



Greater AI Design Control Aids Evolution of Computational Materials

Piper Welch^{1(✉)}, Monica Li³, Shawn Beaulieu¹, Annie Xia³, Dong Wang³, Medha Goyal³, Atoosa Parsa², Corey S. O'Hern³, Rebecca Kramer-Bottiglio³, and Josh Bongard¹

¹ Department of Computer Science, University of Vermont, Burlington, VT 05405, USA

{piper.welch, shawn.beaulieu, josh.bongard}@uvm.edu

² Department of Biology, Tufts University, Medford, MA 02155, USA
atoosa.parsa@tufts.edu

³ Department of Mechanical Engineering and Materials Science, Yale University, New Haven, CT 06520, USA
{monica.s.li, annie.xia, dong.wang, medha.goyal, corey.ohern, rebecca.kramer}@yale.edu

Abstract. Unconventional computing may overcome some of the limitations of traditional silicon-based systems using alternative materials and computational mechanisms. However, due to their complex underlying dynamics and high-dimensional parameter space, the design of these materials such that they perform computation is non-intuitive, making AI-driven design attractive. It has been shown that evolutionary algorithms can tune the structural properties of grains within a granular material such that it computes logical functions. In recent years, programmable granular metamaterials have been developed so that multiple physical properties of individual grains can be altered independently. This raises the question of whether allowing evolutionary algorithms to tune more grain features within a granular material frustrates or facilitates its ability to embed computation. In this work, we show that the latter is the case, when grain sizes and stiffnesses are co-evolved to embed Boolean logic gates, compared to evolving just sizes or stiffnesses alone. We report physical verification of evolved designs, taking a further step toward the provision of alternatives to electronic computing.

Keywords: Granular Metamaterials · Mechanical Computing · Sim2Real

1 Introduction

For over fifty years, Moore's law has governed the rapid development of semiconductors as the density of transistors on silicon chips exponentially increased [11]. But as transistors shrink toward atomic scales, both theoretical and practical barriers are hindering this rate of growth [19]. The impending end of Moore's

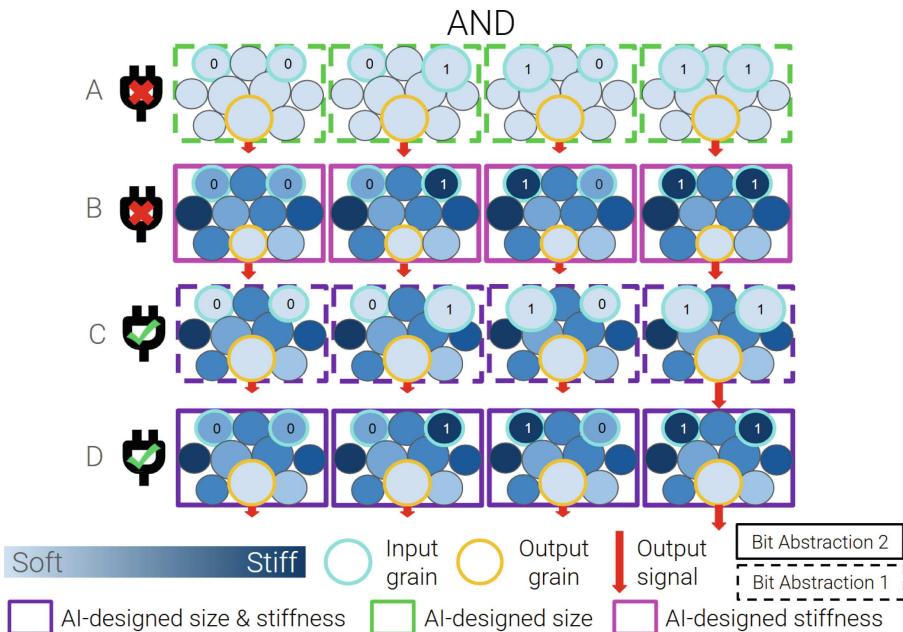


Fig. 1. Increasing AI design control over computational materials. (A): An evolutionary algorithm attempted to tune the sizes of 8 out of 10 grains in an *in silico* granular material to increase its “ANDness”, but failed. Here, manually making the two input grains small or large provided the two input bits; the resulting force at the output grain (red arrow) was interpreted as the output. (B): In another material the evolutionary algorithm could tune only grain stiffnesses, and the two input grains were softened or stiffened to provide the input value of the logic gate. Again, the evolutionary algorithm failed. (C): In a third material the evolutionary algorithm could tune sizes and stiffnesses, and input values were provided via grain size change. Here AND behavior was achieved. (D): AND behavior was also achieved in a fourth material in which sizes and stiffnesses were evolved, but input was provided via grain stiffness change.

law signals an urgent need for alternative computational substrates capable of sustaining progress in computing beyond the limits of conventional silicon architectures. To this end, some investigators have expanded upon traditional semiconductor technologies, such as the work by Jayachandran *et al.*, who demonstrated the successful fabrication of 3D chips using non-silicon materials (MoS₂ and WSe₂) [7]. Meanwhile, others have explored entirely new paradigms, leading to the emergence of mechanical, chemical, neuromorphic, and quantum computing [4, 10, 18, 23].

