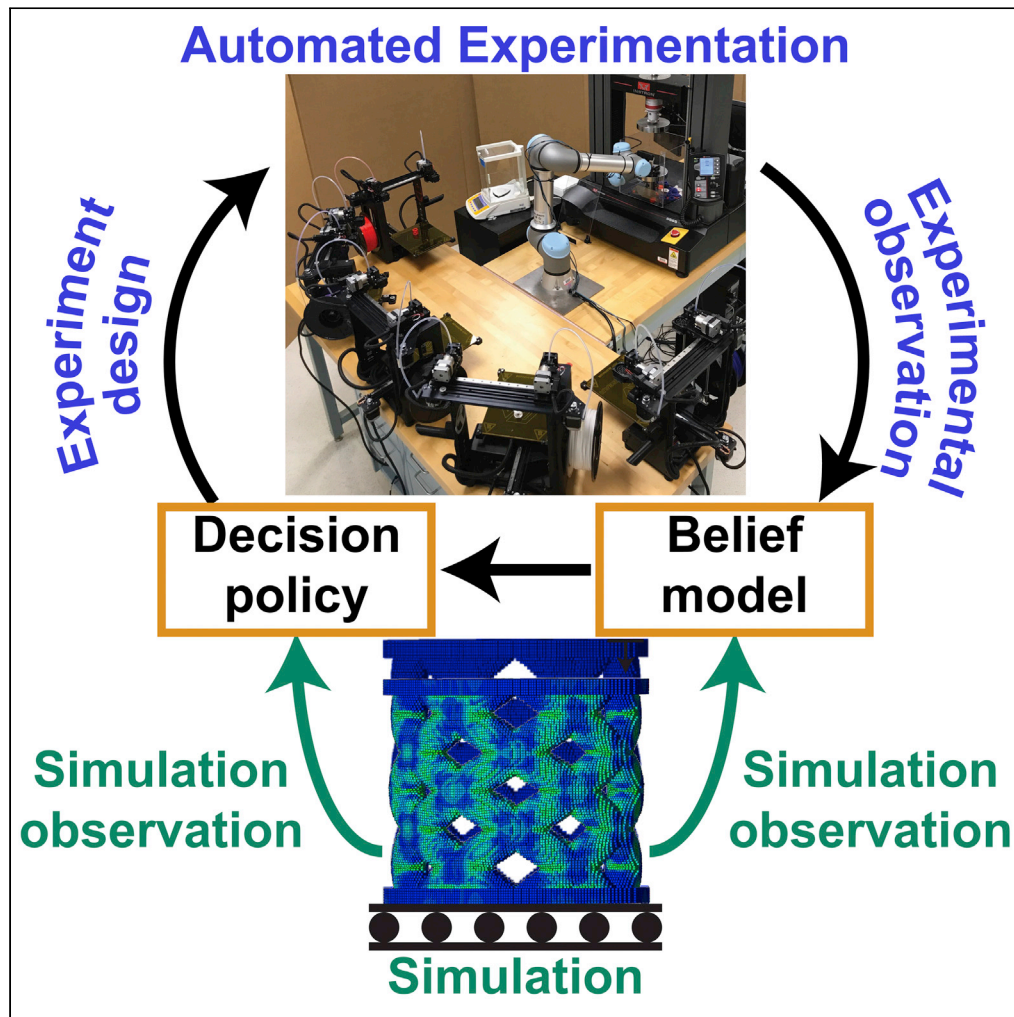CellPress

Article

Using simulation to accelerate autonomous experimentation: A case study using mechanics

Aldair E. Gongora, Kelsey L. Snapp, Emily Whiting, Patrick Riley, Kristofer G. Reyes, Elise F. Morgan, Keith A. Brown

kreyes3@buffalo.edu (K.G.R.)
efmorgan@bu.edu (E.F.M.)
brownka@bu.edu (K.A.B.)

Highlights
Simulation and autonomous experimentation were combined to accelerate research

Resilience, which was accurately simulated, was learned in 10 x fewer experiments

Simulating related properties, i.e. yield force, accelerated learning toughness

Simulation was introduced to an experimental learning loop using transfer learning

Article

Using simulation to accelerate autonomous experimentation: A case study using mechanics

Aldair E. Gongora,[1] Kelsey L. Snapp,[1] Emily Whiting,[2] Patrick Riley,[3] Kristofer G. Reyes,[4,*] Elise F. Morgan,[1,5,6,*] and Keith A. Brown[1,6,7,8,*]

SUMMARY

Autonomous experimentation (AE) accelerates research by combining automation and machine learning to perform experiments intelligently and rapidly in a sequential fashion. While AE systems are most needed to study properties that cannot be predicted analytically or computationally, even imperfect predictions can in principle be useful. Here, we investigate whether imperfect data from simulation can accelerate AE using a case study on the mechanics of additively manufactured structures. Initially, we study resilience, a property that is well-predicted by finite element analysis (FEA), and find that FEA can be used to build a Bayesian prior and experimental data can be integrated using discrepancy modeling to reduce the number of needed experiments ten-fold. Next, we study toughness, a property not well-predicted by FEA and find that FEA can still improve learning by transforming experimental data and guiding experiment selection. These results highlight multiple ways that simulation can improve AE through transfer learning.

INTRODUCTION

Designing materials and structures with optimized properties is a paramount goal of materials science and engineering (Wegst et al., 2015; Yeo et al., 2018). For instance, successes in the study of architected materials have shown that modifying the geometry of lattice-like structures is a powerful method for tuning mechanical properties. Key to the exploration of such intricate structures are advances in high-performance computing and simulation methods, namely finite element analysis (FEA), that have enabled the computation of many facets of mechanical performance (Bar-Sinai et al., 2020; Gao et al., 2003; Kochmann and Bertoldi, 2017). By combining FEA with optimization algorithms, approaches such as topology optimization (Barthelat and Mirkhalaf, 2013; Boddeti et al., 2018; Chen et al., 2018; Jin et al., 2020; Sigmund and Maute, 2013) have led to discovery of intriguing hierarchical structures and composites.

While simulation is powerful, it cannot predict all aspects of mechanical performance, necessitating physical experiments. Mechanics is one of many fields in which experiments can present a bottleneck to progress, a challenge that has motivated the development of autonomous experimentation (AE) systems in numerous fields such as biology (Bryant et al., 2004; King et al., 2009), materials science (MacLeod et al., 2020; Nikolaev et al., 2016; Noack et al., 2019), chemistry (Bédard et al., 2018; Burger et al., 2020; Epps et al., 2020; Porwol et al., 2020), and mechanics (Gongora et al., 2020) to efficiently explore vast and multi-dimensional parameter spaces without human intervention. Ultimately, AE accelerates research by utilizing automation to perform experiments rapidly and using machine learning to select experiments that yield best progress toward the chosen goal. As such, many AE-related advances have involved improved automation (Coley et al., 2019; Li et al., 2018; Nikolaev et al., 2014; Ren et al., 2018; Sun et al., 2019) or algorithms (Wang et al., 2015). Improving collection of experimental data has been critical in this effort, because the premise that simulation is imperfect has led the community to largely proceed in an experimentally data-driven fashion. While this premise is not incorrect, simulation can still in principle provide value for experimental campaigns. A remaining open question, whose answer likely depends upon the relationship between simulation and experiment, is how best to incorporate simulation into AE.

Here, we test the hypothesis that incorporating knowledge from simulation with AE can accelerate the pace of research in the context of mechanics (Figure 1). To explore this concept, we use a robotic system both to 3D-print components and to test them in uniaxial compression (Gongora et al., 2020). When combined with a Bayesian optimization (BO) algorithm to iteratively select experiments that will maximize a performance

[1]Department of Mechanical Engineering, Boston University, Boston, MA 02215, USA

[2]Department of Computer Science, Boston University, Boston, MA 02215, USA

[3]Google Research, Mountain View, CA 94043, USA

[4]Department of Materials Design and Innovation, University at Buffalo, Buffalo, NY 14260, USA

[5]Department of Biomedical Engineering, Boston University, Boston, MA 02215, USA

[6]Division of Materials Science & Engineering, Boston University, Boston, MA 02215, USA

[7]Physics Department, Boston University, Boston, MA 02215, USA

[8]Lead contact

*Correspondence:
kreyes3@buffalo.edu (K.G.R.),
efmorgan@bu.edu (E.F.M.),
brownka@bu.edu (K.A.B.)
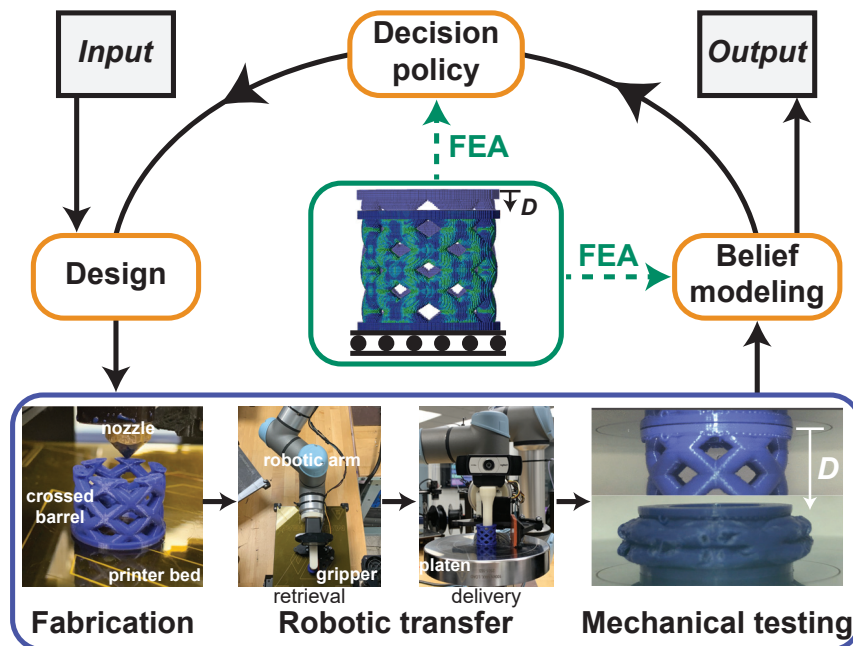https://doi.org/10.1016/j.isci.2021.102262

