

Data-Efficient Discovery of Hyperelastic TPMS Metamaterials with Extreme Energy Dissipation

MAXINE PERRONI-SCHARF*, Massachusetts Institute of Technology, USA

ZACHARY FERGUSON*, Massachusetts Institute of Technology and CLO Virtual Fashion, USA

THOMAS BUTRUILLE, Massachusetts Institute of Technology, USA

CARLOS M. PORTELA, Massachusetts Institute of Technology, USA

MINA KONAKOVIĆ LUKOVIĆ, Massachusetts Institute of Technology, USA

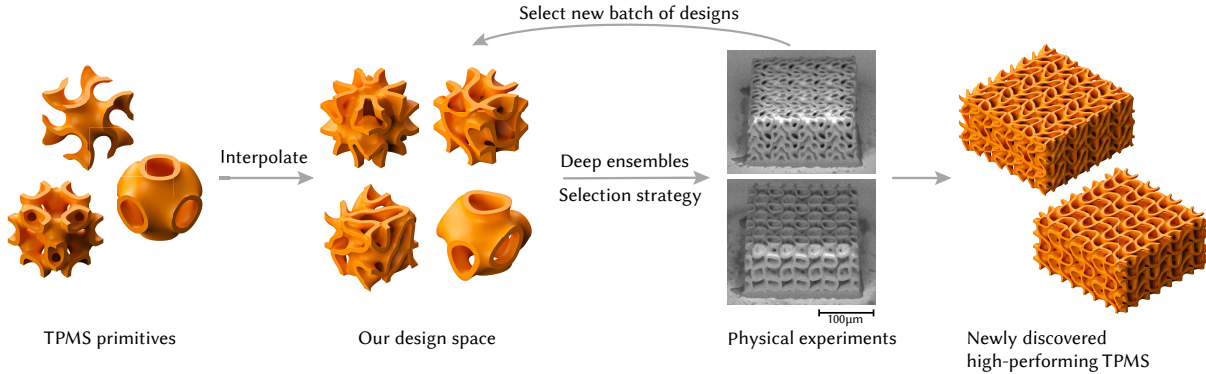


Fig. 1. Starting from a small set of well-studied TPMS metamaterial primitives, we use their implicit representations and interpolate between them to obtain a rich design space of TPMS structures exhibiting a wide variation in mechanical performance. To model the complex space of their behaviors, we physically test their performance and train Deep Ensembles that capture the uncertainty in prediction and use the results to inform our selection strategy for which designs to fabricate and test next. We iteratively select batches of designs for 3D-printing and testing to improve the model and discover new TPMS structures with higher energy dissipation than in known primitive structures.

Triply periodic minimal surfaces (TPMS) are a class of metamaterials with a variety of applications and well-known primitive morphologies. We present a new method for discovering novel microscale TPMS structures with exceptional energy-dissipation capabilities, achieving double the energy absorption of the best existing TPMS primitive structure. Our approach employs a parametric representation, allowing seamless interpolation between structures and representing a rich TPMS design space. As simulations are intractable for efficiently optimizing microscale hyperelastic structures, we propose a sample-efficient computational strategy for rapid discovery with limited empirical data from 3D-printed and tested samples that ensures high-fidelity results. We achieve this by leveraging a predictive uncertainty-aware Deep Ensembles model to identify which structures to fabricate and test next. We iteratively refine our model through batch Bayesian optimization, selecting structures for fabrication that maximize exploration of the performance space and exploitation of our energy-dissipation objective. Using our method, we produce the first open-source dataset of hyperelastic microscale TPMS structures, including a set of novel structures that demonstrate extreme energy dissipation capabilities, and show several potential applications of these structures.

*Joint first authors with equal contributions.



This work is licensed under a Creative Commons Attribution 4.0 International License.

SIGGRAPH Conference Papers '25, August 10–14, 2025, Vancouver, BC, Canada

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1540-2/2025/08

<https://doi.org/10.1145/3721238.3730759>

CCS Concepts: • **Computing methodologies** → **Machine learning**; • **Applied computing** → **Engineering**; *Computer-aided design*.

Additional Key Words and Phrases: 3D-Printed Metamaterials, TPMS Metamaterials, Sim-to-Real Gap

ACM Reference Format:

Maxine Perroni-Scharf, Zachary Ferguson, Thomas Butruille, Carlos M. Portela, and Mina Konaković Luković. 2025. Data-Efficient Discovery of Hyperelastic TPMS Metamaterials with Extreme Energy Dissipation. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25)*, August 10–14, 2025, Vancouver, BC, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3721238.3730759>

1 INTRODUCTION

Metamaterials are artificially engineered materials composed of three-dimensional microstructures, designed to exhibit properties that cannot be achieved by monolithic materials alone [Jiao et al. 2023; Xia et al. 2022]. These materials can achieve unique electromagnetic [Walser 2001], acoustic [Matlack et al. 2016], or mechanical characteristics [Bauer et al. 2017], unlocking innovative applications in a variety of domains. The study of metamaterials is an increasingly active area of research, significantly aided by advances in computer graphics and machine learning, which have enabled new possibilities for optimizing the design of 3D-printable metamaterial structures that exhibit unique or extreme physical properties. These properties include materials with negative Poisson ratios (auxetic

materials) [Shim et al. 2013], negative compressibility [Gatt and Grima 2008; Nicolaou and Motter 2012], ultra-high stiffness [Crook et al. 2020; Zheng et al. 2014], ultra-light weight [Schaedler et al. 2011], and programmable properties and mechanisms [Xia et al. 2022].

Identifying microstructures that can produce such properties poses a significant inverse design challenge. To address this challenge, recent research has increasingly relied on deep neural networks to automate the optimization of microstructures for desired physical responses [Li et al. 2023; Zhang et al. 2022]. A limitation of most of these approaches is that they rely on simulated data for model training—as generating sufficient physical data through 3D printing and testing is prohibitively expensive and time-consuming—and thus offer no guarantee that the resulting models accurately mirror real-world behaviors. As we show in our work, this simulation-to-reality gap is especially prevalent in settings where fabricated structures are very sensitive to the types of constituent materials and conditions in which they are tested. The limited number of studies that incorporate real-world data into training processes typically do not address strategies for efficient data acquisition through 3D-printing.

Our study addresses these challenges by applying a Bayesian framework to optimize hyperelastic triply periodic minimal surface (TPMS) metamaterials. Bayesian optimization iteratively suggests experiments to evaluate that are likely to yield the greatest performance improvement. Our optimization method is uncertainty-aware and based on real-world experimental data. We use Deep Ensembles (DEs) [Lakshminarayanan et al. 2017] to capture the uncertainty in predictions and model the complex space of TPMS behaviors. DEs are iteratively refined through a limited number of real-world experiments. The uncertainty measures allow us to target new regions of the parameter and performance spaces as we print subsequent batches of data. Furthermore, we exploit a specific objective for discovering structures with a high capacity to dissipate energy.

To obtain physical experimental data, we utilize sub-micron resolution 3D printing technology—termed two-photon lithography—to fabricate microscopic metamaterials, of approximately $200 \times 200 \times 100 \mu\text{m}^3$ in size. Fabricating samples through this technique enables high-throughput prototyping and subsequent characterization, while ensuring mechanical properties of the constituent polymer are not scale-dependent (i.e., no size effects are present). Given the widespread use of TPMS architectures in impact absorption applications [Zhang et al. 2023; Zhao et al. 2024], our goal is to discover structures with exceptional energy absorption capabilities when compressed by up to 60% of their original height. TPMS structures inherently possess a high surface area-to-volume ratio, are lightweight, and distribute stress more homogeneously than classical truss-based metamaterials. We choose to focus on large-deformation mechanical responses because the combination of high energy dissipation with lightweight properties can result in robust materials suitable for various applications, including bone-scaffolding and impact absorption.

Figure 2 illustrates three potential uses for our hyperelastic TPMS metamaterial structures. One is lightweight, breathable knee padding,

offering flexibility and reducing heat buildup, suitable for wearables and sports equipment. Another is patient-specific bone scaffolds, which could conform to complex anatomical shapes, support cell growth, and allow nutrient transport, building on prior studies [Dong and Zhao 2021]. Lastly, these structures could serve as impact-absorbing bumpers for robot vacuum cleaners, providing lightweight protection without compromising efficiency.

Our findings suggest broader implications for design workflows. While simulation-based pipelines are prevalent in metamaterial discovery [Lee et al. 2024; Li et al. 2023], they often struggle to capture dissipation-related behaviors driven by contact, buckling, and other nonlinear effects at the microscale. Recent work has addressed these challenges via differentiable simulators [Huang et al. 2024] or physics-informed models [Peng et al. 2023], but these approaches remain limited in capturing emergent behaviors such as jamming or folding. By integrating real-world testing into the optimization loop, our method complements such tools and offers a general strategy for discovering structures whose performance cannot be reliably predicted from simulation alone.

Our key contributions are:

- (1) The first method that utilizes Bayesian optimization to discover extremal hyperelastic TPMS metamaterial structures.
- (2) The first physically validated open-source microscale TPMS metamaterial dataset, which includes a set of novel structures with extreme energy dissipation capabilities.
- (3) A study on the application of traditional finite element method (FEM) simulations to microscale hyperelastic structures.

Figure 1 provides an overview of our method. We first discuss our TPMS design representation in Section 3. Next, we detail our experimental fabrication and testing setup in Section 4. We then address the limitations of simulations in matching experimental data and optimizing microscale TPMS structures in Section 4.1. Finally, in Section 5, we present our primary contribution: an optimization method for discovering novel TPMS structures with exceptional energy absorption capabilities.