One promising novel computing strategy is mechanical computation using granular metamaterials, as there is currently no known upper limit on their computational density. This is due to the fact that these materials report the results of the computation as oscillatory signals at different frequencies simulta-

neously. The substrate of such a system, granular metamaterials, are materials with discrete units, or grains, whose strategic arrangement and local interactions can give rise to emergent properties not present in their constituent components. By tuning the spatial configuration and contact network of these grains, it is possible to engineer novel mechanical behaviors that arise solely from the collective dynamics of the system [3,5,6,8,15,22].

Computational granular metamaterials (CGMMs) represent a class of granular metamaterials that have been designed, *in silico*, to execute logical operations [1,12–14]. However, such materials are notoriously non-intuitive to understand and to design, so in all work to date, machine learning has been employed to tune the features of the grains within the materials such that they instantiate the desired computation. In all previous CGMM work, binary digits have been encoded as the presence ('1') or absence ('0') of mechanical vibration applied to manually specified grains representing the input ports. Then, the presence or absence of resulting vibration at a pre-specified output grain is interpreted as the value of the binary output. The benefit of encoding information as vibration is that one material can simultaneously act like more than one Boolean logic gate, with each gate operating at a different frequency. As the frequency spectrum has, in theory, infinite capacity, there is currently no known upper limit on CGMMs' computational density. The current experimental upper bound is 16 Boolean gates within a single material [1].

1.1 Physical Realization of Computational Materials

A profound limitation of vibrational CGMMs is the difficulty of their physical realization. This stems from, among other things, the sensitivity of the vibrational bit abstraction. Vibrational systems are sensitive to noise, and damping. Even minor perturbations such as deviation in resonance frequency or a slight phase misalignment can disrupt computation, significantly diminishing the robustness and reliability of the system's function.

Given this brittleness, we herein explore metamaterials that encode computation using mechanical phenomena other than vibration, and which do not change during the computing process. We refer to this class of CGMMs as *static* CGMMs. This facilitates the physical realization of AI designed, *in silico* computational materials, which we report below. Specifically, we test bit abstractions that involve changing a grain's bulk properties, such as size and stiffness, to encode binary bits. Computation is instantiated by the particular static force chains that arise within a given material. Static force chains operate by transmitting compressive or tensile forces through stable contact points between grains, forming reliable pathways for information flow [20]. The result of the computation is interpreted as the appearance of high ('1') or low ('0') forces arriving at a pre-specified output grain. By AI designing grain features within the material, different computations can be realized.

Although the abstraction of input and output bits differ, abstraction equality—and thus the cascading of multiple logic gates to compute more complex functions—could be achieved via hardware translation of input forces into

input grain property change [9]. The advantage of our encoding is that static CGMMs are a non-volatile computing system [23]. That is to say, they can sustain their computational states over time without requiring continuous energy input, unlike vibrational systems that rely on ongoing actuation. Unfortunately, in static CGMMs, the ability to operate multiple gates within a substrate at different frequencies is lost. However, we could alternatively utilize several bulk properties, such as force and shear modulus, to operate multiple gates simultaneously. This investigation of polycomputational static CGMMs is beyond the bounds of the present work.

1.2 Inverse Design of Computational Materials

Regardless of what physical phenomena encode information in a CGMM, as mentioned above, the inverse problem of designing a material that embodies a desired computation is non-intuitive and thus is difficult to tackle with hand-designed strategies. For this reason, machine learning has been employed to design CGMMs by allowing the optimization process to tune one grain feature: stiffness. However, new adaptive metamaterials are on the horizon in which multiple features of individual grains may be modified independently [9]. It is unknown whether broadening the design reach of machine learning to these new tunable grain features would frustrate design because of the increased dimensionality of the parameter space, or facilitate design as the larger space would contain more gradients to follow.

Specifically, in the case of static CGMMs, it remains unknown if independently varying multiple grain features could yield better computational results than tuning one feature per grain. To investigate this, here, we use evolutionary algorithms to evolve the properties of granular materials to behave as a variety of logic gates (AND, OR, XOR, and NAND). We find granular materials with AI designed grain sizes and stiffnesses are able to produce more distinguishable binary outputs for all tested logic functions, compared to materials with only AI designed grain sizes or stiffnesses. We then report the successful transfer of a designed logic gates from *in silico* to physical material. We conclude by providing an overview of the potential applications of this new computational platform and discussing future research directions.

2 Simulated Computational Materials

In this section, we first describe how simulated computational materials were designed (Sect. 2.1) and then investigate the quality of those designs (Sect. 2.2).

2.1 Methodologies

Several Boolean gates were AI designed into simulated granular materials with varying degrees of success.¹

¹ This work's code and supplementary materials can be found at [https://github.com/piperwelch/static.logic.gates/..](https://github.com/piperwelch/static.logic.gates/)

CGMM Setup. In this work, our materials are 2-D and comprised of 10 grains placed on a 3-4-3 hexagonal lattice. The lattice is inside a bounding box of fixed size in the x and y directions.

