

This is the accepted manuscript made available via CHORUS. The article has been published as:

Optimizing a reconfigurable material via evolutionary computation

Sam Wilken, Marc Z. Miskin, and Heinrich M. Jaeger

Phys. Rev. E **92**, 022212 — Published 31 August 2015

DOI: [10.1103/PhysRevE.92.022212](https://doi.org/10.1103/PhysRevE.92.022212)

OPTIMIZING A RECONFIGURABLE MATERIAL VIA EVOLUTIONARY COMPUTATION

Sam Wilken,^{*} Marc Z. Miskin,[†] and Heinrich M. Jaeger
*The James Franck Institute and the Department of Physics,
The University of Chicago, Chicago, Illinois 60637*
(Dated: August 10, 2015)

Rapid prototyping by combining evolutionary computation with simulations is becoming a powerful tool for solving complex design problems in materials science. This method of optimization operates in a virtual design space that simulates potential material behaviors and after completion needs to be validated by experiment. However, in principle an evolutionary optimizer can also operate on an actual physical structure or lab experiment directly, provided the relevant material parameters can be accessed by the optimizer and information about the material’s performance can be updated by direct measurements. Here we provide a proof of concept of such direct, physical optimization by showing how a reconfigurable, highly nonlinear material can be tuned to respond to impact. We report on an entirely computer controlled lab experiment in which a 6 x 6 grid of electromagnets creates a magnetic field pattern that tunes the local rigidity of a concentrated suspension of ferrofluid and iron filings. A genetic algorithm is implemented and tasked to find field patterns that minimize the force transmitted through the suspension. Searching within a space of roughly 10^{10} possible configurations, after testing only 1,500 independent trials the algorithm identifies an optimized configuration of layered rigid/compliant regions.

I. INTRODUCTION

In designing materials, one is often faced with a large number of control parameters whose relative importance for reaching optimized performance is a priori unknown [1, 2]. Given that it is typically not feasible to explore the large parameter space through exhaustive experimentation, computational methods have been gaining importance. These methods perform virtual experiments by simulating a given physical situation, and then update the simulation parameters until they get close to the desired result. Among such methods, evolutionary computation has been particularly successful in situations where the searchspace is complex, and they have found application in diverse areas including robotics optimization and learning [3, 4], crystal structure prediction [5–8], directed self-assembly [9, 10] and optimization of granular materials [11–16]. In all of these cases, the standard process is that the optimization occurs within the virtual simulation environment and only once the algorithm arrived at a final solution is this solution translated into a physical object or an experiment to validate performance.

New opportunities for design arise when considering materials whose properties can be changed quickly by external electrical, optical or mechanical stimuli: this enables feedback whereby the optimizer can operate directly on the physical material by measuring its properties and, in response, devising suitable stimuli

to steer those properties towards the desired goal. Circumventing the simulation can have a number of advantages. In particular, while the optimizer can be a black-box algorithm agnostic of the actual physics, the simulation needs to be carefully adjusted and validated to capture all essential aspects of the experiment to be replicated. For complex nonlinear processes this often is a key difficulty, which is avoided by using the actual experiment to generate the signal that is fed back to the optimizer. All of this becomes feasible once the optimizer is sufficiently powerful to require sampling of only a small subset of possible parameter combinations to get close to the desired goal, and as long as the material can cycle through different configurations rapidly.

As a proof of concept, we demonstrate this approach with a reconfigurable material consisting of regions whose stiffness can be switched from rigid to compliant. The particular goal is to have this material automatically adapt and configure itself such that the transmission of impact force is minimized. The material is a field-activated dense suspension of ferrofluid and iron filings. Under the application of an external magnetic field the suspension strongly increases its viscosity, a phenomenon that has been used in a number of applications as smart fluids [17–20]. At the large solids concentration (packing fraction of about 0.5) we are considering here, the suspension behaves like a field-activated granular material that can be driven through a jamming transition and becomes rigid. An array of electromagnets controls this reversible transformation locally across a 6x6 grid pattern of individually addressable regions that cover the volume of the material. The complex, nonlinear dynamic behavior of the material arises both from the field-response of individual regions and from the

^{*} Current Address: Department of Physics, New York University, New York, NY 10003.

