cost information sources with missing physics will exhibit large discrepancies with respect to the 'ground truth', in this work our hypothesis is that every source contains useful information about the 'ground truth' that is expensive/difficult to observe/query. In this specific case this is justified as all the reduced-order models and homogenization schemes use the same inputs and produce the same outputs as the 'ground truth' microstructure-based finite element calculations. More importantly, from an epistemic point of view, all the information sources seek to predict the response of a microstructure as a function of imposed loading using different theoretical frameworks and/or physical assumptions. Regardless of the underlying assumptions (e.g. how stress, strain or work are partitioned among the different phases), all sources are 'causally' connected to the same underlying 'ground truth' and it is thus expected that they will be correlated—to differing degrees—to it.

Since all information sources potentially have information about the 'ground truth', we proceed to fuse them, using following standard approaches for the fusion of normally distributed data (following Winkler [51]), as they are all represented as GPs. According to this fusion method, the fused mean and variance at point $\mathbf{x}$ can be computed as

$$\mu_{\text{Wink}}(\mathbf{x}) = \frac{\mathbf{e}^T \tilde{\Sigma}(\mathbf{x})^{-1} \boldsymbol{\mu}(\mathbf{x})}{\mathbf{e}^T \tilde{\Sigma}(\mathbf{x})^{-1} \mathbf{e}}, \tag{8}$$

$$\sigma_{\text{Wink}}^2(\mathbf{x}) = \frac{1}{\mathbf{e}^T \tilde{\Sigma}(\mathbf{x})^{-1} \mathbf{e}}, \tag{9}$$

where $\mathbf{e} = [1, \ldots, 1]^\top$, $\boldsymbol{\mu}(\mathbf{x}) = [\mu_1(\mathbf{x}), \ldots, \mu_S(\mathbf{x})]^\top$ contains the mean values of S sources at point $\mathbf{x}$, and $\tilde{\Sigma}(\mathbf{x})$ is the covariance matrix between sources,

$$\tilde{\Sigma}(\mathbf{x}) = \begin{bmatrix} \sigma_1^2(\mathbf{x}) & \cdots & \rho_{1S}(\mathbf{x})\,\sigma_1(\mathbf{x})\sigma_S(\mathbf{x}) \\ \rho_{12}(\mathbf{x})\,\sigma_1(\mathbf{x})\sigma_2(\mathbf{x}) & \cdots & \rho_{2S}(\mathbf{x})\,\sigma_2(\mathbf{x})\sigma_S(\mathbf{x}) \\ \vdots & \ddots & \vdots \\ \rho_{1S}(\mathbf{x})\,\sigma_1(\mathbf{x})\sigma_S(\mathbf{x}) & \cdots & \sigma_S^2(\mathbf{x}) \end{bmatrix}, \tag{10}$$

where $\sigma_i^2(\mathbf{x})$ is the total variance of source, i, at a point $\mathbf{x}$ computed in Eq. (7) and $\rho_{ij}(\mathbf{x})$ is the correlation between the deviations of information sources i and j at point $\mathbf{x}$. We use the reification process described in Refs. [22,26], to estimate the correlation between the errors of sources i and j, computed as

$$\rho_{ij}(\mathbf{x}) = \frac{\sigma_j^2(\mathbf{x})}{\sigma_i^2(\mathbf{x}) + \sigma_j^2(\mathbf{x})}\tilde{\rho}_{ij}(\mathbf{x}) + \frac{\sigma_i^2(\mathbf{x})}{\sigma_i^2(\mathbf{x}) + \sigma_j^2(\mathbf{x})}\tilde{\rho}_{ji}(\mathbf{x}), \tag{11}$$

which is the variance weighted average of the correlation coefficients computed by reifying sources i and j respectively as

$$\begin{aligned} \tilde{\rho}_{ij}(\mathbf{x}) &= \frac{\sigma_i(\mathbf{x})}{\sqrt{\left(\mu_i(\mathbf{x}) - \mu_j(\mathbf{x})\right)^2 + \sigma_i^2(\mathbf{x})}}, \\ \tilde{\rho}_{ji}(\mathbf{x}) &= \frac{\sigma_j(\mathbf{x})}{\sqrt{\left(\mu_j(\mathbf{x}) - \mu_i(\mathbf{x})\right)^2 + \sigma_j^2(\mathbf{x})}}, \end{aligned} \tag{12}$$

where $\mu_i(\mathbf{x})$ and $\mu_j(\mathbf{x})$ correspond to the means of sources i and j, and $\sigma_i^2(\mathbf{x})$ and $\sigma_j^2(\mathbf{x})$ are the total variances of information sources i and j at $\mathbf{x}$. The correlations between the errors of two models/sources is estimated using the procedure described in Refs. [22,26]. We used the fused means and variances in the input design space χ to construct a fused GP model. Letting $\boldsymbol{\mu}_{\text{Wink}}(\mathbf{x}_{1:N_f})$ and $\Sigma(\mathbf{x}_{1:N_f}) = \text{diag}(\sigma_{\text{Wink}}^2(\mathbf{x}_1), \ldots, \sigma_{\text{Wink}}^2(\mathbf{x}_{N_f}))$ be the vector of fused means with a diagonal matrix of the fused variances at the sampling set $\mathbf{x}_{1:N_f} \subset \chi$, the posterior predictive distribution of the fused model is given by:

$$\hat{f}^{\text{fused}}(\mathbf{X}) \sim \mathcal{N}\left(\boldsymbol{\mu}^{\text{fused}}(\mathbf{X}), \Sigma^{\text{fused}}(\mathbf{X})\right), \tag{13}$$