Experimental Design. To test the generality of our computational system, we optimize granular materials to perform as one of AND, OR, XOR, or NAND. For each logic gate, we conduct four experiments, testing two input bit abstractions, each under two different AI design conditions. Inputs are passed to a material in one of two ways: in *bit abstraction 1*, binary values were encoded by minimizing ('0') or maximizing ('1') the radial size of two input grains, while in *bit abstraction 2*, they were encoded through the minimization ('0') or maximization ('1') of grain stiffness. For each of the two input bit abstractions, we conducted two AI design conditions: first, optimization of the grain feature used for bit abstraction (only size or stiffness), and second, optimization of both grain features (size and stiffness). In all four cases, the output is measured as the amount of force exerted by the central grain in the bottom lattice row against the bounding wall.

Simulation. To simulate our granular system, we use the Discrete Element Method (DEM) [21]. All grains are frictionless and have the same mass ($m = 1$). Depending on the experimental setup, grains can have different radii or/and elastic moduli. Gravity is not included in our simulations. Particle-particle interactions are modeled as repulsive with linear spring potential. The Fast Inertial Relaxation Engine (FIRE) [2] is used for energy minimization.

Optimization. To design materials that behave as logic gates, we couple our physics simulator with an evolutionary algorithm. Specifically, we use Age-Fitness Pareto Optimization (AFPO) [17]. We selected AFPO as it requires no hyperparameters to promote diversity within the population. It is also well suited for problems with multiple local optima and prevents premature convergence.

Initialization. Each individual in the population represents the configuration of one granular material and is encoded as two float vectors with length 10, one assigning each grain's stiffness and another assigning their diameters. The initial population is instantiated with 200 randomly generated materials with grain properties drawn from a uniform distribution. Grains have variable stiffness ratios ($\in [0.5, 10]$) and diameters ($\in [1.0, 1.04]$). The stiffness ratio is chosen according to previous related work [1], while the diameter ratio was chosen as simulations can become unstable if all grains have a diameter above 1.04. The stiffness and sizes are unitless quantities and derive meaning from the ratios between grains. The locations of the grains chosen as the input and output of the logic gate are fixed to the locations shown in Fig. 1.

Evaluation. To assess a material's fitness, we measure the force, F_n , exerted by the output grain on the bottom wall of the simulation environment under

each input case $n \in \{‘00’, ‘01’, ‘10’, ‘11’\}$. We use fitness functions that measure a given material’s logical behavior (“GATEness”) as a floating point number. This allows gradients to be followed during the evolutionary search. Equations 1–4 present the fitness functions for designing AND, OR, XOR, and NAND gates respectively.

$$\text{ANDness} = \ln \left(\frac{F_{11}}{(F_{01} + F_{10} + F_{00})/3} \right) \quad (1)$$

$$\text{ORness} = \ln \left(\frac{(F_{01} + F_{10} + F_{11})/3}{F_{00}} \right) \quad (2)$$

$$\text{XORness} = \ln \left(\frac{(F_{01} + F_{10})/2}{(F_{00} + F_{11})/2} \right) \quad (3)$$

$$\text{NANDness} = \ln \left(\frac{(F_{01} + F_{10} + F_{00})/3}{F_{11}} \right) \quad (4)$$

Each of these fitness functions is designed to maximize the force under the input conditions where a ‘1’ output is expected, while minimizing the force under the input conditions where a ‘0’ output is expected. In the above equations, \ln denotes the natural logarithm.

Selection. After all 100 materials have been simulated and their behavior assessed, the materials with the lowest fitness values are removed. Survivor selection occurs by iteratively selecting two random individuals from our population, and discarding one if it is Pareto-dominated. Specifically, material i is removed if it is compared to another material j that demonstrates better performance ($f_j > f_i$) and belongs to a lineage l_j that is as young or younger than the lineage of material i ($l_j \leq l_i$). The age of each lineage is defined by the number of generations since it first emerged in the population. In the initial generation of randomly generated materials, each material begins a unique lineage with an age of zero. This process repeats until the population size reaches 100.

Reproduction. The materials allowed to reproduce are selected via tournament selection. In our implementation, two materials are randomly selected and the one with higher fitness is allowed to reproduce. This process repeats until the population contains 100 new materials.

Variation. Material mutation occurs by mutating the size, and/or stiffness of grains within a material. Specifically, we employ uniform mutation with a probability of 0.2 for each grain. The mutation size for grain diameter is $\in \pm 0.005$, while the mutation size for grain stiffness is $\in \pm 0.5$. We do not employ crossover.

Diversity. At each generation, one randomly generated material is injected into the population. This has the effect of introducing diversity into the evolving population of materials.