Figure 1. Incorporating finite element analysis (FEA) into a Bayesian experimental autonomous researcher (BEAR) to study the mechanical behavior of additively manufactured components

metric such as component toughness, this system is termed a ''Bayesian experimental autonomous researcher'' (BEAR). To understand how FEA can improve the operation of the BEAR, we first compare FEA predictions to experimental measurements and determine that while resilience is well predicted by FEA, toughness is not. We then explore the use of the BEAR to optimize resilience by using discrepancy modeling and FEA in the belief model. We evaluate this approach by conducting experimental campaigns and find that, as compared to BO with an uninformative prior, using FEA in this fashion can reduce by a factor of ~10 the number of experiments necessary to find high-performing structures. Finally, we study BEAR campaigns to optimize toughness using a custom method in which the belief model is built on FEA-transformed data and the decision policy is FEA-informed. We find that these FEA-informed experimental learning campaigns resulted in ~15% higher performing structures compared to campaigns using a traditional BO approach. For the evolving field of AE and data-driven research more broadly, this work shows the potential for capitalizing on additional information sources such as simulation to accelerate the pace of research and enable the exploration of more complex parameter spaces.

RESULTS AND DISCUSSION

Comparison of simulated and measured mechanical properties of parametric structures

In order to determine the degree to which FEA could predict the mechanical behavior of additively manufactured structures, we evaluated a ''crossed barrel'' family of components that leverage our previously reported dataset (Gongora et al., 2020). In the previously reported data set, the ''crossed barrel'' family of components is parametrized by n hollow columns with outer radius r and thickness t that are twisted with an angle θ. Thus, the four-dimensional parameter space considered in this work is defined by $x =$ (n, θ, r, t). In a typical experiment, a component was printed out of polylactic acid (PLA) filament and tested in uniaxial, quasi-static compression (Figure 2A). This experiment allowed direct computation of two important metrics of energy absorbed by a component during compression, namely toughness U and resilience U_E. Resilience is defined as the energy stored during the elastic portion of the compressive curve while toughness is defined as the energy absorbed during the entirety of the compression (elastic and failure). Optimizing the former is important for realizing structural components that accommodate a variety of working conditions without damage while optimizing the latter is critical for realizing structures that are safe during catastrophic events. From the design of the ''crossed barrel'' structure used for 3D printing, a mesh generated from hexahedral elements (Figures S1 and S2) can be used in FEA to predict resilience $\tilde{U}_E$ by simulating a uniaxial quasi-static compression test (Figure 2B).
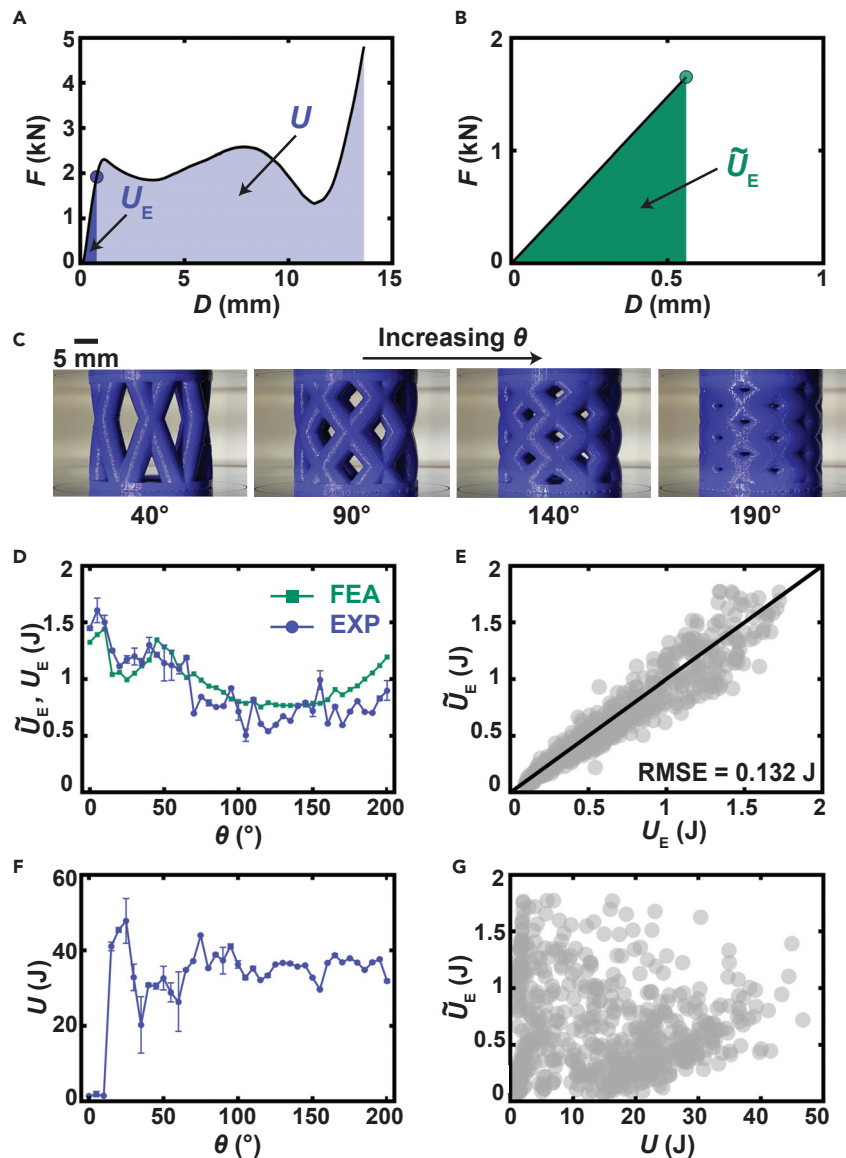
Figure 2. Exploration of resilience and toughness of parametric structures

(A) Force F vs. displacement D for a crossed barrel structure measured using quasi-static compression. Resilience U_E and toughness U are computed as areas under the F-D curve as shown.

(B) F vs. D. computed using finite element analysis (FEA) along with simulated resilience $\tilde{U}_E$ found as the area under the F-D curve.

(C) Photographs of a series of crossed barrel structures that vary in their twist angle θ.

(D) U_E and $\tilde{U}_E$ vs. θ for a series of crossed barrels including those depicted in (C). While this plot highlights the dependence of resilience on θ as a single parameter, it is worth emphasizing that we are exploring a four-dimensional parameter space. Points denote mean with error bars denoting standard deviation.

(E) $\tilde{U}_E$ vs. U_E for 600 crossed barrels selected in a grid across the entire four-dimensional parameter space. Agreement between the quantities is evident based upon the root mean square error (RMSE).

(F) U vs. θ for a series of crossed barrels including those depicted in (C). Points denote mean with error bars denoting standard deviation.

(G) $\tilde{U}_E$ vs. U for 600 crossed barrels selected in a grid across the entire four-dimensional parameter space. These properties are found to be uncorrelated.

To assess the prediction capabilities of FEA, experimental measurements and FEA predictions were obtained for a series of components that varied based upon a single parameter, namely the twist angle θ (Figure 2C). The root mean squared error (RMSE) between U_E and $\tilde{U}_E$ for the series of components was 0.1801 J (Figure 2D). Perhaps more importantly, the FEA predictions reasonably captured the general trends present in the experimental measurements. Not only did U_E and $\tilde{U}_E$ exhibit similar dependencies on θ but they also showed excellent agreement with each other for 600 distinct designs spread across the four-dimensional parameter space, exhibiting RMSE = 0.132 J and a mean squared percentage error of 3.12% (Figure 2E).

Toughness could not be computed due to the need for painstaking management of self-contacts and development of advanced material models (Li et al., 2017; Qiao et al., 2008; Zhu et al., 2010) which limit the throughput of computation. While one could expect that FEA predictions of resilience could provide some value as resilience contributes to toughness by definition (*i.e.* Figure 2A), examining U for this class of structures shows that U and $\tilde{U}_E$ are not similar in magnitude (Figure 2F). This difference is rooted in mechanics with the energy absorbed during plastic deformation not being directly correlated with that stored during elastic deformation. However, there are some structures for which $U \approx U_E$, namely those that exceed the force threshold and therefore never enter the plastic regime during the working range (*i.e.* those for which $\theta \leq 15°$). Nevertheless, the disagreement between these quantities manifests across the entire parameter space with $\tilde{U}_E$ not being correlated with U ($R^2 = 0.0133$) (Figure 2G). It should be noted that the lack of correlation between U_E and U is likely a general property of lattice-based structures as these feature a vast diversity of behaviors that only begin once the structure enters the plastic regime.

