Efficient searching of processing parameter space to enable inverse microstructural design of materials

Dung-Yi Wu[a], Todd C. Hufnagel[a,b,c]

[a]*Department of Materials Science and Engineering, Johns Hopkins University, Baltimore, MD 21218, USA*
[b]*Hopkins Extreme Materials Institute, Johns Hopkins University, Baltimore, MD 21218, USA*
[c]*Department of Mechanical Engineering, Johns Hopkins University, Baltimore, MD 21218, USA*

Abstract

Materials design can be accelerated by the use of physics-based forward models that predict the properties of new materials. In cases where the properties of the material depend on its microstructure, the models can be used as part of an optimization scheme to predict the microstructural features that are required to achieve the design objectives. Producing these microstructures, however, requires that we determine the processing parameters necessary to produce the target microstructure. Here we demonstrate the use of Bayesian optimization using simple analytic forward models to enable this inverse process design in the context of optimization of isothermal heat treatment of a commercial aluminum alloy to achieve (i) a target volume fraction of a specific intermetallic phase, (ii) a specified aluminum grain size distribution, and (iii) both objectives simultaneously. We discuss the use, and limitations, of Bayesian optimization with a scalar desirability function for solving multi-objective problems, and demonstrate Bayesian optimization with expected hypervolume improvement to determine the Pareto front describing the trade-off between multiple objectives. We compare Bayesian optimization with a genetic algorithm (NSGA-II) and show that while the two approaches produce similar Pareto fronts, the Bayesian optimization converges on the solution more quickly. We further illustrate the effect of parametric uncertainty in the forward models on the uncertainty of the Pareto front. Finally, we demonstrate experimentally that the optimized process parameters do indeed allow

---

Email address: `hufnagel@jhu.edu` (Todd C. Hufnagel)

us to make samples with the desired microstructure.

Keywords: Probabilistic materials design, inverse design, microstructure optimization, Bayesian optimization, multi-objective optimization.

1. Introduction

Materials development has traditionally involved an iterative cycle through stages of processing and characterization to arrive at a set of properties that satisfy the design criteria. Such iteration can be time-consuming and expensive; it has frequently been noted that it can take decades to move a material from laboratory to application [1]. One way to accelerate this cycle is through an inverse procedure in which a set of desired properties is defined, and then a series of experiments is conducted with various processing parameters. Guided by a suitable optimization scheme, a final set of processing parameters is identified that optimizes the property (or properties) of interest [2–5]. This approach to materials design and discovery has been applied to the development of porous materials [6], optical materials [7], ion-conductive materials [8], metals [9, 10] , and polymers [11, 12].

Such inverse processes do not directly concern themselves with the effect of microstructure on materials properties, and indeed in many situations, particularly those involving functional properties, the microstructure may not be that important. For situations where microstructure is important, particularly those regarding complex multiscale mechanical behaviors such as plastic deformation and fracture of metals, considerable effort has been devoted to developing physics-based forward models that predict the properties and behavior to be expected from a specified microstructure [13–17]. Such models can then be used as part of an inverse scheme in which various aspects of the microstructure are varied and optimized to produce a desired set of properties [18–21]. This is important because it enables a design process that relies heavily on computation, which may be less time-consuming and less expensive than an experimental campaign that seeks to go directly from processing to properties. Furthermore, the link to microstructure provides essential physical insight into the physical mechanisms that underlie the desired behaviors.

However, this description of microstructural design omits a critical step: Once an "optimal" microstructure has been identified, what processing is required to produce it? This is essentially a second inverse problem; assuming that we have good physics-based forward models that describe the development of microstructure resulting from a given processing treatment, we can use an optimization scheme to determine the appropriate processing conditions. In this way, the entire materials design loop can be closed, with computational models being used to both propose microstructural objectives and determine the necessary processing. In principle, iteration through such a closed loop can be rapid, accelerating the materials design process considerably.

One notable early attempt at such an inverse process design is the work of Johnson and coworkers, who applied a mesh adaptive direct search algorithm to identify a heat treatment schedule for obtaining a specified size distribution of Ni_4Ti_3 precipitates in NiTi shape memory alloys [22]. More recently, several groups have tackled problems in which the goal was optimization of processing conditions for multiple microstructural objectives simultaneously. For example, Tran and coworkers developed an asynchronously parallel Bayesian optimization framework which combined multiple objective functions into a single scalar objective to optimize welding process parameters using a Kinetic Monte Carlo forward model [23]. Honarmandi and coworkers combined composition and process parameters to achieve a desired length scale and volume fraction of phases resulting from spinodal decomposition in TiAlZrN alloys [24]. Dornheim and coworkers constructed a framework combining deep reinforcement learning with an efficient mean-field crystal plasticity model to identify metal-forming process parameters to obtain a specified orientation distribution functions [25]. Liu and coworkers employed Bayesian optimization to explore the optimal initial temperature and cooling rate for a Ti-6Al-4V alloy in metal additive manufacturing, aiming to achieve a targeted dendritic area and level of microsegregation [26].

Here we illustrate the use of inverse design to determine the processing parameters required to achieve multiple microstructural objectives in heat treatment of a commercial aluminum alloy. The two microstructural parameters of interest are the volume fraction of Al_7Fe_2Cu second-phase particles and the aluminum grain size distribution. Extending the prior demonstrations of inverse process design discussed above, we compare the performance of two optimization schemes, Bayesian optimization with expected hypervolume improvement and a genetic algorithm (NSGA-II), in determining the Pareto front (which describes the tradeoff between the two objectives). We further consider the uncertainty in the Pareto front arising from the uncertainty in the values of parameters of our forward models. Finally, we perform validation experiments to show that samples produced using optimized parameters do indeed achieve our microstructural objectives.

2. Background

To illustrate our framework for multi-objective optimization of processing parameters we chose as a specific example the design of aluminum alloys to resist spall, a dynamic failure process which occurs in response to shock loading and involves nucleation, growth, and coalescence of spall voids [27]. Microstructure influences spall strength directly by providing nucleation sites for voids and indirectly through its effect on the dynamic strength, which influences the rate of void growth [28, 29].

In a recent paper we explored the effect of microstructure on the spall strength of 7085 aluminum, manipulating both the volume fraction of intermetallic particles and the aluminum grain size distribution via heat treatment [30]. We observed that spall voids tend to nucleate at Al_7Cu_2Fe second-phase particles and that the spall strength increases with the average aluminum grain size, suggesting that aluminum grain boundaries may also be preferential sites for spall void nucleation.

Ongoing work in this area seeks to develop physics-based computational models to describe spall void nucleation and growth, with the goal of identifying microstructures that might be expected to produce materials with outstanding spall resistance. This naturally raises the question of how to produce such a target microstructure once it has been identified. This problem can be quite complicated due to the large number of potential aspects of microstructure which could be manipulated and which might affect spall strength. Here, we focus on the two microstructural features identified in our earlier work: The volume fraction of Al_7Cu_2Fe second-phase particles and the aluminum grain size distribution. These have the advantages of being easily manipulated via heat treatment, easily measured, and capable of prediction with existing physics-based models of microstructure development. It should be noted that we are not considering the role of smaller (nanometer-scale) precipitates, which influence spall strength through their effect on the dynamic flow stress of the alloy but have little effect on spall void nucleation.

3. Forward models for processing

To develop our optimization scheme we need forward models to predict the microstructure that will result from a given processing schedule. Here we focus on simple isothermal heat treatments to manipulate the nucleation and growth of Al_7Cu_2Fe second-phase particles and the aluminum grain size distribution, although in actual practice more sophisticated thermomechanical processing would probably be used. Our use of simple heat treatments is primarily motivated by the primary focus of this work, which is on optimization and the inverse design process. It also enables use to employ robust yet relatively inexpensive commercially-available modeling software, specifically Pandat [31]. We note this is merely a specific example; the optimization framework described below is general and could be applied to other aspects of microstructure using other models.

The overall scheme of forward modeling is outlined in Fig. 1. Our processing schedules comprise one or two stages of isothermal heat treatment, each specified by a temperature and a time at that temperature. In order to describe the evolution of microstructure, we require forward models for the dissolution, nucleation, and growth of the Al_7Cu_2Fe particles and for coarsening of the aluminum grains, as described below. We note that more sophisticated models of microstructural development are available, but we choose simple models for their computational efficiency and because our primary focus is on the multi-objective optimization and not accurate simulation of the microstructure per se.

3.1. Precipitate growth model

Our starting material is 7085-T711 aluminum, which has pre-existing micron-scale Al_7Cu_2Fe second-phase particles [30]. To calculate the volume fraction of these particles in response to heat treatment we need both a thermodynamic model (to calculate the Gibbs energies of the various phases) and a kinetic model (to describe

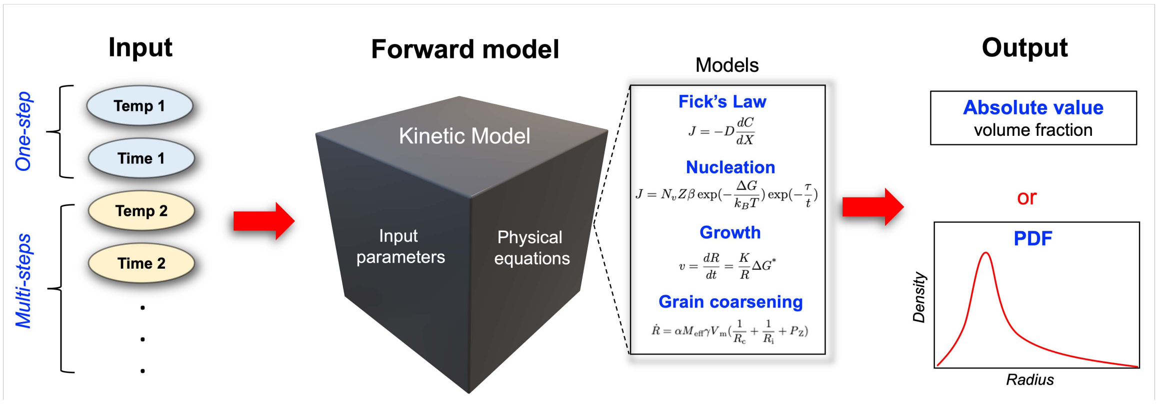

Fig. 1: Flowchart depicting the forward modelling process. The input to the models is a sequence of heat treatment steps, each consisting of a temperature and a holding time at that temperature. The outputs of the models are the volume fraction and size distribution of the Al_7Cu_2Fe second-phase particles, and the aluminum grain size distribution.

nucleation and growth of the particles). Here, we use the Calphad model [32] as implemented in Pandat to calculate the Gibbs energies of the various phases in this alloy system. Nucleation of second-phase particles in Pandat is based on classical nucleation theory [33]. The transient nucleation rate (J) is given by

$$J = NZ\beta \exp\left(-\frac{\Delta G}{k_B T}\right) \exp\left(-\frac{\tau}{t}\right), \tag{1}$$

where N is the nucleation site density, Z is the Zeldovich factor, β is the atomic attachment rate, ΔG is the nucleation barrier, k_{B} is the Boltzmann constant, t is time at temperature T, and τ is the incubation time. The nucleation barrier (ΔG) is calculated from

$$\Delta G = \frac{16\pi}{3} \frac{\sigma_{\alpha\beta}^3}{(\Delta G_V + \Delta G_S)}, \tag{2}$$

where $\sigma_{\alpha\beta}$ is the interfacial energy of the particle/matrix interface, and ΔG_V and ΔG_S are the chemical driving force for nucleation and the elastic strain energy due to lattice misfit, respectively. Assuming that the nuclei are spherical, β is given by

$$\beta = \frac{4\pi R^{*2}}{\alpha^4} D_{\mathrm{eff}} \tag{3}$$

where α is the atomic spacing, D_{eff} is the effective diffusivity, and R^* is the critical radius.

