

2024

Discovering tough and impact-resistant structures using a self-driving lab

<https://hdl.handle.net/2144/49250>

"Downloaded from OpenBU. Boston University's institutional repository."

BOSTON UNIVERSITY
COLLEGE OF ENGINEERING

Dissertation

**DISCOVERING TOUGH AND IMPACT-RESISTANT
STRUCTURES USING A SELF-DRIVING LAB**

by

KELSEY LAWRENCE SNAPP

B.A., University of Oklahoma, 2009
J.D., The University of Texas at Austin, 2012
M.S., Boston University, 2023

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2024

Approved by

First Reader

Keith A. Brown, Ph.D.
Associate Professor of Mechanical Engineering
Associate Professor of Materials Science and Engineering
Associate Professor of Physics

Second Reader

Elise F. Morgan, Ph.D.
Dean *ad interim* of the College of Engineering
Maysarah K. Sukkar Professor of Engineering Design and Innovation
Professor of Mechanical Engineering
Professor of Materials Science and Engineering
Professor of Biomedical Engineering

Third Reader

Emma Lejeune, Ph.D.
Assistant Professor of Mechanical Engineering

Fourth Reader

Timothy J. Lawton, Ph.D.
Research Chemist
Combat Capabilities Development Command
U.S. Army Natick Soldier Research, Development, and Engineering
Center

DEDICATION

To my amazing wife, Marisa.

ACKNOWLEDGMENTS

First, I would like to thank my advisor, Professor Brown. His passion for science is contagious and his wide breadth of knowledge never ceases to amaze me. He has taught me how to analyze data critically, think outside of the box for creative solutions, and to pay attention to small details. Most importantly, I'm grateful for him introducing me to the world of self-driving labs, and I look forward to seeing how he will continue to drive forward that research technique in the years to come.

Secondly, I would like to thank each of my dissertation committee members for their amazing insight and feedback. Dr. Lawton has been a great strategic partner during the years of research that we have collaborated, and his guidance has been invaluable in the work presented here. Professor Lejeune graciously allowed me to do a rotation in her lab to learn more about finite element analysis, and her passion for open-source data influenced my approach to how we shared the BEAR's data and code. Finally, I would like to thank Professor Morgan for her insight into the mechanics underlying some of our high performing parts and for the fantastic feedback during both the prospectus and dissertation review phases.

In addition to committee members, I would like to thank additional faculty that had a profound impact on my work. I would like to thank Dr. Farny first as an amazing lecturer and second as a great supervisor during my graduate student teacher service. Having taken thermodynamics with him three times, I can say that his passion for teaching the next generation of engineers has inspired me to always give my best in whatever I do. Additionally, I would like to thank Professor Andersson who not only was

a great academic advisor for my first year at Boston University, but also did a fantastic job teaching me linear algebra. I would also like to thank him for serving on my qualification exam and prospectus committees.

Next, I would like to thank the excellent support staff at Boston University. Although it is impossible to individually name everyone who helps support the daily operation of the University, I would like to mention a few who were particularly helpful during my time here. I would like to thank Anna Masland and Emery Dutton for their help navigating the administrative requirements to graduate. I would like to thank Kara Mogensen for her help testing on Epic's heavy duty Instron and Sydney Holder for her help taking micro-CT scans. Last, I would like to thank Riley Barrar with the Office of the Ombuds for her tireless work supporting students and staff at Boston University.

I would like to thank all of the great collaborators that I have worked with over the years. In particular, I would like to thank Professor Whiting and her students Ben and Sam who have helped with the development of the python code that generates GCS designs and provided feedback on our work during countless meetings. I would also like to thank Professor Lipton of Northeastern and his two students Brett and Daniel for their work on viscous thread printing and their patience as we brought VTP to the BEAR.

I would also like to thank the amazing students who I have worked with over the years. Specifically, I would like to thank Dr. Gongora for mentoring me when I first joined the KABLlab and for developing the first version of the BEAR. I would like to thank Sebastian, Abin, Rashid, and Thomas for their help keeping the BEAR running during their respective summers in the lab. I would like to thank Neil Frings for his help

taking 4-point bending tests. Finally, I would like to thank all the KABLlab members that I have worked with over the years, including Dr. Alsharif, Dr. Rendos, Dr. Saygin, Harley, Yihong, Adedire, Dylan, and Jiashuo.

Last but not least, I would like to thank my family who have always supported and believed in me. I would like to thank my parents, Cheri and Steve, who have always supported me in my endeavors. I would like to thank my big brother Byron for hanging out with me when I was younger and being a great role model to me. Most importantly, I would like to thank my wife, Marisa. She encouraged me to pursue my dreams of going back to school to study engineering and has provided support during the arduous journey to complete my PhD. I'm looking forward to our next 68 years together!

**DISCOVERING TOUGH AND IMPACT-RESISTANT
STRUCTURES USING A SELF-DRIVING LAB**

KELSEY LAWRENCE SNAPP

Boston University, College of Engineering, 2024

Major Professor: Keith A. Brown, Associate Professor of Mechanical Engineering,
Associate Professor of Materials Science and Engineering, Associate
Professor of Physics

ABSTRACT

Humans depend on energy-absorbing structures constantly during daily life. Crumple zones in cars protect occupants during a crash. Packaging protects sensitive goods during delivery. Sports equipment protects athletes during both expected and unexpected impacts. Unfortunately, the development of energy absorbing structures is slow because phenomena such as high strain and self-self contacts make their performance difficult to model. Thus, physical experiments are often the only way to effectively evaluate the performance of a potential structure.