Optimizing resilience with an FEA-informed BEAR

In this work, we employed BO due to its popularity and previously reported success as an active learning strategy for optimization. Additionally, active learning approaches, such as BO, have been previously reported to outperform one factor at a time (OFAT) and design of experiment (DoEs) approaches (Braham et al., 2020). In particular, OFAT can be slow and inefficient in high-dimensional parameter spaces since a single variable is varied at a time. Further, OFAT is highly sensitive to the initial selection of variables and does not rapidly capture potential correlations between input variables. While DoE approaches address some of the aforementioned shortcomings of OFAT, they depend on an initial round of experiments being conducted before analysis can proceed (Bowden et al., 2019). Additionally, DoE approaches are unable to efficiently capture highly non-linear parameter spaces and when applied iteratively tend to be generally exploitative (Cao et al., 2018). Active learning approaches, such as BO, improve upon OFAT and DoE by using all available experiments to build belief models that can capture complex and highly non-linear parameter spaces. Moreover, active learning approaches enable the iterative selection of subsequent experiments using decision policies that can determine and exploit correlations between input variables, and they can be customized to favor exploration, exploitation, or balance these two goals (Lookman et al., 2019; Rohr et al., 2020).

Given that FEA can reasonably predict resilience, learning or optimizing experimental resilience can be considered a classic transfer learning process (Pan and Yang, 2010). Specifically, both $U_E(x)$ and $\tilde{U}_E(x)$ are defined over the same parameter space – namely the four-dimensional space corresponding to the parametric family of structures given by $x = (n, \theta, r, t)$. To define $\tilde{U}_E(x)$ over the parameter space, we built a surrogate model from FEA predictions selected on a grid (Figures S3 and S4A). Since $\tilde{U}_E(x)$ represents an approximation for $U_E(x)$, the optimization problem can be addressed using relational knowledge transfer, an approach in inductive transfer learning. Specifically, we define a discrepancy model $\delta_E(x) = U_E(x) - \tilde{U}_E(x)$ to represent the difference between FEA and experiment. This approach differs from BO with an uninformative prior in that, rather than selecting experiments based upon a belief model of U_E that is a Gaussian process regression (GPR) against experimental measurements of U_E, the discrepancy model approach uses a GPR trained on $\delta_E = U_E - \tilde{U}_E$. In other words, the FEA-informed approach makes $\tilde{U}_E$ the Bayesian prior for the belief model.

In order to evaluate the incorporation of FEA into a BEAR through discrepancy modeling (Figure 3A), we performed a series of simulated learning campaigns based upon an expected improvement (*EI*) decision policy using either the uninformative prior or the FEA-informed approach. The *EI* decision policy was selected due to its widespread application in the BO community as an improvement-based decision policy that seeks to select subsequent experiments based on the likelihood of exceeding previous observations. Since the simulated learning campaigns are *in silico*, we may evaluate the performance P of these

A

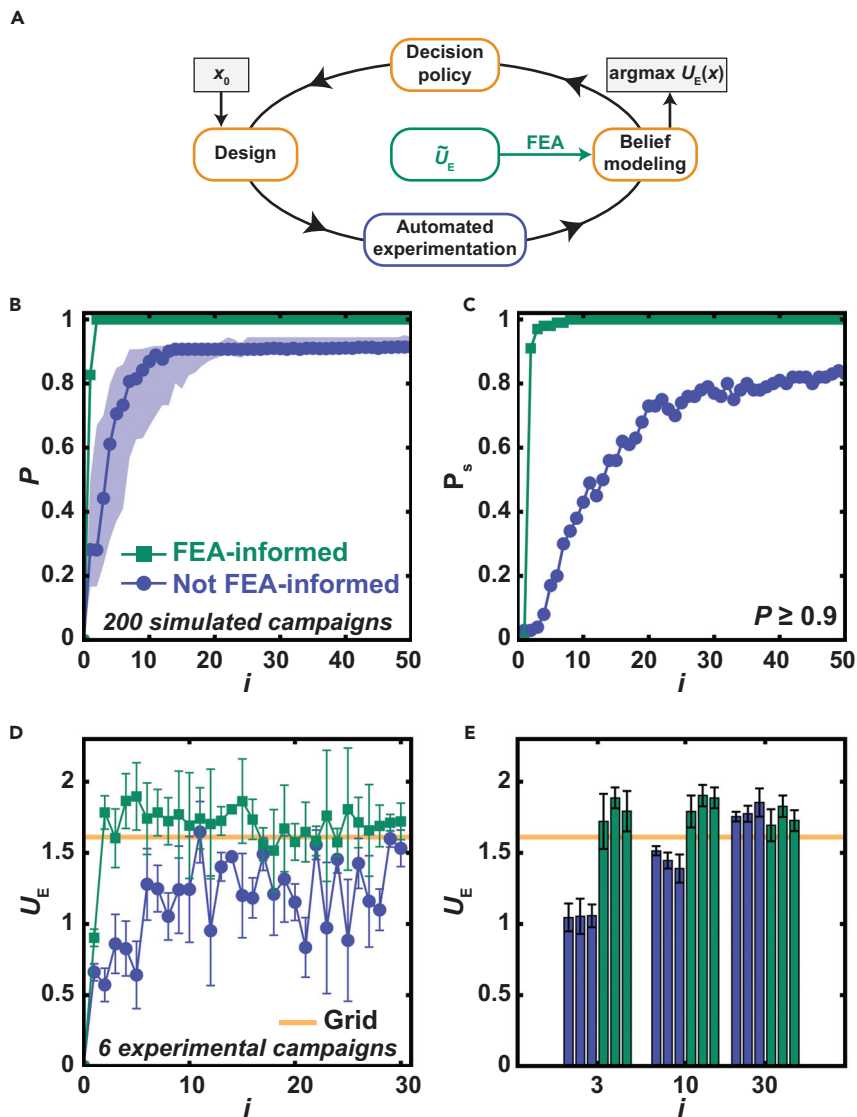

B

C

D

E

Figure 3. Combining AE and simulation to optimize resilience

(A) Scheme showing how FEA was incorporated into the BEAR to find the design x with the maximum U_E. The process begins with a random design x_0.

(B) Simulated performance P at experiment number i for a simulated learning campaign.

(C) Probability of achieving $P \geq 0.90$ at a given i.

(D) Sequence of U_E measured during six experimental learning campaigns. Points denote mean with error bars denoting standard deviation. The horizontal line indicates the maximum value measuring on a grid of 600 points.

(E) U_E measured at the optimum predicted for each of the six campaigns after i experiments. Each value represents the average of five identically prepared samples with the error bars denoting their standard deviation.

approaches using $P = 1$ to indicate that the campaign found the optimum design, or the x that corresponded to the largest U_E in the data set. Campaigns using the FEA-informed approach achieved a median performance $Md(P) = 1$ in only two experiments, whereas campaigns using an uninformative prior only achieved $Md(P) \geq 0.90$ after 13 experiments (Figure 3B). To further assess these approaches, we computed the probability P_s of achieving $P \geq 0.90$ where $P_s = 1$ would indicate that all simulated campaigns had achieved $P \geq 0.90$. Campaigns using an FEA-informed approach achieved $P_s \geq 0.99$ after six experiments, while those using an uninformative prior only achieved $P_s = 0.80$ after 32 experiments (Figure 3C). These results clearly predict that a good prior from simulation will substantially accelerate an experimental campaign.

Table 1. Mean and standard deviation of resilience after 3 experiments for both the uninformative-prior approach (not FEA-informed) and the FEA-informed approach

Approach	Mean U_E (J)	SD (J)
Not FEA-informed	1.05	0.10
	1.05	0.12
	1.06	0.08
FEA-informed	1.72	0.19
	1.89	0.07
	1.79	0.14

To estimate how sensitive the improvement was to the accuracy of the FEA predictions, we first repeated the FEA computations while varying E and σ_y to determine the degree to which these input material properties influenced $\tilde{U}_E$. Interestingly, $\tilde{U}_E$ is more than twice as sensitive to errors in σ_y than errors in E, with overestimations of σ_y increasing $\tilde{U}_E$ and overestimations of E decreasing $\tilde{U}_E$ (Figures S5A and S5B). To determine the degree to which such errors would affect a learning campaign, we performed a series of simulated learning campaigns in which the FEA results were multiplied by a factor $\varphi_{FE} = 0.50$, 0.90, 0.95, 1.05, 1.1, and 1.5. Variations in the FEA results that were ~10% (which correspond to a ~7% error in σ_y or a ~28% error in E), had a minimal effect on convergence (Figures S5C and S5D). Larger errors, however, had a more substantial impact on convergence with the interesting result that overestimations in FEA results were more damaging to learning than underestimations (Figures S5E and S5F). We hypothesize that this asymmetry is due to the goal of the algorithm being to maximize resilience.

Guided by these simulated learning campaigns, we performed six independent experimental learning campaigns to further assess the incorporation of FEA into a BEAR through discrepancy modeling (Table 1). Three campaigns used an FEA-informed approach, and three used an uninformative prior. Each campaign was given an experimental budget of 30 experiments. By comparing the average U_E for each approach as a function of experiment number, the FEA-informed approach on average outperformed the uninformative-prior approach (Figure 3D). Additionally, the average resilience using the FEA-informed approach was mostly larger, after one experiment, than the mean of the predicted maximum resilience of 1.61 ±0.17 J from a GPR surrogate model trained on resilience measurements at 600 distinct design locations specified in a grid search ("Grid").