2 RELATED WORKS

Metamaterial Design. Metamaterials have been found to have mechanical properties unprecedented in bulk material systems. Early work on metamaterials by Deshpande et al. [2001] and Ashby [2006] involving stochastic and periodic beam-based lattices explored the physical reasons supporting their ultra-high stiffness-to-weight ratios, launching broad explorations into how metamaterial design parameters, including relative density, constituent material, and topology, might affect mechanical performance, as categorized by Bauer et al. [2017] and Zhang et al. [2020] among others. Beyond exceptional stiffness up to the theoretical boundary for porous materials [Crook et al. 2020; Wang et al. 2022], metamaterial designs have achieved extraordinary mechanical properties such as negative Poisson’s ratios [Babaee et al. 2013], enhanced fracture toughness [Shaikeea et al. 2022], extreme compliance and resilience at large-strain deformation [Surjadi and Portela 2025], and finally optimized energy dissipation through the formation of compaction zones [Datelbaum et al. 2020; Hawreliak et al. 2016; Lind et al. 2019].

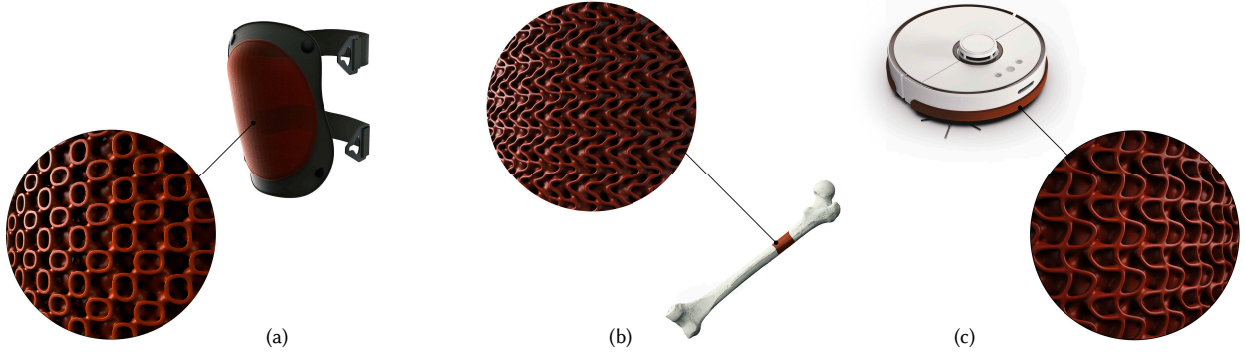


Fig. 2. Three potential applications for our discovered hyperelastic TPMS structures: (a) lightweight, thin, flexible and breathable knee padding, (b) a porous bone implant allowing for cell growth and infiltration, easy to print for patient-specific constraints, and (c) a thin bumper around a robotic vacuum cleaner, providing lightweight protection.

Optimizing energy dissipation in metamaterials for impact mitigation is of special interest due to the wide range of engineering applications where dynamic loading ultimately leads to failure. Deshpande and Fleck [2000] explored how the formation of compaction shocks in stochastic foams enhanced energy dissipation and how the stress state in the compaction shocks was related to the quasistatic behavior of these foams. Experimental studies on the dynamic response in stochastic [Barnes et al. 2014; Deshpande and Fleck 2000], periodic beam-based [Mines et al. 2013; Weeks and Ravichandran 2023] and shell-based [Novak et al. 2023; Tancogne-Dejean et al. 2019] metamaterials, including TPMS structures, have demonstrated the viability of these material design frameworks for the purpose of energy dissipation. However, most of these studies have achieved these results without fully taking advantage of advanced quantitative tools such as machine learning and neural architectures for optimizing metamaterial topologies inside these frameworks.

Several key studies in metamaterial design have made significant contributions without machine learning, including the seminal work in computer graphics of Bickel et al. [2010] introducing computational methods for designing custom materials with specific mechanical properties. Schumacher et al. [2015] developed a framework for creating microstructures that can be optimized to change properties across the given shape. Panetta et al. [2015] explored compliant mechanisms using metamaterials. Tozoni et al. [2020] proposed a parametric representation for flexible planar metamaterials. Martínez et al. [2016, 2018] used Voronoi Diagrams for novel metamaterial designs. Makatura et al. [2023] developed a unified procedural graph representation for cellular metamaterials, facilitating the creation of tileable structures with diverse properties.

Recent advances in metamaterial design have leveraged neural networks to develop encoded implicit parameterizations of metamaterials, primarily trained on simulated data. Lee et al. [2023] used a variational autoencoder to learn latent space representations for 2D pixelated metamaterials. Wang et al. [2020] and Zhang et al. [2022] both employed deep generative models to organize microstructures within structured latent spaces. However, such parameterization approaches provide no assurance that decoded metamaterials will be free from structural anomalies like floating islands, thus limiting their applicability for designing 3D validated datasets.

Several studies have tailored neural network architectures to enhance performance predictions for specific structural subclasses. These models typically simulate thousands of structures rapidly, but cannot ensure high fidelity in real-world fabrication scenarios. Bastek et al. [2022] inverted the property maps of truss structures based on extensive simulated data. Li et al. [2023] developed a neural network model to predict the behavior of extruded isohedral tilings, though this model does not account for out-of-plane behaviors or structures with three-dimensional variability. Fang and Zhan [2019] proposed deep physically-informed neural networks for the design of electromagnetic metamaterials. Meyer et al. [2022] investigated the use of graph neural networks to model complex dependencies within metamaterial structures. Oktay et al. [2023] introduced neuromechanical autoencoders that integrate a differentiable simulator with neural networks to model intricate two-dimensional geometries. Liu et al. [2024] developed programmable neural networks that adapt to various simulation environments.

Our research diverges from these works by integrating data from limited real-world experiments into a robust predictive model, enhancing both the real-world applicability of our designs and discovery efficiency. Unlike the approaches of Ha et al. [2023], who trained their models on a narrow range of strut-based microstructures from a small, 3D-printed dataset, our methodology incorporates a broader experimental base to improve prediction reliability. In alignment with this work, Thakolkaran et al. [2025] fabricate several hundred spinodal metamaterials for the training of a predictive model. However, unlike ours, their structures are not strategically selected and the model leverages physics-based inductive biases, which are challenging to define for many other classes of microstructure parameterizations.

Triply Periodic Minimal Surfaces (TPMS). There are many different subclasses of metamaterials [Jiao et al. 2023; Zadpoor et al. 2023], each of which has particular strengths in different applications. Here we focus on TPMS [Hu et al. 2020], tileable minimal surfaces that divide space into two continuous regions. TPMS structures cover a broad spectrum of physical properties and applications, from heat transfer [Attarzadeh et al. 2022; Fan et al. 2022] to tissue engineering [Feng et al. 2021], offering unique properties that are beneficial across various domains. We adopt the formulation from Al-Ketan

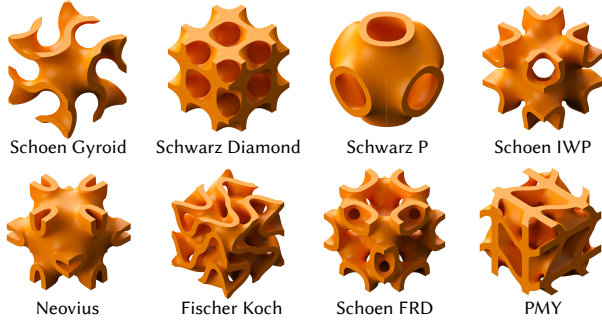


Fig. 3. Unit cells of TPMS primitives used in interpolation to represent the entire design space. Each of these structures shares a periodicity of 2π in each direction.

and Abu Al-Rub [2021], representing a vast TPMS metamaterial design space with eight parameters, enabling smooth interpolation between designs. To our knowledge, ours is the first study to efficiently discover novel TPMS metamaterials with extreme energy dissipation and to obtain insights via high-throughput prototyping and characterization of TPMS metamaterials with hyperelastic constitutive properties at the microscale.

Simulation of Metamaterials. Previous studies rely on simulated data for either building a library of homogenized material properties [Lee et al. 2024; Panetta et al. 2015; Schumacher et al. 2015; Tozoni et al. 2020] or training neural networks to predict metamaterial behavior [Li et al. 2023]. Because of the complexity of simulating full-scale structures, these works often employ *periodic homogenization* [Chen et al. 2021; Nakshatrala et al. 2013]. Furthermore, several studies focus entirely on 2D structures [Lee et al. 2024; Li et al. 2023; Tozoni et al. 2020], which are easier to simulate and fabricate but have limited real-world applicability. We focus on 3D structures, specifically TPMS, which are more challenging to simulate and fabricate but have a broader range of applications. The use of differentiable simulators for metamaterial design has been explored in several works [Lee et al. 2024; Li et al. 2022; Panetta et al. 2017, 2015; Tozoni et al. 2020]. These works utilize a variety of optimization techniques but are limited to specific classes of metamaterials and do not explore the design space of TPMS metamaterials.

Guided Data Collection for Material Discovery. When dealing with experimental design problems for which datasets are not available, collecting new data can be expensive and time-consuming. In this context, Bayesian optimization [Jones et al. 1998; Shahriari et al. 2016] has emerged as an effective method for guiding the search for an optimal solution in various applications. Erps et al. [2021] use Bayesian optimization to guide the discovery of better-performing 3D printing resins. Sharpe et al. [2018] explore the design of lattice structures via Bayesian optimization in simulation. Kuszczak et al. [2023] optimize the thickness of hexagonal honeycomb metamaterial lattices with Bayesian optimization. Zhang et al. [2022] provide a comparative study on Bayesian optimization for material design.