To tackle this problem, we employ a self-driving lab comprised of an automated experimentation system guided by machine learning. This system can in principle provide the acceleration in terms of experimental throughput and information per experiment needed to address the challenge of discovering high-performing structures. Initially, we developed a family of structures called generalized cylindrical shells that are reliable to additively manufacture and that include over trillions of unique designs. Next, we modified a self-driving lab and physically tested tens of thousands of structures in quasistatic compression. During a campaign lasting nearly two years, we discovered a

component with an energy absorbing efficiency of 75.2%, the highest ever reported. Furthermore, using seven different polymer filaments, we found designs that outperformed all prior synthetic structures in energy absorbing efficiency across a broad stress range from 100 Pa to 10 MPa. In addition to this technical outcome, because this campaign lasted orders of magnitude longer than previously published self-driving lab campaigns, we also uncovered insights about the interactions between researchers and self-driving labs, such as how to monitor the progress of the campaign and how to adjust the campaign when problems arise.

Accounting for impact events further complicate the task of developing high-performing structures because of strain-rate dependent material behavior. We developed a physics-informed model to predict optimal impact velocity from a single quasistatic test. We then leveraged limited intermediate strain rate and impact testing to refine the model for strain-rate strengthening and demonstrate the model's extrapolative abilities by applying it to different impactor masses, different designs, and a different material. This model can be used to screen potential designs using our extensive collection of quasistatic experiments or to guide future impact testing.

Finally, we explored the use of a self-driving lab as a community resource by collaborating with outside researchers to study the processing-structure-property relationships of foam-like structures printed on fused filament 3D printers using viscous thread printing. We ran two independent active learning campaigns that tested hundreds of samples, allowing us to develop models that predict their layer height, modulus, and stress-strain curve. These models not only allowed the realization of printed foams that

match the stress-strain response of commercial foams, but it also enabled inverse design of 3D mechanisms by allowing designers to spatially modulate processing parameters. This work not only furthers the mechanical understanding of viscous thread printing, but also serves as a template for increasing the utilization of existing self-driving labs through collaboration.

TABLE OF CONTENTS

DEDICATION	iv
ACKNOWLEDGMENTS	v
ABSTRACT	viii
TABLE OF CONTENTS.....	xi
LIST OF TABLES.....	xv
LIST OF FIGURES	xvi
LIST OF ABBREVIATIONS.....	xxiii
CHAPTER 1: INTRODUCTION.....	1
1.1 Motivation.....	1
1.2 Organization.....	4
CHAPTER 2: DRIVING SCHOOL FOR SELF-DRIVING LABS.....	7
2.1: Introduction.....	7
2.2: Overview of Self-Driving Lab Campaigns.....	9
2.3: Actions to be Taken – the Knobs.....	11
2.3.1: Setting A: The range of parameters considered.....	12
2.3.2: Setting B: The metrics that define performance	13
2.3.3: Setting C: The surrogate model used to approximate the experimental response.....	14
2.3.4: Setting D: The acquisition function used to evaluate the value of sampling at a given location.....	16
2.3.5: Setting E: Finding the maximum of the acquisition function.....	17

2.3.6: Setting F: Which experimental data points are used to condition the surrogate model.....	19
2.4: Sources of Feedback to Adjust the Knobs – the Gauges	20
2.4.1: Response scatter plots.....	20
2.4.2: Parity plots	22
2.4.3: Proximity plots.....	25
2.4.4: Performance gamut	27
2.5: Concluding Remarks	29
 CHAPTER 3: SUPERLATIVE MECHANICAL ENERGY ABSORBING EFFICIENCY DISCOVERED THROUGH SELF-DRIVING LAB-HUMAN PARTNERSHIP.....	
3.1: Introduction.....	30
3.2: Results.....	34
3.2.1: Defining a campaign to study generalized cylindrical shells	34
3.2.2: Discovering high-performing structures.....	39
3.2.3: Material influence on design and performance.....	44
3.2.4: Broader design considerations	48
3.3: Discussion.....	56
3.4: Methods	59
3.4.1: Design of generalized cylindrical shells	59
3.4.2: Development of the Bayesian experimental autonomous researcher	63
3.4.3: The research campaign	67
3.4.3: SHAP analysis	85

CHAPTER 4: A PHYSICS-INFORMED IMPACT MODEL REFINED BY MULTI-FIDELITY TRANSFER LEARNING.....	86
1. Introduction.....	86
2. Theory.....	88
3. Results and Discussion	91
4. Conclusion	100
5. Methods	101
CHAPTER 5: EXPLORING SDLS AS A COMMUNITY RESOURCE.....	104
5.1: Introduction.....	104
5.2: Results and Discussion	106
5.2.1: Modeling and homogeneous subspace	106
5.2.2: Comparison to common foams	110
5.3: Conclusion	112
5.4: Methods	112
5.4.1: Layer height measurements	112
5.4.2: Automated testing	113
5.4.3: Principal component analysis	114
CHAPTER 6: CONCLUSION	116
6.1 Lowering barriers to develop and use SDLs.....	116
6.2 Leveraging dataset for additional mechanical insights.....	117
6.3 Working towards fundamental impact performance metric	119
6.4 Increasing productivity and utilization of SDLs.....	120

6.5 Practical advice for new SDL users	121
BIBLIOGRAPHY	124
CURRICULUM VITAE.....	137

LIST OF TABLES

Table 3.1. Common and superlattice structures and materials: Values of stress threshold σ_t^* and energy absorbing efficiency K_s^* not directly reported were computed based on reported force-displacement or stress-strain plots.	33
Table 3.2. Inputs to the machine learning models used for Bayesian optimization.	71
Table 3.3. Outputs of the machine learning models used for most of the experimental campaign.	71
Table 3.4. Descriptions of decision policies used during campaign.	75
Table 3.5. Filaments studied in this work along with their processing settings.	84

LIST OF FIGURES

Figure 1.1: Compression testing and performance benchmarks. a, Compression testing consists of compressing two platens together at a constant rate while measuring the force and displacement. b, Efficient energy absorbers consist initial elastic region, long and flat plateaus, and a densification region..... 1

Figure 2.1: Schematic showing the interactions between a human researcher and a self-driving lab (SDL). The SDL proceeds autonomously as an iterative process of performing experiments and then conditioning a model to select the next experiment. The human can monitor the progress of the individual experiments and campaigns as a whole using a series of plots to adjust settings that guide the operation of the SDL. Importantly, the human is not in the loop and any adjustments can happen asynchronously without slowing the pace of experimentation..... 11