It is key to note that BO campaigns with *EI* attempt to balance both exploration and exploitation, and thus not all subsequent experiments will yield an increase in the experimental response. While the sequence of experimental responses can provide some insight into performance, it is imperative to assess performance by evaluating the experimental response of the predicted optimum design of the belief model after *i* experiments. To do this, we carried out five repeated experiments at the predicted optimum designs of each of the experimental campaigns after three, 10, and 30 experiments (Figure 3E). After three experiments, the FEA-informed approach outperformed the uninformative-prior approach by ~71% and the Grid result by ~12%. After 10 experiments, the performance of the uninformative prior approach increased but was still ~28% less than the performance of the FEA-informed approach. After 30 experiments, the performance of the uninformative prior approach and the FEA-informed approach were not statistically different ($p = 0.39$), while both approaches outperformed Grid by ~10%. Additionally, the performance of the FEA-informed approach after 30 experiments was not statistically different than the performance of the FEA-informed approach after 10 experiments based on a multiple comparison analysis comparing each campaign ($p > 0.05$). Ultimately, the FEA-informed approach reduced the number of experiments necessary to find a high-performing design by 10-fold relative to the uninformative prior approach and by 600-fold relative to Grid. Notably, campaigns based upon the uninformative prior approach found better designs in 30 experiments than resulted from the 1800 experiments used as part of the Grid. The observed 60-fold reduction in number of experiments recapitulates our previously reported acceleration of BO relative to grid searching when optimizing toughness (Gongora et al., 2020).

Optimizing toughness with an FEA-informed BEAR

While $\tilde{U}_E$ was clearly useful in optimizing U_E, its lack of correlation with U makes its utility in a toughness optimization framework substantially less clear. Indeed, we performed simulated campaigns exploring
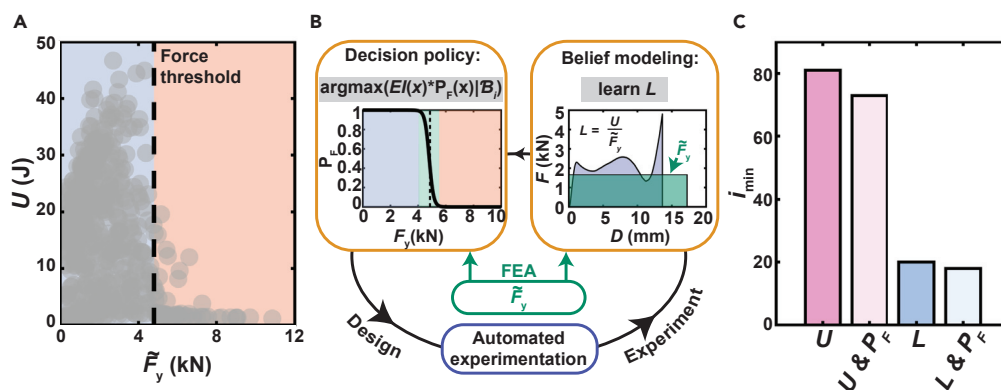
Figure 4. Incorporation of a belief model built on FEA-transformed data and an FEA-informed decision policy into a BEAR

(A) U vs. FEA-predicted yield force $\tilde{F}_y$ for 600 distinct designs selected on a grid. The dashed line corresponds to the experimental force threshold.

(B) Schematic showing how FEA was incorporated into the BEAR. The decision policy was modified using a logistic function P_F. The experimental data was transformed into an effective length L using $\tilde{F}_y$.

(C) Estimate of the minimum number of experiments i_{min} needed to explore parameter space using strategies based on U, U and P_F, L, and L and P_F.

the simplest extension of the resilience studies by building a Bayesian prior with $\tilde{U}_E$ and using discrepancy modeling to find the inelastic component of toughness (*i.e.* $U - U_E$). Unfortunately, simulation suggested that this type of knowledge transfer would not provide any acceleration, indicating that more creative approaches to using FEA are needed. One such approach is motivated by how tough structures are used in practice. Specifically, tough components are often accompanied by a force threshold to allow them to absorb energy before transmitting dangerous reactive forces to other elements in a system. With this force criterion in mind, the FEA-predicted yield force $\tilde{F}_y$ (Figure S4B) becomes a very useful factor as it can differentiate between structures whose design allows them to plastically deform during failure and those that are too strong for the imposed force threshold. For our data set, while a low $\tilde{F}_y$ did not guarantee high performance, structures with high $\tilde{F}_y$ were all low-performing (Figure 4A). To leverage this realization, we constructed a logistic function $P_F(x)$ trained using $\tilde{F}_y$ that biases the system away from regions of parameter space that are too strong, thus shrinking parameter space by ∼9% (Figure 4B). Specifically, P_F was built using $\tilde{F}_y$ computed for 1188 designs selected on a grid, where $P_F = 1$ indicated that $\tilde{F}_y$ would not exceed the force threshold and $P_F = 0$ indicated that it would. The transition between 0 and 1 was given a width of 15%, a number chosen to match the median coefficient of variation of U. Subsequently, P_F was incorporated into the decision policy by selecting the next experiment by finding $argmax(EI(x) \cdot P_F(x))$ (Gelbart et al., 2014). This approach effectively filters the parameter space by removing designs that are predicted to be low-performing.

In parallel, we explored a way to reduce the *effective* size of parameter space by using simulation data to transform experimental data into a format that featured greater correlations in parameter space. Specifically, we defined an effective length $L(x) = U(x)/\tilde{F}_y(x)$ that represents how much compression would be required at constant $\tilde{F}_y$ to produce the same toughness as the experimentally determined value U. A belief model of L was built using a GPR, and then this model was combined with $\tilde{F}_y$ to form the input to the decision policy (Figure 4B).

In order to estimate whether these approaches—one that uses an FEA-informed decision policy and one that uses a belief model build on FEA-transformed experimental data—would reduce the number of experiments needed to explore the parameter space, we hypothesized that examining the correlation lengths of GPRs trained on these datasets would provide insight. In particular, longer correlation lengths would indicate that each data point is providing information relevant to larger regions in parameter space. To explore this systematically, we divided the total volume of parameter space by the product of these correlation lengths to approximate how many experiments would be needed to explore space (Figure 4C). As expected, shrinking the parameter space by ∼9% based on P_F commensurately reduced the number of needed experiments. Strikingly, learning L rather than U was predicted to produce a four-fold reduction
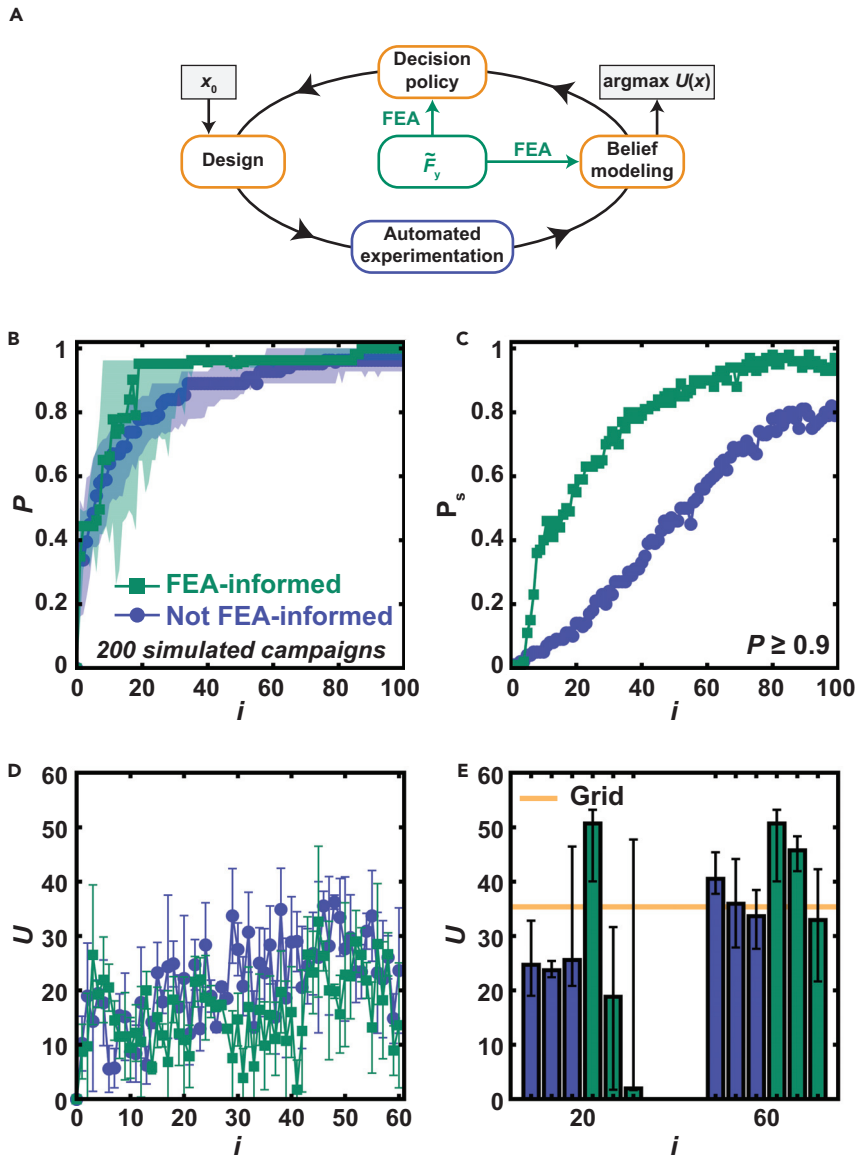
Figure 5. Optimizing toughness of parametric structures

(A) Strategy for incorporating FEA into the BEAR through a belief model built on FEA-transformed data and an FEA-informed decision policy.