where

$$\boldsymbol{\mu}^{\text{fused}}(\mathbf{X}) = K\left(\mathbf{x}_{1:N_f}, \mathbf{X}\right)^T \left[K\left(\mathbf{x}_{1:N_f}, \mathbf{x}_{1:N_f}\right) + \Sigma\left(\mathbf{x}_{1:N_f}\right)\right]^{-1} \boldsymbol{\mu}_{\text{Wink}}\left(\mathbf{x}_{1:N_f}\right),$$
$$\Sigma^{\text{fused}}(\mathbf{X}) = K(\mathbf{X}, \mathbf{X}) - K\left(\mathbf{x}_{1:N_f}, \mathbf{X}\right)^T \left[K\left(\mathbf{x}_{1:N_f}, \mathbf{x}_{1:N_f}\right) + \Sigma\left(\mathbf{x}_{1:N_f}\right)\right]^{-1} K\left(\mathbf{x}_{1:N_f}, \mathbf{X}\right). \tag{14}$$

By constructing the fused GP, in each iteration of our proposed methodology, we determine the next design point—i.e., phase fraction to evaluate—and information source to query—i.e., model connecting microstructure and response—by balancing the cost of the query and the value of such query relative to the objective function, Eq. (2). In order to select the next point to query we first generate a Latin Hypercube experimental design, denoted as $\mathbf{X}_{\text{alt}}$ over the input space. Let $(\mathbf{x}_{1:N}, y_{1:N})$ be the design points and the corresponding objective values, and $i_{1:N}$ be the indices of the queried information sources up to time step N. In a Bayesian optimization framework, the choice of utility function determines the next (unobserved) point in the design space to explore. While there are a number of possible utility functions to choose from, here we propose a two-step look ahead utility which considers the immediate improvement in one step as well as the expected improvement in two steps.

This utility, which is obtained by querying the design point $\mathbf{x} \in \mathbf{X}_{\text{alt}}$ from information source i is defined as

$$U_{\mathbf{x},i} = \mathbb{E}\left[\max_{\mathbf{x}' \in \mathbf{X}_{\text{alt}}} \mu^{\text{fused}}(\mathbf{x}') + \max_{\mathbf{x}'' \in \mathbf{X}_{\text{alt}}} EI_{\mathbf{x},i}(\mathbf{x}'') \middle| \mathbf{x}_{1:N}, y_{1:N}, i_{1:N}, \mathbf{x}_{N+1} = \mathbf{x}, i_{N+1} = i\right], \tag{15}$$

where $EI_{\mathbf{x},i}(\mathbf{x}'')$ is the one-step look-ahead expected increase in the maximum of the fused GP given $\mathbf{x}_{N+1} = \mathbf{x}$ and $i_{N+1} = i$ as

$$EI_{\mathbf{x},i}(\mathbf{x}'') = \mathbb{E}\left[\max_{\mathbf{x}' \in \mathbf{X}_{\text{alt}}} \mu^{\text{fused}}(\mathbf{x}') \middle| \mathbf{x}_{N+2} = \mathbf{x}''\right] - \max_{\mathbf{x}' \in \mathbf{X}_{\text{alt}}} \mu^{\text{fused}}(\mathbf{x}'). \tag{16}$$

We note that to compute the expectation in Eq. (15) we use a Monte Carlo approach, drawing N_q independent samples from the normal distribution of the GP of information source i at a design point $\mathbf{x}$, $f_i^q(\mathbf{x}) \sim \mathcal{N}(\mu_i(\mathbf{x}), \sigma^2_{GP,i}(\mathbf{x}))$, $q = 1, \ldots, N_q$. Then, by temporarily augmenting $(\mathbf{x}, f_i^q(\mathbf{x}))$, one at a time, using the available samples of information source i, the mean of the fused GP is temporarily updated, which is denoted as $\mu^{\text{fused},q}_{\mathbf{x},i}$, and the utility is approximated as

$$U_{\mathbf{x},i} \approx \frac{1}{N_q} \sum_{q=1}^{N_q} \left(\max_{\mathbf{x}' \in \mathbf{X}_{\text{alt}}} \mu^{\text{fused},q}_{\mathbf{x},i}(\mathbf{x}') + \max_{\mathbf{x}'' \in \mathbf{X}_{\text{alt}}} EI^q_{\mathbf{x},i}(\mathbf{x}'')\right), \tag{17}$$

where $EI^q_{\mathbf{x},i}(\mathbf{x}'')$ is the one-step look-ahead expected increase in the maximum of the fused GP upon augmentation of query $(\mathbf{x}, f_i^q(\mathbf{x}))$ to information source i. We compute this expected increase using the Knowledge Gradient metric over the temporary fused GP as discussed in Ref. [52].

Using the Latin Hypercube sampling over the input space, we evaluate the utility function for each of the information sources by removing the previously added sample and augmenting the next sample in the proposed alternatives. The next point to query—among all the candidate points—with the best information source is determined using the following policy:

$$(i_{N+1}, \mathbf{x}_{N+1}) = \underset{i \in [1,\ldots,S], \mathbf{x} \in \mathbf{X}_{alt}}{\text{argmax}} \frac{U_{\mathbf{x},i}}{C_{\mathbf{x},i}}, \tag{18}$$

where $C_{\mathbf{x},i}$ is the cost of querying information source i at a design point $\mathbf{x}$.