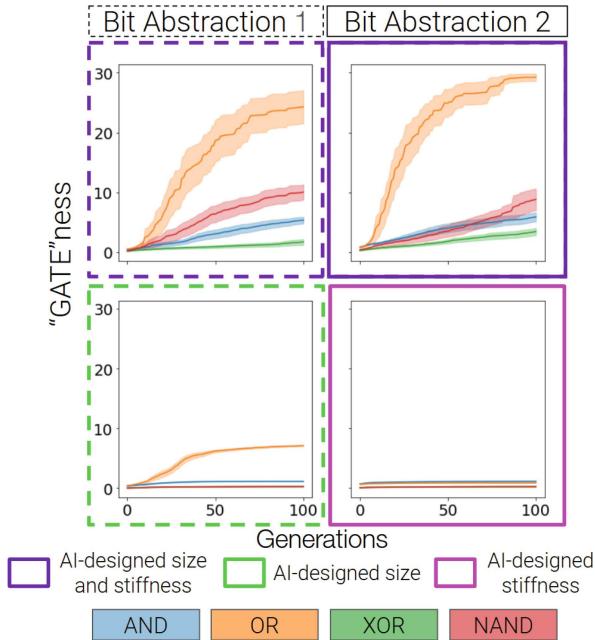


Fig. 2. Fitness over evolutionary time. This figure shows the fitness over evolutionary time for the 4 logic gates we investigated (AND, OR, XOR, NAND) across our 2 input bit abstractions, and 3 AI design conditions. Each line is the average of thirty experimental trials with a 95% confidence interval plotted in a lighter shade.

The process of simulating and evaluating new and freshly generated materials, removing low-fitness materials, duplicating and mutating surviving materials, and introducing a newly generated material is repeated for 100 generations of material evolution. We repeat this evolutionary process thirty times for each experimental condition. Each replicate uses a different random seed, and therefore each initial population is a novel random instantiation. We use the Mann-Whitney U test for all statistical comparisons and Bonferroni correction for multiple pair-wise comparisons.

2.2 Results and Discussions

The maximum fitness of the population over evolutionary time for each input bit abstraction, AI design condition, and logic gate is reported in Fig. 2. Here, each graph reports the average population maximum across 30 trials with the shaded region representing a 95% confidence interval. We find that (+AI designed size, +AI designed stiffness) granular materials have significantly higher fitness than those with only AI design for size or stiffness across both input bit abstractions and across all logic gates ($p < 0.005$ for all comparisons). When comparing the (+AI designed size, +AI designed stiffness) condition across bit abstraction

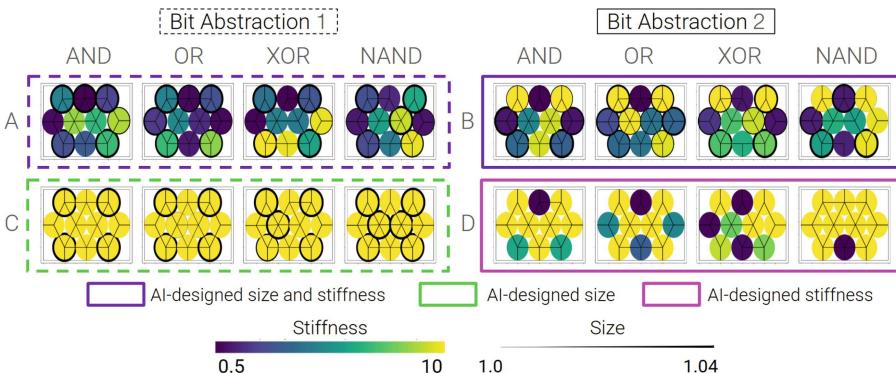


Fig. 3. Highest Fitness Solutions by AI-design Condition and Bit Abstraction Across Logic Gate. This figure shows the highest solution in the '11' input case for each logic gate across each AI-design condition and input bit abstraction.

1 and bit abstraction 2, we find that, after Bonferroni correction, there is no significant difference between any gates, save for XOR ($p < 0.05$), which is significantly higher for bit abstraction 2. In the (+AI designed size, +AI designed stiffness) conditions, for both bit abstraction 1 and bit abstraction 2, we find that the fitness from materials optimized for OR is the highest, followed by NAND, followed by AND, and ending with XOR. This pattern is not seen in the (+AI designed size, -AI designed stiffness) condition, which displays the highest fitness for materials optimized for OR, followed by AND, followed by XOR, and NAND. Similarly, the (-AI designed size, +AI designed stiffness) condition resulted in the materials optimized for AND having the highest fitness, followed by OR, NAND, and XOR. These results indicate that, while it is always beneficial to have AI design in both size and stiffness, some AI interventions are better suited for specific gates.

Top Performing Logic Gates. Figure 3 presents the configurations of the highest-performing materials at generation 100 for each logic gate, organized by AI design condition and bit abstraction. In this visualization, grain stiffness is denoted by color, while grain size is represented by both the plotted size and line width of each grain's border. Several intriguing structural and material property patterns emerge. We proceed with enumerating a few.