(B) Simulated performance P at experiment number i for a simulated learning campaign to estimate toughness.

(C) Probability P_s of achieving $P \geq 0.90$ at a given i.

(D) Sequence of U measured during six experimental learning campaigns. Points denote mean with error bars denoting standard deviation.

(E) U measured at the optimum predicted for each of the six campaigns after i experiments. Each value represents the median of five identically prepared samples with the error bars denoting their total range.

in the number of needed experiments. We hypothesize that this reduction could be explained by the physical processes at play. Specifically, since toughness is both a product of the strength of the component and its ductility, factoring out a representation of its strength could remove one source of variability, allowing experiment to focus on learning one quantity – ductility – more directly.

We next evaluated the incorporation of these two approaches, termed together as an ''FEA-informed approach,'' into a BEAR (Figure 5A). We performed a series of simulated learning campaigns using the uninformative-prior approach and the FEA-informed approach to optimize U defined over the

four-dimensional parameter space $x = (n, \theta, r, t)$. Campaigns using the FEA-informed approach achieved $Md(P) \geq 0.95$ after 19 experiments, outperforming campaigns using an uninformative prior which achieved a $Md(P) \geq 0.95$ after 66 experiments (Figure 5B). Interestingly, the interquartile range of the FEA-informed campaigns was notably reduced after 36 experiments. Additionally, the FEA-informed approach achieved $Md(P) = 1$ after 87 experiments while the uninformed prior approach plateaued at $Md(P) = 0.96$ after 80 experiments. The campaigns based upon the FEA-informed approach reached $P_s \geq 0.85$ after 50 experiments and an average $P_s > 0.90$ after 60 experiments, outperforming campaigns using an uninformative prior which only achieved $P_s \geq 0.58$ after 60 experiments and $P_s = 0.79$ at 100 experiments (Figure 5C). From these simulated campaigns, we concluded that the FEA-informed approach should substantially accelerate optimization of toughness relative to the uninformative-prior approach.

It is worth noting that the experimental burden of the FEA methodology used herein is extremely low. The only input from experiment that is needed is a characterization of the materials. While we performed materials characterization by printing and evaluating cylindrical test coupons, for many relevant materials, high quality tabulated values exist. Further, the general method we developed to convert standard triangle language files appropriate for additive manufacturing to hexahedral meshes for FEA can be used for any general structure beyond the crossed barrel family. Thus, FEA for the simulation of resilience has a low barrier to entry and, with the continued growth of high performance computing resources, is likely to be an increasingly efficient path to gaining mechanical insight.

Based on the results of the simulated campaigns, six independent experimental campaigns were conducted to optimize toughness, each with an experimental budget of 60 experiments (Table 2). Three experimental campaigns used an uninformative prior and the remaining three used the FEA-informed approach. In contrast to what was observed for resilience (Figure 3D), the progression of the experimental response during each campaign was not a clear indicator of progress (Figure 5D), further emphasizing that a campaign's performance must be evaluated by assessing the predicted optimum design of the campaign. To directly evaluate these predicted optima, we carried out five repeated experiments on each predicted optimum after 20 and 60 experiments (Figure 5E). Here, the median and the range are plotted due to the large differences in toughness that arose for designs near the boundary of the imposed force threshold. While the acceleration observed for a particular campaign does in part depend on the location in parameter space of the first randomly selected experiment and the set of experimental responses observed for that particular campaign, the reduction in experiments observed in the experimental campaigns is comparable to reduction in experiments suggested by the simulated campaigns. In both simulated and experimental campaigns, the FEA-informed approach is superior to the uninformative-prior approach suggesting that, more generally, an optimization campaign benefits from the incorporation of FEA in the learning structure.

To compare the performance of the approaches and accounting for the large range in performance for a given design, we computed the probability that a component designed by the FEA-informed approach would be tougher than a component designed by the uninformative prior approach (Figures S6A and S6B). Based on this metric, after 20 experiments, campaigns based on the FEA-informed approach had a 54% chance of producing tougher components that the uninformative-prior approach, which shows the two approaches are effectively equal at this stage. However, after 60 experiments, campaigns based on the FEA-informed approach had a 73% chance of producing tougher components than those produced by campaigns using an uninformative prior. Notably, after 20 experiments, the only campaign that identified a tougher component than the 35.4 $\pm$ 1.5 J mean experimental response of the optimum found using the 600 measurements in a grid (Gongora et al., 2020) was a campaign based on the FEA-informed approach. Interestingly, after 60 experiments, four experimental campaigns, two with uninformative-prior approach and two with the FEA-informed approach, outperformed Grid, with the top performer being the FEA-informed approach outperforming Grid by ~15%. Notably, the top performer discovered by this FEA-informed campaign was tougher than any we had previously identified in any experimental campaign. Ultimately, the FEA-informed approach outperformed the uninformative-prior approach by increasing the probability of finding a high-performing design after 60 experiments and reduced the number of experiments necessary to find a high-performing structure by 30-fold relative to Grid while increasing the performance by ~15%.

Table 2. Median and range of toughness after 60 experiments for both the uninformative-prior approach (not FEA-informed) and the FEA-informed approach

Approach	Median U (J)	Range (J)
Not FEA-informed	40.53	7.66
	35.91	16.25
	33.64	10.81
FEA-informed	50.70	13.18
	45.78	6.37
	32.92	20.60

Conclusion

In this work, we used a case study in mechanics to explore several ways in which simulation data can be inserted into AE and evaluated the degree to which each accelerates research. Two mechanical properties were used: one, resilience, that can be robustly simulated; and one, toughness, that cannot. We found that when a good simulator exists for the property of interest, AE campaigns can be significantly accelerated using simulation knowledge as a Bayesian prior. This was demonstrated for the case of resilience, where a ∼10-fold reduction in the number of experiments was observed when FEA was incorporated in the belief model via discrepancy modeling versus a traditional BO approach. For toughness, in contrast, we developed a custom method for incorporating simulation in a BEAR by using FEA data to transform the space where belief modeling occurs and by using simulation to guide the decision policy. This custom method resulted in a ∼73% chance of outperforming a traditional BO approach with a ∼15% increase in component toughness. While the custom method developed in this work utilized $\tilde{F}_y$ in the AE system, FEA generally presents a valuable addition to the active learning component of an AE system, which is often viewed as purely data-driven. By capitalizing on the ability of FEA to predict certain properties of a system such as stresses and strains under varying loading or boundary conditions, myriad mechanical insights can be extracted and incorporated in the decision policy or the belief model to further accelerate the research process. Collectively, the principles described herein show how knowledge transfer from a simulator to an AE system may increase the pace of research not only in mechanics but also in other domains such as the physical sciences where simulation is ubiquitous but imperfect.

Limitations of the study

The incorporation of simulation into AE in this work was explored using a BEAR in the context of mechanics where FEA was used to predict mechanical properties. To further explore the principles of knowledge transfer from simulation to AE, future studies need to focus on exploring the applicability and utility of these methods in other domains such as chemistry, biology, and materials science where simulation can be accessible.

Resource availability

Lead contact

Further information and requests should be directed to and will be fulfilled by the lead contact, Keith A. Brown (brownka@bu.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

All data needed to evaluate the conclusions in the paper are present in the paper and/or the supplemental information. The raw data can be accessed through www.kablab.org/data. Additional data related to this paper may be requested from the authors.

METHODS

All methods can be found in the accompanying transparent methods supplemental file.

SUPPLEMENTAL INFORMATION

ACKNOWLEDGMENTS

AUTHOR CONTRIBUTIONS

Conceptualization, all authors; Methodology, A.E.G., E.F.M., K.G.R., and K.A.B.; Software development, A.E.G.; Investigation, A.E.G. and K.L.S.; Writing, review, and editing, all authors.

DECLARATION OF INTERESTS

P.R. is an employee of Google, a technology company that sells machine learning services as part of its business. All other authors declare that they have no competing interests.

INCLUSION AND DIVERSITY

One or more of the authors of this paper self-identifies as an underrepresented ethnic minority in science. One or more of the authors of this paper self-identifies as a member of the LGBTQ+ community. One or more of the authors of this paper received support from a program designed to increase minority representation in science.

REFERENCES

Bar-Sinai, Y., Librandi, G., Bertoldi, K., and Moshe, M. (2020). Geometric charges and nonlinear elasticity of two-dimensional elastic metamaterials. Proc. Natl. Acad. Sci. U S A 117, 10195–10202.

Barthelat, F., and Mirkhalaf, M. (2013). The quest for stiff, strong and tough hybrid materials: an exhaustive exploration. J. R. Soc. Interface 10, 20130711.

Bédard, A., Adamo, A., Aroh, K.C., Russell, M.G., Bedermann, A.A., Torosian, J., Yue, B., Jensen, K.F., and Jamison, T.F. (2018). Reconfigurable system for automated optimization of diverse chemical reactions. Science 361, 1220–1225.

Boddeti, N., Ding, Z., Kaijima, S., Maute, K., and Dunn, M.L. (2018). Simultaneous digital design and additive manufacture of structures and materials. Sci. Rep. 8, 1–10.