Once the selected design point has been identified and a query has been made, the corresponding fused GP is updated. This loop is continued until the objective of the optimization problem has been met or the budget allocated to this optimization task has been exhausted. The final design, $\mathbf{x}^*$, is then chosen as:

$$\mathbf{x}^* = \underset{\mathbf{x} \in \chi}{\text{argmax}}\ \mu^{\text{fused}}(\mathbf{x}) \tag{19}$$

4. Designing dual-phase microstructures for enhanced mechanical performance

The framework, Fig. 4, is now demonstrated against the problem of identifying the phase constitution —represented in this simplified instance as just the volume fraction of the constituents in the dual-phase microstructure— that results in the maximum formability metric. We consider the microstructure-based finite element calculations as 'ground truth' and use the three reduced-order models and three micromechanical homogenization schemes as the cheap information sources available to elicit the behavior of the composite dual-phase microstructure.

The optimization process starts by querying (once) the cheap and the 'ground truth' sources. These data are then used to construct the initial GPs used to emulate each of the information sources. The framework is then used to identify which information source to use next, and where in the input space to use it.

In this first demonstration of the framework, after five queries to any of the cheap sources the framework makes the recommendation for the next point to query using the expensive 'ground truth'. We note that the framework allows for the querying of any of the cheap sources multiple times (or not at all) as long as the five-query budget has not been exhausted. We also note that while the cost is considered in the construction of the utility function, the total expenditure (i.e. total computational resources used) is not accumulated. Implicitly, this set up thus considers that there is no real hard budget constraining the optimization. This problem setup will be contrasted with a case in which there is a hard budget constrain for the total computational expenditures before a query to the 'ground truth' is made, as described later in Section 5.

Fig. 5 shows the fused model obtained by our approach following seven queries from the 'ground truth', with mean represented by the smooth red line and 99.7% confidence interval represented by the red shaded area. We compare these results to the 'ground truth', represented with the jagged green line and the green shaded area. Recall, from Section 2.4, that for the 'ground truth' the variation in the values of $(1/\tau)(d\tau/d\varepsilon_{pl})$ for a fixed value of f_{hard} is due to different realizations of the RVE representation of the dual-phase microstructure with identical phase constitutions. As can be seen from Fig. 5, the fused model obtained by our approach represents the 'ground truth' well in the region of the optimal design.

Table 3 shows the progression of the optimization procedure as the framework identifies the design point that corresponds to maximum formability. The table compares the predictions from the fuse model evaluated at the best design point thus far, $\mathbf{x}^*_{fused}$, with the value of the formability parameter evaluated using the 'ground truth' at the same best design point. The table shows that it takes seven queries of the 'ground truth' to produce a close-to-optimal solution. While not shown here, in previous work [27] we have shown how this multi-information source framework is superior to the use of the 'ground truth' as the only query to sample.

Fig. 6 shows the number of times that any of the information sources, including the expensive 'ground truth', are queried. Between each of these expensive queries, the different cheap information sources are queried depending on their utility (including consideration of their cost). The figure shows that all the cheap sources are being queried, albeit at different frequencies, in order to identify the optimal design. As described above, the querying policy is controlled through Eq. (18), which balances improvement in objective with cost. The figure shows, for example, how in the early stages of the optimization the 'isostress' and 'secant method-1' sources tend to be queried more extensively than the other four available sources. In later stages of the process it is evident that many more sources are being queried and this serves as an indication of the ability of the present framework to optimally select the most cost effective information sources. We also point out here that very cheap information sources, such as the isostress and isostrain models, could potentially be queried exhaustively and still follow the policy determined by Eq. (18). In this case, the cost of querying is so negligible that any potential for objective improvement would lead to these sources being selected. This is an excellent feature of the policy, since in general, if we have access to a very inexpensive information source, it makes intuitive sense to quickly gather all information we can from it.

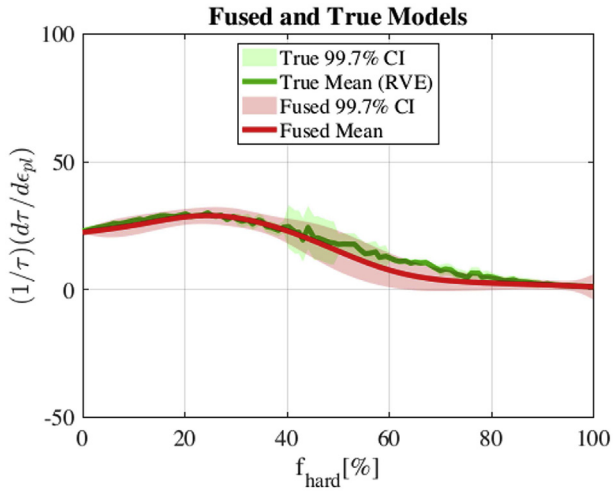

Fig. 5. The mean and 99.7% CI of the fused model obtained by the proposed approach in comparison with the 'ground truth' (RVE).

Table 3
Performance of the optimization framework. The true optimal solution as obtained by the microstructure-based finite element calculation (RVE) is $(\mathbf{x}^*, f^*) = (21.6080, 29.3123)$.