Figure 2.2: Selection methods for sampling points. a, Grid based sampling (orange open circles) will sample the same points each time. The target maximum (black asterisk) and the closest point (solid orange circle) are also shown. b, Random sampling is statistically likely to cluster in local regions, leaving other regions with no points. c, Latin hypercube sampling (LHS) points are selected by dividing the space into a grid and selecting one point randomly in each domain of the resulting grid. d, A second round of LHS points (dark orange open circles) can be used to further search in a promising region, finding a better point (solid orange) than the originally selection points..... 18

Figure 2.3: Response scatter points. a, Observed response y vs. x_1 . Points colored based on the order that they were tested (light gray to black), with the most recent test data shown as a blue star. In this panel, the peak of y is firmly within the boundaries of x_1 , indicating that the range is likely appropriate. b, Observed y vs. x_2 . In contrast with x_1 which is a design variable that can be chosen continuously, x_2 corresponds to a material property and therefore only discrete values corresponding to known materials can be selected. In this panel, the peak of y is on the boundary that cannot be expanded due to material constraints, pointing to potential material development goals..... 21

Figure 2.4: Parity plots. a, Predicted response \tilde{y} vs. y built from 13,250 data points. Red line represent $\tilde{y} = y$ and pink lines represent estimated measurement error σ_y . Blue star represents the latest experimental prediction and the subsequent result. b, Same model as a, but zoomed in on the upper section of the data to highlight the flat plateau that occurs at $\tilde{y} \approx 0.7$, leading to significant underprediction of performance for the latest experiment. c, Model built from 4,060 data points near max (y) shows improved performance at high values, which are the most important in a maximization problem. Prediction of σ_y is also decreased. d, Model built from 410

data points illustrating deteriorated performance and a very small estimated σ_y due to little data. 23

Figure 2.5: Proximity plots. Observed y vs. distance to selected point p . Black points represent previous experiments, while the blue star represents the selected experiment and its predicted performance y . The distance of the point closest to the star is an indicator of the exploration/exploitation tradeoff. In this panels, this tradeoff is illustrated by selecting a point based on upper confidence bound with varied hyperparameter λ . a, $\lambda = 0$, pure exploitation. b, $\lambda = 1$, balance of exploration/exploitation. c, $\lambda = 5$, focus on exploration. d, $\lambda = 25$, nearly pure exploration. 26

Figure 2.6: Performance gamut. a, Experimental data (black dots) are represented by two dimensions: y_1 and y_2 . b, To find the maximum of the acquisition function, LHS is used to evaluate sampling points (orange) using the zoomed-in GPR model from Figure 2.4c. c, A second pass of LHS (dark orange) is done with the same GPR model but in a smaller region to more closely locate the maximum of the model. ... 28

Figure 3.1: Common and superlative structures and materials. Values of critical stress threshold σ_t^* and critical energy absorbing efficiency K_s^* for synthetic structures (blue diamond) and natural materials (green triangle) gathered from literature, with superlative plastic and hyperelastic generalized cylindrical shells (GCS-this study) components (red circle)..... 32

Figure 3.2: Challenge of designing energy absorbing structures. a, Force F vs. displacement D and effective medium stress σ vs. compressive effective medium strain ϵ for an additively manufactured cylindrical shell made of thermoplastic polyurethane (TPU). Maximum energy absorbing efficiency K_s^* is calculated at an optimum threshold stress σ_t^* (dashed line) by dividing the energy absorbed while $\sigma \leq \sigma_t^*$ (blue region) by the theoretical maximum amount absorbed (red rectangle). b, Eleven independent geometric parameters including diameter d , height h , wall thickness t , and eight other parameter x_{1-8} that together define a generalized cylindrical shell (GCS). When combined, at least trillions of unique designs are possible. c, Elastic modulus E of the seven polymers studied as determined by compression tests. Error bars represent one standard deviation. d, Schematic showing an autonomous research system in which five 3D printers are used to fabricate polymeric structures that are automatically weighed, imaged, and tested using quasistatic compression. The output of this testing is automatically interpreted and used to select subsequent designs to test..... 36

Figure 3.3 Details of the human/machine collaboration. a, K_s^* of each successful test (gray dots), along with the highest K_s^* to date (blue line). Key changes to the processing and sampling space are marked. b, Modulus of each experiment's filament roll plotted in semi-log (right axis) and colored according to the Figure 3.2c, along with

the highest K_s^* to date (blue line – left axis). c, Decision policy of each experiment (right axis), along with the highest K_s^* to date (blue line – left axis). Decision policies are listed with descriptions in Table 3.4..... 41

Figure 3.4: Research campaign to find highly efficient structures. a, Each maximum energy absorbing efficiency K_s^* measured over the first ~21,500 experiments performed. Pictures highlight noteworthy components (black stars) and the highest performing structure (red star). The color of the pictured components is indicative of the material used, with Green indicating PLA, Blue indicating PETG, and Red/Gray indicating different types of TPU. The solid blue line denotes the running best K_s^* observed. b, effective medium stress σ vs. effective medium strain ϵ for experiment 21,285, named Palm, which resulted in $K_s^* = 75.2\%$. Inset photographs show the state of the component at various points indicated on the curve (images enhanced to improve clarity). Shading denotes regions used to compute K_s^* as described in Figure 3.2a..... 44

Figure 3.5: Exploration of high-performing designs discovered in elastic and plastic materials. a, Rendering of Willow, a high-performing design discovered using the plastic polymer polylactic acid (PLA) together with effective medium stress σ vs. effective medium strain ϵ for 15 identically prepared PLA Willow components. b, Rendering of Iroko, a high-performing design discovered using the hyperelastic polymer TPU-2, together with σ vs. ϵ for 15 identically prepared TPU-2 Iroko components. c, Measured maximum energy absorbing efficiency K_s^* vs. polymer elastic modulus E for Iroko and Willow components made from one of five polymers. Dashed lines show a sigmoidal fit to guide the eye. Error bars represent one standard deviation. Marker colors denote component composition, as shown in Figure 3.2c. 46