The precipitation dissolution and growth rates are based on a simple model proposed by Morral and coworkers [34], which assumes that growth (or dissolution) are controlled by diffusion. The rate of motion of the curved particle interface is

$$v = \frac{dR}{dt} = \frac{K}{R}\Delta G^*, \tag{4}$$

where K is a kinetic parameter related to the concentration difference between the particles and matrix, R is the radius of curvature of the interface, and ΔG^* is a thermodynamic driving force.

Most of the parameters used in these models are calculated using values built into Pandat, but several require special attention. The molar volume and atomic spacing used are values for pure aluminum ($V_m = 9 \times 10^{-5}\,\mathrm{m}^3$ and $\alpha = 6.4 \times 10^{-10}\,\mathrm{m}$, respectively). Most important is the interfacial energy $\sigma_{\alpha\beta}$, which has a profound effect on the nucleation rate. To find an appropriate value for the interfacial energy we used Bayesian calibration based on a series of heat treatment experiments (see Supplementary Material for details). In the as-received 7085-T711 material the volume fraction of $\mathrm{Al_7Cu_2Fe}$ second-phase particles is $(3.2 \pm 1.7) \times 10^{-3}$. We performed heat treatments consisting of a solutionizing step at 600 °C for up to 30 h optionally followed by annealing at 300 °C for up to 70 h. Heat treatments were performed in air followed by a water quench. The measured precipitate volume fractures were used to calibrate the value of $\sigma_{\alpha\beta}$. The results are shown in Supplementary Information Fig. S2, with a final value of $\sigma_{\alpha\beta} = 2.3 \times 10^{-2}\,\mathrm{J\,m^{-2}}$.

3.2. Grain growth model

The evolution of the aluminum grain size distribution is calculated based on the Kampmann/Wagner Numerical (KWN) model [35]. In this model the continuous probability density function (PDF) representing the grain size distribution is divided into a large number of classes, and the evolution of each class is evaluated for each time step. The grain growth model is

$$\dot{R}_i = \alpha M_{\mathrm{eff}} \gamma V_m \left(\frac{1}{R_c} - \frac{1}{R_i} + Z \right), \tag{5}$$

where α is a shape constant, γ is the interfacial energy of the grain boundary (assumed the same for all boundaries), R_i is the radius of grains in the ith class, R_c a radius parameter related to both grain size and shape, and Z is a parameter describing the Zener drag effect [36, 37]. Finally, M_eff is the effective grain boundary mobility,

$$M_\mathrm{eff} = \frac{A_2}{\delta}\overline{M}, \tag{6}$$

where A_2 is the mobility coefficient, δ is the width of the grain boundary, and $\overline{M}$ is an effective atomic mobility. To use Eqns. 5 and 6 we require values for the various parameters. Several of them (R_c, V_m for aluminum, and $\overline{M}$) we draw from the Pandat thermodynamics and mobility databases. We take the grain boundary width δ to be twice the Al nearest-neighbor distance ($\delta = 3.68 \times 10^{-10}\,\mathrm{m}$) and assume the Zener parameter Z to be zero.

The remaining parameters (γ and A_2) we obtain by calibrating the model against suitable experimental data. To select the kinetic parameters and validate the grain growth model, we performed single-stage heat treatments and characterized the aluminum grain structure via EBSD; details of the characterization are provided in the Supplementary Information. We report grain sizes as effective radii (that is, the radius of the circle with the same area as a given grain). Cumulative grain size distributions for the starting material and material annealed at 500 °C for up to 192 h are shown in Supplementary Information Fig. S3, together with the predictions of the model using the final values of the parameters ($\gamma = 0.65\,\mathrm{J\,m^{-2}}$ and $A_2 = 0.9$).

From Supplementary Information Fig. S3 we can see that the model captures the mean of the grain size distribution well, but does a poorer job replicating the shape of the distribution. This is probably partly due to the relatively complex starting aluminum grain size distribution, which is both heterogeneous and anisotropic [30]. Furthermore, the grain growth model is a simple one, which does not fully capture all aspects of grain structure evolution, such as nucleation of new grains via recrystallization. One could imagine incorporating a more realistic model (such as a Potts model [38–40] or a phase-field model [41–43]) to better capture the evolution of the grain size distribution, at the cost of greater computational complexity and introducing additional parameters into the optimization. Here, our focus is on the optimization problem so we elected to use the simpler model, which allows us to rapidly iterate potential solutions.

4. Single objective inverse problem

In this section we describe inverse process design for the two single objectives (precipitate volume fraction and aluminum grain size distribution) separately, before tackling the more complex problem of multiple objectives in the following section.

4.1. Bayesian optimization for process design

The inverse problem of determining the optimal process parameters to achieve a particular target microstructure, given physics-based forward models of microstructural development, is basically one of optimization. Essentially, we need to define an objective function describing the agreement between the target microstructure and various candidate microstructures predicted by the forward model, and then optimize the process parameters to make the agreement as good as possible. In practice, this is complicated because the parameter space is continuous and multidimensional, and the forward model may be time-consuming and expensive to run. To efficiently explore this space we use Bayesian optimization [44]. We note that Bayesian optimization has been used by others in the context of selecting process parameters to optimize materials properties [45–48]. Here, we seek to use it to identify process parameters necessary to produce a specified microstructure.

Bayesian optimization uses a surrogate model to describe the behavior of the objective function in the parameter space, which is initially unknown, based on a relatively small amount of training data. In our case, we use Gaussian process regression to generate this surrogate model. Because the training data set is small there are uncertainties associated with the predictions of the model, which may vary from point to point in the parameter space. In Gaussian process regression knowledge of the uncertainties allows us to calculate an acquisition function that identifies places (in parameter space) where the forward model should be run to reduce the uncertainty in the surrogate function and locate the optimal value of an objective function. Here we use the expected improvement function as our acquisition function to select the next point to test [49]. The algorithm continues until either the expected improvement reaches a selected small value (implying that little more improvement is anticipated), or a pre-defined number of iterations have been performed. In practice, we have found that running one hundred iterations results in satisfactory convergence even if the expected improvement function is still decreasing. We discuss convergence issues in more detail below.

Figure 2 is a schematic overview of the Bayesian optimization framework for the specific case of a one-step thermal treatment (time t at temperature T) of our aluminum alloy. The starting point is a target microstructure that we aim to achieve. To allow quantitative comparison with the results of the forward models this microstructure is described by a few parameters (for example, a volume fraction of a specific phase of interest, or the aluminum grain size distribution). The optimization process is initiated by random selection of (T, t) points, and the forward models are run for each combination to produce candidate microstructures (again represented by statistical descriptors). The candidate microstructures are compared with the target microstructure through calculation of an objective function which quantifies the differences. This allows creation of a surrogate model which describes the shape of the objective function surface across the parameter space. Next, the expected improvement acquisition function is used to determine a new set of points in parameter space at which to run the forward model. When a specified stopping criterion is met, the algorithm terminates and the optimum values of the process parameters are determined from the global minimum in the objective function.

4.2. Optimizing the volume fraction of second-phase particles

As a demonstration of the Bayesian optimization framework, we first describe its use for a simple case: Finding the heat treatment conditions (time at temperature) necessary to produce a specified volume fraction of an intermetallic phase, Al_7Cu_2Fe. Our as-received starting material (aluminum alloy 7085-T711) has a volume fraction

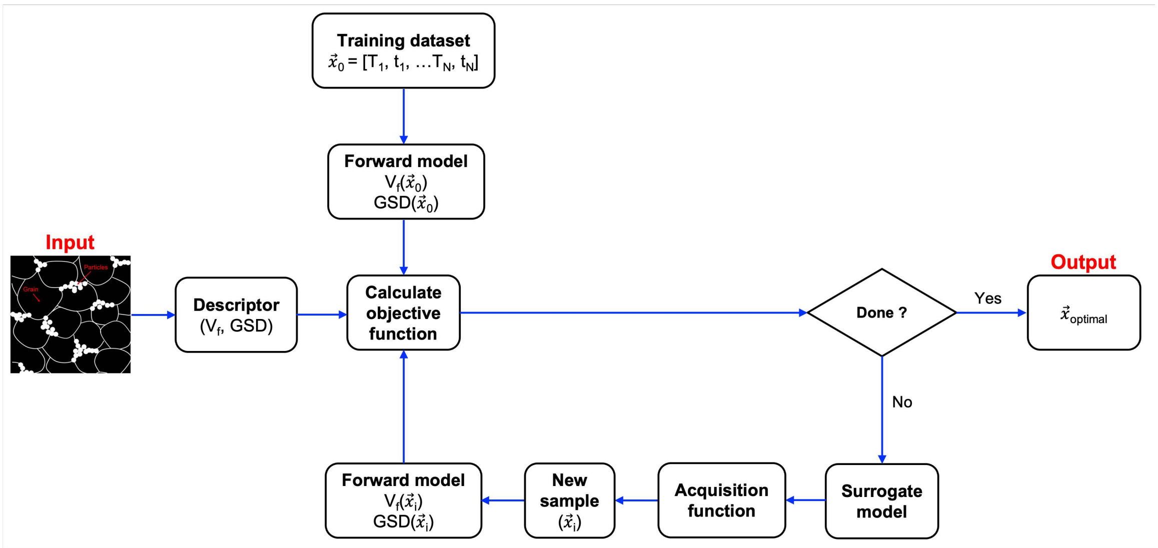

Fig. 2: An overview of the Bayesian optimization framework for determination of process parameters to achieve a specified target microstructure. A surrogate function describing the differences in microstructure derived from various points in process parameter space (calculated from the forward models) is constructed, and additional points are calculated until a defined stopping criterion is met. At this point the optimal process parameters are determined from the global minimum in the surrogate function and returned as the output of the process.

of approximately 0.003 for this phase, and as a target for demonstration purposes we arbitrarily chose to reduce this to 0.001. In principle, this can be achieved by a single-stage heat treatment at a temperature where Al_7Cu_2Fe is unstable, monotonically reducing its volume fraction as it dissolves. Alternatively, we can imagine a two-stage heat treatment consisting of a solutionizing step at high temperature, followed by re-precipitation and growth to the target volume fraction at a lower temperature. We illustrate both in this section.

4.2.1. Single-step heat treatment

Some microstructural features, such as the volume fraction of second-phase particles, can be described by a single number. Others, such as the grain size distribution, require more complex descriptors. If a single-valued descriptor is appropriate we can quantify the difference between the candidate and target microstructures as the absolute value of the difference between the single-value descriptors,

$$y(\vec{x}_i) = |S_C(x_i) - S_o|, \tag{7}$$

where $y(x_i)$ is the value of the objective function at point x_i in process parameter space, $S_C(x_i)$ is the value of the microstructural descriptor of the candidate microstructure at point x_i, and S_o is the descriptor of the target microstructure. For the simple case of precipitation of a second phase considered in this section, S is simply the volume fraction of that phase. The goal of the optimization is to find the point in parameter space x that minimizes the value of $y(x)$.