Experiment	$\mathbf{x}^*_{fused}$	f^*_{fused}	$f_{true}(\mathbf{x}^*_{fused})$
2	36.1809	18.8470	25.4273
3	36.6834	20.7475	25.1903
4	32.1608	22.5232	27.1322
5	31.6583	28.9564	27.3177
6	26.1307	29.8824	28.8633
7	22.6181	29.7564	29.2840

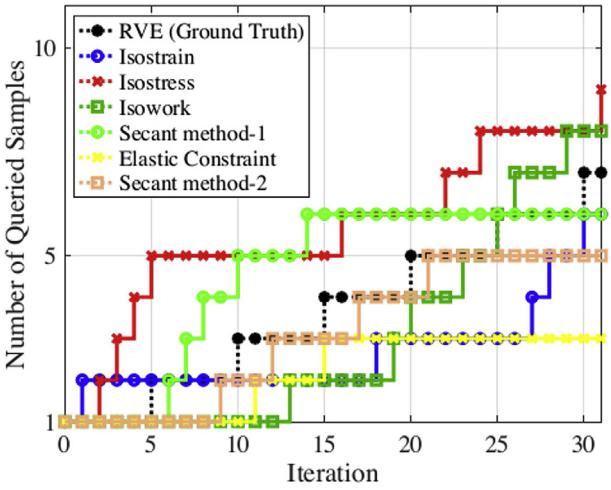

Fig. 6. Number of queries as a function of the experimental sequence for all information sources.

Fig. 7 shows the GPs of the information sources. The black dots show the samples queried from the information sources and the black lines represent the mean function of these sources. The solid shaded regions for each of the GPs correspond to the intrinsic uncertainty of the GPs themselves and originates from the lack of information about the response of that particular source in regions not explored yet. In addition to this uncertainty, a comparison with the response of the 'ground truth' is used to compute the discrepancy of the information sources. The intrinsic uncertainty of the GP and the discrepancy between the information source and the 'ground truth' are added into the total uncertainty, shown as a lighter shaded region.

The figure shows that none of the information sources performs well over the entire design domain. Some sources overstimate their predictions relative to the 'ground truth' (black solid line) while some sources underestimate it. None of the sources provides even a qualitative indication of where the optimal value of the 'ground truth' should be. The application of the reification-based information fusion, however, results in a fused model that reproduces, with high fidelity, the response of the 'ground truth'.

Fig. 6 shows that different sources are queried at different stages of the optimization sequence. Moreover, Fig. 7 shows that while all sources have a significant discrepancy with the 'ground truth', some sources tend to be closer to it in some regions of the problem space. To understand how different sources were correlated to the 'ground truth', we present the effective independent information sources index in Fig. 8. This index was introduced first time in our previous work [27], and can be understood as a metric that

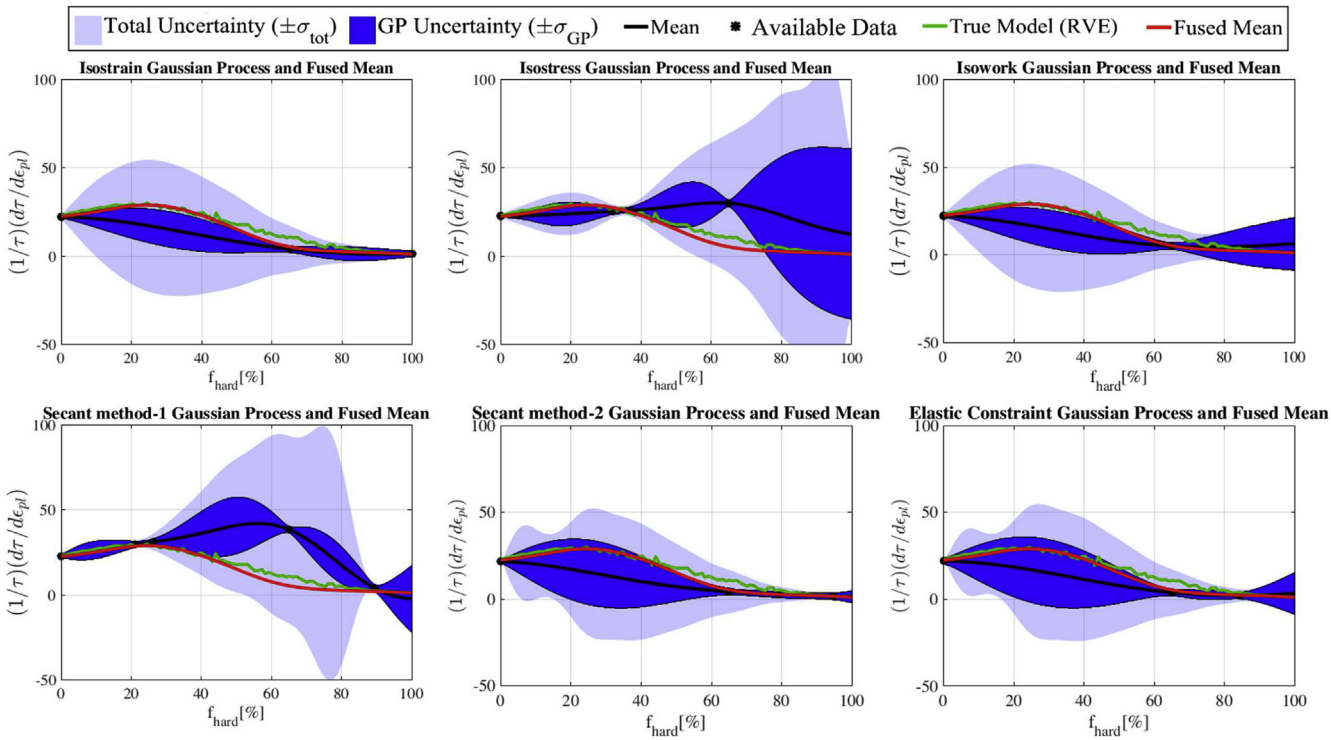

Fig. 7. Comparison between the 'ground truth', fused model, and individual information sources.