Recently, Snapp et al. [2024] employed a self-driving lab using Gaussian Processes (GPs) to discover high-performing FDM printed structures. While both studies focus on data-driven optimization of mechanical properties, we address a different domain by modeling

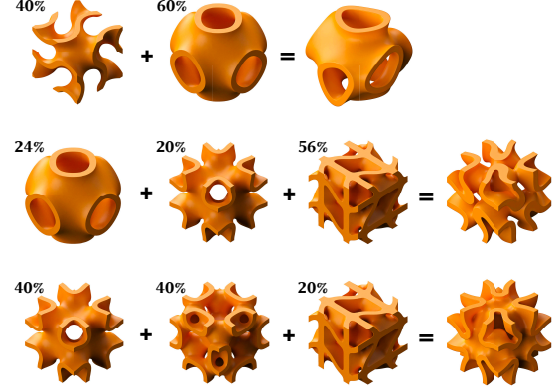


Fig. 4. Our designs are obtained as a weighted sum of eight TPMS primitives. Each design cell is fully parameterized with a set of weights. For example, the cell in the first row is a combination of 40% of Schoen Gyroid and 60% of Schwarz P.

hyperelastic microscale structures. Additionally, we utilize a neural network-based surrogate model in place of GPs. This contrast in modeling approaches and targeted applications offers complementary insights into the application of Bayesian optimization in structure design.

Recent advances in the mechanics and materials science communities further highlight the potential of such data-driven strategies. Serles et al. [2025] applied Bayesian optimization to identify carbon nanolattices with record-breaking specific strength, leveraging simulations and constrained experimental fabrication. Similarly, Peng et al. [2023] introduced a constrained multi-objective framework to explore trade-offs between mechanical performance and manufacturability in architected materials. These approaches demonstrate how machine learning and optimization can transform structural material discovery when coupled with domain-specific priors and physical experimentation.

Our work contributes to this growing space by targeting a distinct objective—maximizing energy dissipation in 3D-printed hyperelastic TPMS lattices—under experimental and computational constraints where high-fidelity simulations are infeasible. Unlike the above approaches, which typically use Gaussian Processes, we use Deep Ensembles to represent the surrogate model, enabling uncertainty-aware predictions in complex performance spaces. We also adapt the acquisition function for large batch selection under fabrication budget limits, demonstrating the versatility of Bayesian frameworks across diverse structural material domains.

3 TPMS DESIGN SPACE

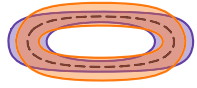
The most popular and well-studied TPMS structures, which we call TPMS primitives, are shown in Figure 3. To represent a wider range of TPMS structures, we model our design space after the one proposed by Al-Ketan and Abu Al-Rub [2021]. That is, we take these eight TPMS primitives each expressed as an implicit function (we refer to Al-Ketan and Abu Al-Rub [2021] for the definitions) then interpolate between them using barycentric coordinates:

$$F(x, y, z) := \sum_{i=1}^8 w_i f_i(x, y, z),$$

where w_i are the weights for each primitive, f_i is the corresponding implicit function, and F is the resulting interpolated structure’s implicit function. Examples of three such interpolations are shown in Figure 4. This ability to seamlessly interpolate between structures enables potential downstream applications where one could place structures from our dataset in various sections of a larger lattice and smoothly transition between them, creating spatially varying physical properties within a single object [Schumacher et al. 2015].

The zero level set of the implicit function represents the *sheet-networks* of the structure. To fabricate these structures, we need to assign a thickness to the sheets. This is commonly done by adding an offset $\pm \frac{1}{2}t$ to F [Al-Ketan and Abu Al-Rub 2021]:

$$\left(F + \frac{1}{2}t\right) \left(F - \frac{1}{2}t\right) = 0.$$



This offset, however, is not in standard units of distance and, as we show in purple in the inset figure, results in uneven wall thicknesses. To address this issue, we instead propose a spatially varying thickness and scale the offset by $1/\|\nabla F\|$ to ensure the walls have a uniform thickness (orange curves):

$$\left(F + \frac{1}{2} \frac{t}{\|\nabla F\|}\right) \left(F - \frac{1}{2} \frac{t}{\|\nabla F\|}\right) = 0.$$

We use a thickness of $t = 0.5 \mu\text{m}$ in all our experiments.

We mesh these structures using CGAL [The CGAL Project 2024] as a $4 \times 4 \times 2$ tiling with side lengths of $200 \times 200 \times 100 \mu\text{m}^3$. We choose CGAL over marching cubes because CGAL provides a tetrahedral mesh directly from our implicit function. These tetrahedral meshes are necessary for the finite element analysis in Section 4.1. Additionally, we use CGAL to produce periodic tetrahedral meshes for periodic simulations. We discard parametric combinations that produce cavities and keep only the largest disconnected component (i.e., discarding isolated material regions).

4 EXPERIMENTAL SETUP

4.1 Limitations of Simulation-Based Approaches

While prior work has successfully used simulated datasets—particularly with periodic boundary conditions—to train predictive models [Li et al. 2023; Panetta et al. 2015; Schumacher et al. 2015], the assumptions in those studies do not fully align with our experimental setup.

We tested both periodic simulations of single unit cells using PolyFEM [Schneider et al. 2019] and full-scale simulations of $4 \times 4 \times 2$ structures with rigid plate contact and substrate adhesion. However, calibrating simulation parameters—such as friction coefficients, geometric tolerances, and material properties—is particularly challenging due to inherent uncertainties introduced by our fabrication and development processes, causing a mismatch between simulated and real curves. This is further discussed in the supplementary material.

4.2 Fabrication

Benchmark, microscale experiments involving the uniaxial compression of 3D printed TPMS metamaterials are performed to capture

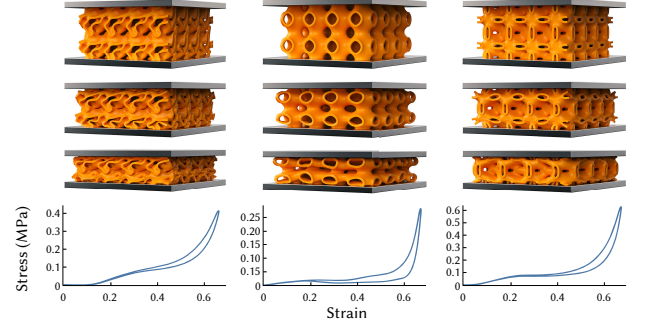


Fig. 5. Renders of three TPMS structures being compressed, along with their experimental stress-strain curves. Rows present 0, 25%, and 50% compression of each structure respectively.

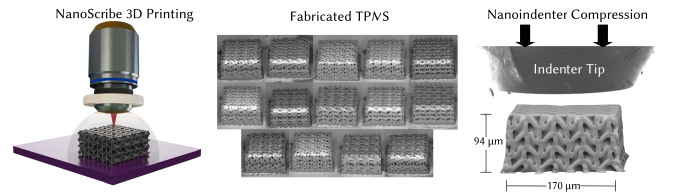


Fig. 6. Fourteen fabricated micro-TPMS structures printed on a Nanoscribe 3D printer and tested using an Alemnis nanoindenter.

sufficient data for Bayesian optimization of their properties, illustrated in Figure 5. Here, TPMS metamaterials are first fabricated out of a hyperelastic photosensitive resin (IP-PDMS) using two-photon lithography (Nanoscribe Professional GT2), where a femtosecond laser is focused through an optical objective and rastered in three dimensions to produce the TPMS architectures in a layer-by-layer manner, on top of a silicon substrate (Figure 6). To ensure high-fidelity prints, we employ a $25\times$ magnification objective with a slicing distance (direction normal to the substrate) set to $1.0 \mu\text{m}$ and a hatching distance of $0.5 \mu\text{m}$. The laser speed during printing is set to $13,500 \mu\text{m/s}$ using a laser power of 60 mW . A polymerization reaction occurs within a microscale ellipsoidal volume (the voxel) where the probability of multi-photon absorption is highest. This fabrication process differs from the more common fused deposition modeling (FDM) 3D printers in its ability to create structures with sub-micron features that lack interfaces between layers (due to voxel overlap between layers) and complex interior features and overhangs without the need for supports. The resulting 3D printed metamaterials are thus composed of an isotropic, hyperelastic polymer that exhibits properties similar to PDMS at the macroscale. It is important to note that polymers printed through two-photon lithography are isotropic and insensitive to print hatching direction, since polymerization occurs stochastically and slicing and hatching distance are chosen so each voxel pass overlaps with former voxel paths. After printing, the TPMS metamaterials are developed and the uncured resin washed away first with a 10-minute wash of isopropyl alcohol followed by a 1-minute rinse in clean isopropyl alcohol.

Each TPMS lattice is a $4 \times 4 \times 2$ tessellation of an individual cubic unit cell in x , y , and z directions respectively. The unit cells have a nominal side length of $50 \mu\text{m}$ and wall thickness of $4 \mu\text{m}$. The

as-fabricated dimensions of the TPMS lattices are measured with scanning electron microscopy (SEM) micrographs of gold-coated, uncompressed samples, having an average width and depth of 170 ± 2 μm and an average height of 94 ± 2 μm . These dimensions are used in calculating the stress-strain relations following uniaxial compression for each TPMS lattice sample. Prior to compression, every fabricated TPMS lattice sample is checked for print errors with an SEM to ensure experiments for each lattice are representative of their architectures.

4.3 Mechanical Characterization

Each metamaterial is uniaxially compressed *ex situ* using a displacement-controlled nanoindenter (Alemnis ASA) with a 400 μm diameter flat-punch tip. The compression is performed at a strain rate of 1 s^{-1} to a maximum strain of approximately 55%. At the beginning of each uniaxial compression experiment, the tip begins displacement 5 μm from a TPMS metamaterial sample such that it can accelerate to the desired tip velocity by the time contact is initiated with the sample. Some variation in maximum strain for each sample is expected due to slight variations in sample height and tip distance from sample. During uniaxial compression, displacement is recorded through the piezoelectric actuator that drives the compression, and force is measured through a load cell to which the tip is attached. The force-displacement data—along with the sample dimensions—are then used to calculate the uniaxial engineering stress-strain response for each sample.