We describe each stage of heat treatment by two processing variables, temperature T and time t, which together form a vector. For instance, a one-step heat treatment is defined as $\vec{x} = [T_1, t_1]$, and two-step heat treatment is defined as $\vec{x} = [T_1, t_1, T_2, t_2]$. For the single-stage heat treatment we constrained the range of temperatures searched to be from $T = 0$ °C up to the melting temperature of the alloy, $T = 630$ °C [50], knowing that only the upper part of this range is useful for reducing the volume fraction. For the two-stage treatment we decided to reduce the size of the parameter space to be searched by constraining the temperature range of the first step (solutionizing) step to be between 400 °C (above the nose of the TTT curve for Al_7Cu2Fe predicted by Pandat) and the melting point of the alloy. For the second step (precipitation and growth) we constrained the temperature range to be between 100 °C and 400 °C. For similar reasons, we constrained the duration of the stages to be 0 h to 30 h for solutionizing and 0 h to 100 h for precipitation and growth.

Using the UQpy package developed by Olivier and coworkers [51], we initiated the Bayesian optimization framework by selecting twenty random points in processing parameter space, which we designate $\vec{x}_0$. These are used as inputs to the forward model, which produces as output the volume fraction of particles V_f at each $\vec{x}_0$. From these we calculate the objective function $y(\vec{x}_0)$ at each point; these then represent the training data for the model (Fig. 3(a)). The surrogate model is constructed by Gaussian process regression, together with an estimate of the uncertainty in the model at each point $\vec{x}$ in process parameter space. The acquisition function then

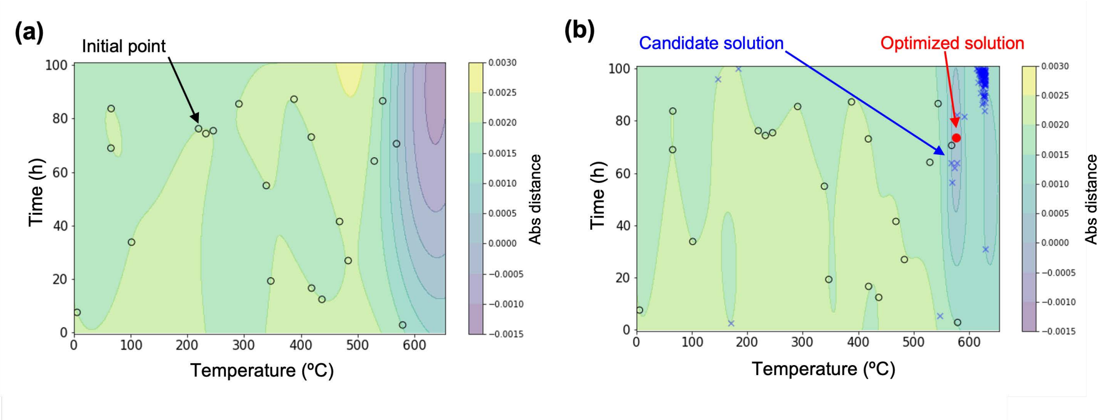

Fig. 3: Contour plots of the Gaussian process regression surrogate model, with contours of constant objective function (Eqn. 7). (a) Initial model with 20 training data points indicated. Note the minimum in the objective function at high temperatures and long times (upper-right corner). (b) Final surrogate model obtained by Bayesian optimization, with the candidate solutions indicated. Note the refinement of the shape and position of the global minimum.

determines new points at which to run the forward model. This process iterates until the stopping criterion is satisfied; as discussed below, in practice we found that running one hundred iterations was sufficient for the objective function to converge on a final value. The final surrogate model is shown in Fig. 3(b), with the minimum identified at ($\vec{x}_{i=12} = \vec{x}_{\text{opt}} = [576\,^\circ\text{C}, 74\,\text{h}]$).

To check this result we prepared a physical sample according to the identified optimum heat treatment. EDS maps of randomly-chosen areas on this sample are provided in the Supplementary Information Fig. S4. The volume fraction determined by manual point count method was $V_f = 0.0011$, close to the target volume fraction of 0.001.

4.2.2. Two-step heat treatment

Next, we move on to a more complicated case of a two-step heat treatment, consisting of a solutionizing stage to dissolve the Al_7Cu_2Fe second-phase particles followed by a nucleation and growth step to re-form the particles, which is described by a vector $\vec{x}_i = [T_1, t_1, T_2, t_2]$ as discussed above. The bounds for the search space are set up according to the TTT diagram, also described above. The initial training points, candidate solutions, and final solution are shown in Fig. 4. Because

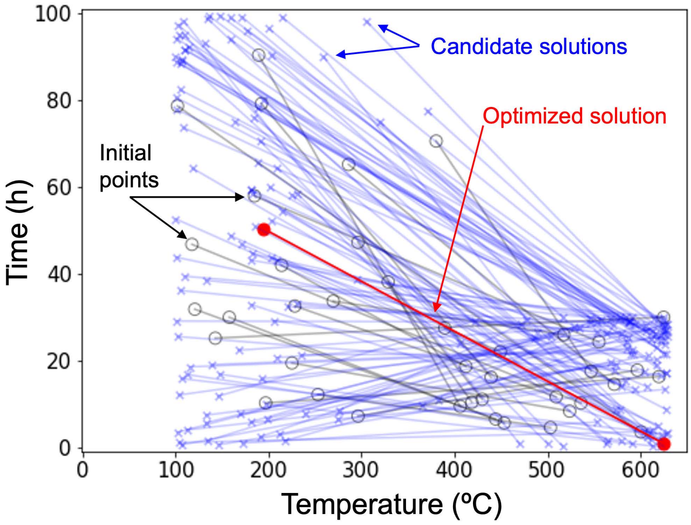

Fig. 4: Two-step heat treatment schedules for producing volume fraction $V_f = 0.001$ of Al_7Cu_2Fe second-phase particles. Each treatment consists of a solutionizing step at high temperature for a given time (t_1, T_1) followed by a nucleation and growth step at lower temperature (t_2, T_2). Each treatment is represented by these two points connected by a line. The initial points used for training are shown as gray open circles, candidate solutions as blue crosses, and the final solution as red solid circles. In each case the first and second stages of the heat treatment are connected by a line of the same color.

the search space is four-dimensional the shape of the objective function surface cannot be easily represented in a manner similar to that for the single-stage treatment. Again we ran the optimization for 100 iterations, but in fact the optimal solution ($\vec{x}_{i=16} = \vec{x}_{\text{opt}} = [625\,^\circ\text{C}, 1\,\text{h}, 194\,^\circ\text{C}, 50\,\text{h}]$) was identified after only 16 iterations.

To verify this result we measured the volume fraction of Al_7Cu_2Fe second-phase particles produced by the optimized heat treatment using the systematic manual point count method on an EDS map (shown in Supplementary Information Fig. S5). The measured volume fraction of $(8 \pm 4) \times 10^{-4}$ is within the experimental uncertainly of the target value, 1×10^{-3}.

4.3. Optimizing the aluminum grain size distribution

Unlike the volume fraction of particles which is described by a single scalar, the grain size distribution of the aluminum is described by a probability density function

(PDF). To compare distributions we use the Kolmogorov–Smirnov (KS) test, which is a commonly-used non-parametric test defined as the absolute maximum difference between two continuous distribution functions (CDFs) [52, 53]. After transforming the grain size distribution PDF to a normalized CDF, we calculate the KS distance as

$$y(\vec{x}_i) = \sup \left| F_{\text{candidate}}(\vec{x}_i) - F_{\text{target}} \right| , \tag{8}$$

where “sup” indicates the largest difference between the candidate distribution $F_{\text{candidate}}(\vec{x}_i,\ j)$ for heat treatment $\vec{x}_i$ and the target distribution $F_{\text{target}}(j)$ across all of the grain sizes in the distribution. This process is illustrated in Fig. 5 for the comparison between the initial microstructure (prior to any heat treatment) and the desired target microstructure.

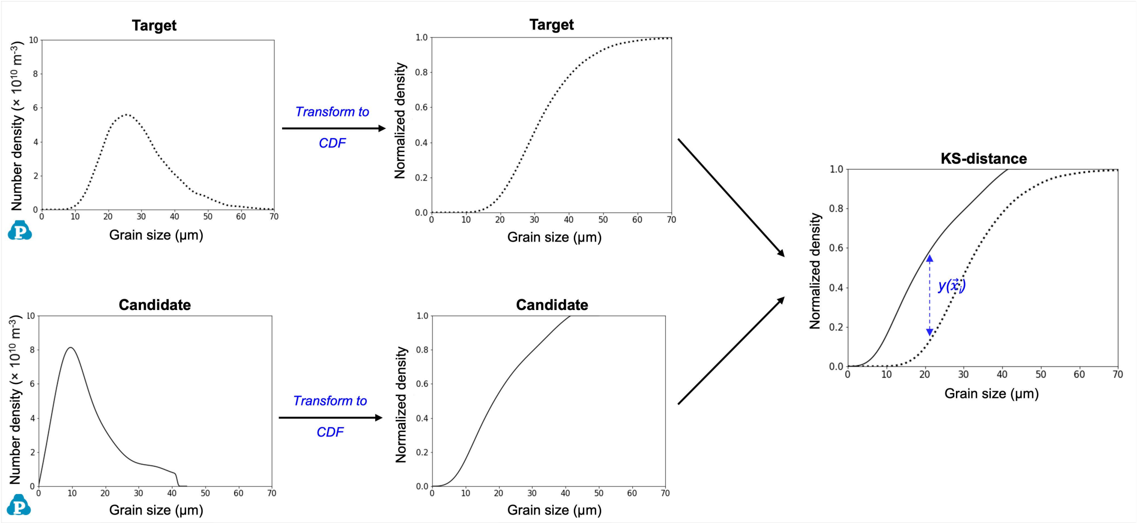

Fig. 5: Calculation of the KS distance $y(\vec{x}_i)$ between a target grain size distribution F_{target} and a candidate distribution $F_{\text{candidate}}(\vec{x}_i)$ corresponding to heat treatment $\vec{x}_j$. Top: Target grain size distribution as a probability density function (PDF) at left, which is cumulatively summed to produce the cumulative density function (CDF) at center. The same procedure is applied to the candidate distribution (bottom). At right the two CDFs are compared; the KS distance $y(\vec{x}_i)$ is the largest absolute difference between the two CDFs (at any grain size).

To test Bayesian optimization of the processing parameters for a target grain size distribution, we arbitrarily chose as our target a log-normal size distribution with a mean of 25 µm and standard deviation 10 µm. In what follows we describe the grain size distribution in terms of unweighted (instead of log-weighted) grain sizes, which for our target grain size distribution is 29 µm. Our search space spanned a two-stage heat treatment with the same constraints as for the particle volume fractions (described above). We note, however, that very little grain growth happens during the second stage of the heat treatment, which is at a lower temperature than the first stage. Fig. 6(a) shows the search space with the initial (training) points, candidate solutions, and the final optimal solution. The optimal solution, $\vec{x}_{i=90} = \vec{x}_{\text{opt}} = [530\,^{\circ}\mathrm{C}, 11\,\mathrm{h}, 208\,^{\circ}\mathrm{C}, 80\,\mathrm{h}]$, was reached after ninety iterations. The initial and target grain size CDFs, along with the CDF of the optimal solution, are shown in Fig. 6(b). Parts (c) and (d) of that figure are contour plots showing the surrogate function before and after optimization.