[†] Current Address: Department of Physics, Cornell University, Ithaca, NY 14853.

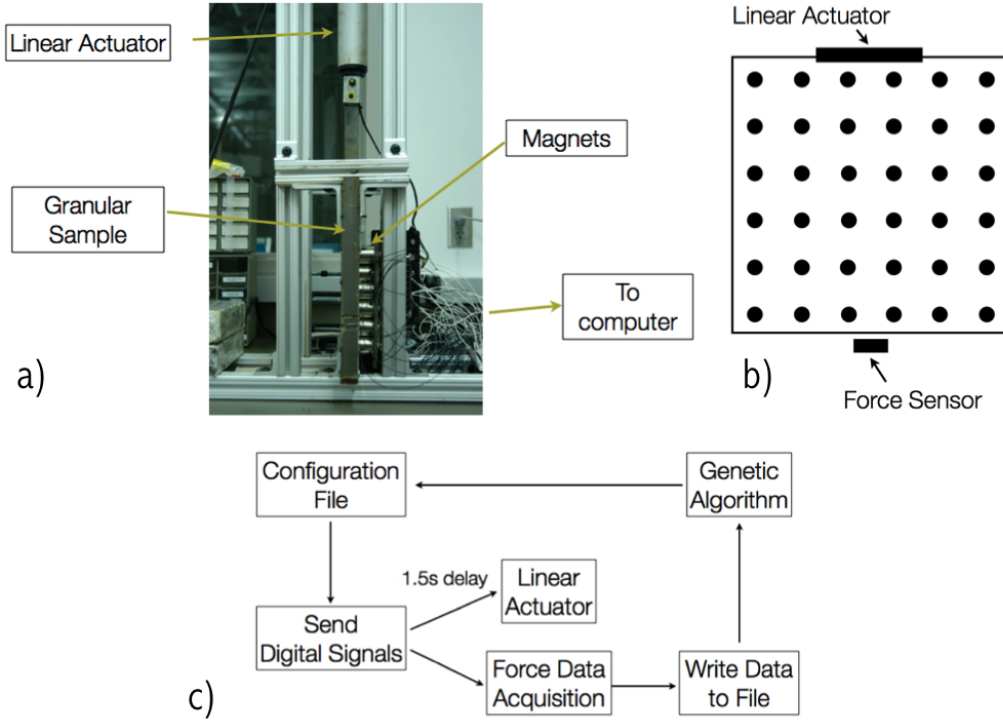


FIG. 1. a) Picture of the experimental setup, showing the primary components: the linear actuator, the container holding the suspension, and the electromagnet array. The force sensor sits at the bottom of the container. As seen from the side, the container is 1" wide and 12" tall. b) Sketch of the side wall of the container, indicating the layout of the magnet array, the size of the impactor at the top and the position of the force sensor at the bottom. c) Diagram of the automation process that is run for each configuration tested.

interactions of neighboring regions. On one side of the suspension volume is a linear actuator that impacts the material and on the opposite side we mount a force sensor. The task we are setting ourselves is to configure the field-induced pattern of spatial rigidity variation inside the material such that it minimizes the signal measured by the force sensor.

Computer control of this material allows us to create any specific field-induced pattern, via the electromagnets, and reliably evaluate its performance. However, already in the simplest case, just using each electromagnet in binary ON/OFF mode, there are 2^{36} or roughly 10^{10} possible configurations. Furthermore, the relationship between the control parameter—in this case the current to the electromagnet—and the local degree of rigidity inside the material are highly nonlinear. Nevertheless, the essentially digital, binary nature of the material with its grid of rigid/compliant cells, makes optimizing the impact response a natural problem for a genetic algorithm to tackle.

Artificial evolution achieves its power by comparing members of a population of trial solutions that comprise a generation, selecting the best-performing member(s), and then applying modifications (mutations) to generate the next generation of trials in an effort to get closer

to the desired performance. Typically, these trial solutions are all generated in parallel; for example, by performing a separate simulation for each member of the population. In principle, the same could be done with a large number of parallel experiments. For the proof-of-concept approach described in this paper, we reduce this to a single experimental setup and use it to cycle through all trial solutions per generation before moving on to the next generation. We show that such an algorithm can explore the enormous configuration space efficiently, operating on the material directly, and can find optimized configurations.

