

Supplement to 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

ACM Reference Format:

Maxine Perroni-Scharf, Zachary Ferguson, Thomas Butruille, Carlos M. Portela, and Mina Konaković Luković. 2025. Supplement to 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, 5 pages. <https://doi.org/10.1145/3721238.3730759>

1 CHALLENGES WITH SIMULATIONS

Previous studies in metamaterial design and exploration rely on simulated data for training neural networks to predict metamaterial behavior [Li et al. 2023]. These works have employed fast simulations with periodic boundary conditions and slower full-scale simulations with aperiodic boundary conditions to generate training datasets of metamaterial stress-strain behavior. To test the validity of this approach in our application, we performed extensive simulation tests to match simulations to reality. We also attempt to calibrate the simulation using the differentiable simulator of Huang et al. [2024]. However, due to inaccuracies and long simulation times, the process is intractable for modeling the proposed triply periodic minimal surface (TPMS) structures to large deformations to perform a high-fidelity exploration of the design space. We provide more insights below.

Periodic Boundary Conditions. A common approach for simulating and generating data for metamaterials is using periodic boundary conditions [Lee et al. 2024; Li et al. 2023; Panetta et al. 2015; Schumacher et al. 2015]. These impose that the displacement on one side (typically axis aligned) of the mesh should equal that on the opposite side, taking advantage of the periodicity of the TPMS shapes to generalize deformations in a single cell to a tiled pattern.

Using the open-source implementation of periodic boundary conditions in PolyFEM [Schneider et al. 2019], we simulate a single unit cell of a TPMS structure and apply a constant strain rate to it, measuring the normal stress as described by Li et al. [2023]. As expected, the deformations predicted by the periodic boundary conditions do not match the experimental data (see Figure S.1). The discrepancy

is due to the periodic boundary conditions not matching our experimental setup—where the finite TPMS structure is placed between two rigid plates and the top plate is displaced downwards. This results in contact forces between the TPMS structure and the rigid plates being unaccounted for any periodic boundary conditions, in addition to free-boundary effects not being captured. Additionally, simulations of a single unit cell fail to capture the aperiodic buckling or deformation localization across multiple cells (see Figure S.1).

Full-Scale Simulations. Because periodic boundary conditions do not match the experimental setup, we perform full-scale simulations with contact boundary conditions. Matching the fabricated structures, we place a $4 \times 4 \times 2$ tiling between two rigid plates (see Figure 8). The bottom plate is fixed, and the top plate is displaced downwards at a constant rate. Because the printed structures adhere to the bottom silicon substrate, we apply a zero-displacement boundary condition to the bottom surface of the structure. This is a good fit because the silicon substrate does not deform giving the loading conditions and we limit the compression to amounts that do not cause de-bonding. We then measure the contact force applied to the top plate by the metamaterial.

However, full-scale simulations are computationally expensive, as meshing these TPMS structures requires a large number of vertices. On average, a single $4 \times 4 \times 2$ TPMS with two tetrahedral elements in the structure’s thickness requires 123K–329K vertices with an average of 217K. By comparison, a single periodic cell requires 20K vertices on average. Furthermore, to accurately capture large deformation of the structure we need to handle nonlinear behavior and self-contact. Because of these constraints, full-scale simulations take on average ~ 14 hours to simulate a single structure, with the longest-running simulation taking over 43 hours, whereas running the periodic simulations takes an average ~ 2.3 hours (max: 6.6 hours). This difference in timing highlights why the majority of prior works build datasets from periodic rather than full-scale simulations.

Figure 8 shows the stress-strain curve for three full-scale simulations. The stress-strain curve’s linear regime (up to ~ 0.2 strain) matches well with the experimental data, but the nonlinear regime shows a large gap between the simulation and the experimental data. Note that this is after performing the differentiable calibration of bulk material. This gap is due to various factors, including material properties, mesh resolution, and geometry differences due to the fabrication process. Performing the differentiable calibration of the full-scale simulations to close the gap is intractable, as 20 iterations of material optimization would require up to 36 days of computation.