Figure 3.6: Willow/Iroko by material. a, Stress σ –strain ϵ curves for components made using the Willow design printed in TPU-1, TPU-2, TPU-3, PETG, and PLA. PLA, the original Willow material, has 15 tests, while the other materials have three each. b, σ – ϵ curves for components made using the Iroko design printed in the same five materials. TPU-2, the original Iroko material, has 15 tests while the other materials have three each. Colors depict the material as in Figure 3.2c..... 48

Figure 3.7: Convex hulls for seven materials. All tests for each of the seven materials studied, with their final σ_{tp} marked. The materials are TPE (a), TPU-1-3 (b-d), nylon (e), PETG (f), and PLA (g)..... 49

Figure 3.8: Design insights that emerge from mechanical dataset. a, σ_{tp} vs. polymer plateau stress σ_p for seven polymers tested during the campaign together with a power law fit shown as a dashed line. Marker color indicates the material type. Error bars indicate one standard deviation of σ_{tp} found throughout the campaign. b, K_s^* vs.

relative density ρ_r for all components tested during the campaign with point color denoting K_s^* . c, Normalized height ω vs. d/t for all components tested during the campaign in which point color denotes K_s^* as in b. 50

Figure 3.9: Shapley additive explanations (SHAP) analysis of Willow as the superlative design. a, Parity plot of the neural network built on all data taken using PLA with Willow highlighted. b, SHAP waterfall plot for the Willow design tested in PLA relative to a PLA cylindrical shell with the same height, diameter, and thickness. These values show the cumulative effect of positive (red) or negative (blue) contributions of individual feature values to model predictions. 53

Figure 3.10: Simulation of efficiency for combinations of tested willow components. a, Two or more experimental F-separation curves (red) can be combined to create a simulated test of in-plane tiled components (blue). Curves are added in platen separation space to account for possible variations in component height. b, Mean and one standard deviation in K_s^* for original components (red) and combined prediction (blue) for varying number of components. 54

Figure 3.11: Examining toughness metrics beyond energy absorbing efficiency. a, Toughness per mass U_m vs. component effective medium density for all successful experiments. The color of each dot corresponds to the material used, as designated in Figure 3.2c. b, Toughness per volume U_v vs. component effective medium density for all successful experiments. c, U_v vs σ_t for the original Willow component with the point representing K_s^* marked by a black dot and the parity line showing the maximum possible U_v . d, U_v vs σ_t for the original Iroko component. 56

Figure 3.12: Generalized cylindrical shells. a, Generalized cylindrical shells (GCS) are realized by transforming a cylindrical shell of height h and wall thickness t to create interesting shapes that preserve the topology of the shell. b, The perimeter P varies linearly along the height z of the shell based on an average perimeter P_0 and a perimeter difference x_1 . c, The cross sections of each layer are deformed in a z -dependent manner using a summed cosine function with 4-period amplitude C_4 and 8-period amplitude C_8 . These are defined at the top and bottom by four variables $x_2, x_3, x_4,$ and $x_5,$ and linearly interpolated to determine the cross section at any z . d, The cross sections of the design are rotated about the cylinder axis in a z -dependent manner by rotation angle ϕ_0 using both linear and sinusoidal twists as defined by linear twist $x_6,$ sinusoidal twist amplitude $x_7,$ and sinusoidal twist period $x_8.$ 61

Figure 3.13: Picture of the Bayesian experimental autonomous researcher (BEAR), consisting of five fused filament fabrication 3D printers, a six-axis robot arm, a scale, and a universal testing machine. 64

Figure 3.14: Hardware and software Organization of the Bayesian experimental autonomous researcher (BEAR). 64

Figure 3.15: Flowchart of the software loops executed in the four computers running as part of the BEAR. Colors on the panels correspond to systems in Figure 3.14. Order of Main Loop actions can be adjusted by researchers to maximize throughput by prioritizing potential bottlenecks. 66

Figure 3.16: Experiments carried out by the BEAR. Experiments performed during the campaign, which are defined by eleven GCS parameter values. The color of each dot corresponds to the material used, as designated in Figure 3.2c. Black and red stars correspond to breakthrough experiments, as designated in Figure 3.4a. 68

Figure 3.17: Mass calibration through extrusion multiplier adjustments. a, A single component printed five times with different extrusion multiplier shows that component mass is linearly correlated with extrusion multiplier. b, Applying integral tuning to the extrusion multiplier was ineffective because of variations in slicing complex curved structures, as seen in print number < 60. However, when integral tuning was applied to the slicer filament length by adjusting the extrusion multiplier rapidly, consistent mass was obtained (print number > 60). Mass is normalized by target mass and slicer filament length is normalized by the initial slicer filament length (print number = 1). 77

Figure 3.18: Calculating the effective area of GCS designs. a-b, To calculate the effective cross-sectional area of a design, it is fit into a cylinder based on its maximum radius. c, This cylinder is then enclosed in a hexagonal prism. d, The hexagonal prism can be tiled infinitely in a plane. Thus, the effective cross-sectional area of a design is estimated as the area necessary per design to tile it in a plane without collisions.... 81

Figure 3.19: Material characterization of polymers studied. a, Plateau stress σ_p vs. elastic modulus E for seven materials used in this campaign. b, Rebound fraction vs. E. c, Rebound fraction vs. σ_p . Error bars represent one standard deviation. Here, σ_p is the stress at 25% strain. Rebound fraction is the height after 1 minute relaxation divided by the initial height. 84