We expect the majority of contributions to the stress response to be from material deformation. The marked increase in the stress response at high strain is caused by contact between walls within the TPMS samples, resulting in a compacted state. Frictional contributions at a state of high strain are difficult to quantify due to their highly nonlinear nature but may slightly hasten the onset of compaction by halting the relative movement between walls. Frictional energy dissipation from the TPMS walls sliding against each other is expected to be minimal in comparison to material deformation.

A minimum of two experiments for each sample architecture are performed to ensure measurement repeatability, with the experiments averaged together and denoised using a Savitzky-Galoy filter [Schafer 2011], from which 120 evenly spaced strain-stress values are selected to pass to the machine learning model for training and model performance verification.

5 OPTIMIZATION FRAMEWORK

Due to the intractability of simulations and the need for high-fidelity results, we use only 3D-printed and tested data. Our goal is to efficiently discover novel TPMS structures with superior energy dissipation, given fixed, limited 3D printing costs. Energy dissipation (MJ/m^3), measured by the area between the loading and unloading portions of a stress-strain curve (see inset), quantifies a structure's energy absorption ability.

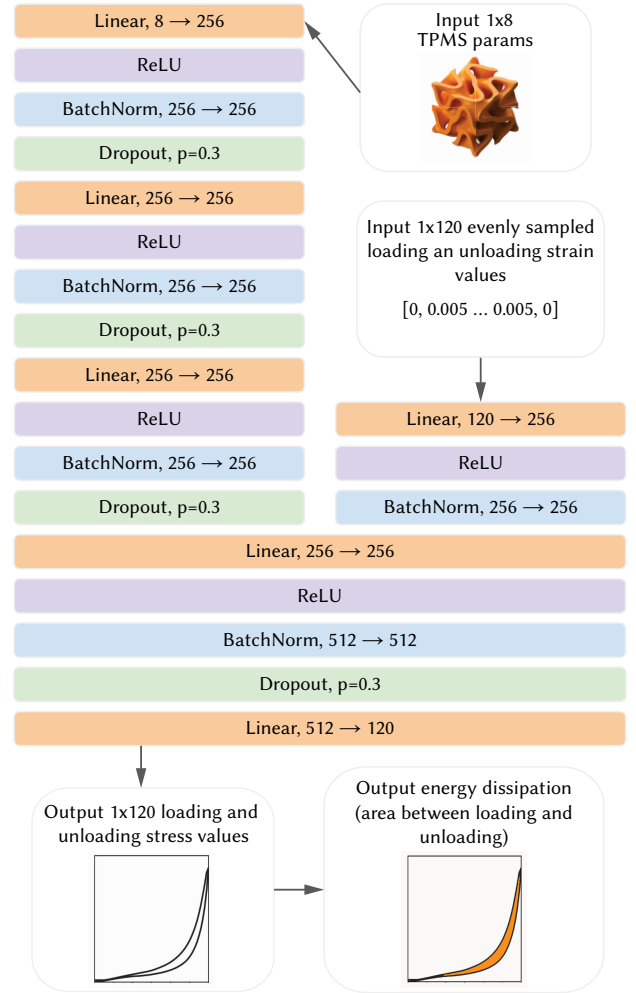
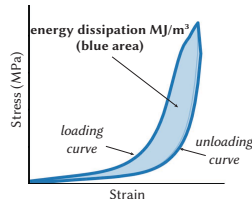


Fig. 7. Architecture of a single MLP model in the deep ensemble. One head takes in TPMS parameters, the other head takes in strain values. The output is a stress-strain curve, from which we can calculate energy dissipation.

To efficiently identify high-performing structures, we avoid the naive approach of uniformly or randomly sampling and testing points in the design space, which would be prohibitively costly. Instead, we propose a data-efficient method using an uncertainty-aware surrogate model to iteratively guide the sampling process.

We explore the design space by fabricating and testing batches of structures. The surrogate model then helps us decide on what structures to include in the next batch by balancing two objectives: exploring the design space to enhance accuracy and exploiting promising high-energy dissipation regions that we have already identified. This approach allows us to quickly and efficiently identify the best-performing designs, with our best structure exceeding the average uniform sample by more than a factor 20 and achieving twice the energy density of the best primitive TPMS structure.

Below, we outline the steps in our Bayesian optimization-based pipeline, from modeling the performance space to using our model to select the next batch to 3D-print and test.

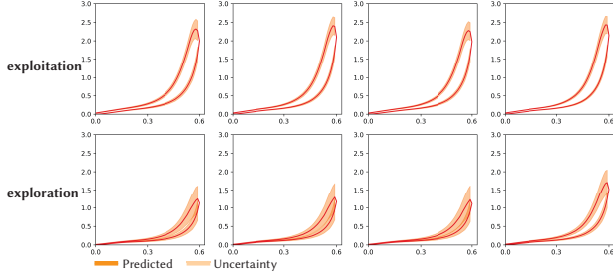


Fig. 8. Study on the effect of κ for batch acquisition. We choose a batch of size four using both $\kappa = 0$ and $\kappa = 10$. The top row corresponds to the top four scoring samples when $\kappa = 0$, whereas the bottom row is the top four selected points when $\kappa = 10$. Larger κ values encourage exploration by maximizing uncertainty, whereas smaller κ values encourage exploitation by maximizing the energy dissipation of our structures.

5.1 Neural Network Surrogate Model

Architecture. To capture the complex range of behaviors exhibited by our TPMS structures, we use Deep Ensembles [Lakshminarayanan et al. 2017] as a surrogate to predict the structure behavior based on their parameters, as prior research has demonstrated that ensembles outperform Gaussian Processes in terms of convergence time to optimal solutions and handling complex, non-convex performance spaces [Lim et al. 2021; Tian et al. 2024]. Our deep ensemble comprises 30 independently trained neural networks, enabling us to obtain both mean predictions and measures of prediction uncertainty that can inform our sampling strategy and improve accuracy.

A key challenge in Bayesian optimization is finalizing the model architecture without access to a large dataset, as collecting extensive experimental data upfront wastes fabrication resources and time. To address this, we used a proven two-branch architecture for predicting stress-strain curves in metamaterials [Li et al. 2023]. Preliminary tests on simulated structures and initial experimental data informed dropout values and hidden layer sizes, reducing overfitting on small datasets.

Each model in our ensemble is a dual-headed MLP (details in Figure 7) adapted from the neural architecture in [Li et al. 2023] to predict stress-strain behavior. One head processes the eight TPMS parameters, while the other handles an array of evenly applied strain values ranging from 0 to the maximum applied strain for each structure. Treating strain values as a model’s input rather than a constant is necessary due to slight variations in the maximum strain applied during testing. This ensures unbiased exploration when sampling for our next batch, preventing undue favoritism toward structures subjected to slightly extra strain.

As we are exploring the uncharted territory of the performance of microscale hyperelastic TPMS structures, we lack data to design a more complex model architecture that performs well with our specific type of structure. Therefore, we keep our model as agnostic as possible and exclude any inductive biases that might limit our model’s ability to explore and interpret this new performance space.

We first train our model on an initial batch of 23 structures. Given the very small size of this initial training set, we take careful steps to mitigate overfitting, which would hinder the algorithm’s ability to explore new regions of the performance space. To address this, we

add dropout [Srivastava et al. 2014] and batch normalization layers [Santurkar et al. 2018] to the model and implement early stopping at convergence during training [Caruana et al. 2000].

Training. During each phase of the optimization process, we train each network in our ensemble model using all available physically validated data, including input TPMS parameters and the full range of applied strain values, paired with the corresponding stress values obtained from experimental testing. In the inference phase, the model receives specific metamaterial parameters along with a pre-defined array of strain values, uniformly sampled from 0 to the maximum applied strain (approximately 0.6). Throughout the training phase, we use the Adam optimizer and the mean squared error loss function between predicted and real stress values. We use a learning rate of 0.001 and train each single multilayer perceptron for up to 2,000 epochs with early stopping patience of 100.

Inference. To obtain predicted stress-strain curves, each model M_i in the ensemble is used to predict an array of stress values S from input strain array E and structure parameters w . Our final prediction is then the mean of stress-strain predictions across all the N models in the ensemble, $\mu_S(w, E) = \frac{1}{N} \sum_i^N M_i(w, E)$. To quantify the uncertainty in stress, we calculate the variance across the predicted stress curves from all the ensemble models as $\sigma_S(w, E) = \frac{1}{N-1} \sum_i^N (M_i(w, E) - \mu_S(w, E))^2$.

For each structure, we analogously calculate the mean energy dissipation score D , our main outcome of interest, by quantifying the area enclosed between the loading and unloading portions of the stress-strain curve using the trapezoidal rule, which we denote as f_d . The average and uncertainty in energy dissipation are then assessed via the mean, $\mu_D(w, E) = \frac{1}{N} \sum_i^N f_d(M_i(w, E))$, of this energy dissipation area across all predictions and the corresponding variance, $\sigma_D(w, E) = \frac{1}{N-1} \sum_i^N (f_d(M_i(w, E)) - \mu_D(w, E))^2$.

5.2 Energy-dissipation Optimization

Overview. Our optimization process works as follows. We first acquire a batch of uniformly selected structures in design space, and 3D print and test these structures to get our first training dataset. Our initial batch is uniformly selected, and does not include the TPMS primitives, as is standard in Bayesian optimization to prevent biases in our exploration. We train our neural network and query it with an *acquisition function* to tell us which structures to 3D-print and test next. We then train the network on all available data once again, and repeat the process, iteratively acquiring a total of ten batches of approximately 25 TPMS structures at a time. Each time we query the network for the next batch, we aim to strike a balance between exploring the TPMS performance space (which is more heavily emphasized in our first three batches) and maximizing energy dissipation (which we focus on in the last three batches). For batch selection, we follow the Bayesian-optimization pipeline [González et al. 2016; Shahriari et al. 2016].