II. EXPERIMENTAL SETUP

Figure 1a shows the experimental setup. In the center is the rectangular container (dimensions 12"x12"x1", with walls made from 430-type stainless steel of 0.035" thickness), seen here from its side, that holds the volume of dense suspension—about 1.5 liters. The two-dimensional array of electromagnets is mounted on the outside of one of the large faces of the container (Fig. 1b). Each magnet in its on-state can apply a field around 0.5T at a depth of $\frac{1}{2}$ " into the suspension. A Gauss meter measured the magnetic field as $0.2 \pm 0.1T$,




FIG. 2. Magnetic field dependence of the viscosity for a 6 cSt ferrofluid, a suspension of 50% by volume 250-400 micron iron filings dispersed in 100 cSt silicone oil, and a suspension of 50% by volume of 250-400 micron iron filings suspended in 6 cSt ferrofluid. The red box indicates the magnetic field $\frac{1}{2}$ " into the suspension applied by one magnet in the ON state during optimization experiments. The blue box corresponds to the field with the magnet OFF. The nonzero field corresponds to contributions from neighboring magnets turned ON.

$\frac{1}{2}$ " away from the magnet. The relative permeability of the ferrofluid of 2.6 (from the manufacturer) is then used to arrive at the quoted value. The force sensor is mounted at the center of the bottom container wall, in line with the linear actuator that drives a T-shaped piston vertically into the suspension above. The piston cross-section is 3" wide x 1" deep. For each trial, when the piston impacts the sample, the penetration corresponds to roughly two magnet rows.

To maximize the difference between rigid and compliant states in the suspension material, 250-400 micron iron filings were dispersed in a ferromagnetic fluid composed of colloidal magnetite particles in a proprietary blend of light hydrocarbons manufactured by Ferro Tec (Type EFH1), at a packing fraction $\phi = 0.50$ iron filings by volume. Figure 2 shows the field response of this mixture and compares it with the response of the ferrofluid by itself as well as iron filings in 100cSt silicone oil at the same ϕ but without ferrofluid. All data were taken in an Anton Paar MC301 rheometer at constant shear rate of $2s^{-1}$. Given the field strengths available with our magnet array, the iron filings/ferrofluid mixture can transform from a viscosity around 10 Pas when the magnet is off to values in excess of 500Pas with the magnet turned on. By contrast, the ferrofluid alone reaches only 1Pas even at the highest fields, and the iron filings in oil exhibit an off-state with significantly larger viscosity (around 100 Pas) and thus a smaller dynamic range.

The electromagnets, the linear actuator and the force sensor are computer-controlled to allow for automated testing of configurations. Each electromagnet is activated by a data acquisition device (NI USB-6008/9). Since the device can only supply 5mA at 5V, it drives a small circuit with power transistors capable of handling the 0.35A at 12V required for fully turning on each magnet. The linear actuator is activated in a similar way, except that the larger current (5A at 12V) required to propel it forward is handled by a transistor circuit controlling a double-pole, double-throw relay. The force sensor (Omega DLC101-10, $(1.64 \pm 0.05) \times 10^{-2}$ Volts per Newton output) continuously measures the force, while the computer collects a fixed number of voltage values as the impact is taking place. Figure 3a shows the measured force signal.

During operation, the genetic algorithm sends a file designating a field configuration to the magnet control portion of the software, which turns on the right combination of electromagnets and triggers the force sensor to begin collecting data (Fig. 1c). After a 1.5s delay, a signal is sent to the linear actuator to begin impact. The data acquisition software then writes the force data to a file for the algorithm to interpret. In our experiment, the algorithm focuses on the peak force value (Fig. 3a). Then the cycle begins again as the algorithm sends out another configuration to the apparatus. Each configuration is tested in approximately 3s. As the configurations approach the optimized pattern the peak force is seen to decrease.