First, the AND and OR logic gates within the (+AI designed size, -AI designed stiffness) condition for bit abstraction 1 (Fig. 3C) share notable similarities in structure, though the OR gate exhibits a slight size increase of certain grains that are not found in the AND gate. The differences between the AND and OR gates for all other experimental conditions are more pronounced. We do not have any intuition regarding why this similarity has only appeared in the (+AI designed size, -AI designed stiffness) condition.

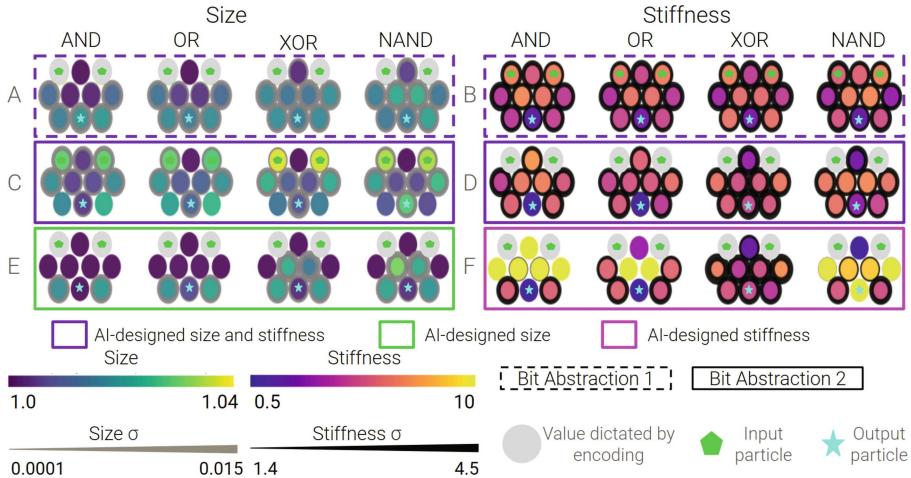


Fig. 4. Heat maps of mean size and stiffness across logic gates, input bit abstractions, and AI-design condition. This figure shows the mean size and stiffness across each input bit abstraction. **(A)**: Heat maps for mean size in materials with AI designed size and stiffness. **(B)**: Heat maps for mean stiffness in a materials with AI designed size and stiffness. **(C)**: Heat maps for mean size in materials with AI designed size and stiffness. **(D)**: Heat maps for mean stiffness in materials with AI designed size and stiffness. **(E)**: Heat maps for mean size in materials with only AI designed size. **(F)**: Heat maps for mean stiffness in materials with only AI designed stiffness.

Furthermore, when comparing the solutions in the (+AI designed size, -AI designed stiffness) condition (Fig. 3C) and the (-AI designed size, +AI designed stiffness) condition (Fig. 3D), grain property symmetry is observable along the central longitudinal axis in all materials. This pattern holds for all logic gates, except for the XOR gate, which remains distinctly asymmetric. This divergence in XOR may indicate a necessary structural asymmetry for achieving this logic function.

Another pattern we observe is that when stiffness is controlled by AI (Fig. 3A-B,D), the grain located between the input nodes consistently displays a low stiffness. The XOR in the (-AI designed size, +AI designed stiffness) (Fig. 3D) condition is the sole exception. This anomaly might stem from an alternative stress distribution pattern required to achieve the XOR functionality.

General Anatomy of a Logic Gate. Figure 4A-F presents heatmaps depicting the mean stiffness and size of individual grains extracted from the top-performing material across each of our thirty evolutionary runs. Each grain's color represents its relative mean stiffness or size, while the edge width indicates the standard deviation. Similar to Fig. 3, several interesting patterns have emerged across different input bit abstractions, AI design conditions, and logic gates. We proceed with enumerating a few.

The variation in mean stiffness, size, and standard deviation across different conditions highlights the complex ways in which the level of AI design control and choice of input bit abstractions shape the functionality of the granular materials. For instance, in the (+AI designed size, +AI designed stiffness) condition for bit abstraction 1 (Fig. 4A,B), the stiffness distributions for each logic gate appear to share a similar pattern, while the mean sizes display similarities between the AND and OR gates, and likewise between the XOR and NAND gates. This statement is correct for both the mean values and standard deviations. Similarly, under the (+AI designed size, -AI designed stiffness) condition (Fig. 4E), the AND and OR gates exhibit comparable configurations, while the XOR and NAND exhibit comparable configurations.

In the (-AI designed size, +AI designed stiffness) condition (Fig. 4F), the mean stiffness values for several grains are maximized to the upper limit (10) for the AND, OR, and NAND gates. Interestingly, many of these high-stiffness grains exhibit very low standard deviation. We can also see in this panel that the grains horizontally adjacent to the output grain maintain a relatively high standard deviation across all gates. In the case of XOR, the overall standard deviation across grains is higher than in other gates, indicating divergent solutions to this particular logic gate.

The distinct patterns under each level of AI design control condition emphasize that different combinations AI design of size and stiffness, along with input bit abstractions, lead to divergent internal structures. This exemplified the nuanced interplay between material properties and their computational outcomes.

3 Physical Computational Materials

In this section, we proceed with the physical validation of a granular material evolved *in silico* to act as an AND gate.