To verify that this works, we prepared a sample following the optimized heat treatment schedule. The mean aluminum grain size obtained from EBSD on that specimen (shown in Supplementary Information Fig. S6) is 31.0 ± 1.3 µm, which is reasonably close to the target value of 29 µm.

5. Multi-objective inverse problem

In the prior section we demonstrated how Bayesian optimization can be used to determine processing parameters corresponding to a single target microstructural objective, either particle volume fracture or grain size distribution. However, in practice the need to achieve multiple properties simultaneously may produce a requirement to meet more than one microstructural objective at the same time. This is, in principle, significantly more challenging than optimizing for a single objective because a change to the processing parameters to improve one feature of the microstructure may come at the expense of another. In this section we demonstrate several approaches to this problem, optimizing both the volume fraction of Al_7Cu_2Fe second-phase particles and the aluminum grain size distribution simultaneously.

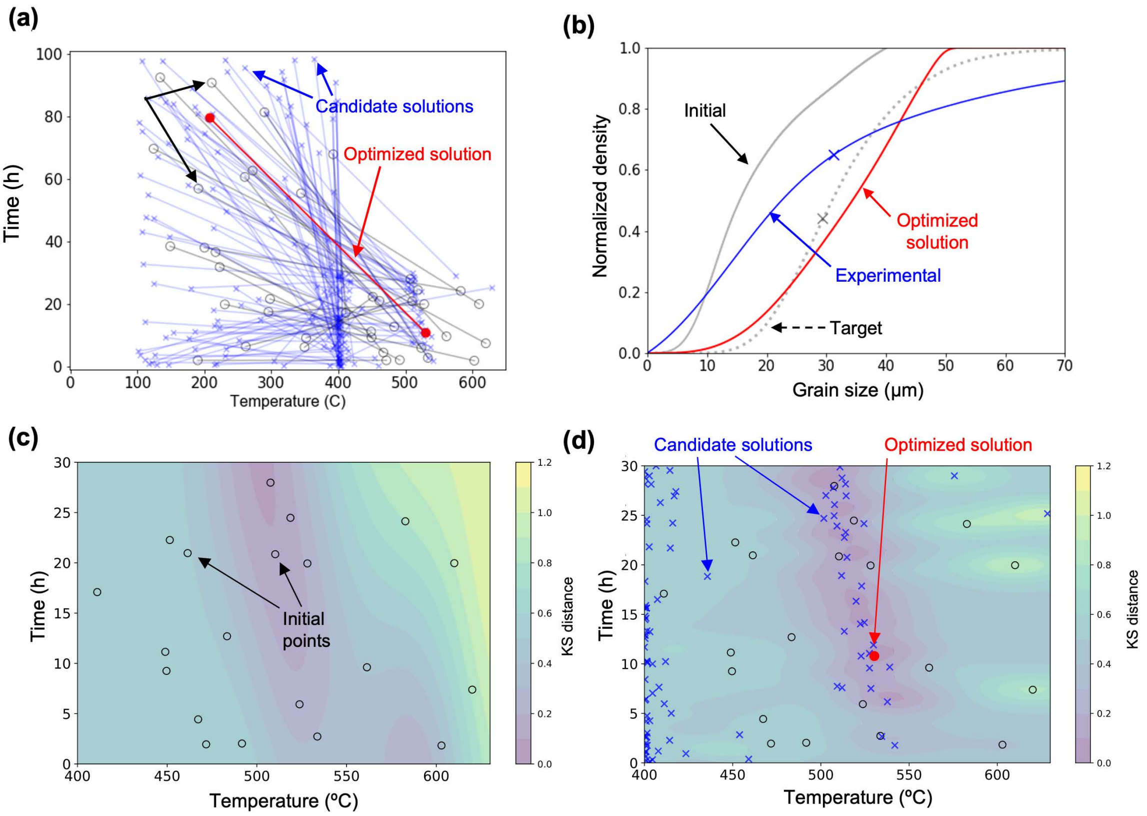

Fig. 6: Two-stage heat treatment for achieving a target grain size distribution. (a) shows the heat treatment conditions, where each solution represents a combination of a time at high temperature followed by a time at a lower temperature (similar to Fig. 4 above). (b) Comparison of CDFs for the initial grain size distribution, target distribution, and distribution corresponding to the optimized heat treatment. (c) Contour plots of the initial surrogate model, with contours of constant objective function (Eqn. 8). Note that although the search space is four-dimensional, this plot shows only the higher-temperature stage ($[T_1,\, t_1]$) because the second stage (which occurs at lower temperature) has almost no effect on the aluminum grain size distribution. (d) Final surrogate model obtained by Bayesian optimization, with the candidate solutions indicated. Note the refinement of the shape and position of the global minimum.

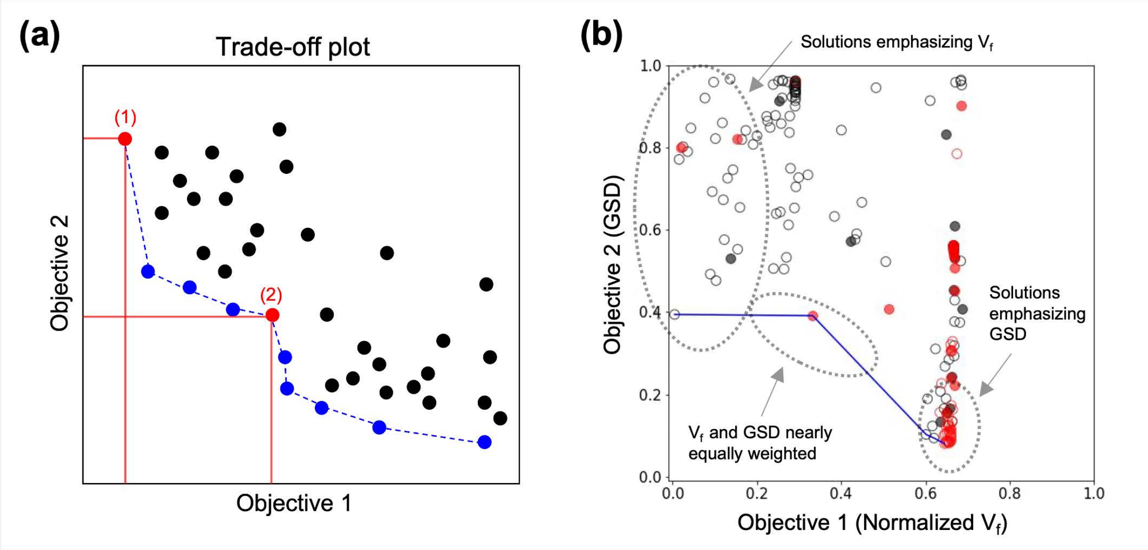

Fig. 7: (a) Schematic Pareto trade-off plot. The points represent candidate solutions, and the points defining the Pareto front are shown in blue. The two points shown in red illustrate two optimal solutions with different relative importance of the two objectives. (b) Trade-off plot for the particle volume fraction (black) and grain size (red) objectives optimized independently. The initial points and candidate solutions are represented as filled and open circles, respectively. The Pareto front is in blue.

5.1. Pareto front methods

To begin we need some way to assess the trade-offs inherent in optimizing against multiple objectives. This is the idea behind the Pareto front, illustrated schematically in Fig. 7(a). The value of the objective function for each goal is plotted against the other; the Pareto front is the set of solutions for which no other solution improves on both objectives (*i.e.* is closer to the origin in Fig. 7(a)). For example, points (1) and (2) on Fig. 7(a) represent different trade-offs between the two objectives. Thus the Pareto front comprises a set of solutions which are all in some sense optimal but which give different weights to the two objective functions.

A simple approach to optimizing against multiple objectives would be to optimize for each separately, combine all the candidate solutions into a single set, construct the Pareto front, and then choose a solution on the front that reflects the relative importance of the objectives. To illustrate the present case, Fig. 7(b) is a Pareto-front plot for the candidate solutions for the single-objective volume fraction and grain size distribution optimizations discussed above. It is apparent that the candidate solutions in the vicinity of the Pareto front are sparsely distributed, particularly in the regions of the plot where the two objectives are nearly equally weighted. This is not surprising, given that in the two optimizations the objectives were considered independently. It thus seems likely that there may be other, potentially better solutions to be found in this region, and it becomes important to optimize against both objectives simultaneously.

As a first attempt at multiple-objective optimization we used an implementation of the non-dominated sorting genetic algorithm II (NSGA-II) [54, 55] as implemented in the pymoo package [56]. Here, we used the simulated binary crossover (SBX) [57] and polynomial mutation (PM) [58] genetic operators. The NSGA-II parameters were gene length of four, crossover rate of 0.6, mutation rate of 0.05 and an initial population size of twenty. The Pareto front evolves with an increasing number of generations; in practice we found that the Pareto front evolved to a stable position after about forty generations.

The operation of the NSGA-II algorithm is illustrated in Fig. 8. Starting from an initial population of candidate solutions (the first generation), the forward models are run to determine the volume fraction of Al_7Fe_2Cu second-phase particles and the aluminum grain size distribution, and the values of the objective functions (relative to the targets) are calculated. The individual candidate solutions are then ranked (by the fast non-dominated sorting approach and the crowding distance [59]) and then the top individuals (ten here) are selected to be parents of the next generation. From those parents, offspring are generated by crossover and mutation of the four variables (two times and two temperatures) and the fitness of each offspring is calculated using the forward models and the values of the objective functions. The algorithm iterates until the specified number of generations (or some convergence criterion) is reached.

The results are shown in Fig. 9, which compares the solutions produced by the NSGA-II algorithm after forty generations with the naive Pareto front derived from the separate single-objective Bayesian optimizations. Notice that the smallest achievable values of the two objective functions are about the same as in Fig. 7(b), nearly

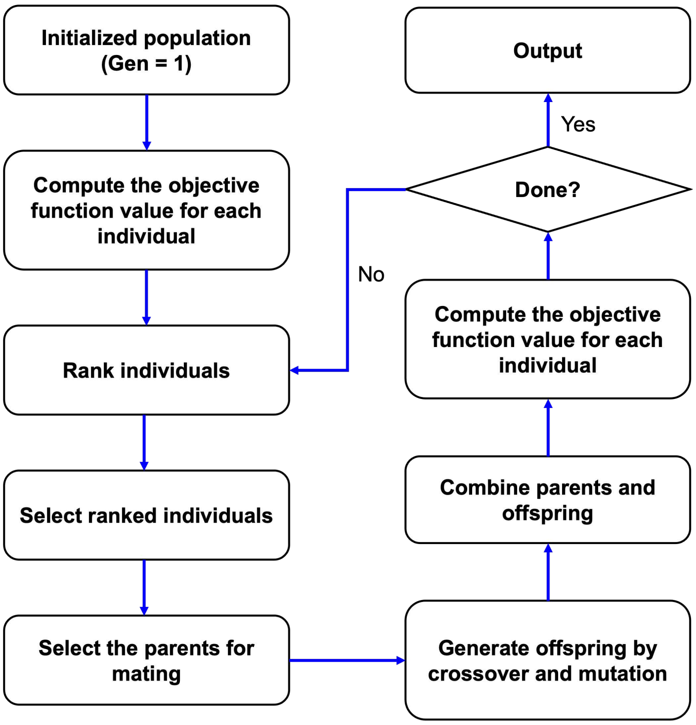

Fig. 8: An overview of the NSGA-II framework for determination of process parameters to achieve a specified target microstructure.

zero for the volume fraction objective and about 0.1 for the grain size distribution objective. But with the genetic algorithm the objective function for both objectives can be reduced simultaneously in the region between these extremes, reflecting more optimal solution for various trade-offs between the two objectives. Notice also that the solutions identified are more uniformly and densely distributed along the Pareto front.