Because we are testing the material under conditions of severe impact by ramming the linear actuator into the suspension, a method is required to minimize mechanical compaction of the material after repeated tests and to refresh the initial state of the material before testing the result of magnet configuration changes. As Fig. 3b shows, without such refreshing there is a monotonic decay of the peak force when the same configuration is impacted repeatedly. To reduce this effect, the material is vibrated by a large speaker for 30 seconds between each impact. The speaker is mounted on the side of the container, opposite the magnets (Fig. 1), and outputs a 55 Hz wave train with 4 pulses per second (Fig. 3b, inset). This vibration before each trial significantly increases the reproducibility of the initial state, which now fluctuates around a fixed value by $\pm 2N$ and no longer decreases continually (blue trace in Fig. 3b).

The optimization is performed on the computer by a population-based incremental learning algorithm [21, 22], a type of genetic algorithm. The algorithm first selects ten random magnet configurations, constituting the initial generation. The configurations are randomly sampled assuming a fixed probability for the magnet at site (i,j) to be P_{ij} . Each of these configurations is then

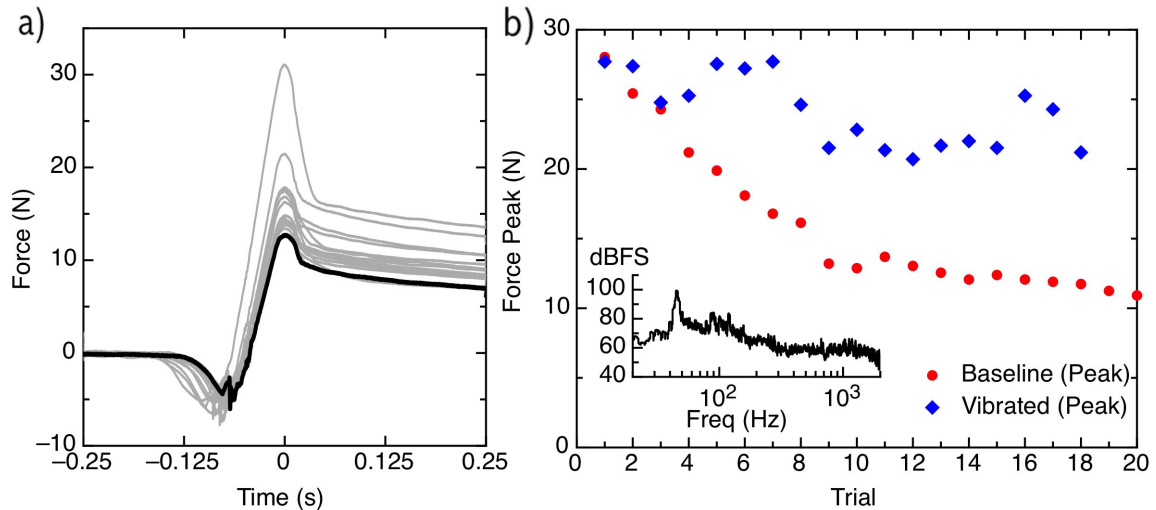


FIG. 3. a) Temporal force profile during impact as measured by the sensor at the bottom of the container. The different traces indicate the evolution of the profile during optimization (gray) and in the optimized state (black). The piston impacts the top of the suspension at about -0.125s. b) Peak force as a function of the number of impacts performed on one fixed configuration without (red) and with (blue) vibration between subsequent impacts. Inset) Power spectrum of the audio clip used to vibrate the suspension for 30s between trials.

evaluated by performing an impact experiment and the algorithm assigns a fitness value to each, in this case the peak force measured. Once all configurations in one generation have been tested, they are ranked by performance and used to update the probabilities for magnet activation. The new probabilities P'_{ij} are set according to $P'_{ij} = P_{ij} + \tau \text{Cov}(R, M)$ where τ is the learning rate, which determines how much the probability can jump in one generation, set to $\tau = 0.1$ for all optimization runs; R is the relative fraction of other configurations in the generation worse than the configuration in question and M_{ij} is a binary representation of the magnets where $ON = 1$ and $OFF = 0$. The covariance matrix is implemented to allow for correlated mutations. Configurations for the next generation are sampled using these updated probabilities and the process is iterated. We consider the optimization complete when we see that most of the probabilities converge to 0 or 1 (some seem to stay around 0.5 because it is not guaranteed that every magnet will influence the impact behavior).