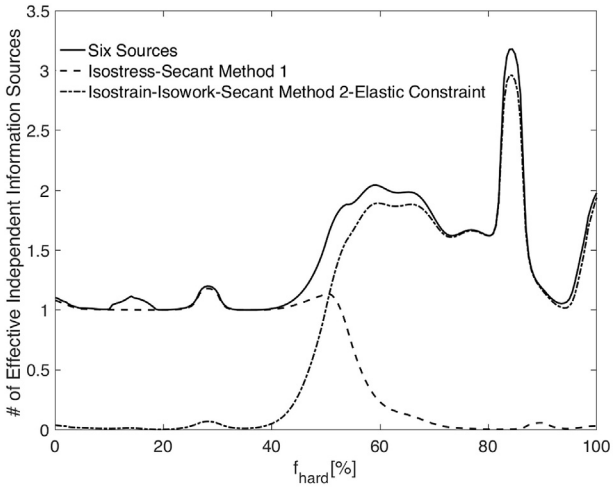

Fig. 8. Number of effective independent information sources (EIS), I_{eff} as a function of f_{hard}.

indicates the number of effective number of independent information sources used to estimate the fused model at different regions of the domain. If we define $\sigma_*^2(\mathbf{x})$ as the variance of the current best information source at point $\mathbf{x}$, the number of effective independent information sources with variance $\sigma_*^2(\mathbf{x})$ at the point $\mathbf{x}$ is given as

$$I_{\text{eff}} = \sigma_*^2(\mathbf{x})\mathbf{e}^\top \tilde{\Sigma}(\mathbf{x})^{-1}\mathbf{e}. \tag{20}$$

If we have S total information sources, each with variance σ_*^2, then this index will take the value S. Otherwise, if any information source has a greater variance than the best current information source (the source that is closest to the 'ground truth' in this region), then the value of I_{eff} will be less than S. This is typically the case since information sources differ in their fidelity. This index also takes into account the correlation between information sources, which often results in redundant information that reduces the I_{eff} value. The reification approach used in this work takes this into account by detecting the correlation and discounting this effect appropriately. We note that if these correlated sources were to be treated as statistically independent, they would have undue evidence on the overall estimation of the fused model. The value of I_{eff} thus counts the number of information sources that provide useful information.

Fig. 8 includes the I_{eff} for all the cheap sources simultaneously and compares this index against that of the isostress + secant method-1 pair and the combination of the isostrain + isowork + secant method-2 + elastic constraint. As can be seen in Fig. 8, I_{eff} is not large over the input space. For this particular problem, it is often the case that only a few of the sources are reasonable estimates of 'ground truth' at any given location in the domain. This renders the contribution of the other more inaccurate sources to I_{eff} to be very small. In this figure, it is clear that initially, until $f_{\text{hard}} \approx 40\%$, the isostress + secant method-1 pair is driving the fused approximation, which is clear from the fact that the six-source index and the isostress + secant method-1 index nearly overlap in this region. This means that the isostress + secant method-1 are capable of explaining most of the variance with regards to the 'ground truth'. Examining Fig. 7 one can see that these two sources are closer to the 'ground truth' than the other four sources within this region of the input space. At the other end of the domain, the other four sources are contributing more to the prediction, which can again be seen from the near overlap with the six-source index and the drop-off of the isostress + secant method-1 pair. Note, that the proposed framework indicates that there is no instance in which there are six effective information sources over the entire domain. This results from two different factors: (i) there are strong correlations between two or more sources at different

regions of the domain — the correlated sources are not necessarily the same everywhere; and (ii) some sources are very poor at predicting the 'ground truth' in some regions of the domain having very large variance that renders their contribution negligible. This is exemplified by the isostress + secant method-1 pair, which according to Fig. 7 exhibit very large variance at intermediate-to-high volume fractions of martensite. Accordingly, their contribution to I_{eff} in this region of the domain is negligible.

This result is significant as I_{eff} essentially indicates what information source (or model) operates in a specific region of the problem domain. Thus, it is possible to use I_{eff} to uncover the relevant sources and, indirectly, the relevant physics governing the behavior of a system in different regimes. In this case, for example, I_{eff} suggests that the assumption of equal partitioning of stress in the regime at low volume fraction of the hard particle seems to be operationally correct. At higher volume fractions of the hard phase, other theories may be more in agreement with the 'ground truth'.

5. The impact of cost on optimal decision making under budget constraints

We note that the set up of the test problem presented in Section 4 was somewhat arbitrary as we defined the acquisition protocol (five queries to the cheap sources before querying the 'ground truth') ahead of carrying out the exploration/exploitation of the design space. In principle, we could have chosen to query the cheap sources any number of times in between queries to the 'ground truth'. More importantly, while cost was used to compute the utility metric of the design space, there was no explicit consideration of budget available during the optimization exercise.

To resolve this issue, we considered a more realistic situation that is more relevant to the process of optimal experimental design for materials optimization:

- There is a finite set of information sources that can be queried at any given time in order to learn more about the problem space.
- Each of the sources has different degree of fidelity with regards to the 'ground truth' and, most importantly, different cost.
- There is a finite budget, expressed in terms of total cost (in this case in computational time) available to explore and exploit the domain space before the information source considered as the 'ground truth' is queried.

In order to demonstrate the effect of cost in our decision-making process, we compare the average results obtained over 100 independent simulations in two conditions—labeled as 'with cost consideration' and 'without cost consideration'—in Fig. 9. In both conditions, we consider a fixed budget (computational cost) of 100 s to be spent among all the cheap information sources before querying the 'ground truth'.

In the case identified as 'with cost consideration', we consider the actual cost of information sources in the selection criterion in Eq. (18), however in the case of 'without cost consideration', the cost of all information sources are assumed to be the same during the selection process. After the budget of 100 s is spent, then the 'ground truth' is queried in both cases.