*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>

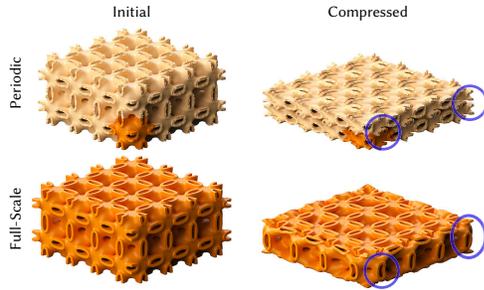


Fig. S.1. A comparison between periodic boundary conditions and full-scale simulation. The periodic simulation (top row) only simulates a single cell (shaded darker). However, this does not capture the correct dynamics (circled in blue) of the compression of a full-scale $4 \times 4 \times 2$ simulation (bottom row).

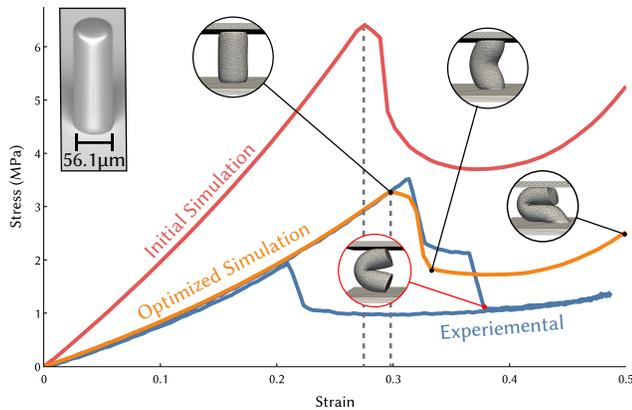


Fig. S.2. Results of material optimization on a micro-scale cylinder of bulk elastic material to determine a suitable set of parameters for use in simulation. Key points of the dynamics are highlighted. Note that the simulation strain-stress curve accuracy improves in both magnitude and time of buckling (grey dashed lines) as we optimize Young’s modulus and Poisson’s ratio. We see the stress increases linearly with strain until the cylinder buckles (around 0.3 strain) it plateaus for a small time until dipping further as the cylinder tears away from the silicon substrate base (as verified in simulation with free boundary conditions; circled in red). Because we simulate the bottom adhesion with a hard constraint of zero displacement, our simulation fails to capture this second dip in the stress.

We provide a comparison of periodic and full-scale simulations with experimental measurements on 42 different TPMS structures in Figure S.3.

1.1 Calibrating Simulations

Recent work has shown the promise of using differentiable simulation to reduce the gap between simulation and reality [Arnavaz and Erleben 2024; Huang et al. 2024]. We use the differentiable simulator of Huang et al. [2024] to predict a good set of material parameters for our TPMS structures. The results of this calibration are shown in Figure S.2.

We start by printing 16 cylinders of bulk IP-PDMS elastic material with a measured diameter of $56.1 \mu\text{m}$ and height of $153.9 \mu\text{m}$

after developing. We then measure the stress-strain curve of these cylinders experimentally and average the results to a single curve.

We model the printed material using a neo-Hookean material model. We find this to be the simplest model capable of capturing the nonlinear elasticity exhibited. Alternatively, using a more parameter-rich model (e.g., Ogden or Mooney-Rivlin) would increase the complexity of calibration.

Further, while the metamaterials exhibit a significant energy dissipation upon release, this was not found in the bulk material, which indicates that the property under study is uniquely characterized by the microstructure design and how it bulks and **dissipates energy through friction**. This also supports our use of a hyperelastic material model and not explicitly modeling visco-elasticity/plasticity.

We initialize the neo-Hookean material in our simulation using material parameters provided by the IP-PDMS dataset: Young’s modulus $E = 9.5 \text{ MPa}$ and Poisson’s ratio $\nu = 0.3$. As shown in Figure S.2, these material parameters do not accurately capture the behavior of the material.

To improve the accuracy of our simulation, we perform material optimization on the cylinder of bulk elastic material. We perform a joint optimization over E and ν simultaneously to minimize the integrated difference between the simulation and experimental stress-strain curves. Our optimization converges to a value of $E = 3.572 \text{ MPa}$ and $\nu = 0.475$ in 20 iterations. With these parameters, the simulation stress-strain curve closely matches the experimental data.