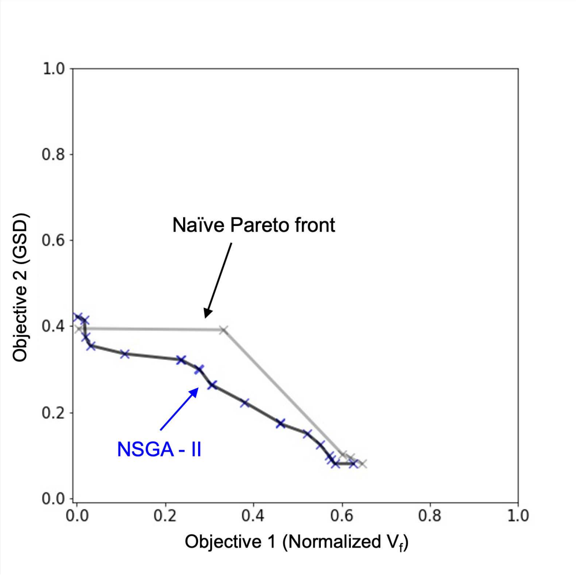

Fig. 9: The optimized Pareto front determined by the genetic algorithm (blue) compared with the naive Pareto front determined from the single-objective optimizations (black).

5.2. Bayesian optimization with desirability function-based scalarization

Conceptually, a simple way to approach Bayesian optimization for multi-objective problems is to combine the objectives into a single, scalar function. For example, we can simply calculate a weighted sum of the individual objection functions,

$$Y(x) = \sum_{i=1}^{k} w_i \, y_i(x), \tag{9}$$

where w_i is the weight for objective function y_i, x represents the independent variables, and $Y(x)$ is the overall objective function.

One complication in applying Eqn. 9 is that the various individual objectives $y_i(x)$ may represent quite different values depending on the nature of the objective (for example, the particle volume fraction and the KS distance of the grain size distributions calculated above). This issue is addressed by the use of "desirability functions," which are essentially normalized versions of individual objective functions [60]. For example, if $y_i(x)$ is the distance between a value and the target then the desirability can be described by

$$d_i[y_i(x)]\begin{cases} 1, & \text{for } y_i(x) < \text{LSL} \\ \left(\frac{y_i(x)-\text{USL}}{\text{LSL}-\text{USL}}\right)^s & \text{for LSL} \leq y_i(x) \leq \text{USL} \\ 0, & \text{for } y_i(x) > \text{USL}, \end{cases}$$

where LSL and USL are the lower and upper specification limits respectively and s is a geometrical parameter. The desirability scales from 1 (if the value of $y_i(x)$ is within the LSL of its target) to zero (if $y_i(x)$ is more than the USL above the target. The geometrical parameter s describes the shape of the transition between the two extremes of $d_i = 1$ close to the target and $d_i = 0$ far away. We note that the desirability may be calculated in slightly different ways, depending on the nature of the objective [60].

To optimize multiple objectives simultaneously, Derringer and coworkers constructed an overall desirability function Y as a weighted product of the individual desirability functions,

$$Y = \left[\prod_{i=1}^{m} d_i(y_i)^{w_i}\right]^{\frac{1}{\sum w_i}} \tag{10}$$

where w_i is the weighting factor of the ith objective function [61]. The weights are chosen based on the assessed relative importance of the individual objectives, with the constraint that $0 \leq w_i \leq 1$ and normalized such that $\sum_i w_i = 1$. In essence, the optimal solutions form a Pareto front as above, and the chosen weights determine which of the solutions on that front is chosen. By transforming the multi-objective problem into a single objective problem in this way we can apply the Bayesian optimization framework to its solution.

To this approach, we optimize the processing conditions for producing a particular Al_7Cu_2Fe volume fraction and aluminum grain size distribution (using the same targets and limits as above), assigning different weights to the two objectives (Fig. 10).

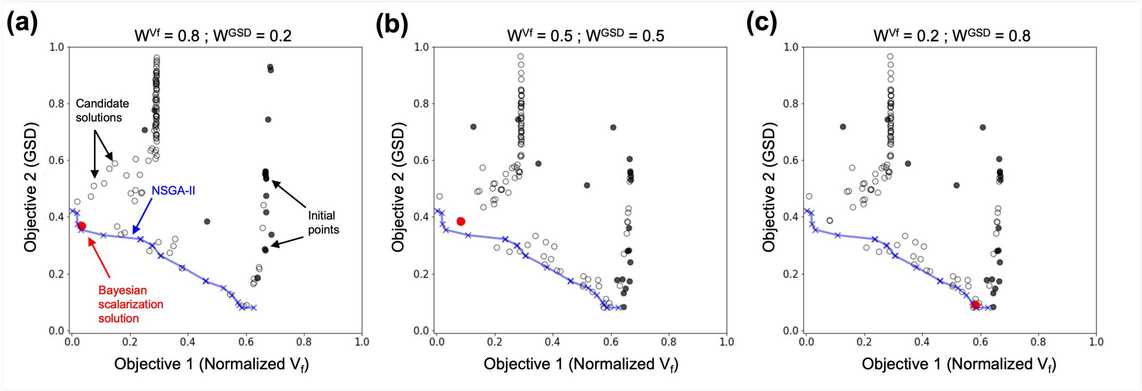

Fig. 10: Multi-objective problem (particle volume fraction and aluminum grain size distribution) solved using overall desirability and Bayesian optimization. Trade-off plots showing the two objective functions by applying different weights for each objective function with (a) $w^{\mathrm{VF}} = 0.8$, $w^{\mathrm{GSD}} = 0.5$ (b) $w^{\mathrm{VF}} = 0.5$, $w^{\mathrm{GSD}} = 0.5$ (c) $w^{\mathrm{VF}} = 0.2$, $w^{\mathrm{GSD}} = 0.8$. The initial points (filled circles), candidate solutions (open circles), and the final optimized solution (red). Also shown is the Pareto front determined by the genetic algorithm NSGA-II described above.

Figure 11(a) illustrates the search space, including the initial (training) points, candidate solutions, and the optimized solution. The optimization process applies a weight of 0.2 to objective function 1 (volume fraction) and a weight of 0.8 to objective function 2 (grain size distribution). The optimal solution was $\vec{x}_{\mathrm{opt}} = [626\,^{\circ}\mathrm{C}, 0.2\,\mathrm{h}, 169\,^{\circ}\mathrm{C}, 60\,\mathrm{h}]$, reached after 86 iterations.

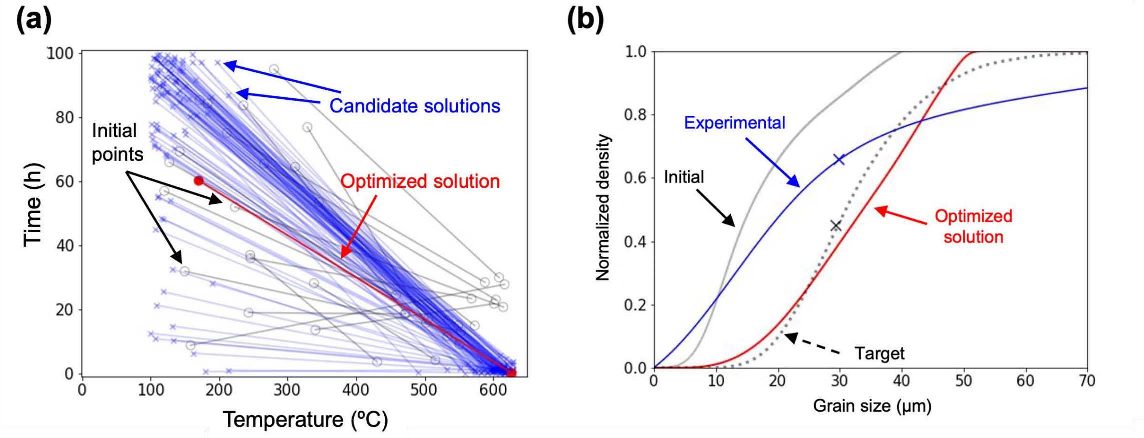

Fig. 11: Multi-objective problem (particle volume fraction and aluminum grain size distribution) solved using overall desirability and Bayesian optimization. (a) Initial, candidate, and optimized two-stage heat treatments, similar to Fig. 4 above). (b) Grain size distributions: Initial (gray), target (dashed), optimized solution (red), and experimentally-determined final (blue).

The initial and target grain size CDFs, along with that for the optimal solution are shown in Fig. 11(b). To verify this result, we heated-treated a sample following the optimized solution. The experimentally-determined CDF of the grain size distribution obtained from EBSD on that specimen (Supplementary Information Fig. S7) is shown in blue in Fig. 11(b). As with the single-objective optimization above, the agreement between the mean of the experimentally-determined distribution (30.0 ± 1.4 μm) is in good agreement with the target 29 μm, although the shape of the experimentally-determined distribution varies significantly from that of the optimized solution. The optimized volume fraction of Al_7Fe_2Cu second-phase particles was 0.0015, compared to a target value of 0.001. The experimental value of the volume fraction of Al_7Fe_2Cu second-phase particles was (0.0010 ± 0.0005), and the corresponding EDS map is shown in Supplementary Information Fig. S8, in good agreement with both the target and optimized solution. We note that the solutions determined by Bayesian optimization lie close to the Pareto front found by the NSGA-II algorithm, as expected (Fig. 10 (c)).

5.3. Bayesian optimization using expected hypervolume improvement

A limitation of the desirability function approach to Bayesian optimization of multi-objective problems is that it only locates a single point on the Pareto front, reflecting the weighting of the individual objectives selected by the user. This is problematic because it does not provide a complete view of the trade-offs involved, and because preferences (*i.e.* weighting of the objectives) may change with time.

An alternative approach which allows us to identify the entire Pareto front is Bayesian optimization using expected hypervolume improvement (EHVI) as the acquisition function [62]. The hypervolume is the size of the dominated subspace bounded by a provisional Pareto front (P) and a reference point (r), as illustrated schematically in Fig. 12(a). As new points y are calculated, the improvement in the Pareto front is quantified by the increase in the hypervolume, as shown in Fig. 12(b). The hypervolume improvement (HVI) is given by

$$\mathrm{HVI} = \mathrm{HV}(P_{\mathrm{new}}) - \mathrm{HV}(P) \tag{11}$$

where P_{new} is the Pareto front determined with the additional candidate points y. To choose these new points, we calculate the expected hypervolume improvement,

$$\mathrm{EHVI}(\mu, \sigma, P, r) = \int \mathrm{HVI}(P, y) \cdot \mathrm{PDF}_{\mu,\sigma}(y)\, dy, \tag{12}$$

where $\mathrm{PDF}_{\mu,\sigma}(y)$ is the probability density function associated with the new points y. Here we use single-task Gaussian process regression to generate surrogate models, with a multivariate independent normal distribution with mean μ and standard deviation σ [63].