Bowden, G.D., Pichler, B.J., and Maurer, A. (2019). A design of experiments (DoE) approach accelerates the optimization of copper-mediated 18F-fluorination reactions of arylstannanes. Sci. Rep. 9, 1–10.

Braham, E.J., Davidson, R.D., Al-Hashimi, M., Arróyave, R., and Banerjee, S. (2020). Navigating the design space of inorganic materials synthesis using statistical methods and machine learning. Dalt. Trans. 49, 11480–11488.

Bryant, C.H., Jones, F.M., Kell, D.B., King, R.D., Muggleton, S.H., Oliver, S.G., Reiser, P.G.K., and Whelan, K.E. (2004). Functional genomic hypothesis generation and experimentation by a robot scientist. Nature 427, 247–252.

Burger, B., Maffettone, P.M., Gusev, V.V., Aitchison, C.M., Bai, Y., Wang, X., Li, X., Alston, B.M., Li, B., Clowes, R., et al. (2020). A mobile robotic chemist. Nature 583, 237–241.

Cao, B., Adutwum, L.A., Oliynyk, A.O., Luber, E.J., Olsen, B.C., Mar, A., and Buriak, J.M. (2018). How to optimize materials and devices via design of experiments and machine learning: demonstration using organic photovoltaics. ACS Nano 12, 7434–7444.

Chen, D., Skouras, M., Zhu, B., and Matusik, W. (2018). Computational discovery of extremal microstructure families. Sci. Adv. 4, eaao7005.

Coley, C.W., Thomas, D.A., Lummiss, J.A.M., Jaworski, J.N., Breen, C.P., Schultz, V., Hart, T., Fishman, J.S., Rogers, L., Gao, H., et al. (2019). A robotic platform for flow synthesis of organic compounds informed by AI planning. Science 365, eaax1566.

Epps, R.W., Bowen, M.S., Volk, A.A., Abdel-Latif, K., Han, S., Reyes, K.G., Amassian, A., and Abolhasani, M. (2020). Artificial chemist: an autonomous quantum dot synthesis bot. Adv. Mater. 32, e2001626.

Gao, H., Ji, B., Jager, I.L., Arzt, E., and Fratzl, P. (2003). Materials become insensitive to flaws at nanoscale: lessons from nature. Proc. Natl. Acad. Sci. U S A 100, 5597–5600.

Gelbart, M.A., Snoek, J., and Adams, R.P. (2014). Bayesian optimization with unknown constraints. Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence, UAI'14 (AUAI Press), pp. 250–259.

Gongora, A.E., Xu, B., Perry, W., Okoye, C., Riley, P., Reyes, K.G., Morgan, E.F., and Brown, K.A. (2020). A Bayesian experimental autonomous researcher for mechanical design. Sci. Adv. 6, eaaz1708.

Jin, Z., Zhang, Z., Demir, K., and Gu, G.X. (2020). Machine learning for advanced additive manufacturing. Matter 3, 1541–1556.

King, R.D., Oliver, S.G., Rowland, J., Soldatova, L.N., Whelan, K.E., Young, M., Aubrey, W., Byrne, E., Liakata, M., Markham, M., et al. (2009). The automation of science. Science 324, 85–89.

Kochmann, D.M., and Bertoldi, K. (2017). Exploiting microstructural instabilities in solids and structures: from metamaterials to structural transitions. Appl. Mech. Rev. 69, 050801.

Li, J., Lu, Y., Xu, Y., Liu, C., Tu, Y., Ye, S., Liu, H., Xie, Y., Qian, H., and Zhu, X. (2018). AIR-chem: authentic intelligent robotics for chemistry. J. Phys. Chem. A 122, 9142–9148.

Li, L., Zhang, G., and Khandelwal, K. (2017). Topology optimization of energy absorbing structures with maximum damage constraint. Int. J. Numer. Methods Eng. *112*, 737–775.

Lookman, T., Balachandran, P.V., Xue, D., and Yuan, R. (2019). Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. npj Comput. Mater. *5*, 21.

MacLeod, B.P., Parlane, F.G.L., Morrissey, T.D., Häse, F., Roch, L.M., Dettelbach, K.E., Moreira, R., Yunker, L.P.E., Rooney, M.B., Deeth, J.R., et al. (2020). Self-driving laboratory for accelerated discovery of thin-film materials. Sci. Adv. *6*, eaaz8867.

Nikolaev, P., Hooper, D., Perea-López, N., Terrones, M., and Maruyama, B. (2014). Discovery of wall-selective carbon nanotube growth conditions via automated experimentation. ACS Nano *8*, 10214–10222.

Nikolaev, P., Hooper, D., Webber, F., Rao, R., Decker, K., Krein, M., Poleski, J., Barto, R., and Maruyama, B. (2016). Autonomy in materials research: a case study in carbon nanotube growth. npj Comput. Mater. *2*, 16031.

Noack, M.M., Yager, K.G., Fukuto, M., Doerk, G.S., Li, R., and Sethian, J.A. (2019). A kriging-based approach to autonomous experimentation with applications to X-ray scattering. Sci. Rep. *9*, 1–19.

Pan, S.J., and Yang, Q. (2010). A survey on transfer learning. IEEE Trans. Knowl. Data Eng. *22*, 1345–1359.

Porwol, L., Kowalski, D.J., Henson, A., Long, D.L., Bell, N.L., and Cronin, L. (2020). An autonomous chemical robot discovers the rules of inorganic coordination chemistry without prior knowledge. Angew. Chem. Int. Ed. *59*, 11256–11261.

Qiao, P., Yang, M., and Bobaru, F. (2008). Impact mechanics and high-energy absorbing materials: review. J. Aerosp. Eng. *21*, 235–248.

Ren, F., Ward, L., Williams, T., Laws, K.J., Wolverton, C., Hattrick-Simpers, J., and Mehta, A. (2018). Accelerated discovery of metallic glasses through iteration of machine learning and high-throughput experiments. Sci. Adv. *4*, eaaq1566.

Rohr, B., Stein, H.S., Guevarra, D., Wang, Y., Haber, J.A., Aykol, M., Suram, S.K., and Gregoire, J.M. (2020). Benchmarking the acceleration of materials discovery by sequential learning. Chem. Sci. *11*, 2696–2706.

Sigmund, O., and Maute, K. (2013). Topology optimization approaches: a comparative review. Struct. Multidiscip. Optim. *48*, 1031–1055.

Sun, S., Hartono, N.T.P., Ren, Z.D., Oviedo, F., Buscemi, A.M., Layurova, M., Chen, D.X., Ogunfunmi, T., Thapa, J., Ramasamy, S., et al. (2019). Accelerated development of perovskite-inspired materials via high-throughput synthesis and machine-learning diagnosis. Joule *3*, 1437–1451.

Wang, Y., Reyes, K.G., Brown, K.A., Mirkin, C.A., and Powell, W.B. (2015). Nested-batch-mode learning and stochastic optimization with an application to sequential multistage testing in materials science. SIAM J. Sci. Comput. *37*, 361–381.

Wegst, U.G.K., Bai, H., Saiz, E., Tomsia, A.P., and Ritchie, R.O. (2015). Bioinspired structural materials. Nat. Mater. *14*, 23–36.

Yeo, J., Jung, G.S., Martín-Martínez, F.J., Ling, S., Gu, G.X., Qin, Z., and Buehler, M.J. (2018). Materials-by-design: computation, synthesis, and characterization from atoms to structures. Phys. Scr. *93*, 053003.

Zhu, F., Lu, G., Ruan, D., and Wang, Z. (2010). Plastic deformation, failure and energy absorption of sandwich structures with metallic cellular cores. Int. J. Prot. Struct. *1*, 507–541.

Supplemental information

Using simulation to accelerate autonomous experimentation:

A case study using mechanics

Aldair E. Gongora, Kelsey L. Snapp, Emily Whiting, Patrick Riley, Kristofer G. Reyes, Elise F. Morgan, and Keith A. Brown

Supplemental Information

Transparent Methods:

Fabrication and Experimental Testing of Crossed Barrel Structures:

A crossed barrel structure with design parameters $x = (n, \theta, r, t)$, where n is the number of struts, θ is the angular displacement of the strut, r is the outer strut radius, and t is the thickness of the strut, was converted to a standard triangle language (STL) file using OpenSCAD. Using Slic3r, the generated STL file was converted to g-code and uploaded for fabrication to the MakerGear M3 FDM printers using OctoPrint. The crossed barrel structures were printed using polylactic acid (PLA) filament. The diameter of the printer nozzle was 0.35 mm and the diameter of the PLA filament was 1.75 mm. The nozzle temperature was 215 °C, and the printer bed temperature was set to 75 °C for all layers after the first, which was printed with a 85°C bed temperature. The structures were retrieved when the bed temperature was below 40 °C, and then tested using uniaxial quasi-static compression on a universal testing machine (5965, Instron Inc.) at a speed of 3 mm/min with a force threshold of 4.8 kN. Toughness U was measured as the area under the force F-displacement D curve. The F-D curve was truncated if, after an initial displacement of 2 mm, the force was below 50 N. The height of the crossed barrel was 23 mm, and the outer and inner radius of the top disk was 13 and 8 mm, respectively. The fabrication process and testing protocol described up to this point was based upon our prior work (Gongora et al., 2020). Resilience U_E of the crossed barrel structure was measured as the area under the F-D up to the yield force, which was calculated from the 0.2% offset yield stress. When the yield force of the crossed barrel structure was larger than the force threshold, U_E was measured as the area under the entire F-D curve.