However, the curves do not match exactly, due to differences in how the adhesion between the cylinder and silicon substrate is modeled. We simulate the bottom adhesion with a zero-displacement boundary condition, while in reality the cylinder tears away from the silicon substrate base. This is verified in simulation with free boundary conditions. Further work is needed to accurately characterize and model this behavior to fully capture all aspects of the stress-strain curve.

We apply these parameters to the full-scale simulations of TPMS structures. While the differentiable calibration improves the accuracy of the linear regime, the nonlinear regime still shows a large gap between the simulation and experimental data. This is because the calibration is performed on a cylinder of bulk material, which does not capture the complex geometry and **frictional contact** of the TPMS structures. Calibrating for these effects requires simulating the TPMS structures themselves, but this would be computationally too costly as performing the differentiable calibration of the full-scale simulations would require up to 36 days of computation.

2 BEST AND WORST STRUCTURES OF EACH BATCH

We present the best and worst performing structures of each batch in Figure S.4.

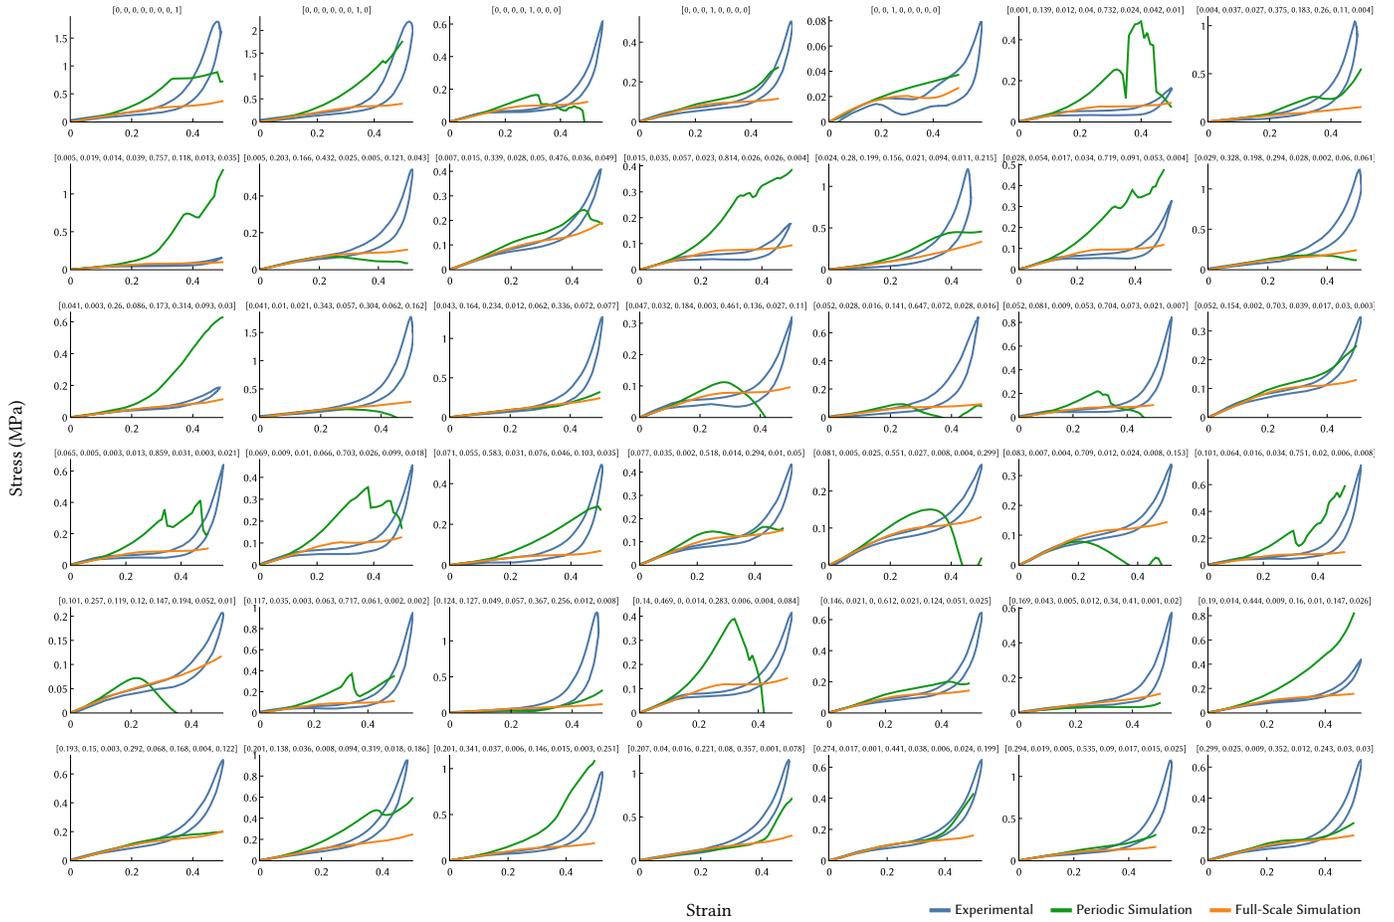


Fig. S.3. A comparison of our periodic simulation (green), full-scale simulation (orange) and experimental measurements (blue) on 42 different TPMS structures. The periodic simulations differ wildly from the experimental measurements. The full-scale simulations match well in the linear regime of the material (up to ~ 0.2), but diverge in the nonlinear response. The simulations exhibit more buckling and therefore lower stress response than the real TPMS.