3.1 Experimental Setup

To design and construct a physical implementation of a computational granular material, we created a simple hardware with six grain types. The grain types are differentiated by two stiffnesses and three sizes, shown in Fig. 5A. These grains are made by casting silicone elastomer in molds laser-cut from a 0.25 in thick acrylic sheet. The grain diameters are 1.00 in (blue), 1.02 in (green), and 1.04 in (red), which span the 4% size change analogous to the simulations. The grain's material dictates its stiffness, which either has an elastic modulus at 100% strain of 10 psi (light, Smooth-On Ecoflex 30) or 86 psi (dark, Smooth-On Dragon Skin 30). We constructed a rudimentary force-sensing box experimental setup for these grains, shown in Supplementary Fig. 1. The grain packing was rigidly restricted from out-of-plane motion on both faces. The exterior in-plane bounding box was laser-cut from a 0.25 in thick clear acrylic sheet, similar to the negative molds for the grains. We designed the box's inner dimension to account for the laser beam width and to rigidly constrain the packed configuration of the

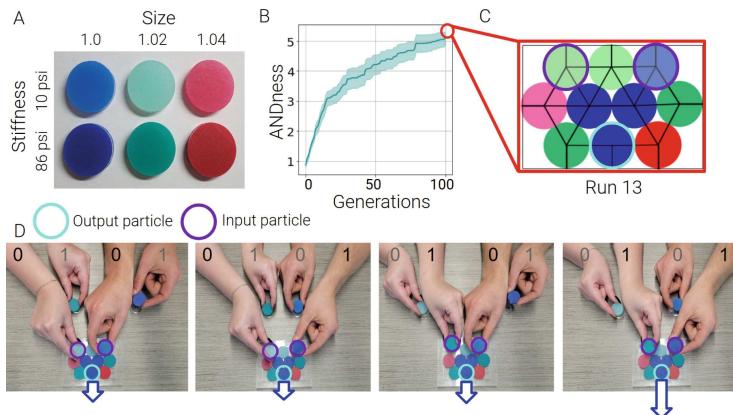


Fig. 5. Physical validation of an AND gate. (A): This panel shows the 6 grain types that we used to build the hardware implementation of our logic gates. Grains are available in 3 sizes and 2 stiffnesses. (B): The evolution of ANDness over generations across 30 replicates with 95% confidence interval. In these evolutionary runs, the packings are AI-designed for size and stiffness and the logical input control policy is guided by stiffness. (C): This panel shows the packing with the highest fitness value found by evolution. (D): This panel shows how the 4 input cases ['11', '10', '01', '00'] are passed to the material. Here, both individuals have a '0' in their right hand (soft grain) and a '1' in their left hand (stiff grain). Because this gate was evolved to behave as logical AND, we expect a high force on the output grain in the '11' input case, and a low force on the output grain in the '10', '01', and '00' input cases.

smallest grain size. The stiffness of cast acrylic is 10,000 psi, so we assume the 0.5 in wide walls are rigid. To measure force output, the wall in contact with the output grain was rigidly linked to the force sensor (Vernier Go Direct®Sensor Cart) and unattached from the rest of the wall. This section of the rigid wall was a T-shape to keep the piece from slipping out completely and was free to move 0.02 in perpendicular to the wall face. Force measurements were sampled at 50 Hz and averaged over a 2 s collection interval.

To transfer the optimized designs from simulation to the reduced-degree-of-freedom physical material, we re-execute evolution using the aforementioned experimental setup. However, we constrain the grain properties to discrete values. Specifically, we use the non-dimensional stiffness values $\in \{0.11627, 1\}$ and sizes $\in \{1.0, 1.02, 1.04\}$, which are based on the physical grain ratios. We conduct 30 independent replicates for each logic gate, under both input bit abstractions. Here, we only test the (+AI designed size, +AI designed stiffness) condition as the prior simulation results revealed its superiority. Figure 5B shows the fitness over time for 30 replicates for bit abstraction 2, where logic is passed to the material based on the stiffness of input grains. Figure 5C shows the highest fitness evolved design across all replicates at generation 100. Figure 5D shows how the different input cases are passed into the material. In this visual metaphor, there are two individuals, each of whom has a soft ('0') and stiff ('1') grain in

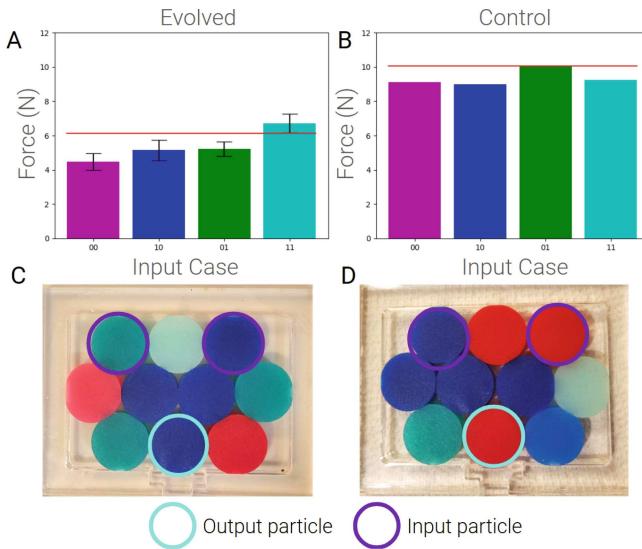


Fig. 6. Physical validation results. (A): Mean force recorded for each input case in the AND gate hardware implementation. The red line represents the magnitude of the threshold above which a force value is outputting a '1'. (B): Mean force recorded for each input case for a control configuration. (C): The configuration of the evolved AND gate when supplied with the '11' input case. (D): The configuration of the control material when supplied with the '11' input case.