Batch Selection. We begin by choosing 1,000,000 new metamaterial parameter configurations utilizing the Dirichlet distribution [Lin 2016] with $\alpha = 1$ to ensure all parameters sum to unity and are sampled uniformly across the 8-dimensional hyper-simplex. We then apply the surrogate model to each set of parameters to predict

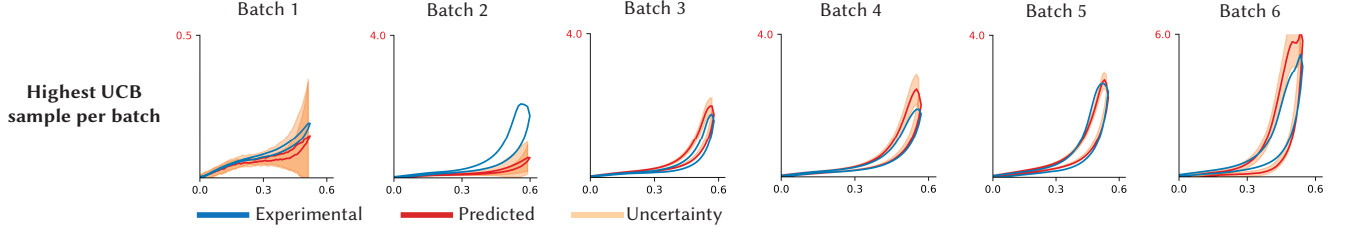


Fig. 9. We present the experimental (blue), predicted (red), and uncertainty (light orange) stress-strain curves for the highest scoring UCB sample in successive batches. This sample represents the optimal choice suggested by our optimization algorithm for each batch. The predicted curves are generated by a surrogate model trained on all data up to but excluding the current batch. It is important to note that the y -axis scale of the plots increases to accommodate the rise in maximum stress as optimization progresses. Initially, our algorithm selects points with high uncertainty to encourage exploration, transitioning to more exploitative and optimistic points in the later stages.

stress-strain curves, energy dissipation scores, and associated uncertainty scores. From these 1,000,000 potential designs, we iteratively select the top 40 structures based on their performance as evaluated by the *acquisition function* (details below) and, of these, we fabricate and test an average of 25 structures per batch (as many as time allows per 3D printing slot).

Acquisition Function. Our acquisition function informs our selection of the structures to be printed and tested based on a specified objective: either exploration of design space or exploitation of our energy-dissipation performance goal. We base our acquisition on the Upper Confidence Bound Acquisition Function [Srinivas et al. 2010], whereby each sample is scored according to a weighted sum of the uncertainty scores and the mean objective scores:

$$\text{UCB}(\mathbf{w}) = \mu(\mathbf{w}) + \kappa\sigma(\mathbf{w}),$$

where $\mu(\mathbf{w})$ is the predicted mean energy dissipation from our model for parameters \mathbf{w} , $\sigma(\mathbf{w})$ is the variance energy dissipation, and κ is a trade-off parameter between exploration and exploitation.

We introduce a localization penalty to reduce excessive clustering when selecting new sample points. During batch selection, we consider a sphere with radius r around all previously selected points from earlier and current batches. If a new sample falls within this radius, it is disallowed. Departing from the theoretical localization penalty suggested in [González et al. 2016], which in our case would give $r \approx 0.01$, we use a larger radius of $r = 0.2$, as our limited number of larger batches requires a stronger deterrent against clustering.

We evaluate $\text{UCB}(\mathbf{w})$ for all \mathbf{w} out of our 1,000,000 samples, iteratively selecting structures for our batch by choosing the one with the highest UCB score that is at least $r = 0.2$ away from all other selected structures in previous and current batches. This process continues until we acquire 40 points to 3D print and test in priority order. We repeat this process for all the batches of data.

To balance exploration and exploitation, we adjust the exploration parameter κ across batches, starting at $\kappa = 2$ for batches 2 and 3 to encourage exploration, progressively reducing it to 1, 0.75, 0.5, and finally 0 for the fourth to tenth batches to focus on optimizing energy dissipation. As shown in Figure 8, high κ values favor uncertainty, while low κ values favor energy dissipation.

Figure 9 shows batch-wise experimental and predicted stress-strain curves, where predictions are made by a model trained on all data up to but not including the current batch. As the process moves

Table 1. Energy Dissipation Values (kJ/m³) by Batch

Batch	Minimum	Maximum	Mean	Median
primitives (baseline)	6.80	165.96	54.13	15.51
1 (initial samples)	5.81	45.77	18.81	17.92
2	4.74	44.89	15.60	12.63
3	13.14	218.99	63.70	41.25
4	16.48	175.21	82.64	76.12
5	24.88	196.87	81.65	77.59
6	49.78	312.94	138.91	123.7
7	54.96	339.18	182.30	179.69
8	104.51	226.45	168.76	170.36
9	56.29	320.47	151.88	132.99
10	92.74	435.98	232.39	239.55

from exploration to exploitation, the uncertainty in the selected samples decreases, following the typical behavior of batch Bayesian optimization. This shift is also reflected in the improved predictive accuracy with each successive batch, as seen in the figure.

6 MATERIAL DISCOVERY RESULTS

Material Discovery. Our optimization effectively discovers novel TPMS structures with significantly enhanced energy dissipation. Figure 10 shows the best and worst performing structures from each batch, with experimental stress-strain curves for loading and unloading phases. Our goal is to maximize the area between these curves to optimize energy dissipation. Both the best and worst structures improve throughout the optimization.

Table 1 reports the minimum, maximum, median, and mean energy dissipation of each batch. The initial uniformly sampled batch had a mean energy dissipation of 18.81 kJ/m³, while the final optimized batch achieved a mean of 232.39 kJ/m³. Our approach identifies structures with energy dissipation twenty times higher than the average for randomly selected structures.

We identified 70 structures with energy dissipation exceeding 165.96 kJ/m³, outperforming the best primitive, Fischer-Koch. These materials often resemble Fischer-Koch but include significant proportions of other primitives, enhancing their performance.

Dataset Analysis. Our final hyperelastic TPMS dataset includes ten batches and TPMS primitives, labeled with stress-strain curves. Energy dissipation ranges from 4.74 kJ/m³ to 435.98 kJ/m³, over