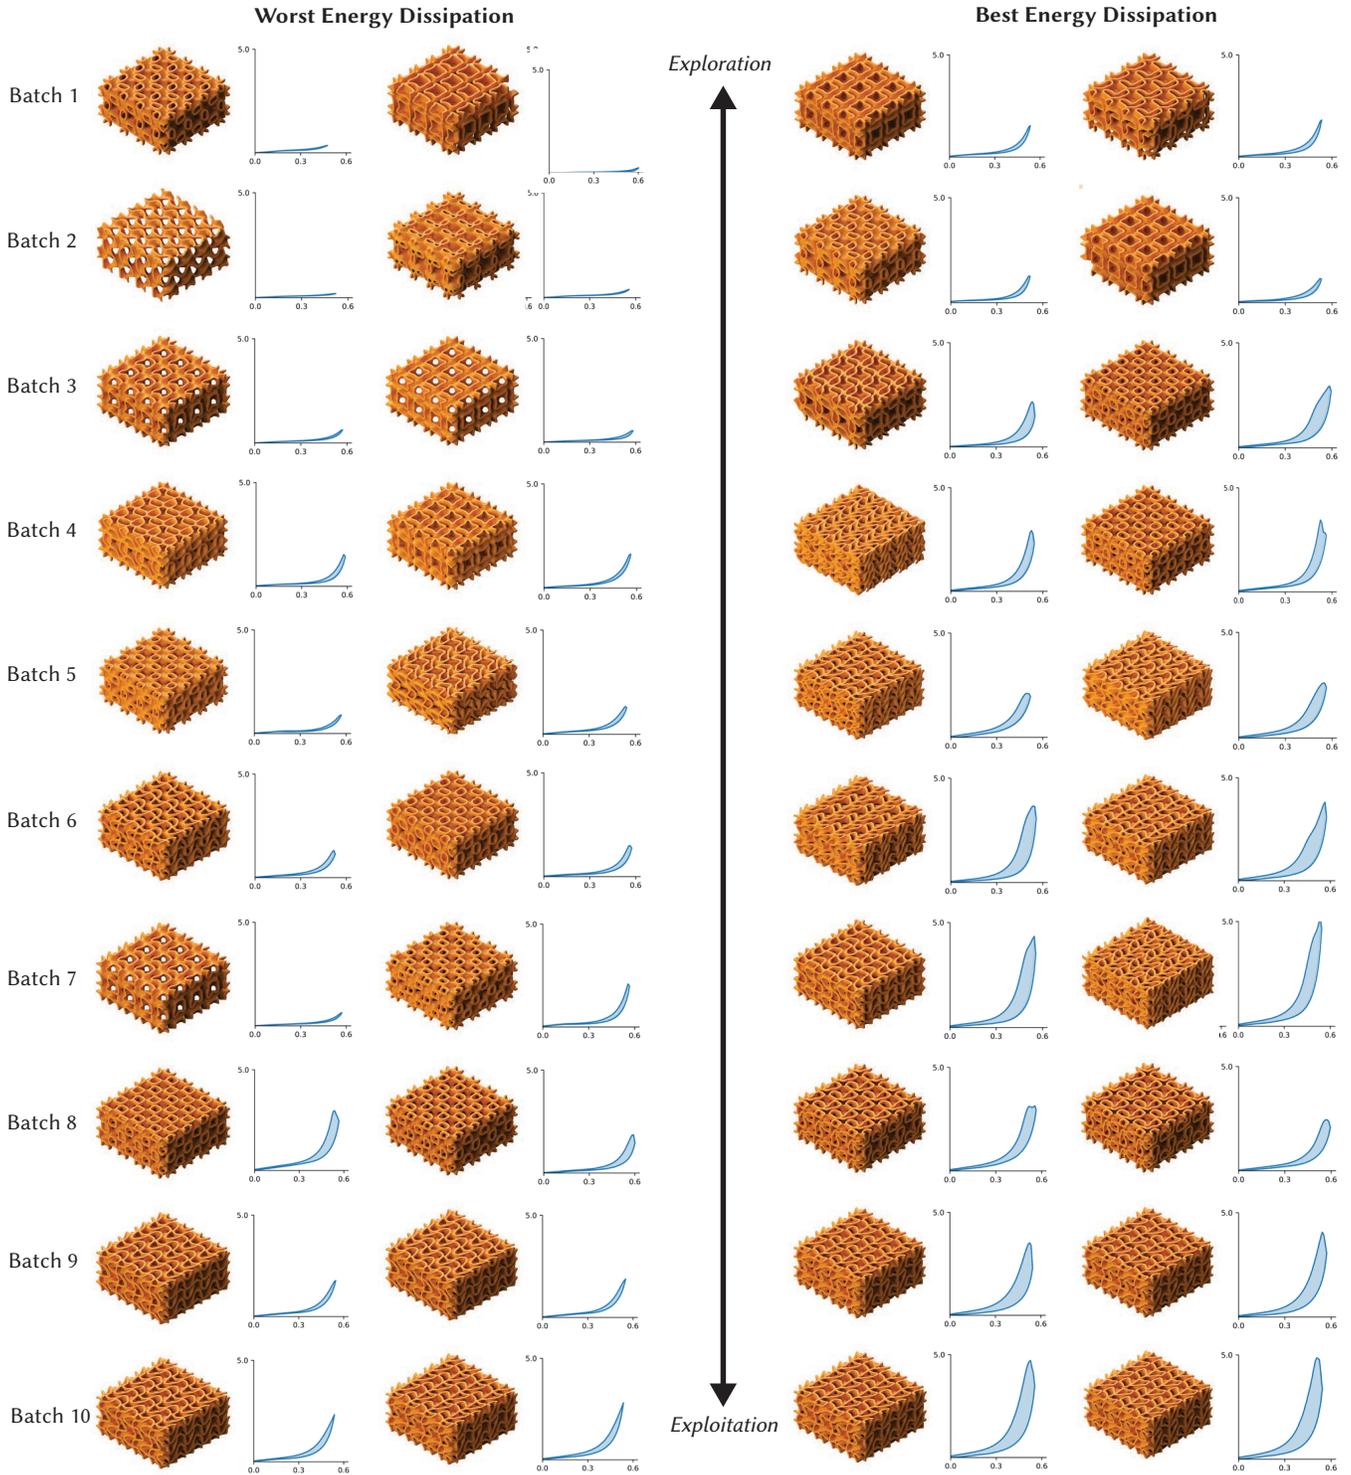


Fig. S.4. We plot the experimental strain (x-axis) and stress (y-axis, MPa) curves for all 10 iterative batches. We display the two samples for each batch with the worst (left two columns) and best (right two columns) energy dissipation (area between loading and unloading stress-strain curves, $\text{MJ}^{-100}/\text{m}^3$, 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. The highest-energy dissipation structure is that at the bottom right of the plot.

REFERENCES

- Kasra Arnavaz and Kenny Erleben. 2024. Differentiable Rendering as a Way to Program Cable-Driven Soft Robots. arXiv:2404.07590 [cs.RO]
- 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>
- 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>
- 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.
- 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.
- 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>