Figure 4.1: Predicting impact performance from quasistatic tests. a, Photograph showing a generalized cylindrical shell (GCS) component made from thermoplastic polyurethane (TPU). B, Schematic of an impact test with impactor of mass m hitting the component of height h at an initial velocity V_0 . c, Acceleration a vs. displacement D for impact tests of the same component at three different V_0 and $m = 1.57$ kg. The max acceleration a_m for each V_0 is denoted as a black dot. The units of a are the gravitational acceleration g. d, The a_m at various V_0 with $m = 1.57$ kg. e, Janssen factor J vs. V_0 . The test with the lowest J is designated as the point (V_0^*, J^*) . f, Schematic of a fixed speed test in which a platen is lowered at constant velocity V. g, Force F vs. D for quasistatic (QS) test ($V = 2$ mm/min). This data is used to compute the most efficient operating conditions in terms of the critical force threshold F_t^* and maximum mechanical energy absorbing efficiency K_s^* , which is

defined by dividing the amount of energy absorbed (light blue area) by the maximum amount of energy that could be absorbed below F_t^* (light red area). h, Photograph showing five additional GCS designs made from TPU. i, QS model uses a single QS test to compute the predicted optimal impactor velocity \tilde{V}_0^* . j, Parity plot showing \tilde{V}_0^* vs. V_0^* for the six shown designs using the QS model for $m = 1.57$ kg. Error bars denote one standard deviation..... 91

Figure 4.2: Predicting impact performance from intermediate strain rate tests. a, Velocity V vs. D during impact and intermediate strain rate (ISR) tests. While impact tests start at $V = V_0$ and decrease as energy is absorbed, ISR tests have a fixed V for the entire test. b, F vs. D for ISR tests of copies of the component shown in Figure 4.1a. c, K_s^* vs. V for the ISR tests shown in b. d, F_t^* vs. V for the ISR tests shown in b. The red line shows a fit to Equation (4.6) with key fitting constants $F_t^*(0)$ and V_c marked by gray and black dashed lines respectively. e, ISR model uses both QS and ISR tests of the target design to predict \tilde{V}_0^* . f, Parity plot showing \tilde{V}_0^* vs. V_0^* for the six original designs using the ISR model for $m = 1.57$ kg. Error bars denote one standard deviation. In panels a, b, c, and d, shades of blue indicate V 94

Figure 4.3: Predicting impact performance using transfer learning. a, Values from ISR fitting each of the six designs to Equation (4.6) with the mean value shown as a dashed red line. b, Plot of QS $F_t^*(2 \text{ mm/min})$ vs. fitting constant $F_t^*(0)$. c, Transfer learning (TL) ISR model that uses ISR tests of other designs to calculate V_c and α , allowing the prediction of \tilde{V}_0^* with a single QS test for each target design. d, Parity plot showing \tilde{V}_0^* vs. V_0^* for the six designs using the TL ISR model for impactor mass $m = 1.57$ kg using leave one out cross validation (LOOCV) to calculate V_c and α for each prediction. e, F_t^* vs. V for ISR tests with theoretical impact point (V_0^*, F_t^*) (pink square) calculated using Equation (4.7). f, Values from ISR and impact fitting each of the six designs to Equation (4.6) with the mean value shown as dashed red line. g, TL impact model uses ISR tests and impact F_t^* prediction from other designs to calculate parameters V_c and α , allowing the prediction of \tilde{V}_0^* with just a single QS test for each new target design. h, Parity plot of \tilde{V}_0^* vs. V_0^* for the six original designs using the TL impact model for $m = 1.57$ kg using LOOCV. i, New design printed from TPU. j, QS test for new design showing F vs. D with red dashed line indicating F_t^* . k, J vs. V_0 for the design shown in i with $m = 3.1$ kg. Dashed lines represent \tilde{V}_0^* using the four models. l, Second new design printed from TPU. m, QS test for second new design showing F vs. D with red dashed line indicating F_t^* . n, Values of J vs. V_0 for design shown in l for $m = 3.1$ kg. Dashed lines represent \tilde{V}_0^* using the four models. Error bars throughout represent one standard deviation..... 96

Figure 4.4: Extending transfer learning models to an additional material. a, Stress σ vs. strain ϵ for QS tests of solid cylinders 8 mm in diameter with $h = 16$ mm used in material testing. b, Picture of seven of the 16 additional designs fabricated with a foaming TPU filament. c, The h vs. component mass for original six designs and the

new 16 foaming TPU designs. d, The K_s^* vs. F_t^* for original six designs and new 16 foaming TPU designs. e, Root mean squared error (RMSE) for predictions of \tilde{V}_0^* of new foaming TPU designs. Both TL models use fitting constants V_c and α trained solely on data from original six TPU designs. f, Parity plot showing \tilde{V}_0^* vs. V_0^* for the 16 foaming TPU designs using the TL Impact model for $m = 1.57$ kg. 99

Figure 5.1: Parameters for viscous thread printing (VTP). a, In VTP, the nozzle speed V_n travels slower than the thread exit velocity V_t , while at a height above the print H significantly higher than the thickness of the extruded thread Dt , causing the thread to coil in patterns that can be predictable. b, To create 3D objects, rectilinear toolpath (blue) can be used while varying the spacing between lines in the x-y plane by ΔL and spacing layers by ΔZ . A cross hatch is used to ensure continuity from layer to layer. 105

Figure 5.2: Mapping homogenous subspace for PLA filament. a, H^* vs. V^* for all 300 experiments performed in PLA with color indicating ΔL and size indicating ΔZ . b, Parity plot of predicted layer height \hat{H}_L vs. H_L for Gaussian process regression (GPR) model using leave one out cross validation (LOOCV). c, Log-log parity plot for predicted effective modulus \hat{E} vs. effective modulus E for the GPR model using LOOCV. d, Slice in 2D of H^* and V^* where ΔL and ΔZ are equal to 1.5 mm. Color indicates H_L and the red line indicates predicted homogenous subspace where $\hat{H}_L = \Delta Z$. Black dots indicate five equally spaced samples selected for validation testing. e, GPR model's \hat{E} (red line) vs. results of five validation experiments (black). Error bars represent one standard deviation in semi-log space from each condition being tested in triplicate. f, Layer height H_L vs. layer number for three PLA cubes with different H^* . After several layers, H_L for each cube stabilizes to 2.2 mm, which is equal to ΔZ 106