their hands. The varying hands that are extended represent four different input cases passed into the material.

To account for variability in grain positioning and contact networks, we systematically 1) unpacked, 2) repacked, and 3) remeasured the force, 10 times for each grain configuration. Consequently, we can assess the robustness and stability of the material's functional output under minor perturbations in grain arrangement. As the control experiment, we measured five randomized packing configurations. Each grain was independently and randomly chosen from the six size and stiffness combinations, shown in Fig. 5A. Figure 6D shows one of the randomized configurations tested, and images of the others are available in the supplementary material. It is worth noting that each random material was measured once. All data collected is also available in the supplementary material.

3.2 Design Transfer

We selected the highest-performing configuration for AND, with the input logic coded in the grain stiffness, for hardware deployment. To evaluate the performance of the physical implementation, we packed the optimized configuration into the bounding box and took force measurements on the output grain for the input cases $\in \{00, 01, 10, 11\}$.

Figure 6A shows the mean force measured across each input case for the evolved configuration from Fig. 5B. Figure 6C shows the randomized reconfiguration of the evolved material. Figure 6B shows the mean force measured across each input case for the randomized configurations.

During evolution, we used a continuous metric termed “GATEness” for $\text{GATE} \in \{\text{AND}, \text{OR}, \text{NAND}, \text{XOR}\}$, to guide the optimization for materials that exhibit the desired logical behavior. However, post-evolution, this quantitative measure loses its meaning, as a gate either functions correctly or it does not. Therefore, we implement discrete thresholding to quantify the gate’s logical behavior. Specifically, we define an observed output force above a given threshold as a ‘1’, and an observed output force below a given threshold as a ‘0’. This approach is similar to that used in digital electronics [16]. We define a threshold based on the highest mean recorded for an input configuration that should produce a ‘0’ output. For example, in the physical instantiation of an AND gate, we observe a mean maximum force of 5.965 N where a ‘0’ output is expected. Consequently, this value serves as our threshold, meaning any output force greater than 5.965 N is interpreted as a ‘1’. For each randomized control, we set the threshold to the value of the force when the input is ‘00’. This value then becomes the control threshold for categorizing outputs.

Table 1 presents the observed frequencies of logical AND behavior across various materials, with the aforementioned universal threshold set for each material type. To evaluate whether the evolved configurations show a statistically higher proportion of AND behavior than random configurations, we use Fisher’s exact test. We receive a p-value < 0.005 , providing evidence that the evolved materials produce AND logic more frequently than their randomized counterparts.

Table 1. Evaluation of physical realization of the evolved and random configurations.

| Metric | Evolved Configuration | Random Configuration |
|------------------------|-----------------------|----------------------|
| % trials acting as AND | 90% | 0% |

4 Conclusions and Future Work

The shift away from traditional digital electronic computing will become increasingly salient in the coming years. To continue the pace of technological innovation, we must explore alternatives using novel computational substrates and paradigms. In this work, we introduced a new approach for using computational granular metamaterials as logic gates by leveraging static force chains that emerge within their internal contact network. By encoding binary inputs as physical properties like grain size and stiffness, our system translates these variations into force outputs, allowing the implementation of fundamental logic gates such as AND, OR, XOR, and NAND. We discovered that granular materials with AI design in both size and stiffness exhibit significantly more computational behavior compared to materials with only AI designed size or stiffness.

Our findings highlight the importance of material diversity in achieving emergent behaviors that are critical for logic operations. To ground our results in physical reality, we transferred an evolved AND gate configuration from simulation, demonstrating the first functional physical realization of a CGMM. Our sim2real transfer confirms the feasibility of using granular materials as next-generation computing technologies.

We plan to continue this research by improving upon our *in silico* design and optimization pipeline. Specifically, we hope to optimize our materials to display more complex computational behaviors. For example, one granular assembly could be designed to function as multiple logic gates simultaneously by using different bulk properties as inputs and outputs. Multiple gates could also operate simultaneously by using two sets of input grains and 2 output grains. Alternatively, one pair of input grains could have size and stiffness consecutively act as bit abstractions while outputs could be measured at different output grains, or different force angles on the same output grain (i.e. one in the x-direction and the other in the y-direction). We also plan to explore shape as an additional design parameter, allowing the evolution to optimize not only the size and stiffness of grains but also their geometry. It remains to be seen whether shape-varying CGMMs can further enhance the functionality of a material acting as a logic gate, but this direction holds promising potential for expanding the CGMM computational capacity.