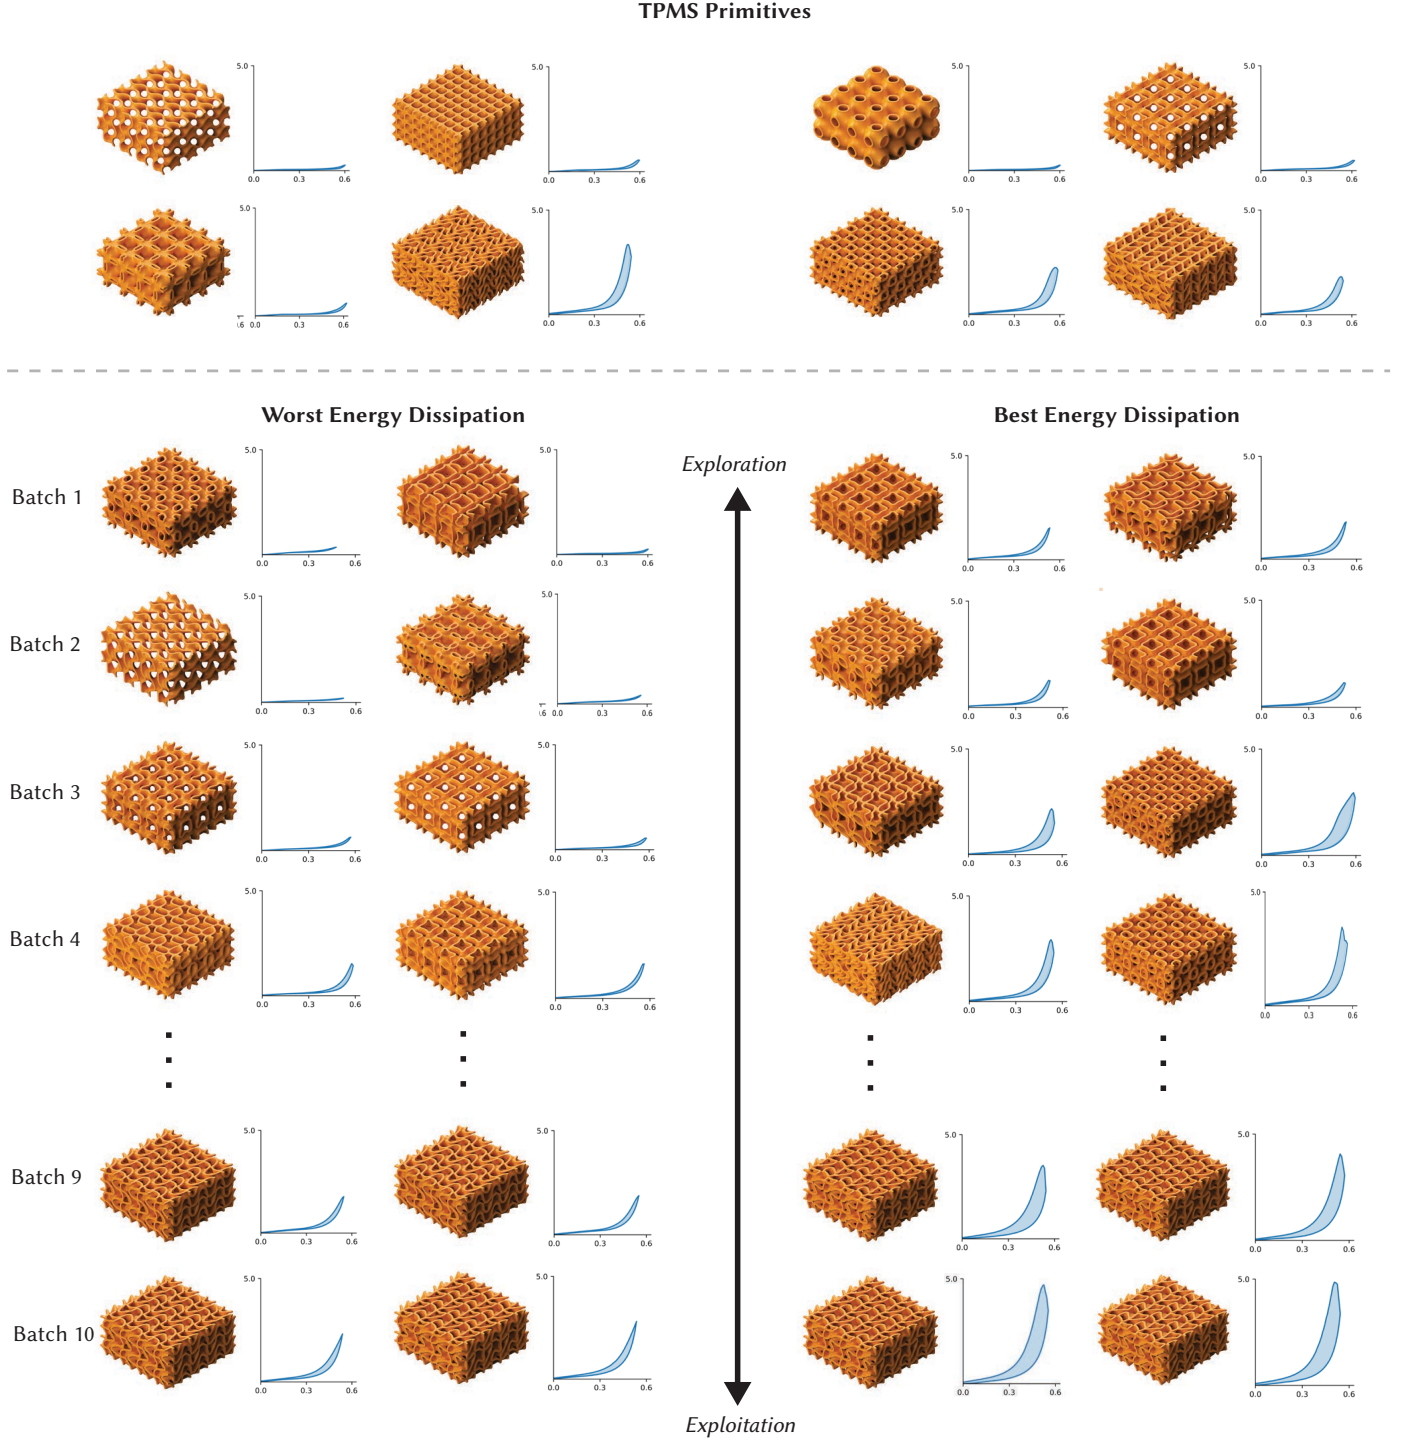


Fig. 10. We plot the experimental strain (x -axis) and stress (y -axis, MPa) curves for the TPMS primitive structures (top) and our batched data (bottom). We display the two samples for each batch with the worst (left two columns) and best (right two columns) energy dissipation (area between loading and unloading stress-strain curves, kJ/m^3 , highlighted in light blue). We can see the effects of our optimization procedure on the batched data. As we transition from exploration in batch 2 towards exploitation in batch 10, both the best and worst performing structures of each batch improve, and we are ultimately able to find structures with energy dissipation significantly higher than the initial uniform samples and the TPMS primitives. The highest-energy dissipation structure is that at the bottom right of the plot, with energy dissipation of 435.98 kJ/m^3 .

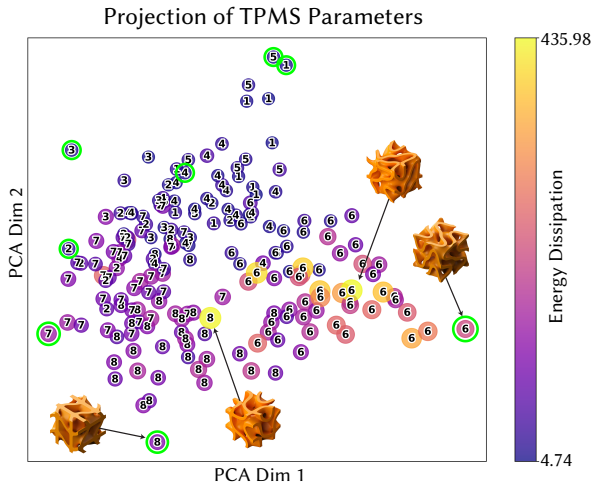


Fig. 11. Principal component analysis projection of metamaterial parameters, with color indicating energy dissipation and labels showing the highest-percentage TPMS primitive for each structure. Green circles highlight the single TPMS primitives. Structures with the highest percentages of primitives 6 and 8 exhibit the highest energy dissipation, surpassing the performance of those primitives on their own. We highlight a few of our highest-performing mixture structures and their corresponding highest percentage TPMS primitives.

double that of the energy dissipation of Fischer-Koch, which has an energy density of 165.96 kJ/m^3 .

We use Principal Component Analysis (PCA) [Abdi and Williams 2010] to visualize the dataset's performance, as shown in Figure 11. Each circle's color and size represent energy dissipation, with numerical labels indicating the predominant primitive. Structures 6 and 8 account for most high energy-dissipation instances, while primitives 1, 4, and 5 generally have lower energy dissipation.

Applications. TPMS structures have a wide range of potential applications due to their unique properties. They are lightweight, porous, flexible, and, as we have shown, can be designed to have high energy dissipation. We illustrate three potential applications of our discovered hyperelastic TPMS structures in Figure 2: a knee pad, a bone filling, and a robot vacuum cleaner bumper.

7 CONCLUSION

Limitations and Future Work. Our method has limitations worth exploring. While effective with TPMS structures, it could be extended to other metamaterials and objectives. Future work should develop a multi-class database of validated hyperelastic metamaterials. Using $4 \times 4 \times 2$ lattices may miss behaviors like aperiodic buckling in larger, nonuniform lattices. Additionally, the IP-PDMS photoresist used here, while hyperelastic, dissipates energy less effectively than metals or ceramics because it lacks deformation mechanisms such as plasticity and fracture. Testing optimized TPMS lattices made from stiffer hyperelastic materials could validate applicability for impact mitigation, focusing on material recoverability, peak stress reduction, and energy dissipation. Lastly, while we optimized our

model using efficient microscale fabrication techniques and experimental testing methods, to replicate our findings in a performance setting, future work should explore larger length scales and also identify the physical basis of the dissipation mechanisms.

Summary. In ten 3-hour 3D-printing sessions, our Bayesian optimization framework effectively discovered microscale TPMS structures that dissipate over twice the energy of the best-performing Fischer-Koch TPMS. We present the first physically validated microscale hyperelastic TPMS dataset, including 274 structures, 70 of which exhibit properties superior to those of the best primitive structure.

ACKNOWLEDGMENTS

We would like to thank Stefanie Mueller for her support throughout this work, Ane Zuñiga for her valuable insights, and Pavle Konaković for his help with modeling and rendering. We are also grateful to the MIT Morningside Academy for Design, the Andrew (1956) and Erna Viterbi Fellowship, and the Mathworks Fellowship for their support. Financial support from the DEVCOM Army Research Laboratory Army Research Office through the Massachusetts Institute of Technology (MIT) Institute for Soldier Nanotechnologies (ISN) under Cooperative Agreement Number W911NF-23-2-0121 is gratefully acknowledged. This work was carried out in part through the use of the MIT.nano facilities.

REFERENCES

- Hervé Abdi and Lynne J Williams. 2010. Principal component analysis. *Wiley Interdisciplinary Reviews: Computational Statistics* 2, 4 (2010), 433–459.
- Oraib Al-Ketan and Rashid K. Abu Al-Rub. 2021. MSLattice: A free software for generating uniform and graded lattices based on triply periodic minimal surfaces. *Material Design & Processing Communications* 3, 6 (2021), e205. e205 MDPC-2020-042.R1.
- M.F Ashby. 2006. The properties of foams and lattices. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 364, 1838 (Jan. 2006), 15–30. <https://doi.org/10.1098/rsta.2005.1678>
- Reza Attarzadeh, Seyed-Hosein Attarzadeh-Niaki, and Christophe Duwig. 2022. Multi-objective optimization of TPMS-based heat exchangers for low-temperature waste heat recovery. *Applied Thermal Engineering* 212 (2022), 118448.
- Sahab Babaee, Jongmin Shim, James C. Weaver, Elizabeth R. Chen, Nikita Patel, and Katia Bertoldi. 2013. 3D Soft Metamaterials with Negative Poisson's Ratio. *Advanced Materials* 25, 36 (Sept. 2013), 5044–5049. <https://doi.org/10.1002/adma.201301986>
- A.T. Barnes, K. Ravi-Chandar, S. Kyriakides, and S. Gaitanaros. 2014. Dynamic crushing of aluminum foams: Part I – Experiments. *International Journal of Solids and Structures* 51, 9 (May 2014), 1631–1645. <https://doi.org/10.1016/j.jis.2013.11.019>
- Jan-Hendrik Bastek, Siddhant Kumar, Bastian Telgen, Raphaël N Glaesener, and Dennis M Kochmann. 2022. Inverting the structure–property map of truss metamaterials by deep learning. *Proceedings of the National Academy of Sciences* 119, 1 (2022), e2111505119.
- Jens Bauer, Lucas R. Meza, Tobias A. Schaedler, Ruth Schwaiger, Xiaoyu Zheng, and Lorenzo Valdevit. 2017. Nanolattices: An emerging class of mechanical metamaterials. *Advanced Materials* 29, 40 (2017), 1701850. <https://doi.org/10.1002/adma.201701850> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.201701850>
- Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. 2010. Design and Fabrication of Materials with Desired Deformation Behavior. In *ACM SIGGRAPH 2010 Papers*. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/1833349.1778800>
- Rich Caruana, Steve Lawrence, and C Giles. 2000. Overfitting in neural nets: Backpropagation, conjugate gradient, and early stopping. *Advances in Neural Information Processing Systems* 13 (2000), 381–387.
- Tian Chen, Julian Panetta, Max Schnaubelt, and Mark Pauly. 2021. Bistable auxetic surface structures. *ACM Trans. Graph.* 40, 4, Article 39 (jul 2021), 9 pages.
- Cameron Crook, Jens Bauer, Anna Guell Izard, Cristine Santos de Oliveira, Juliana Martins de Souza e Silva, Jonathan B. Berger, and Lorenzo Valdevit. 2020. Platenan lattices at the theoretical limit of stiffness and strength. *Nature Communications* 11, 1 (2020), 1579. <https://doi.org/10.1038/s41467-020-15434-2>