III. RESULTS

Tasking the algorithm with minimizing the peak force value corresponds to finding the configuration that functions as the best shock absorber. One might, at first glance, guess that the all ON and all OFF configurations would set the upper and lower performance limits, given that the average rigidity of the whole system is lowest in the all OFF state and highest with all magnets ON. However, this assumes a homogeneous material that is compressed uniformly. On our experiment, the

situation is deliberately made more complex by letting the impact occur over only a portion of the top free surface of the dense suspension. Because this surface can deform, the material in its low-viscosity OFF state can let the impactor penetrate considerably and thus create a sizable force at the bottom of the container. By turning all magnets ON a more rigid connection to the bottom is established and while the impactor does not penetrate as much, it can now transmit even larger force. One might also consider a random pattern of ON/OFF cells. However, in this case there is a high chance, on average, for rigid cells to percolate across the material and thereby channel impact momentum vertically rather than distributing it more widely in the horizontal direction.

In Fig. 4a we show the evolution of the material toward its optimized configuration by plotting the median peak force for each generation. In generation 1, the optimizer starts with a random pattern, which transmits larger forces than either the all ON or OFF configurations (red and blue lines). After about 100 generations the optimizer has settled into an asymptotic configuration that absorbs peak forces roughly twice as efficiently than the initial random state and 20% better than the all OFF state. In Fig. 4b we plot the evolution of the total error in the peak force, given by $\sqrt{(\Delta F)^2 + (\delta F)^2}$, where ΔF is the spread in peak force values among configurations belonging to a given generation, and δF is the experimental uncertainty in determining each peak force. This error reaches a low, asymptotic level after a similar number of generations, N , as the peak force itself. The solid is a fit to an exponential decay, $F \propto \exp(-N/N_0)$ from which we