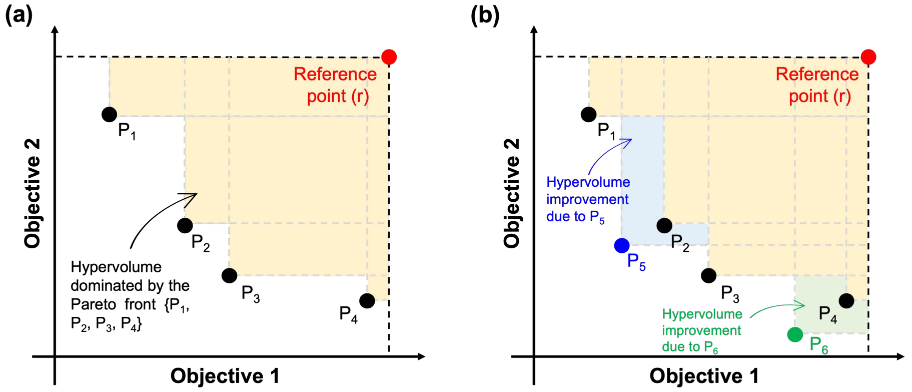

Fig. 12: (a) Schematic hypervolume (shaded yellow) between a reference point r and a provisional Pareto front $P = \{P_1, P_2, P_3, P_4\}$. (b) Improvement of the hypervolume by introduction of two new points (P_5 (shaded blue), P_6 (shaded green)).

We performed Bayesian optimization using the EHVI acquisition function using the Botorch package [64]. We chose an initial population of twenty points, and introduced a new candidate point at every iteration. The optimization process was terminated after eight hundred iterations. The results are shown in Fig. 13, which compares the Pareto fronts produced by Bayesian optimization with EHVI with that from NSGA-II. We see that the Pareto fronts determined by the two algorithms are similar.

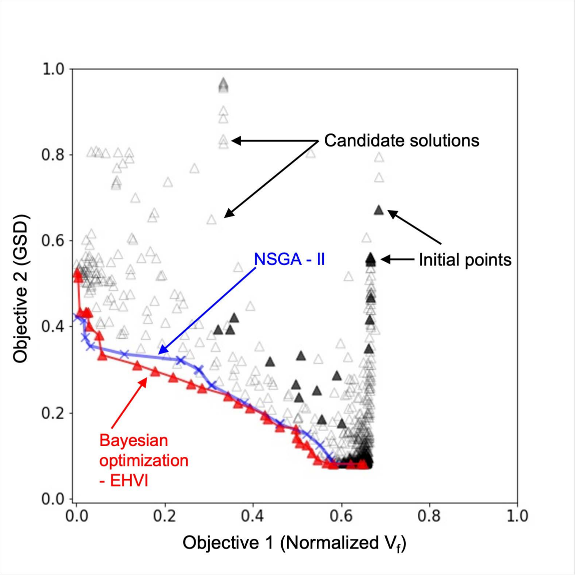

Fig. 13: The optimized Pareto front determined by Bayesian optimization with expected hypervolume improvement (red) and NSGA-II (blue). The filled black triangles are the initial points for the Bayesian optimization, and candidate solutions are shown as open triangles.

6. Discussion

Our results show that both Bayesian optimization with expected hypervolume improvement and a genetic algorithm (NSGA-II) can be used for inverse process design to achieve a particular target microstructure, and that for our test case the two approaches produce similar Pareto fronts. Two key questions regarding this work are the relative efficiencies of the two algorithms (which may become especially important when more realistic and therefore more computationally-intensive physics-based forward models are used) and the uncertainty in the Pareto front determined by this approach.

The relative efficiency of Bayesian optimization with expected hypervolume improvement and the NSGA-II genetic algorithm is illustrated in Fig. 14, which shows the increase of the hypervolume as a function of the number of candidate points evaluated for each approach. A total of fourteen runs were performed (seven each for Bayesian optimization and NSGA-II), and in the plot the hypervolumes are normalized by the value of the largest hypervolume after 800 evaluation among these fourteen runs. We see that although both algorithms converge to similar hypervolumes (and thus similar Pareto fronts), Bayesian optimization improves more quickly as a function of the number of evaluations. In practice, one would not pre-determine the number of evaluations, but would instead choose a convergence criterion based on the rate of increase of the hypervolume [65]. In this light, Bayesian optimization would be preferred because it converges more quickly, at least for the problem considered here.

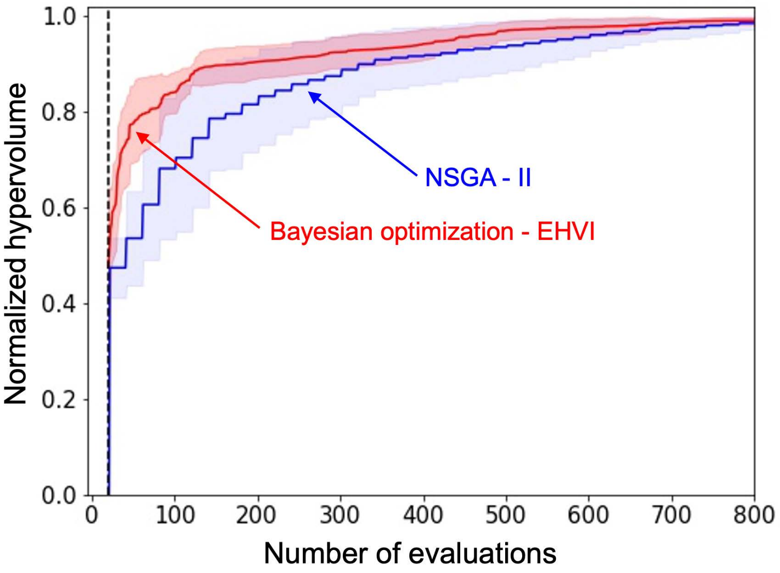

Fig. 14: Hypervolume improvement for the Bayesian optimization and generic algorithms, as a function of the number of points evaluated. The solid line represents the mean of the hypervolume (averaged over seven runs for each algorithm) and the shaded area illustrates the standard deviation. The hypervolumes are normalized to the value of the largest final hypervolume from the fourteen total runs.

Naturally, the question of relative efficiency becomes more important as the computational cost of the forward models increases. For this work we deliberately chose relatively simple forward models because their low computational cost allowed for rapid iteration through the optimization process. More accurate results could be obtained with more sophisticated models. For example, Monte Carlo Potts models [38, 39] phase field models [41, 42], and hybrid models [66] would likely predict the grain size distribution more accurately by accounting for complex effects such as anisotropy, abnormal grain growth, and coupling of grain boundary migration with local chemistry. Similarly, in this study we chose to examine only simple heat treatments; consideration of more complex thermomechanical processing would require appropriate forward models.

We note that although Bayesian optimization works well for this fairly simple problem with two microstructural objectives, this may not be true in other situations. In particular, in a situation with a larger number of microstructural objectives (which may be correlated with each other) the relative efficiency of Bayesian optimization may suffer as the computational complexity of constructing the surrogate model increases. For problems of higher complexity, more efficient and scalable optimization algorithms, such as reinforcement learning [67, 68], may be required.

A second important question relates to the uncertainties in the results. The systematic (epistemic) uncertainty of the predicted optimal processing sequence arises from two sources. First is the choice of forward models; different models naturally produce different results and will result in a different optimization [69]. Second is uncertainty in the parameters of the models — most notably in our case, that of the interfacial energy $\sigma_{\alpha\beta}$ in Eqn. 2 and the grain boundary energy in Eqn. 5. Because the choice of forward models is not a focus of this study (as discussed above, we deliberately chose simple models of low computational complexity), here we discuss only the parametric uncertainty.

We built the surrogate model with a fixed set of parameters of the forward model, obtained through a Bayesian calibration process using the results of a limited number of experiments (as discussed in Supplemental Material). The primary objective of Bayesian calibration is to reduce the difference between the experimental observations and simulation results based on an absolute distance metric (as shown in Supplementary Information Fig. S2 and S3). An advantage of Bayesian calibration is that it returns the posterior distribution of the model parameters, which we can use in assessing the parametric uncertainty in the Pareto front. To do so, for each point along the Pareto front (Fig 13) we sampled the posterior distributions of both interfacial energy and grain boundary energy, and used these values to compute the value of the two objective functions around each point on the Pareto front. The results are shown in Fig. 15.

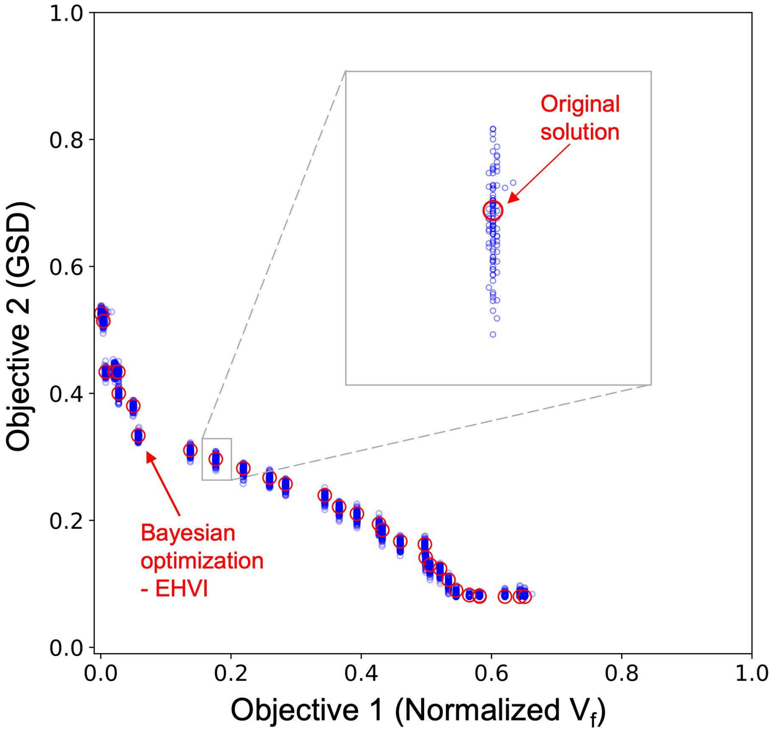

Fig. 15: The effect of uncertainty in model parameters on the Pareto front. The red circles are the Pareto front positions determined by Bayesian optimization with expected hypervolume improvement (Fig 13). Each blue point represents a new combination of objective functions calculated from the forward models using model parameters sampled from the posterior distributions obtained from Bayesian calibration.