Figure 5.3: Predicting full stress-strain curves to replicate commercial foams. a, Force (F)-displacement (D) of a TPU cube to the 4.5 kN force limit. b, Log stress-strain curve (blue) converted from F-D curve is down sampled to 100 points (red) equally spaced to 50% strain. c, Principal component analysis breaks down sampled curves into 100 components (only top five shown for clarity). d, Cumulative variance captured by number of components. The vast majority of variance is captured by the first five components. e-g, Stress-Strain curves for three foam samples (black curves), the sample predicted to most closely match their performance (red dashed), and the tested performance of that predicted sample (blue dotted). 110

LIST OF ABBREVIATIONS

a	Acceleration
a_f	Acquisition function
a_m	Maximum acceleration
a_t	Theoretical lowest maximum acceleration
BEAR	Bayesian experimental autonomous researcher
BO	Bayesian optimization
D	Displacement
d	Diameter
d_m	Dimension
D_t	Displacement threshold
E	Modulus
\hat{E}	Predicted modulus
EI	Expected improvement
ε	Strain
ε_d	Densification strain
F	Force
F_t	Force threshold
FEA	Finite element analysis
FFF	Fused filament fabrication
GCS	Generalized cylindrical shells
GPR	Gaussian process regression

h	Component height
H_L	Layer height
\hat{H}_L	Predicted layer height
ISR	Intermediate strain rate
J	Janssen factor
J^*	Minimum Janssen factor
KE	Kinetic energy
K_S	Energy absorbing efficiency
K_S^*	Critical energy absorbing efficiency
LHS	Latin hypercube sampling
ΔL	Lateral toolpath separation
m	Mass
ω	Normalized height
P	Perimeter
p	Printability
P_0	Average perimeter
PCA	Principal component analysis
PETG	Polyethylene terephthalate glycol
PLA	Polylactic acid
ρ_r	Relative density
QS	Quasistatic
RMSE	Root mean square error

SDL	Self-driving lab
σ	Stress
σ_p	Plateau stress
σ_t	Stress threshold
σ_t^*	Critical stress threshold
σ_{tp}	Material peak stress threshold
t	Wall thickness
TPE	Thermoplastic elastomer
TPU	Thermoplastic polyurethane
U_m	Toughness per mass
U_v	Toughness per volume
UCB	Upper confidence bound
UTM	Universal testing machine
V_0	Initial impact velocity
V_0^*	Ideal initial impact velocity
\tilde{V}_0^*	Predicted ideal initial impact velocity
V_n	Nozzle velocity
V_t	Thread velocity
VTI	Viscous thread instability
VTP	Viscous thread printing
ΔZ	Vertical toolpath separation

CHAPTER 1: INTRODUCTION

1.1 Motivation

Structures and materials that absorb energy are critical to modern society. Every day as we drive to and from work, crumple zones in cars and crash attenuators mitigate the danger to occupants in the event of a crash^{1,2}. Padding in helmets protect not only cyclists and football players, but also soldiers in combat zones³⁻⁵. Sports equipment like boxing gloves or hockey pads prevent injuries from repeated blows between players⁶. Even mundane events like receiving an order from online shopping is facilitated by cheap and effective packaging that protects the item in transit^{7,8}. Because of its many uses, research continues to strive for energy absorbing materials and structures that occupy a smaller volume, are lighter, or are safer. One approach is to improve the energy absorbing efficiency in compression, which is characterized by a force-displacement response with a flat plateau region that persists into densification at a high value of strain⁷ (Figure 1.1).




Figure 1.1: Compression testing and performance benchmarks. a, Compression testing consists of compressing two platens together at a constant rate while measuring the force and displacement. b, Efficient energy absorbers consist initial elastic region, long and flat plateaus, and a densification region.

Unfortunately, discovering high-performing energy absorbing structures and materials is a slow process. Optimizing for a flat plateau region is often in conflict with

having a late densification, which requires low relative density. Compounding this issue, there are a near infinite number of possible structures that can be made. In our own generalized cylindrical shell family described in Chapter 3, we estimate that trillions of unique designs are possible from a mere 11 design parameters. This is both the blessing and the curse of dimensionality. The vast design space makes it more likely that interesting or useful designs exist but also makes them harder to find, preventing the efficient use of brute force methods to discover optimal structures when there are even just a handful of tuning parameters with more than a couple of choices each. Of course, this challenge is only intensified by the adoption of additive manufacturing, which, while enabling the rapid prototyping of new designs, also gives unparalleled diversity of potential structures to study. The hierarchical nature of materials design, where properties and structures at different length scales contributes to the performance of the whole, also adds additional challenges.

Simulation has been successfully used to overcome this bottleneck in certain research domains. For example, finite element analysis (FEA) can be used to screen potential designs, and constrained optimization techniques such as topology optimization can limit design to promising candidates. These techniques have been fruitfully employed to study mechanical properties such as stiffness or strength⁹⁻¹¹. However, they struggle with other properties, such as toughness or energy absorbing efficiency, because they are difficult to simulate. The challenges arise primarily because it is necessary to simulate high strain compression, which often involves buckling events and self-self contacts within the structure. These difficulties are only compounded when modeling impact

events, where strain-rate effects come into play and where the speed of the impact changes as energy is absorbed. In addition, simulation often fails to capture manufacturing processes or defects that affect performance. For example, in fused filament fabrication (FFF) where material is deposited layer by layer, the side of the component has grooves marking the layers and the bonds between layers are usually weaker than intralayer bonds. These characteristics can affect performance, leading to a discrepancy between simulated and experimental results.

To accelerate the discovery of superlative energy absorbing structures, we utilized a self-driving lab (SDL). SDLs combine automation of experimentation with active learning techniques from the machine learning community. Although SDLs can take a variety of forms and even go by different names, a key requirement is that they are able to form a hypothesis, propose an experiment, conduct the experiment, and record the results, all without human intervention¹²⁻¹⁶. SDLs have been successfully employed to study a variety of disciplines, including biology¹⁴, chemistry¹⁷⁻²¹, material science^{15,22-25}, and mechanics²⁶⁻²⁸.