- Dana M. Dattelbaum, Axinte Ionita, Brian M. Patterson, Brittany A. Branch, and Lindsey Kuettner. 2020. Shockwave dissipation by interface-dominated porous structures. *AIP Advances* 10, 7 (July 2020), 075016. <https://doi.org/10.1063/1.50015179>
- V.S. Deshpande, M.F. Ashby, and N.A. Fleck. 2001. Foam topology: Bending versus stretching dominated architectures. *Acta Materialia* 49, 6 (April 2001), 1035–1040. [https://doi.org/10.1016/S1359-6454\(00\)00379-7](https://doi.org/10.1016/S1359-6454(00)00379-7)
- V.S. Deshpande and N.A. Fleck. 2000. High strain rate compressive behaviour of aluminium alloy foams. *International Journal of Impact Engineering* 24, 3 (March 2000), 277–298. [https://doi.org/10.1016/S0734-743X\(99\)00153-0](https://doi.org/10.1016/S0734-743X(99)00153-0)
- Zhifei Dong and Xin Zhao. 2021. Application of TPMS structure in bone regeneration. *Engineered Regeneration* 2 (2021), 154–162.
- Timothy Erps, Michael Foshey, Mina Konaković Luković, Wan Shou, Hanns Hagen Goetzke, Herve Dietsch, Klaus Stoll, Bernhard von Vacano, and Wojciech Matusik. 2021. Accelerated discovery of 3D printing materials using data-driven multiobjective optimization. *Science Advances* 7, 42 (2021), eabf7435. <https://doi.org/10.1126/sciadv.abf7435> arXiv:<https://www.science.org/doi/pdf/10.1126/sciadv.abf7435>
- Zhaohui Fan, Renjing Gao, and Shutian Liu. 2022. Thermal conductivity enhancement and thermal saturation elimination designs of battery thermal management system for phase change materials based on triply periodic minimal surface. *Energy* 259 (2022), 125091.
- Zhiwei Fang and Justin Zhan. 2019. Deep physical informed neural networks for metamaterial design. *IEEE Access* 8 (2019), 24506–24513.
- Jiawei Feng, Bo Liu, Zhiwei Lin, and Jianzhong Fu. 2021. Isotropic porous structure design methods based on triply periodic minimal surfaces. *Materials & Design* 210 (2021), 110050.
- Ruben Gatt and Joseph N Grima. 2008. Negative compressibility. *physica status solidi (RRL)—Rapid Research Letters* 2, 5 (2008), 236–238.
- Javier González, Zhenwen Dai, Philipp Hennig, and Neil Lawrence. 2016. Batch Bayesian optimization via local penalization. In *Artificial Intelligence and Statistics*. PMLR, 648–657.
- Chan Soo Ha, Desheng Yao, Zhenpeng Xu, Chenang Liu, Han Liu, Daniel Elkins, Matthew Kile, Vikram Deshpande, Zhenyu Kong, Mathieu Bauchy, et al. 2023. Rapid inverse design of metamaterials based on prescribed mechanical behavior through machine learning. *Nature Communications* 14, 1 (2023), 5765.
- J. A. Hawreliak, J. Lind, B. Maddox, M. Barham, M. Messner, N. Barton, B. J. Jensen, and M. Kumar. 2016. Dynamic Behavior of Engineered Lattice Materials. *Scientific Reports* 6, 1 (June 2016), 28094. <https://doi.org/10.1038/srep28094>
- Jiangbei Hu, Shengfa Wang, Baojun Li, Fengqi Li, Zhongxuan Luo, and Ligang Liu. 2020. Efficient representation and optimization for TPMS-based porous structures. *IEEE Transactions on Visualization and Computer Graphics* 28, 7 (2020), 2615–2627.
- Zizhou Huang, Davi Colli Tozoni, Arvi Gjoka, Zachary Ferguson, Teseo Schneider, Daniele Panozzo, and Denis Zorin. 2024. Differentiable solver for time-dependent deformation problems with contact. *ACM Transactions on Graphics* 43, 3 (May 2024), 1–30. <https://doi.org/10.1145/3657648>
- Pengcheng Jiao, Jochen Mueller, Jordan R Raney, Xiaoyu Zheng, and Amir H Alavi. 2023. Mechanical metamaterials and beyond. *Nature Communications* 14, 1 (2023), 6004.
- Donald R Jones, Matthias Schonlau, and William J Welch. 1998. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization* 13, 4 (1998), 455.
- I. Kuszczak, F.I. Azam, M.A. Bessa, P.J. Tan, and F. Bosi. 2023. Bayesian optimisation of hexagonal honeycomb metamaterial. *Extreme Mechanics Letters* 64 (2023), 102078. <https://doi.org/10.1016/j.eml.2023.102078>
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. 2017. Simple and S scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems* 30 (NeurIPS 2017). Curran Associates, Inc., 6402–6413. https://proceedings.neurips.cc/paper_files/paper/2017/file/9ef2ed4b7fd2c810847ffa5fa85bce38-Paper.pdf
- Chan-Lock A. Lee, Zizhou Huang, Daniele Panozzo, and Denis Zorin. 2024. Computational design of flexible planar microstructures. *ACM Transactions on Graphics* 43, 6 (2024), 1–21. <https://doi.org/10.1145/3687765>
- Doksoo Lee, Yu-Chin Chan, Wei Chen, Liwei Wang, Anton van Beek, and Wei Chen. 2023. t-METASET: Task-aware acquisition of metamaterial datasets through diversity-based active learning. *Journal of Mechanical Design* 145, 3 (March 2023), 031704. <https://doi.org/10.1115/1.4055925>
- Weichen Li, Fengwen Wang, Ole Sigmund, and Xiaojia Shelly Zhang. 2022. Digital synthesis of free-form multimaterial structures for realization of arbitrary programmed mechanical responses. *Proceedings of the National Academy of Sciences* 119, 10 (2022), e2120563119. <https://doi.org/10.1073/pnas.2120563119> arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.2120563119>
- Yue Li, Stelian Coros, and Bernhard Thomaszewski. 2023. Neural metamaterial networks for nonlinear material design. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–13.
- Yee-Fun Lim, Chee Koon Ng, US Vaiteswar, and Kedar Hippalgaonkar. 2021. Extrapolative Bayesian optimization with Gaussian process and neural network ensemble surrogate models. *Advanced Intelligent Systems* 3, 11 (2021), 2100101.
- Jiayu Lin. 2016. On the Dirichlet distribution. *Department of Mathematics and Statistics, Queens University* 40 (2016).
- Jonathan Lind, Brian J. Jensen, Matthew Barham, and Mukul Kumar. 2019. In situ dynamic compression wave behavior in additively manufactured lattice materials. *Journal of Materials Research* 34, 1 (Jan. 2019), 2–19. <https://doi.org/10.1557/jmr.2018.351>
- Chenyang Liu, Xi Zhang, Jiahui Chang, You Lyu, Jianan Zhao, and Song Qiu. 2024. Programmable mechanical metamaterials: Basic concepts, types, construction strategies—a review. *Frontiers in Materials* 11 (2024), 1361408.
- Liane Makatura, Bohan Wang, Yi-Lu Chen, Bolei Deng, Chris Wojtan, Bernd Bickel, and Wojciech Matusik. 2023. Procedural metamaterials: A unified procedural graph for metamaterial design. *ACM Transactions on Graphics* 42, 5 (2023), 1–19.
- Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. 2016. Procedural Voronoi foams for additive manufacturing. *ACM Trans. Graph.* 35, 4, Article 44 (jul 2016), 12 pages. <https://doi.org/10.1145/2897824.2925922>
- Jonàs Martínez, Samuel Hornus, Haichuan Song, and Sylvain Lefebvre. 2018. Polyhedral Voronoi diagrams for additive manufacturing. *ACM Trans. Graph.* 37, 4, Article 129 (jul 2018), 15 pages. <https://doi.org/10.1145/3197517.3201343>
- Kathryn H. Matlack, Anton Bauhofer, Sebastian Krödel, Antonio Palermo, and Chiara Daraio. 2016. Composite 3D-printed metastructures for low-frequency and broadband vibration absorption. *Proceedings of the National Academy of Sciences* 113, 30 (2016), 8386–8390. <https://doi.org/10.1073/pnas.1600171113> arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.1600171113>
- Paul P Meyer, Colin Bonatti, Thomas Tancogne-Dejean, and Dirk Mohr. 2022. Graph-based metamaterials: Deep learning of structure-property relations. *Materials & Design* 223 (2022), 111175.
- R.A.W. Mines, S. Tsopanos, Y. Shen, R. Hasan, and S.T. McKown. 2013. Drop weight impact behaviour of sandwich panels with metallic micro lattice cores. *International Journal of Impact Engineering* 60 (Oct. 2013), 120–132. <https://doi.org/10.1016/j.ijimpeng.2013.04.007>
- Praveen Babu Nakshatrala, Daniel A Tortorelli, and KB3069875 Nakshatrala. 2013. Nonlinear structural design using multiscale topology optimization. Part I: Static formulation. *Computer Methods in Applied Mechanics and Engineering* 261 (2013), 167–176.
- Zachary G Nicolaou and Adilson E Motter. 2012. Mechanical metamaterials with negative compressibility transitions. *Nature Materials* 11, 7 (2012), 608–613.
- Nejc Novak, Oraib Al-Ketan, Anja Mauko, Yunus Emre Yilmaz, Lovre Krstulović-Opara, Shigeru Tanaka, Kazuyuki Hokamoto, Reza Rowshan, Rashid Abu Al-Rub, Matej Vesenjak, and Zoran Ren. 2023. Impact loading of additively manufactured metallic stochastic sheet-based cellular material. *International Journal of Impact Engineering* 174 (April 2023), 104527. <https://doi.org/10.1016/j.ijimpeng.2023.104527>
- Deniz Oktay, Mehran Mirramezani, Eder Medina, and Ryan P Adams. 2023. Neuro-mechanical autoencoders: Learning to couple elastic and neural network non-linearity. In *The Eleventh International Conference on Learning Representations*. https://openreview.net/forum?id=QubsmJT_A0
- Julian Panetta, Abtin Rahimian, and Denis Zorin. 2017. Worst-case stress relief for microstructures. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–16.
- Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. 2015. Elastic textures for additive fabrication. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–12.
- Bo Peng, Ye Wei, Yu Qin, Jiabao Dai, Yue Li, Aobo Liu, Yun Tian, Liulu Han, Yufeng Zheng, and Peng Wen. 2023. Machine learning-enabled constrained multi-objective design of architected materials. *Nature Communications* 14, 1 (2023), 6630.
- Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. 2018. How does batch normalization help optimization?. In *Advances in Neural Information Processing Systems* 31 (NeurIPS 2018). Curran Associates, Inc., 2488–2498. https://proceedings.neurips.cc/paper_files/paper/2018/file/90556c1c1a5d7c2c4f5f7f1e1c1e1e-Paper.pdf
- Tobias A Schaedler, Alan J Jacobsen, Anna Torrents, Adam E Sorensen, Jie Lian, Julia R Greer, Lorenzo Valdevit, and Wiliam B Carter. 2011. Ultralight metallic microlattices. *Science* 334, 6058 (2011), 962–965.
- Ronald W Schafer. 2011. What is a Savitzky-Golay filter? [Lecture notes]. *IEEE Signal Processing Magazine* 28, 4 (2011), 111–117.
- Teseo Schneider, Jérémie Dumas, Xifeng Gao, Denis Zorin, and Daniele Panozzo. 2019. PolyFEM. <https://polyfem.github.io/>
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to control elasticity in 3D printing. *ACM Trans. Graph.* 34, 4, Article 136 (jul 2015), 13 pages. <https://doi.org/10.1145/2766926>
- Peter Serles, Jinwook Yeo, Michel Haché, Pedro Guerra Demingos, Jonathan Kong, Pascal Kiefer, Somayajulu Dhulipala, Boran Kumral, Katherine Jia, Shuo Yang, et al. 2025. Ultrahigh specific strength by Bayesian optimization of carbon nanolattices. *Advanced Materials* 37, 14 (2025), 2570108.
- B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. 2016. Taking the human out of the loop: A review of Bayesian Optimization. *Proc. IEEE* 104, 1 (2016), 148–175.