The new points (blue) in Fig. 15 represent the uncertainty in the Pareto front position. We note that the position of the front is not strongly affected by the uncertainty in the model parameters. This is an interesting observation, particularly because the nucleate rate of the second-phase particles is extremely sensitive to the value of the interfacial energy, which appears as a cubic term in an exponential (Eqns. 1 and 2). The lack of sensitivity of the Pareto front to this parameter is likely because the objective function is associated with the volume fraction of the second phase, which is associated with thermodynamic phase stability, not necessarily the kinetics of precipitation. It may well be that if the objective function was based on the size distribution of the second-phase particles, rather than their volume fraction, the uncertainty in the results would be larger.

We note that this is a preliminary approach to understanding uncertainty in the inverse design process, providing a basic understanding of the effect of parametric uncertainty on the Pareto front. However, it simplifies the complexity of the problem by assuming that the posteriors of the calibrated parameters are normally distributed [70]. A more sophisticated approach would be to apply a stochastic simulation technique, such as the Markov Chain Monte Carlo (MCMC) method, during Bayesian calibration. This method applies a flexible probabilistic framework that better captures the full scope of potential variability of the parameters [71]. The MCMC method generates samples of the model parameters using a Markov chain, which is designed such that its stationary distribution is expected to converge to the posterior distribution of the model parameters. This produces parameter values without assumptions about the shape of the posterior distribution. Finally, we note that a more complete analysis of uncertainty would include the choice of the forward models. As mentioned above other models such as phase field or Potts could provide more sophisticated predictions of grain growth, but are computationally much more intense than the simple analytical model used here. The effect of these kinds of choices on inverse process design is an important subject for future research.

7. Conclusion

In this study, we described the development of an inverse “process-structure” framework that identifies the optimal processing parameters required to generate a target microstructure. We applied isothermal heat treatments to a commercial aluminum alloy in order to modify specific aspects of its microstructure, including the volume fraction of Al_7Cu_2Fe second-phase particles and the aluminum grain size distribution, both individually and separately. For the multiple objective optimization problem we find that Bayesian optimization with expected hypervolume improvement and a genetic algorithm (NSGA-II) produce similar Pareto fronts describing the trade-off between the two objectives, but that Bayesian optimization converges to the solution more quickly. We also considered the effect of uncertainty in the parameters of the forward model (specifically, interfacial energy of the Al_7Cu_2Fe particles and the aluminum grain boundary energy) on the Pareto front. Interestingly, even though these parameters (and particularly the interfacial energy) can have a profound effect on the kinetics of the microstructural evolution, parametric uncertainty has only a modest effect in the position of the Pareto front.

Funding

This material is based upon work supported by the National Science Foundation under Grant No. 1921959.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Author contributions

Dung-Yi Wu: Conceptualization, methodology, investigation, writing – original draft, software, formal analysis. **Todd C. Hufnagel**: Conceptualization, review and editing, supervision, funding acquisition, project administration.

Acknowledgement

We gratefully acknowledge Michael Shields, Dimitris Giovanis and Dimitrios Tsapetis for assistance with UQpy and discussions regarding uncertainty quantification; Margaret Eminizer for help in automating the forward models; and valuable technical suggestions regarding Pandat from Weisheng Cao. Portions of the experimental work were performed at the Johns Hopkins Materials Characterization and Processing (MCP) facility.

Supplementary materials

Supplementary material related to this article can be found in the online version.

References

[1] J. Correa-Baena, K. Hippalgaonkar, J. van Duren, S. Jaffer, V. Chandrasekhar, V. Stevanovic, C. Wadia, S. Guha, T. Buonassisi, Accelerating materials development via automation, machine learning, and high-performance computing, Joule 2 (8) (2018) 1410–1420. doi:10.1016/j.joule.2018.05.009.

[2] J. Wang, X. Chen, L. Yang, Reliability-based multi-objective optimization incorporating process–property–performance relationship of double-pulse MIG welding using hybrid optimization strategy, Struct. Multidiscipl. Optim. 65 (5) (2022) 148. doi:10.1007/s00158-021-03103-x.

[3] C. Lu, X. Jia, J. Lee, J. Shi, Knowledge transfer using bayesian learning for predicting the process-property relationship of inconel alloys obtained by laser powder bed fusion, Virtual Phys. Prototyp. 17 (4) (2022) 787–805. doi:10.1080/17452759.2022.2068447.

[4] J. Wang, Y. Wang, Y. Chen, Inverse design of materials by machine learning, J. Mater. 15 (5) (2022) 1811. doi:10.3390/ma15051811.

[5] A. Sardeshmukh, S. Reddy, B. Gautham, Bayesian framework for inverse inference in manufacturing process chains, Integr. Mater. Manuf. 8 (2019) 95–106. doi:10.1007/s40192-019-00140-9.

[6] B. Kim, S. Lee, J. Kim, Inverse design of porous materials using artificial neural networks, Sci. Adv. 6 (1) (2020) eaax9324. doi:10.1126/sciadv.aax9324.

[7] R. Dong, Y. Dan, X. Li, J. Hu, Inverse design of composite metal oxide optical materials based on deep transfer learning and global optimization, Comput. Mater. Sci. 188 (2021) 110166. `doi:10.1016/j.commatsci.2020.110166`.

[8] M. Karasuyama, H. Kasugai, T. Tamura, K. Shitara, Computational design of stable and highly ion-conductive materials using multi-objective Bayesian optimization: Case studies on diffusion of oxygen and lithium, Comput. Mater. Sci. 184 (2020) 109927. `doi:10.1016/j.commatsci.2020.109927`.

[9] M. A. Mahmood, A. U. Rehman, B. Karakaş, A. Sever, R. U. Rehman, M. U. Salamci, M. Khraisheh, Printability for additive manufacturing with machine learning: Hybrid intelligent gaussian process surrogate-based neural network model for Co-Cr alloy, J Mech Behav Biomed Mater 135 (2022) 105428. `doi:10.1016/j.jmbbm.2022.105428`.

[10] F. Yu, Y. Zhao, Z. Lin, Y. Miao, F. Zhao, Y. Xie, Prediction of mechanical properties and optimization of friction stir welded 2195 aluminum alloy based on BP neural network, J. Met. 13 (2) (2023) 267. `doi:10.3390/met13020267`.

[11] K. Sattari, Y. Xie, J. Lin, Data-driven algorithms for inverse design of polymers, Soft Matter 17 (33) (2021) 7607–7622. `doi:10.1039/D1SM00725D`.

[12] T. K. Patra, T. D. Loeffler, S. K. Sankaranarayanan, Accelerating copolymer inverse design using Monte Carlo tree search, Nanoscale 12 (46) (2020) 23653–23662. `doi:10.1039/D0NR06091G`.

[13] J. Shen, V. Venkatesh, R. Noraas, S. Ghosh, Parametrically upscaled crack nucleation model (PUCNM) for fatigue nucleation in titanium alloys containing micro-texture regions (MTR), Acta Mater. 252 (2023) 118929. `doi:10.1016/j.actamat.2023.118929`.

[14] B. D. Snow, S. G. Baird, C. Kurniawan, D. T. Fullwood, E. R. Homer, O. K. Johnson, Grain boundary structure-property model inference using polycrystals:

The underdetermined case, Acta Mater. 209 (2021) 116769. doi:10.1016/j.actamat.2021.116769.

[15] H. Tang, H. Huang, C. Liu, Z. Liu, W. Yan, Multi-scale modelling of structure-property relationship in additively manufactured metallic materials, Int. J. Mech. Sci. 194 (2021) 106185. doi:10.1016/j.ijmecsci.2020.106185.

[16] D. Uribe-Suárez, P. O. Bouchard, M. Delbo, D. Pino-Muñoz, Numerical modeling of crack propagation with dynamic insertion of cohesive elements, Eng. Fract. Mech. 227 (2020) 106918. doi:10.1016/j.engfracmech.2020.106918.

[17] S. P. Ong, Accelerating materials science with high-throughput computations and machine learning, Comput. Mater. Sci. 161 (2019) 143–150. doi:10.1016/j.commatsci.2019.01.013.

[18] L. Morand, N. Link, T. Iraki, J. Dornheim, D. Helm, Efficient exploration of microstructure-property spaces via active learning, Front. Mater. Sci. 8 (2022) 628. doi:10.3389/fmats.2021.824441.

[19] X. Y. Lee, J. R. Waite, C.-H. Yang, B. S. S. Pokuri, A. Joshi, A. Balu, C. Hegde, B. Ganapathysubramanian, S. Sarkar, Fast inverse design of microstructures via generative invariance networks, Nat. Comput. Sci 1 (3) (2021) 229–238. doi:10.1038/s43588-021-00045-8.

[20] M. Hasan, P. Acar, Microstructure-sensitive stochastic design of polycrystalline materials for quasi-isotropic properties, AIAA Journal 60 (12) (2022) 6869–6880. doi:10.2514/1.J061788.

[21] R. Zhu, S. Li, I. Karaman, R. Arroyave, T. Niendorf, H. Maier, Multi-phase microstructure design of a low-alloy TRIP-assisted steel through a combined computational and experimental methodology, Acta Mater. 60 (6) (2012) 3022–3033. doi:10.1016/j.actamat.2012.02.007.

[22] L. Johnson, R. Arróyave, An inverse design framework for prescribing precipitation heat treatments from a target microstructure, Mater. Des. 107 (2016) 7–17. doi:10.1016/j.matdes.2016.06.009.

[23] A. Tran, J. A. Mitchell, L. P. Swiler, T. Wildey, An active learning high-throughput microstructure calibration framework for solving inverse structure–process problems in materials informatics, Acta Mater. 194 (2020) 80–92. doi:10.1016/j.actamat.2020.04.054.

[24] P. Honarmandi, V. Attari, R. Arroyave, Accelerated materials design using batch Bayesian optimization: A case study for solving the inverse problem from materials microstructure to process specification, Comput. Mater. Sci. 210 (2022) 111417. doi:10.1016/j.commatsci.2022.111417.

[25] J. Dornheim, L. Morand, S. Zeitvogel, T. Iraki, N. Link, D. Helm, Deep reinforcement learning methods for structure-guided processing path optimization, J. Intell. Manuf. 33 (2022) 333–352. doi:10.1007/s10845-021-01805-z.

[26] D. Liu, Y. Wang, Metal additive manufacturing process design based on physics constrained neural networks and multi-objective Bayesian optimization, Manuf. Lett. 33 (2022) 817–827. doi:10.1016/j.mfglet.2022.07.101.

[27] Y. Wang, M. Qi, H. He, L. Wang, Spall failure of aluminum materials with different microstructures, Mech. Mater. 69 (1) (2014) 270–279. doi:10.1016/j.mechmat.2013.11.005.

[28] V. K. Euser, D. R. Jones, D. T. Martinez, J. A. Valdez, C. P. Trujillo, C. Cady, S. J. Fensin, The effect of microstructure on the dynamic shock response of 1045 steel, Acta Mater. 250 (2023) 118874. doi:10.1016/j.actamat.2023.118874.

[29] D. Mallick, J. Parker, J. Wilkerson, K. Ramesh, Estimating void nucleation statistics in laser-driven spall, J. Dyn. Behav. Mater. 6 (3) (2020) 268–277. doi:10.1007/s40870-020-00248-6.

[30] D.-Y. Wu, C. Miao, C. S. DiMarco, K. Ramesh, T. C. Hufnagel, Microstructural effects on the spall failure of 7085 aluminum alloy, Mater. Sci. Eng. A 866 (12) (2023) 144674. doi:10.1016/j.msea.2023.144674.