SDLs have many strategic advantages. First, they increase the throughput of experiments by eliminating the need for a researcher to be present. During our two-year campaign, we ran our system nearly continuously, including nights and weekends, with an uptime of ~60%. Second, SDLs can document an assortment of relevant metadata, such as recording the mass of the component, taking photos of each part, filming the compression test, and recording ambient conditions such as humidity and temperature. Third, SDLs allow improved consistency and replicability in experiments. Recording this

level of detail day in and day out without error would have been unreasonably tedious for human researchers and could have led to errors due to fatigue. Finally, SDLs have the ability to increase the knowledge gained from each experiment. Traditional methods of high-throughput experimentation, such as grid-based searching or design of experiments, select all experiments before testing begins. This means that they are not using results to inform experiment selection, decreasing the insight gained from each individual experiment²⁷. However, active learning also outperforms human experts, especially in high dimensional spaces, because humans tend to default to one factor at a time optimization, which is inefficient and prone to finding local, rather than global, maxima²⁹.

Here, we used an SDL to find superlative structures for both quasistatic and impact energy absorption. With these data, we gain knowledge about material properties, design motifs, and strain-rate effects. In addition, during our unprecedented two-year SDL campaign, we gained insights on how to monitor and adjust SDL settings during an open-ended campaign. Finally, we demonstrated the utility of an SDL as a community resource by partnering with outside researchers to study the interesting mechanical properties of viscous thread printing (VTP) foams.

1.2 Organization

Chapter 2 explores the challenges that arise when running an SDL for an open-ended, long-term campaign. Because of the extended duration, it is likely that researchers will need to monitor and adjust various settings for the SDL. Taking inspiration from automobiles, we present six knobs that can be tuned to guide SDL priorities and four

gauges that can be monitored to inform the researcher about progress and potential adjustments. This guidance seeks to push forward the understanding of how SDLs can benefit from human-machine collaboration, where researcher expertise can be integrated into the SDL framework to further accelerate research.

In Chapter 3, we present the results of our two-year campaign exploring generalized cylindrical shells (GCS). We study this rich 11-dimensional design space with trillions of unique experiments using seven different filament types. After >25,000 experiments, we found a design that achieved a record breaking 75.2% efficiency at absorbing mechanical energy. In addition, we explored how material affects design and performance, as well as examining the rich dataset for design motifs that could transcend the GCS space.

Chapter 4 shifts focus to predicting performance of GCS parts in impact. Because strain-rate effects can significantly alter the stress-strain response, quasistatic measurements do not directly predict impact performance. However, we created a physics-informed model that employs quasistatic data to predict the optimal impact velocity of a design for a given impact mass. We further refine this model using limited intermediate strain rate and impact tests. This transfer learning model is able to predict performance of new impactor masses, new designs, and new materials with an RMSE of 0.23 m/s.

Chapter 5 transitions to our work opening our SDL up as a community resource. In this chapter, we highlight our collaboration with outside researchers to study viscous thread printing (VTP). Our system ran two independent campaigns to map the response

space for both the layer height and effective modulus of the four-parameter input space. Models built on the data from these campaigns allow targeted performance of homogenous foam parts. Additionally, using principal component analysis, prediction of full stress-strain curves is possible, which allows replication of the performance of common commercial foams using VTP foams.

Finally, Chapter 6 reflects on the lessons learned during the five years working with the SDL and contemplates on the future of SDLs. It concludes with advice for researchers interested in developing and running an SDL for the first time.

CHAPTER 2: DRIVING SCHOOL FOR SELF-DRIVING LABS

This chapter is adapted from work published in *Digital Discovery*³⁰.

2.1: Introduction

To fully take advantage of the effectively infinite number of possible materials and processing conditions to study, it is necessary to leverage every opportunity to accelerate the pace of progress. Advances in machine learning and artificial intelligence have shined a spotlight on the possibility of using these advanced computational methods to transform research. Unfortunately, many material properties can only be reliably determined by experiment, making experiments the gold standard and often the only method for generating high-fidelity data. These considerations have led to the development of self-driving labs (SDLs), which are research systems that iteratively select physical experiments using machine learning and automatically performs these experiments without human intervention.³¹ While systems with closed-loop control over experimental conditions date back decades to self-optimizing chemical reactors,³² the past five years has seen the rapid introduction of increasingly sophisticated SDLs that operate with a wide variety of material systems and form factors.^{31,33–36} These SDLs have included systems that study carbon nanotube synthesis,¹⁵ yeast genetics,¹⁴ catalyst composition,¹⁷ nanoparticle synthesis,^{18–21} properties of thin films,^{22–25} and the mechanics of additively manufactured structures.^{27,28} SDLs have been found to reduce by 10-600 fold the number of experiments needed to reach a given performance level relative to grid-based searching,^{27,28,37–39} to say nothing of their ability to perform experiments at a faster pace than what is possible with human experimentation. The value

of these systems has been further proved through their discovery of new multi-component electrolytes,⁴⁰ previously unreported chemical compounds,⁴¹ and structures with superlative mechanical performance.⁴²

The value proposition of SDLs is that they can accelerate research progress, but the fact that they are self driving does not mean that they cannot have input from people. Indeed, the construction and programming of SDLs is inherently a human process in which people have provided input, codifying their priorities and goals. This process itself can be complex and important as it involves turning the human domain knowledge into an algorithmic format that can be utilized by the SDL.⁴³ Nevertheless, for most SDL campaigns published to date, this initialization is the end of the meaningful human interaction outside of restocking reagents or other feedstock materials. In particular, the SDL iteratively selects and performs experiments based on the programming until the allotted number of experiments has been reached. While this “hands-off” approach is reasonable when considering SDL campaigns that collect data over a few days, it is increasingly unrealistic as the timeframe of campaigns stretches into weeks and months. Under these conditions, it is very natural for the human experimenters to monitor the SDL and make adjustments. Indeed, there have been recent reports highlighting the importance of the human-machine collaboration and how this collaboration can be used as part of otherwise autonomous systems.⁴⁴ However, given that we are only now seeing the widespread adoption of SDLs, there do not yet exist resources to help experimenters know what to monitor and what to adjust.