In the first case (with cost consideration), the cheap information sources are queried exhaustively allowing the policy to obtain all available information from the very inexpensive sources with almost negligible cost. After this, the more expensive sources are queried according to the balance between cost and utility, in a fully automated fashion. While in the second case (without cost consideration), as cost accumulation does not play a role in the selection process, the expensive information sources are queried in the early iterations. Therefore, fewer queries are performed,

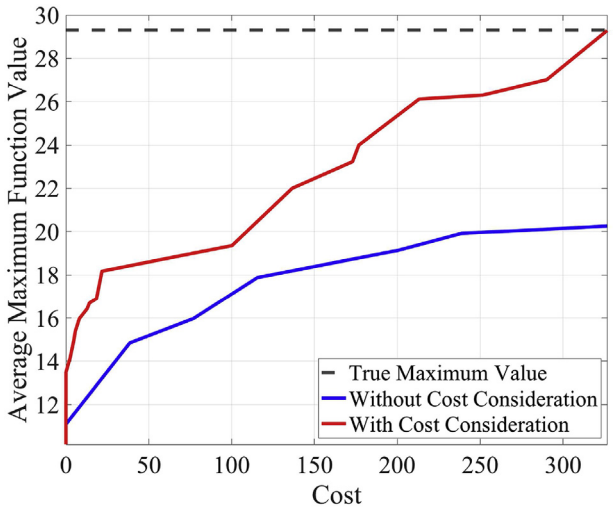

Fig. 9. Average maximum function value per cost for cases of decision-making with and without cost consideration.

particularly on the cheaper information sources, in the limited budget of 100 s. The result is that the total budget is exhausted before reaching the optimum. However, for the 'with cost consideration case', the optimum is obtained.

Fig. 10 demonstrates the average number of queries to each information source and the 'ground truth' over average cost for two conditions of with and without cost consideration. As can be seen, in the case of 'with cost consideration', the cheaper information sources are exhaustively queried, which helps to find the optimum with much less average cost in comparison to the case of 'without cost consideration' as seen in Fig. 9. Essentially, the policy in place has the ability to quickly gather all information from the negligible expense sources prior to making a query to a more expensive source. This provides for more informed queries to the more expensive sources, and also matches the intuitive strategy consisting of exhaustively evaluating the very inexpensive sources immediately.

6. Concluding remarks

While there has been real progress in the development of ICME-based frameworks for materials design, major issues still remain. The challenges are closely associated to the resource-intensive nature of the computational and/or experimental approaches to exploring PSPP relationships as well as to the fact that most ICME frameworks tend to assume that there is a single model or experiment available at a given scale or PSPP linkage. A further limitation of traditional ICME approaches is the fact that in most cases there is no prescribed way to correct for discrepancies between models and 'ground truth'.

In this work, we have presented a framework capable of utilizing multiple available sources of information to carry out a materials optimization task. By exploiting statistical correlations among the different information sources as well as between the sources and the 'ground truth' it is possible to develop a fused model that incorporates the available (useful) information originating from all of the sources. The fused model is represented as a GP and is thus amenable to Bayesian Optimization. Here we have used the Knowledge Gradient metric to efficiently balance the need to explore and exploit knowledge of the materials design space in order to find an optimal solution in as efficient manner as possible. Our framework, however, goes beyond optimal experimental

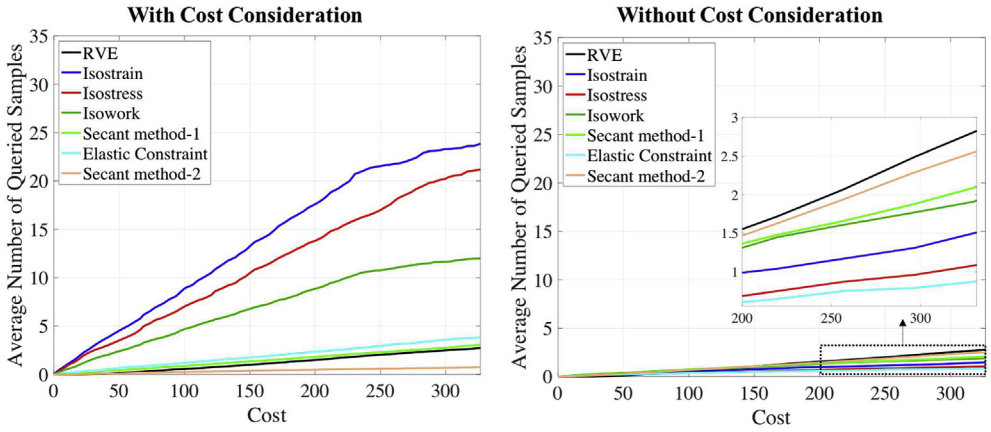

Fig. 10. Average Number of samples queried from the 'ground truth' (RVE) model and the information sources over average cost for the cases of decision-making with and without cost consideration.

design, as it enables the selection of the most cost-effective information source to use every time that we need to query the problem space, while accounting for a total available budget.

While the proposed framework has been demonstrated by using a combination of relatively inexpensive computational models, it is important to note that each information source (including the 'ground truth') is represented as a stochastic model (i.e., a GP) and thus there is no fundamental limitation to using experimental or any other type of information as an independent information source. One could, for example, combine mechanistic models, machine learning derived non-parametric models, experiments at different degrees of resolution, and even expert opinion as long as each source is modeled as a GP or any other suitable stochastic representation. The framework proposed thus provides a natural approach to seamlessly combining experiments and simulations, and hopefully will inspire and instigate future works in this direction.