[31] F. Li, Z. Li, C. Tang, L. Zhang, Q. Tan, C. Chen, M. Zhang, K. Zhou, Design high-strength Al–Mg–Si alloy fabricated by laser powder bed fusion: Cracking suppression and strengthening mechanism, Mater. Sci. Eng. A (2023) 144591.doi:10.1016/j.msea.2023.144591.

[32] G. Yi, B. Sun, J. D. Poplawsky, Y. Zhu, M. L. Free, Investigation of pre-existing particles in Al 5083 alloys, J. Alloys and Compd. 740 (2018) 461–469. doi:10.1016/j.jallcom.2017.12.329.

[33] B. Miesenberger, E. Kozeschnik, B. Milkereit, P. Warczok, E. Povoden-Karadeniz, Computational analysis of heterogeneous nucleation and precipitation in AA6005 Al-alloy during continuous cooling DSC experiments, Materialia 25 (2022) 101538. doi:10.1016/j.mtla.2022.101538.

[34] J. Morral, G. Purdy, Particle coarsening in binary and multicomponent alloys, Scr. Mater. 30 (7) (1994) 905–908. doi:10.1016/0956-716X(94)90413-8.

[35] Q. Du, M. Chen, J. Xie, Modelling grain growth with the generalized Kampmann-Wagner numerical model, Comput. Mater. Sci. 186 (2021) 110066. doi:10.1016/j.commatsci.2020.110066.

[36] R. S. Varanasi, B. Gault, D. Ponge, Effect of Nb micro-alloying on austenite nucleation and growth in a medium manganese steel during intercritical annealing, Acta Mater. 229 (2022) 117786. doi:10.1016/j.actamat.2022.117786.

[37] E. Nes, N. Ryum, O. Hunderi, On the Zener drag, Acta Metall. 33 (1) (1984) 11–22. doi:10.1016/0001-6160(85)90214-7.

[38] M. Anderson, D. Srolovitz, G. Grest, P. Sahni, Computer simulation of grain growth—I. Kinetics, Acta Metall. 32 (5) (1984) 783–791. doi:10.1016/0001-6160(84)90151-2.

[39] W. E. Frazier, S. Hu, V. V. Joshi, A Potts model parameter study of particle size, Monte Carlo temperature, and “Particle-Assisted Abnormal Grain Growth”, Comput. Mater. Sci. 185 (2020) 109945. doi:10.1016/j.commatsci.2020.109945.

[40] W. E. Frazier, G. S. Rohrer, A. D. Rollett, Abnormal grain growth in the potts model incorporating grain boundary complexion transitions that increase the mobility of individual boundaries, Acta Mater. 96 (2015) 390–398. doi:10.1016/j.actamat.2015.06.033.

[41] L.-Q. Chen, Y. Wang, The continuum field approach to modeling microstructural evolution, JOM 48 (12) (1996) 13–18. doi:10.1007/bf03223259.

[42] E. Miyoshi, T. Takaki, S. Sakane, M. Ohno, Y. Shibuta, T. Aoki, Large-scale phase-field study of anisotropic grain growth: Effects of misorientation-dependent grain boundary energy and mobility, Comput. Mater. Sci. 186 (2021) 109992. doi:10.1016/j.commatsci.2020.109992.

[43] T. Takaki, Y. Tomita, Static recrystallization simulations starting from predicted deformation microstructure by coupling multi-phase-field method and finite element method based on crystal plasticity, Int. J. Mech. Sci. 52 (2) (2010) 320–328. doi:10.1016/j.ijmecsci.2009.09.037.

[44] D. Khatamsaz, B. Vela, P. Singh, D. D. Johnson, D. Allaire, R. Arróyave, Multi-objective materials bayesian optimization with active learning of design constraints: Design of ductile refractory multi-principal-element alloys, Acta Mater. 236 (2022) 118133. doi:10.1016/j.actamat.2022.118133.

[45] S. R. Kalidindi, D. B. Brough, S. Li, A. Cecen, A. L. Blekh, F. Y. P. Congo, C. Campbell, Role of materials data science and informatics in accelerated materials innovation, MRS Bull. 41 (8) (2016) 596–602. doi:10.1557/mrs.2016.164.

[46] P. Fernandez-Zelaia, S. N. Melkote, Process–structure–property relationships in

bimodal machined microstructures using robust structure descriptors, J. Mater. Process Technol. 273 (2019) 116251. doi:10.1016/j.jmatprotec.2019.116251.

[47] A. Tran, J. Sun, J. M. Furlan, K. V. Pagalthivarthi, R. J. Visintainer, Y. Wang, pBO-2GP-3B: A batch parallel known/unknown constrained bayesian optimization with feasibility classification and its applications in computational fluid dynamics, Comput. Methods Appl. Mech. Eng. 347 (2019) 827–852. doi:10.1016/j.cma.2018.12.033.

[48] S. Pfeifer, B. S. S. Pokuri, P. Du, B. Ganapathysubramanian, Process optimization for microstructure-dependent properties in thin film organic electronics, Mater. Discov. 11 (2018) 6–13. doi:10.1016/j.md.2018.06.002.

[49] D. R. Jones, M. Schonlau, W. J. Welch, Efficient global optimization of expensive black-box functions, J. Glob. Optim. 13 (4) (1998) 455–492. doi:10.1023/A:1008306431147.

[50] Y.-L. Zheng, C.-B. Li, S.-D. Liu, Y.-L. Deng, X.-M. Zhang, Effect of homogenization time on quench sensitivity of 7085 aluminum alloy, T. Nonferr. Metal Soc. 24 (7) (2014) 2275–2281. doi:10.1016/S1003-6326(14)63344-6.

[51] A. Olivier, D. G. Giovanis, B. Aakash, M. Chauhan, L. Vandanapu, M. D. Shields, Uqpy: A general purpose python package and development environment for uncertainty quantification, J. Comput. Sci. 47 (2020) 101204. doi:10.1016/j.jocs.2020.101204.

[52] N. D. Alexopoulos, C. Charalampidou, P. Skarvelis, S. K. Kourkoulis, Synergy of corrosion-induced micro-cracking and hydrogen embrittlement on the structural integrity of aluminium alloy (Al-Cu-Mg) 2024, Corros. Sci 121 (2017) 32–42. doi:10.1016/j.corsci.2017.03.001.

[53] X. Tu, A. Shahba, J. Shen, S. Ghosh, Microstructure and property based statistically equivalent RVEs for polycrystalline-polyphase aluminum alloys, Int. J. Plast. 115 (2019) 268–292. doi:10.1016/j.ijplas.2018.12.002.

[54] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, IEEE Trans. Evol. 6 (2) (2002) 182–197. doi:10.1109/4235.996017.

[55] H. Zhu, D. Li, H. Nie, X. Wei, Y. Wei, Multiobjective optimization of a staggered-rotor octocopter design based on a surrogate model, Aerosp. Sci. Technol. 139 (2023) 108387. doi:10.1016/j.ast.2023.108387.

[56] J. Blank, K. Deb, Pymoo: Multi-objective optimization in python, IEEE Access 8 (2020) 89497–89509. doi:10.1109/ACCESS.2020.2990567.

[57] Y. Chang, Z. Bouzarkouna, D. Devegowda, Multi-objective optimization for rapid and robust optimal oilfield development under geological uncertainty, Comput. Geosci 19 (4) (2015) 933–950. doi:10.1007/s10596-015-9507-6.

[58] K. Deb, D. Deb, et al., Analysing mutation schemes for real-parameter genetic algorithms., Int. J. Artif. Intell. Soft Comput. 4 (1) (2014) 1–28. doi:10.1504/IJAISC.2014.059280.

[59] R. Wang, An improved nondominated sorting genetic algorithm for multiobjective problem, Math. Probl. Eng. 103 (2016) 1–7. doi:10.1155/2016/1519542.

[60] L. V. Candioti, M. M. De Zan, M. S. Cámara, H. C. Goicoechea, Experimental design and multiple response optimization. using the desirability function in analytical methods development, Talanta 124 (2014) 123–138. doi:10.1016/j.talanta.2014.01.034.

[61] G. Derringer, R. Suich, Simultaneous optimization of several response variables, J. Qual. Technol. 12 (4) (1980) 214–219. doi:10.1080/00224065.1980.11980968.

[62] M. Abdolshah, A. Shilton, S. Rana, S. Gupta, S. Venkatesh, Expected hypervolume improvement with constraints, in: 2018 24th International Conference on Pattern Recognition (ICPR), IEEE, 2018, pp. 3238–3243. doi:10.1109/ICPR.2018.8545387.

[63] K. Yang, M. Emmerich, A. Deutz, T. Bäck, Multi-objective Bayesian global optimization using expected hypervolume improvement gradient, Swarm Evol. Comput. 44 (2019) 945–956. doi:10.1016/j.swevo.2018.10.007.

[64] M. Balandat, B. Karrer, D. Jiang, S. Daulton, B. Letham, A. G. Wilson, E. Bakshy, Botorch: A framework for efficient Monte-Carlo Bayesian optimization, Adv. Neural Inf. Process 33 (2020) 21524–21538. doi:10.48550/arXiv.1910.06403.

[65] A. Mathern, O. S. Steinholtz, A. Sjöberg, M. Önnheim, K. Ek, R. Rempling, E. Gustavsson, M. Jirstrand, Multi-objective constrained Bayesian optimization for structural design, Struct. Multidiscip. Optim. 63 (2021) 689–701. doi:10.1007/s00158-020-02720-2.

[66] E. R. Homer, V. Tikare, E. A. Holm, Hybrid Potts-phase field model for coupled microstructural–compositional evolution, Comput. Mater. Sci. 69 (2013) 414–423. doi:10.1016/j.commatsci.2012.11.056.

[67] P. C. Nguyen, N. N. Vlassis, B. Bahmani, W. Sun, H. Udaykumar, S. S. Baek, Synthesizing controlled microstructures of porous media using generative adversarial networks and reinforcement learning, Sci. Rep. 12 (1) (2022) 9034. doi:10.1038/s41598-022-12845-7.

[68] X. Hui, H. Wang, W. Li, J. Bai, F. Qin, G. He, Multi-object aerodynamic design optimization using deep reinforcement learning, AIP Adv. 11 (8) (2021). doi:10.1063/5.0058088.

[69] E. Rigdon, M. Sarstedt, O.-I. Moisescu, Quantifying model selection uncertainty via bootstrapping and Akaike weights, Int. J. Consum. Stud. 37 (2023) 1596–1608. doi:10.1111/ijcs.12906.

[70] K. C. Bledsoe, J. Hite, M. A. Jessee, J. P. Lefebvre, Application of Markov chain Monte Carlo methods for uncertainty quantification in inverse transport problems, IEEE Trans. Nucl. Sci. 68 (8) (2021) 2210–2219. doi:10.1109/TNS.2021.3089018.

[71] Q.-B. Cheng, X. Chen, C.-Y. Xu, C. Reinhardt-Imjela, A. Schulte, Improvement and comparison of likelihood functions for model calibration and parameter uncertainty analysis within a Markov chain Monte Carlo scheme, J. Hydrol. 519 (2014) 2202–2214. doi:10.1016/j.jhydrol.2014.10.008.