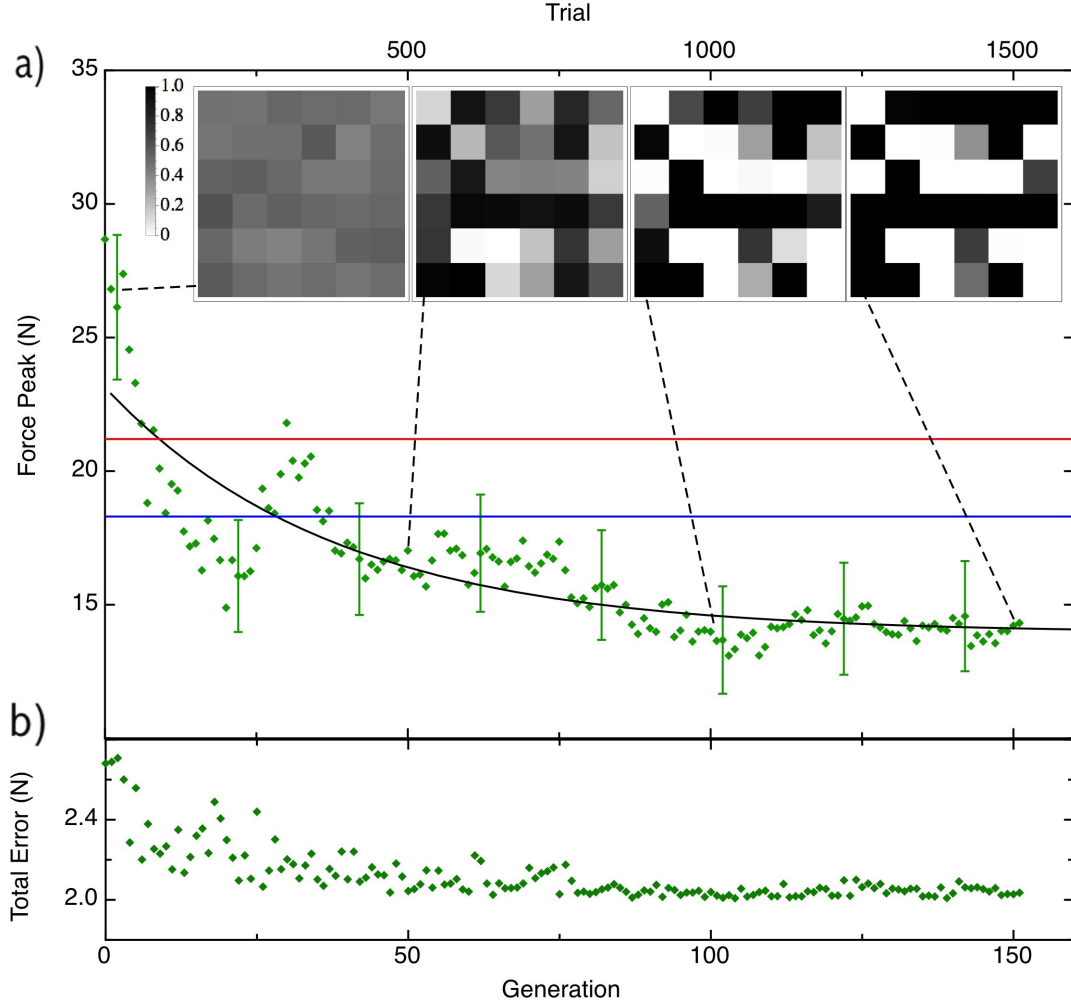


FIG. 4. a) Median peak force for each generation during a particular optimization run where ten configurations (trials) are tested per generation. The black line is a fit to a decaying exponential, while the red and blue lines correspond to the limit of all magnets ON or OFF, respectively. The pictures indicate the average magnet configuration for a given generation, with color indicating the probability to be ON (black = 100%). Initially all probabilities are about 0.5 but as the optimizer finds better solutions the configurations converge to a particular state. b) Total error in the median peak force as a function of generation.

extract a characteristic optimization time of $N_0 = 42 \pm 9$ generations. While the median peak force decreases as the algorithm narrows in on promising configurations, the shape of the impact profile as a function of N still remains the same (Fig. 3a). This demonstrates that the qualitative behavior of the force transmission during impact is independent of the configuration. Performing several independent optimization runs with the same fitness metric and material configuration we find that they all show similar behavior and the data shown in Fig. 4 is representative for the whole group.

The call-outs in Figure 4a give a visual image of the 6 x 6 magnet pattern after a certain number of generations have elapsed. Black corresponds to ON, white to OFF. Each picture is an average over the patterns that populate a particular generation. Early in the evolution, the various patterns exhibited by members of a population

appear to be uncorrelated because the optimizer is still exploring a wide swath of configuration space. As a result, the average patterns appear almost uniformly gray. Further into the evolution, the algorithm fine-tunes patterns based on the better performing configurations from previous generations. As a result, the probability of a cell to be either rigid or compliant converges and the average patterns become more well-defined.

Close inspection of the evolving configurations indicates that some cells converge very quickly, indicating that they are clearly important to decreasing the transmitted force. For example, the third row of cells from the bottom stays always ON, starting from just the fiftieth generation, which is approximately the optimization decay time. However, some cells do not converge to a fixed state and continue to fluctuate between ON and OFF. We interpret this as a sign that these cells do not

affect the force transmission significantly. Which cells are contributing strongly, or minimally, to an optimization will depend on the task at hand. In the specific situation considered here, the details of the impactor geometry, the place at which impact occurs with respect to the boundaries of the suspension volume, and the place at which the transmitted force is measured all are likely to play key roles. However, when minimizing force transmission through the material, we always find one dominant characteristic of the optimized configuration: a horizontal structuring with layers of alternating rigid and compliant cells.

Layering of materials of different compliance is widely used in energy dissipation applications. Without having been programmed to explore layering (or, in fact, any particular aspect of the underlying physics), the algorithm independently identified layering as a promising approach. In our realization of the algorithm with a population of 10 members per generation this solution is realized after no more than 100-150 generations, i.e., in no more than 1,500 trials. The total time for finding the optimized state was a little less than 14h. This includes 30s of vibration to refresh the initial state before each trial; the actual testing time was on the order of seconds per trial. Without the need to reset the initial state, optimizations could be performed in one hour. The total time could also be reduced by at least a factor of 10 if all trials per generation were performed in parallel instead of using a single experimental setup sequentially as we did for this proof of concept.