Finite-Element Analysis (FEA) of Crossed Barrel Structures:

To predict the resilience $\widetilde{U}_E$ and the yield force $\widetilde{F}_y$ of the crossed barrel structures, simulations of uniaxial quasi-static compression were performed with the finite element analysis (FEA) software package ABAQUS/Standard on the full-size model. The generated STL structure of a crossed barrel structure with design parameters $x = (n, \theta, r, t)$ was converted to a binary voxelized volume representation with voxels of length 0.25 mm where each voxel was then converted to an 8-node hexahedral finite element (C3D8 brick element) (**Figure S1**). The element size was selected to balance computational time and accuracy. To capture the response of PLA in FEA, an isotropic elastic, perfectly-plastic material model was used with Young's modulus $E = 1.66$ GPa, yield strength $\sigma_y = 56.62$ MPa, and Poisson's ratio $\nu = 0.36$. The E and σ_y of PLA used in the simulation were determined from stress σ-strain ε curves from uniaxial quasi-static compression tests of cylindrical samples of 8 mm diameter and 16 mm height, where E was measured from the slope of the σ-ε curve and σ_y was measured as the stress at the 0.2% offset strain (**Figure S2 A**). Based on 12 compression tests on identically prepared samples (**Figure S2 B**), the mean measurements of E and σ_y were used in FEA where the coefficient of variation (CV) of E and σ_y was 8.9% and 10.5%, respectively (**Figure S2 C**). To increase the throughput of

simulations, the simulations were conducted using a custom MATLAB script that automated the aforementioned meshing and analysis and ran on Boston University's shared computing cluster (SCC). To capture FEA predictions throughout parameter space, we built a surrogate model using Gaussian process regression trained using FEA calculations selected on a grid (**Figure S3**). The surrogate model for resilience had a root mean squared error (RMSE) of 0.0242 J and mean squared percentage error (MSPE) of 0.07% (**Figure S4 A**). The trained surrogate model for yield force had a RMSE of 0.0382 kN and MSPE of 0.02% (**Figure S4 B**).

Bayesian Optimization Formulation:

In the Bayesian optimization (BO) framework, we used a Gaussian process (GP) to build a belief model $\mathcal{B}_i = GP\left(\mu_i(x), \sigma_i^2(x)\right)$ with mean $\mu_i(x)$ and variance $\sigma_i^2(x)$ from i experimental observations $\{x_{1:i}, y_{1:i}\}$ to model the property of interest in the design space $x = (n, \theta, r, t)$. In particular,

$$\sigma_i^2(x) = \Sigma(x, x) - \Sigma(x, x_{1:i})(\Sigma(x_{1:i}, x_{1:i}) + \lambda^2 I_i)^{-1}\Sigma(x_{1:i}, x), \qquad (1)$$

where λ defines the homoscedastic noise and I_i is the identity matrix. The covariance kernel $\Sigma(x, x')$ was a squared exponential, specifically,

$$\Sigma(x, x') = \alpha^2 \exp\left(-\frac{1}{2}\sum_{j=1}^{d}\left(\frac{(x_j - x'_j)^2}{\beta_j^2}\right)\right). \qquad (2)$$

The kernel was parametrized by $d + 1$ parameters (specifically α and β_j, the latter of which comprises d values) where the design space dimensionality $d = 4$. The parameters of the kernel and the noise were optimized using maximum likelihood estimation after every subsequent experiment. Additionally, the hyperparameters were bounded to be greater than or less than their initialization values by at most a factor of 100 to avoid extremal hyperparameters or over-fitting.

Learning was performed using experimental resilience U_E, experimental toughness U, FEA-calculated resilience $\tilde{U}_E(x)$, and FEA-calculated yield force $\tilde{F}_y(x)$. In the case of using an FEA-prior approach to optimize U_E, we defined a discrepancy model $\delta(x) = U_E(x) - \tilde{U}_E(x)$ to explicitly learn the difference between experiment and FEA. In the case of using an FEA-informed approach to optimize U, we defined an effective length $L(x) = \frac{U(x)}{\tilde{F}_y(x)}$ to conceptually separate effects from strength and ductility. In campaigns that were not informed by FEA, a zero-mean prior was used. For optimizing toughness, we defined a logistic function P_F to filter parameter space. Specifically, $\mathrm{P}_F = 1$ indicated that $\tilde{F}_y(x)$ would not exceed the force threshold $F_t = 4.8$ kN and $\mathrm{P}_F = 0$ indicated that $\tilde{F}_y(x)$ would exceed F_t. Additionally, a transition region with a width of 15%, a number based on the coefficient of variation (CV) of U, was defined to account for structures with $\tilde{F}_y(x)$ near F_t. The logistic regression function was built from FEA and the force threshold as,

$$\mathrm{P}_F(x) = \frac{1}{1 + \exp\left(k(\tilde{F}_y(x) - F_t)\right)}, \qquad (3)$$

with $k = 6.396$ to produce a 15% width.

Experimental campaigns began with an experiment that was selected uniformly at random in parameter space. Subsequent experiments were selected based upon an expected improvement (EI) decision policy. Specifically, after i experiments, the mean and variance of the belief model were calculated based on Table S1 and Equation (1), respectively. The decision policy used for each campaign is given in Table S1.

Table S1. Mean and decision policy used in each of the experimental and simulated campaigns, related to Figures 3 and 5.

Campaign	Mean $\mu_i(x)$	Decision policy x_{i+1}
FEA-informed resilience	$\widetilde{U}_E(x) + \Sigma(x, x_{1:i})(\Sigma(x_{1:i}, x_{1:i}) + \lambda^2 I_i)^{-1}\left(U_E(x_{1:i}) - \widetilde{U}_E(x_{1:i})\right)$	$argmax(EI(x)\|\mathcal{B}_i)$
Resilience without FEA	$\Sigma(x, x_{1:i})(\Sigma(x_{1:i}, x_{1:i}) + \lambda^2 I_i)^{-1}(U_E(x_{1:i}))$	$argmax(EI(x)\|\mathcal{B}_i)$
FEA-informed toughness	$\widetilde{F}_y(x) * \Sigma(x, x_{1:i})(\Sigma(x_{1:i}, x_{1:i}) + \lambda^2 I_i)^{-1}\left(\dfrac{U(x_{1:i})}{\widetilde{F}_y(x_{1:i})}\right)$	$argmax(EI(x) \cdot P_F(x)\|\mathcal{B}_i)$
Toughness without FEA	$\Sigma(x, x_{1:i})(\Sigma(x_{1:i}, x_{1:i}) + \lambda^2 I_i)^{-1}(U(x_{1:i}))$	$argmax(EI(x)\|\mathcal{B}_i)$

Simulated learning campaigns were conducted to assess the performance P of the various approaches to optimize U_E or U. The experimental observations used in the simulations were drawn from a dataset y_{600} comprised of the mean U_E or U for 600 distinct designs where each of the 600 designs was fabricated and tested in triplicate for a total of 1,800 experiments. To test a given combination of learning approach and campaign goal, a total of 100 independent simulated learning campaigns were conducted where the initial experiment for each simulated campaign was selected uniformly at random from the dataset. In the simulated learning campaigns, the subsequent experiments were selected from y_{600} using EI as described in Table S1. The performance of a given campaign after i experiments was given by its predicted optimum x_i^* and was defined as $P_i = \dfrac{y_{600}(x_i^*)}{\max(y_{600})}$. Zero-mean Gaussian noise was also added to the experimental observations with SD equal to the median SD of the dataset, 0.09 J for U_E and 1.66 J for U. Based on the P of 100 independent simulations, the probability of success as a function of i was calculated as $P_S = \dfrac{\sum_{j=1}^{100} \phi(P_{i,j})}{100}$ where $\phi(P_{i,j}) = 1$ if $P_{i,j} \geq 0.90$ and $\phi(P_{i,j}) = 0$ otherwise. These simulations were conducted in MATLAB using the Statistics and Machine Learning Toolbox.

To evaluate the degree to which the FEA simulation would deviate if the material properties were set to incorrect values, we predicted resilience $\widetilde{U}_E$ with varying E and σ_y. Specifically, we increased and decreased E by 5%, 10%, and 50%, while keeping σ_y constant. Additionally, we repeated this process for σ_y while keeping E constant. From this study, we found that both properties linearly affected $\widetilde{U}_E$ with a 1% increase in E resulting in a 0.36% decrease in $\widetilde{U}_E$ and a 1% increase in σ_y resulting in a 1.36% increase in $\widetilde{U}_E$. To investigate the performance of

the FEA-informed approach in the presence of variable input parameters that affect the FEA predictions, we conducted simulating learning campaigns in which FEA predictions were multiplied by a factor ϕ_{FE} = 0.50, 0.90, 0.95, 1.05, 1.1, and 1.5 (**Figure S5**). The range of ϕ_{FE} was selected to reflect the variability introduced when incorrect material properties are employed that would skew the FEA predictions underestimating or overestimating the experimental resilience U_E. From these simulations, we observed that within 10% (ϕ_{FE} = 0.90, 0.95, 1.05, and 1.1), the performance P of the FEA-informed approach varied slightly from the FEA-informed approach used in the study (ϕ_{FE} = 1). For ϕ_{FE} = 0.5, a small deviation was observed in P while for ϕ_{FE} = 0.90 a large deviation was observed in P. In totality, these results support that while variable material properties can affect FEA predictions, the FEA-informed approach still outperforms the uninformed-prior approach even when very inaccurate material properties are used.