Acknowledgements. We would like to thank Patrick Charron for lending the Vernier Cart used in the physical instantiation experiments. We would also like to thank Kameron Bielawski and Amanda Bertschinger for their assistance in creating Fig. 5. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 2235204. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. We would like to acknowledge financial support from the National Science Foundation under the DMREF program (award number: 2118810).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Beaulieu, S., Welch, P., Parsa, A., O'Hern, C., Kramer-Bottligio, R., Bongard, J.: Refractive Computation: parallelizing logic gates across driving frequencies in a mechanical polycomputer. In: ALIFE 2024: Proceedings of the 2024 Artificial Life Conference. MIT Press (2024)
2. Bitzek, E., Koskinen, P., Gähler, F., Moseler, M., Gumbsch, P.: Structural relaxation made simple. *Phys. Rev. Lett.* **97**(17), 170201 (2006)
3. Ciliz, D., O'Hern, C.S.: Granular metamaterials for soft robotic applications. *J. Next Front. Life Sci. AI* p. 48 (2021)
4. De Leon, N.P., et al.: Materials challenges and opportunities for quantum computing hardware. *Science* **372**(6539), eabb2823 (2021)

5. Fu, K., Zhao, Z., Jin, L.: Programmable granular metamaterials for reusable energy absorption. *Adv. Func. Mater.* **29**(32), 1901258 (2019)
6. Gantzounis, G., Serra-Garcia, M., Homma, K., Mendoza, J.M., Daraio, C.: Granular metamaterials for vibration mitigation. *J. Appl. Phys.* **114**(9) (2013)
7. Jayachandran, D., et al.: Three-dimensional integration of two-dimensional field-effect transistors. *Nature* **625**(7994), 276–281 (2024)
8. Kim, E., Yang, J.: Wave propagation in granular metamaterials. *Funct. Composites Struct.* **1**(1), 012002 (2019)
9. Li, M.S., Do, B.H., Le, C.L., O'Hern, C., Kramer-Bottiglio, R.: Variable stiffness and variable size particles for reconfigurable granular metamaterials. In: 2025 8th IEEE-RAS International Conference on Soft Robotics (RoboSoft) (2025)
10. McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S.C., Yuan, X.: Quantum computational chemistry. *Rev. Mod. Phys.* **92**(1), 015003 (2020)
11. Moore, G.E., et al.: Progress in Digital Integrated Electronics, vol. 21, pp. 11–13. Washington, DC (1975)
12. Parsa, A., Wang, D., O'Hern, C.S., Shattuck, M.D., Kramer-Bottiglio, R., Bongard, J.: Evolving programmable computational metamaterials. In: Proceedings of the Genetic and Evolutionary Computation Conference, pp. 122–129 (2022)
13. Parsa, A., Wang, D., O'Hern, C.S., Shattuck, M.D., Kramer-Bottiglio, R., Bongard, J.: Evolution of acoustic logic gates in granular metamaterials. In: International Conference on the Applications of Evolutionary Computation (Part of EvoStar), pp. 93–109. Springer (2022)
14. Parsa, A., Witthaus, S., Pashine, N., O'Hern, C., Kramer-Bottiglio, R., Bongard, J.: Universal mechanical polycomputation in granular matter. In: Proceedings of the Genetic and Evolutionary Computation Conference, pp. 193–201 (2023)
15. Pashine, N., et al.: Tessellated granular metamaterials with tunable elastic moduli. *Extreme Mech. Lett.* **63**, 102055 (2023)
16. Pedroni, V.A.: Digital electronics and design with VHDL. Morgan Kaufmann (2008)
17. Schmidt, M.D., Lipson, H.: Age-fitness pareto optimization. In: Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation, pp. 543–544 (2010)
18. Shrestha, A., Fang, H., Mei, Z., Rider, D.P., Wu, Q., Qiu, Q.: A survey on neuromorphic computing: models and hardware. *IEEE Circuits Syst. Mag.* **22**(2), 6–35 (2022)
19. Theis, T.N., Wong, H.: The End of Moore's Law: a new beginning for information technology. *Comput. Sci. Eng.* **19**(2), 41–50 (2017)
20. Tordesillas, A., Walker, D.M., Lin, Q.: Force cycles and force chains. *Phys. Rev. E-Stat. Nonlinear Soft Matter Phys.* **81**(1), 011302 (2010)
21. Wu, Q., Cui, C., Bertrand, T., Shattuck, M.D., O'Hern, C.S.: Active acoustic switches using two-dimensional granular crystals. *Phys. Rev. E* **99**(6), 062901 (2019)
22. Xia, A., Wang, D., Zhang, J., Shattuck, M.D., O'Hern, C.: Anisotropic shear response of 3D tessellated granular metamaterials. *Bull. Am. Phys. Soc.* (2024)
23. Yasuda, H., et al.: Mechanical computing. *Nature* **598**(7879), 39–48 (2021)