IV. CONCLUSIONS

We have shown that automated material design is possible without simulation by coupling a nonlinear, reconfigurable material directly to a genetic algorithm. In a specific example of a highly nonlinear material consisting of a planar array of 6 x 6 controllable cells, and thus a search space of 2^{36} possible configurations, individual configurations could be tested in seconds and solutions for the best shock-absorbing configurations were found on the timescale of hours. This example of a dense, field-activated suspension that can be magnetic field-activated locally was introduced here to serve as proof-of-principle, but the approach we introduced is generally applicable to a wide range of different material optimization tasks. Because the optimizer operates on the actual physical material or structure, issues that hamper simulation of complex dynamical processes, such as proper validation, are automatically taken care of. Furthermore, if the optimizer is not just used in a single-shot fashion to identify optimized parameters for a fixed performance target, but remains active and coupled to the material, self-learning responses can be envisioned by which the material adapts to changing targets.

ACKNOWLEDGMENTS

We thank Carlos Orellana for rheology measurements. This work was supported by the National Science Foundation through grant CBET 1334426. We acknowledge additional support through award 70NANB14H012 from the U.S. Department of Commerce, National Institute of Standards and Technology as part of the Center for Hierarchical Material Design (CHiMaD).

-
- [1] A. Jain, J. A. Bollinger, and T. M. Truskett, *American Institute of Chemical Engineers Journal* **60**, 2732 (2014).
 - [2] F. S. Bates, M. A. Hillmyer, T. P. Lodge, C. M. Bates, K. T. Delaney, and G. H. Fredrickson, *Science* **336**, 434 (2012).
 - [3] H. Lipson and J. B. Pollack, *Nature* **406**, 974 (2000).
 - [4] M. T. Tolley and H. Lipson, *International Journal of Robotics Research* **30**, 1566 (2011).
 - [5] A. O. Lyakhov and A. R. Oganov, *Physical Review B* **84**, 092103 (2011).
 - [6] A. R. Oganov, A. O. Lyakhov, and M. Valle, *Accounts of Chemical Research* **44**, 227 (2011).
 - [7] E. Bianchi, G. Doppelbauer, L. Filion, M. Dijkstra, and G. Kahl, *Journal of Chemical Physics* **136** (2012).
 - [8] L. Filion and M. Dijkstra, *Physical Review E* **79**, 046714 (2009).
 - [9] J. Qin, G. S. Khaira, Y. Su, G. P. Garner, M. Miskin, H. M. Jaeger, and J. J. de Pablo, *Soft Matter* **9**, 11467 (2013).
 - [10] G. S. Khaira, J. Qin, G. P. Garner, S. Xiong, L. Wan, R. Ruiz, H. M. Jaeger, P. F. Nealey, and J. J. de Pablo, *ACS Macro Letters* **3**, 747 (2014).
 - [11] M. Z. Miskin and H. M. Jaeger, *Nature Materials* **12**, 326 (2013).
 - [12] M. Z. Miskin and H. M. Jaeger, *Soft Matter* **10**, 3708 (2014).
 - [13] H. M. Jaeger, *Soft Matter* **11**, 12 (2015).
 - [14] R. Zalewski and M. Pyrz, *Mechanics of Materials* **57**, 75 (2013).
 - [15] K. Sobolev and A. Amirjanov, *Construction and Building Materials* **24**, 1449 (2010).
 - [16] T. Zohdi, *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **361**, 1021 (2003).
 - [17] R. Stanway, *Materials Science and Technology* **20**, 931 (2004).
 - [18] J. de Vicente, D. J. Klingenberg, and R. Hidalgo-Alvarez, *Soft Matter* **7**, 3701 (2011).

- [19] R. H. Ewoldt, P. Tourkine, G. H. McKinley, and A. Hosoi, *Physics of Fluids* (1994-present) **23**, 073104 (2011).
- [20] C. S. Orellana and H. M. Jaeger, *Soft Matter* **9**, 8519 (2013).
- [21] S. Baluja, *Population-based incremental learning. a method for integrating genetic search based function optimization and competitive learning*, Tech. Rep. (DTIC Document, 1994).
- [22] A. E. Eiben and J. E. Smith, *Introduction to evolutionary computing* (Springer, New York, 2003).