Six independent experimental campaigns were conducted using the BEAR for U_E or U with three campaigns using the uninformative-prior approach and three campaigns using an FEA-based approach, namely the FEA-prior for U_E and the FEA-informed approach for U. The first experiment was selected uniformly at random, and the campaigns ran for the corresponding allotted experiment budget without a human in the loop. In the formulation of the decision policy, subsequent experiments were selected from a uniformly random, finite number of candidate designs drawn from parameter space. In the uninformed-prior approach for optimizing U_E and U, the predicted optimum after i experiments was selected as $x_i^* = \text{argmax}(\mu_i(x))$. In the FEA-prior approach to optimize U_E, the optimum structure also selected as $x_i^* = \text{argmax}(\mu_i(x))$. In the FEA-informed approach to optimize U, the predicted optimum structure at i experiments was selected as $x_i^* = argmax(\mu_i(x_{P_F \geq 0.90}))$, where $x_{P_F \geq 0.90}$ was the subset of designs in parameter space $P_F \geq 0.90$. This was done to account for areas in parameter space near the force threshold, where there is a sharp decrease in U. The probability of an FEA-informed structure outperforming the uninformative approach was computed by calculating the difference between the FEA evaluation measurement and the uninformative-prior measurement and assigning 1 if the difference was positive and 0 otherwise. The probability was then computed for each structure, and the average probability was calculated (**Figure S6 A and B**).

Supplemental Information:

**STL**   **Voxelized representation**   **Hexahedral mesh**

Figure S1: The standard triangle language (STL) file of each crossed barrel structure was first converted to a voxelized representation and then the individual voxels were converted to 8-node hexahedral finite-elements, related to Figure 2.

Figure S2: Material properties of polylactic acid (PLA), related to Figure 2.

A. Stress σ-strain ε curve showing uniaxial compression testing of a cylindrical specimen that was 8 mm in diameter and 16 mm tall. This curve was used to determine Young's modulus E and yield strength σ_y.

B. Measurements of E and σ_y obtained from the compression tests of 12 cylindrical specimens.

C. From the measurements of E and σ_y, the mean (height of the bar) and standard deviation (error bars) were determined. The individual measurements are shown as grey ticks.

Figure S3: Markers indicate the 1,188 design locations sampled by grid-search to obtain finite-element analysis (FEA) predictions from which to build a surrogate model using gaussian process regression, related to Figures 3 and 5.

A

B

Figure S4: Parity plots of surrogate model predictions versus FEA observations from grid-search, related to Figures 3 and 5.

A. Surrogate model predictions of resilience $\widetilde{U}_E$ vs. FEA observations from grid-search $\widetilde{U}_{E,grid}$.

B. Surrogate model predictions of yield force $\widetilde{F}_y$ vs. FEA observations from grid-search $F_{y,grid}$.
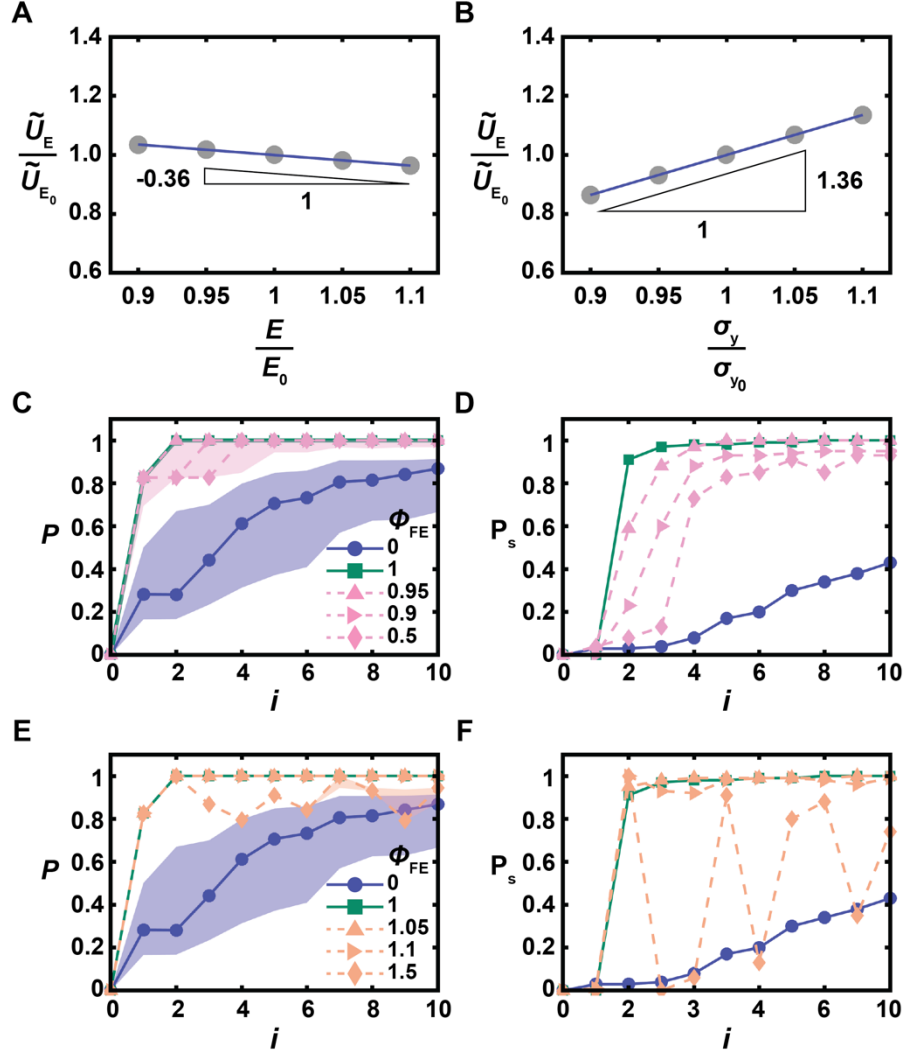
Figure S5: Performance of simulated learning campaigns for resilience U_E with varying finite element analysis (FEA) prediction quality, related to Figure 3.

A. **FEA predicted resilience $\widetilde{U}_E$ vs. Young's modulus E where $\widetilde{U}_{E_0}$ is the FEA prediction using the experimentally determined Young's modulus E_0 and yield strength σ_{y_0}.**

B. **$\widetilde{U}_E$ vs. yield strength σ_y.**

C. **Simulated performance P at experiment number i for a simulated learning campaign to optimize resilience using FEA predictions varied with multiplication factor ϕ_{FE} where $\phi_{FE} = 0$ corresponds to the uninformative prior approach, $\phi_{FE} = 1$ corresponds to the FEA-informed approach used in the study, and $\phi_{FE} = 0.95$, 0.9, and 0.5 correspond to decreasing the FEA predictions.**

D. **Probability P_s of achieving $P \geq 0.90$ at a given i for $\phi_{FE} = 0$, 0.5, 0.9, 0.95, and 1.**

E. **P at i for a simulated learning campaign to optimize resilience using FEA predictions varied with ϕ_{FE} where $\phi_{FE} = 1.05$, 1.1, and 1.5 correspond to increasing the FEA predictions.**

F. **P_s of achieving $P \geq 0.90$ at a given i for $\phi_{FE} = 0$, 1, 1.05, 1.1, and 1.5.**
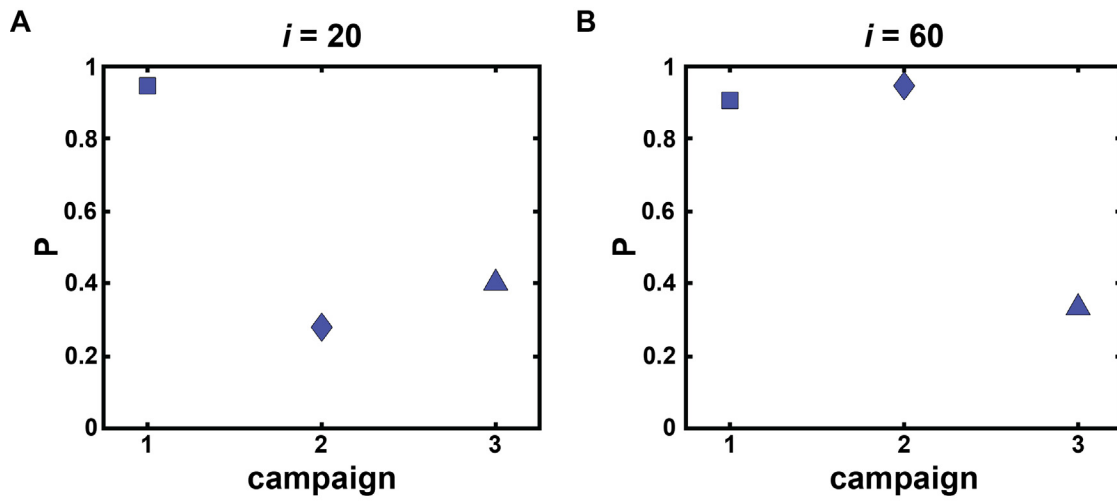
Figure S6: Experimentally determined probability that an FEA-informed campaign will identify a design with higher toughness U than that identified by an uninformative-prior campaign after $i = 20$ experiments (A) and $i = 60$ experiments (B), related to Figure 5.