- Angkur Jyoti Dipanka Shaikhe, Huachen Cui, Mark O'Masta, Xiaoyu Rayne Zheng, and Vikram Sudhir Deshpande. 2022. The toughness of mechanical metamaterials. *Nature Materials* 21, 3 (March 2022), 297–304. <https://doi.org/10.1038/s41563-021-01182-1>
- C. Sharpe, C. Seepersad, D. A. Tortorelli, and S. E. Watts. 2018. Design of mechanical metamaterials via constrained Bayesian optimization. In *Proceedings of the ASME 2018 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE 2018)*. <https://doi.org/10.1115/DETC2018-85270>
- Jongmin Shim, Sicong Shan, Andrej Košmrlj, Sung H. Kang, Elizabeth R. Chen, James C. Weaver, and Katia Bertoldi. 2013. Harnessing instabilities for design of soft reconfigurable auxetic/chiral materials. *Soft Matter* 9 (2013), 8198–8202. Issue 34. <https://doi.org/10.1039/C3SM51148K>
- Kelsey L. Snapp, Benjamin Verdier, Aldair E. Gongora, Samuel Silverman, Adedire D. Adesiji, Elise F. Morgan, Timothy J. Lawton, Emily Whiting, and Keith A. Brown. 2024. Superlative mechanical energy absorbing efficiency discovered through self-driving lab-human partnership. *Nature Communications* 15, 1 (2024), 4290.
- Niranjana Srinivas, Andreas Krause, Sham Kakade, and Matthias W. Seeger. 2010. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*. Omnipress, 1015–1022. <http://www.icml2010.org/papers/422.pdf>
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- James Utama Surjadi and Carlos M. Portela. 2025. Enabling three-dimensional architected materials across length scales and timescales. *Nature Materials* 24, 4 (April 2025), 493–505. <https://doi.org/10.1038/s41563-025-02119-8>
- T. Tancogne-Dejean, X. Li, M. Diamantopoulou, C. C. Roth, and D. Mohr. 2019. High strain rate response of additively-manufactured plate-lattices: Experiments and modeling. *Journal of Dynamic Behavior of Materials* 5, 3 (Sept. 2019), 361–375. <https://doi.org/10.1007/s40870-019-00219-6>
- Prakash Thakolkaran, Michael Espinal, Somayajulu Dhulipala, Siddhant Kumar, and Carlos M. Portela. 2025. Experiment-informed finite-strain inverse design of spinodal metamaterials. *Extreme Mechanics Letters* 74 (Jan. 2025), 102274. <https://doi.org/10.1016/j.eml.2024.102274>
- The CGAL Project. 2024. CGAL (5.6.1 ed.). <https://www.cgal.org/>
- Yunsheng Tian, Ane Zuniga, Xinwei Zhang, Johannes P. Dürholt, Payel Das, Jie Chen, Wojciech Matusik, and Mina Konaković Luković. 2024. Boundary exploration for Bayesian optimization with unknown physical constraints. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*. PMLR, 48295–48320. <https://arxiv.org/abs/2402.07692>
- Davi Colli Tozoni, Jérémie Dumas, Zhongshi Jiang, Julian Panetta, Daniele Panozzo, and Denis Zorin. 2020. A low-parametric rhombic microstructure family for irregular lattices. *ACM Trans. Graph.* 39, 4, Article 101 (aug 2020), 20 pages.
- Rodger M. Walser. 2001. Electromagnetic metamaterials. In *Complex Mediums II: Beyond Linear Isotropic Dielectrics*, Vol. 4467. SPIE, 1–15.
- Liwei Wang, Yu-Chin Chan, Faez Ahmed, Zhao Liu, Ping Zhu, and Wei Chen. 2020. Deep generative modeling for mechanistic-based learning and design of metamaterial systems. *Computer Methods in Applied Mechanics and Engineering* 372 (2020), 113377.
- Yujia Wang, Xuan Zhang, Zihe Li, Huajian Gao, and Xiaoyan Li. 2022. Achieving the theoretical limit of strength in shell-based carbon nanolattices. *Proceedings of the National Academy of Sciences* 119, 34 (Aug. 2022), e2119536119. <https://doi.org/10.1073/pnas.2119536119>
- J. S. Weeks and G. Ravichandran. 2023. Effect of topology on transient dynamic and shock response of polymeric lattice structures. *Journal of Dynamic Behavior of Materials* 9, 1 (March 2023), 44–64. <https://doi.org/10.1007/s40870-022-00359-2>
- Xiaoxing Xia, Christopher M. Spadaccini, and Julia R. Greer. 2022. Responsive materials architected in space and time. *Nature Reviews Materials* 7, 9 (2022), 683–701. <https://doi.org/10.1038/s41578-022-00450-z>
- Amir A. Zadpoor, Mohammad J. Mirzaali, Lorenzo Valdevit, and Jonathan B. Hopkins. 2023. Design, material, function, and fabrication of metamaterials. *APL Materials* 11, 2 (2023), 020401. <https://doi.org/10.1063/5.0144454>
- Hengrui Zhang, Wei Chen, Akshay Iyer, Daniel W. Apley, and Wei Chen. 2022. Uncertainty-aware mixed-variable machine learning for materials design. *Scientific Reports* 12, 1 (2022), 19760.
- Jing Zhang, Suchao Xie, Tao Li, Zinan Liu, Shiwei Zheng, and Hui Zhou. 2023. A study of multi-stage energy absorption characteristics of hybrid sheet TPMS lattices. *Thin-Walled Structures* 190 (2023), 110989.
- Xuan Zhang, Yujia Wang, Bin Ding, and Xiaoyan Li. 2020. Design, fabrication, and mechanics of 3D micro-/nanolattices. *Small* 16, 15 (2020), e1902842. <https://doi.org/10.1002/sml.201902842>
- Xuejin Zhao, Zhenzong Li, Yupeng Zou, and Xiaoyu Zhao. 2024. Compressive characteristics and energy absorption capacity of automobile energy-absorbing box with filled porous TPMS structures. *Applied Sciences* 14, 9 (2024), 3790.
- Xiaoyu Zheng, Howon Lee, Todd H. Weisgraber, Maxim Shusteff, Joshua DeOtte, Eric B. Duoss, Joshua D. Kuntz, Monika M. Biener, Qi Ge, Julie A. Jackson, Sergei O. Kucheyev, Nicholas X. Fang, and Christopher M.
- Spadaccini. 2014. Ultralight, ultrastiff mechanical metamaterials. *Science* 344, 6190 (2014), 1373–1377. <https://doi.org/10.1126/science.1252291> arXiv:<https://www.science.org/doi/pdf/10.1126/science.1252291>