Acknowledgements

The authors acknowledge the U.S. National Science Foundation through grant No. NSF-CMMI-1663130, *DEMS: Multi-Information Source Value of Information Based Design of Multiphase Structural Materials*. DA and RA acknowledge grant No. NSF-DGE-1545403. RA would also like to acknowledge grant No. NSF-CMMI-1534534.

References

[1] J. Allison, Integrated computational materials engineering: a perspective on progress and future steps, JOM 63 (4) (2011) 15–18.
[2] N.R. Council, et al., Integrated Computational Materials Engineering: a Transformational Discipline for Improved Competitiveness and National Security, National Academies Press, 2008.
[3] P. Voorhees, G. Spanos, et al., Modeling across Scales: a Roadmapping Study for Connecting Materials Models and Simulations across Length and Time Scales, Tech. rep., Tech. rep., the Minerals, Metals & Materials Society (TMS), 2015.
[4] S. Reddy, B. Gautham, P. Das, R.R. Yeddula, S. Vale, C. Malhotra, An ontological framework for integrated computational materials engineering, in: Proceedings of the 4th World Congress on Integrated Computational Materials Engineering (ICME 2017), Springer, 2017, pp. 69–77.
[5] V. Savic, L. Hector, U. Basu, A. Basudhar, I. Gandikota, N. Stander, T. Park, F. Pourboghrat, K.S. Choi, X. Sun, et al., Integrated Computational Materials Engineering (ICME) Multi-Scale Model Development for Advanced High Strength Steels, Tech. Rep., SAE Technical Paper, 2017.
[6] M. Diehl, M. Groeber, C. Haase, D.A. Molodov, F. Roters, D. Raabe, Identifying structure–property relationships through dream. 3D representative volume elements and DAMASK crystal plasticity simulations: an integrated computational materials engineering approach, JOM 69 (5) (2017) 848–855.
[7] L. Lin, W. Ren, An implementation of icme in materials information exchanging interfaces, Mater. Discov. 12 (2018) 9–19.
[8] V. Balabanov, R. Haftka, B. Grossman, W. Mason, L. Watson, Multifidelity response surface model for HSCT wing bending material weight, in: 7th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, St. Louis, MO, 1998. AIAA 1998-4804.
[9] T.W. Simpson, T.M. Mauery, J.J. Korte, F. Mistree, Kriging models for global approximation in simulation-based multidisciplinary design optimization, AIAA J. 39 (12) (2001) 2233–2241.
[10] V. Balabanov, G. Venter, Multi-fidelity optimization with high-fidelity analysis and low-fidelity gradients, in: 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Albany, New York, 2004. AIAA 2004-4459.
[11] R.A. Moore, C.J. Paredis, Variable fidelity modeling as applied to trajectory optimization for a hydraulic backhoe, in: ASME 2009 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers, 2009, pp. 79–90.
[12] A.J. Keane, Wing optimization using design of experiment, response surface, and data fusion methods, J. Aircr. 40 (4) (2003) 741–750.
[13] S. Chen, Z. Jiang, S. Yang, W. Chen, Multimodel fusion based sequential optimization, AIAA J. 55 (1) (2016) 241–254.
[14] S. Choi, J. Alonso, I. Kroo, Multi-fidelity design optimization of low-boom supersonic business jets, in: 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Albany, New York, 2004. AIAA 2004-4371.
[15] S. Choi, J. Alonso, I. Kroo, Two-level multifidelity design optimization studies for supersonic jets, J. Aircr. 46 (2) (2009) 776–790.
[16] N. Alexandrov, J. Dennis, R. Lewis, V. Torczon, A Trust Region Framework for Managing the Use of Approximation Models in Optimization, Tech. Rep. CR-201745, NASA, October 1997.
[17] N. Alexandrov, R. Lewis, C. Gumbert, L. Green, P. Newman, Optimization with Variable-Fidelity Models Applied to Wing Design, Tech. Rep. CR-209826, NASA, December 1999.
[18] N. Alexandrov, R. Lewis, C. Gumbert, L. Green, P. Newman, Approximation and model management in aerodynamic optimization with variable-fidelity models, AIAA J. 38 (6) (2001) 1093–1101.
[19] D. Jones, M. Schonlau, W. Welch, Efficient global optimization of expensive black-box functions, J. Glob. Optim. 13 (1998) 455–492.
[20] D. Huang, T.T. Allen, W.I. Notz, R.A. Miller, Sequential kriging optimization using multiple-fidelity evaluations, Struct. Multidiscip. Optim. 32 (5) (2006) 369–382.
[21] P.I. Frazier, W.B. Powell, S. Dayanik, A knowledge-gradient policy for sequential information collection, SIAM J. Control Optim. 47 (5) (2008) 2410–2439.
[22] D. Allaire, K. Willcox, Fusing information from multifidelity computer models of physical systems, in: Information Fusion (FUSION), 2012 15th International Conference on, IEEE, 2012, pp. 2458–2465.
[23] D. Allaire, K. Willcox, O. Toupet, A bayesian-based approach to multifidelity multidisciplinary design optimization, in: 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference, 2010, p. 9183.
[24] D. Allaire, K. Willcox, A mathematical and computational framework for multifidelity design and analysis with computer models, Int. J. Uncertain. Quantification 4 (1).
[25] R. Lam, D. Allaire, K. Willcox, Multifidelity optimization using statistical surrogate modeling for non-hierarchical information sources, in: Proceedings of the AIAA SciTech Multidisciplinary Design Optimization Conference, 2015. AIAA-0143.
[26] W.D. Thomison, D.L. Allaire, A model reification approach to fusing information from multifidelity information sources, in: 19th AIAA Non-deterministic Approaches Conference, 2017, p. 1949.
[27] S.F. Ghoreishi, A. Molkeri, A. Srivastava, R. Arroyave, D. Allaire, Multi-

information source fusion and optimization to realize icme: application to dual-phase materials, J. Mech. Des. 140 (11) (2018) 111409.
[28] B. Peherstorfer, K. Willcox, M. Gunzburger, Survey of multifidelity methods in uncertainty propagation, inference, and optimization, SIAM Rev. 60 (3) (2018) 550–591.
[29] M. G. Fernández-Godino, C. Park, N.-H. Kim, R. T. Haftka, Review of Multi-Fidelity Models, arXiv preprint arXiv:1609.07196.
[30] R. Potyrailo, K. Rajan, K. Stoewe, I. Takeuchi, B. Chisholm, H. Lam, Combinatorial and high-throughput screening of materials libraries: review of state of the art, ACS Comb. Sci. 13 (6) (2011) 579–633.
[31] S.K. Suram, J.A. Haber, J. Jin, J.M. Gregoire, Generating information-rich high-throughput experimental materials genomes using functional clustering via multitree genetic programming and information theory, ACS Comb. Sci. 17 (4) (2015) 224–233.
[32] M.L. Green, C. Choi, J. Hattrick-Simpers, A. Joshi, I. Takeuchi, S. Barron, E. Campo, T. Chiang, S. Empedocles, J. Gregoire, et al., Fulfilling the promise of the materials genome initiative with high-throughput experimental methodologies, Appl. Phys. Rev. 4 (1) (2017) 011105.
[33] S. Curtarolo, G.L. Hart, M.B. Nardelli, N. Mingo, S. Sanvito, O. Levy, The high-throughput highway to computational materials design, Nat. Mater. 12 (3) (2013) 191–201.
[34] P. V. Balachandran, D. Xue, J. Theiler, J. Hogden, T. Lookman, Adaptive strategies for materials design using uncertainties, Sci. Rep. 6.
[35] A. Talapatra, S. Boluki, T. Duong, X. Qian, E. Dougherty, R. Arroyave, Towards an Autonomous Efficient Materials Discovery Framework: an Example of Optimal Experiment Design under Model Uncertainty, arXiv preprint arXiv: 1803.05460.
[36] A. Talapatra, S. Boluki, T. Duong, X. Qian, E. Dougherty, R. Arróyave, Autonomous efficient experiment design for materials discovery with bayesian model averaging, Phys. Rev. Mater. 2 (11) (2018) 113803.
[37] M. Rashid, Dual phase steels, Annu. Rev. Mater. Sci. 11 (1) (1981) 245–266.
[38] P. Chen, H. Ghassemi-Armaki, S. Kumar, A. Bower, S. Bhat, S. Sadagopan, Microscale-calibrated modeling of the deformation response of dual-phase steels, Acta Mater. 65 (2014) 133–149.
[39] A. Srivastava, A. Bower, L. Hector Jr., J. Carsley, L. Zhang, F. Abu-Farha, A multiscale approach to modeling formability of dual-phase steels, Model. Simul. Mater. Sci. Eng. 24 (2) (2016) 025011.
[40] A. Srivastava, H. Ghassemi-Armaki, H. Sung, P. Chen, S. Kumar, A.F. Bower, Micromechanics of plastic deformation and phase transformation in a three-phase trip-assisted advanced high strength steel: experiments and modeling, J. Mech. Phys. Solids 78 (2015) 46–69.
[41] D. Gerbig, A. Srivastava, S. Osovski, L.G. Hector, A. Bower, Analysis and design of dual-phase steel microstructure for enhanced ductile fracture resistance, Int. J. Fract. (2017) 1–24.
[42] F. Feyel, Multiscale fe2 elastoviscoplastic analysis of composite structures, Comput. Mater. Sci. 16 (1–4) (1999) 344–354.
[43] E. Weinan, B. Engquist, X. Li, W. Ren, E. Vanden-Eijnden, Heterogeneous multiscale methods: a review, Commun. Comput. Phys. 2 (3) (2007) 367–450.
[44] W. Voigt, On the relation between the elasticity constants of isotropic bodies, Ann. Phys. Chem 274 (1889) 573–587.
[45] A. Reuss, Berechnung der fließgrenze von mischkristallen auf grund der plastizitätsbedingung für einkristalle, ZAMM J. Appl. Math. Mech. Z. Angew. Math. Mech. 9 (1) (1929) 49–58.
[46] O. Bouaziz, P. Buessler, Mechanical behaviour of multiphase materials: an intermediate mixture law without fitting parameter, Rev. Métall. Int. J. Metall. 99 (1) (2002) 71–77.
[47] G. Weng, The overall elastoplastic stress-strain relations of dual-phase metals, J. Mech. Phys. Solids 38 (3) (1990) 419–441.
[48] Dassault Systemes, ABAQUS User's Mannual, 2017.
[49] S. Nemat-Nasser, M. Hori, Micromechanics: Overall Properties of Heterogeneous Materials, vol. 37, Elsevier, 2013.
[50] C. K. Williams, C. E. Rasmussen, Gaussian Processes for Machine Learning, The MIT Press.
[51] R.L. Winkler, Combining probability distributions from dependent information sources, Manag. Sci. 27 (4) (1981) 479–488.
[52] P. Frazier, W. Powell, S. Dayanik, The knowledge-gradient policy for correlated normal beliefs, Inf. J. Comput. 21 (4) (2009) 599–613.