Based on our experience running the Bayesian experimental autonomous researcher (BEAR), we detail our key learnings from this process and codify two important aspects, namely the choices an experimenter needs to make and the information that they should be monitoring to make these choices. We narrow these items down into six settings that must be chosen and four plots that should be monitored periodically throughout the campaign to adjust these settings. Despite our insights coming from a campaign based on studying the mechanical properties of additively manufactured structures, the lessons learned are largely system agnostic. We discuss these lessons using the commonly adopted framework of Bayesian optimization to aid in their adoption by others. While the heuristics presented herein cannot account for every situation faced when operating SDLs, we hope that the language, principles, and processes can provide experimenters with greater intuition and increase the adoption and effectiveness of SDLs across the materials space.

2.2: Overview of Self-Driving Lab Campaigns

While the details of materials research campaigns performed by SDLs vary considerably with the SDL architecture, materials system, and campaign goals, they share the need to sequentially select experiments based on the currently available data (Figure 2.1). There are a number of algorithms that can perform this selection process. However, we focus our discussion on Bayesian optimization (BO) as it is commonly used in SDLs and because BO separates the steps of amalgamating knowledge and choosing experiments in a manner that is heuristically useful.⁴⁵ In particular, BO generally proceeds as a two-step process (Figure 2.1). For the sake of defining a consistent

language, we consider the task of maximizing some experimentally measured property y across some multi-dimensional parameter space \vec{x} . The objective or ground truth function $y = f(\vec{x})$ is unknown and heteroskedastic (variance is not constant across parameter space). The first step of BO is to use the available knowledge to condition a surrogate model $\tilde{f}(\vec{x}; \vec{\theta})$, or a probabilistic approximation of the objective function that depends on hyperparameters $\vec{\theta}$. Available knowledge can include, for example, experimental data, physical knowledge about the system such as expected symmetries, or related prior knowledge such as simulation results. Second, the surrogate model is used as an input to an acquisition function $a_f(\vec{x}; \tilde{f})$ which quantifies the expected benefit of performing an experiment at any location in \vec{x} . The maximum of a_f across \vec{x} is generally selected as the next experiment.

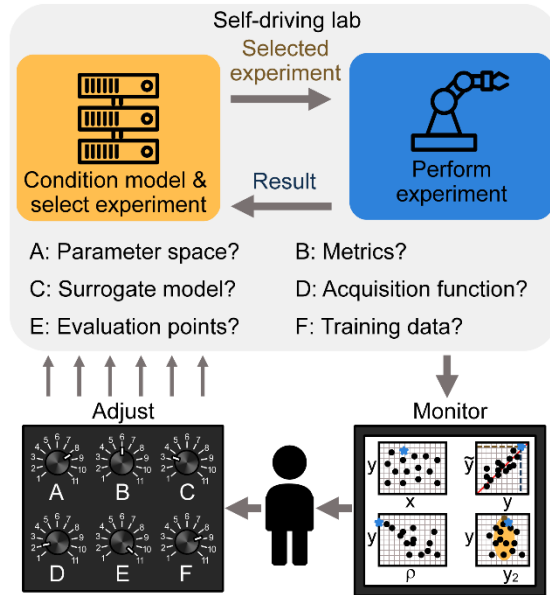


Figure 2.1: Schematic showing the interactions between a human researcher and a self-driving lab (SDL). The SDL proceeds autonomously as an iterative process of performing experiments and then conditioning a model to select the next experiment. The human can monitor the progress of the individual experiments and campaigns as a whole using a series of plots to adjust settings that guide the operation of the SDL. Importantly, the human is not in the loop and any adjustments can happen asynchronously without slowing the pace of experimentation.

Each step of BO requires the experimenter to set a number of parameters whose values may not be obvious or physically intuitive. Crucially, the BO process, and the operation of the SDL more generally, need not be static, and the experimenter should continually revisit settings based on recent results to improve the chances of converging towards the best outcome. Below, we highlight the actions that the experimenter can take (Section 2.3) and then detail the feedback that forms the basis for how to choose which actions to perform (Section 2.4).

2.3: Actions to be Taken – the Knobs

Despite the use of the term self driving, in practice there are a number of deliberate choices that the human experimenter must make to productively guide an SDL

to make efficient progress. In this section, we list six decisions that we have found to be most significant when running an SDL campaign.

2.3.1: Setting A: The range of parameters considered

The range of \vec{x} that are considered as the parameter space is a non-trivial choice that can be crucial to the success of an SDL campaign. In physical experiments, each component of \vec{x} corresponds to a tangible physical quantity such as the identity of a compound, a processing temperature, or the length of a geometric feature. As such, not all permutations of \vec{x} correspond to valid experiments and there is no guarantee that the space of valid experiments is a convex region. In the simplest case, when each component of \vec{x} is independent, the full parameter space can be considered a hypercube. Even in this case, determining the maximum and minimum values for each parameter is often a judgement call based on experimenter intuition. More generally, there exists some geometrically complicated domain in \vec{x} that corresponds to valid experiments. If the boundary of valid experiments is known at the beginning of the campaign, invalid experiments can be avoided. In contrast, if the boundary is unknown but the system is such that invalid experiments can be attempted, the boundary of this domain can itself be learned during the campaign. In addition to the range of parameters considered, it can be useful to transform the parameter space such that correlations are easier to learn. For instance, monotonic transformations such as taking the logarithm of a parameter can facilitate learning if the parameter is positive definite (always greater than zero) and varies across several orders of magnitude.

2.3.2: *Setting B: The metrics that define performance*

As stated above, we are considering an SDL campaign whose goal is to maximize a metric y . Even when stated so simply, there are two additional considerations to mention. First, it is rare to conceive of a material that is only evaluated on a single axis. More often, multiple properties are considered important.^{25,46} For example, a structural material can be evaluated based on its strength, stiffness, density, cost, or toughness. When multiple metrics are important, these metrics must be distilled into a single value to guide automated experiment selection. There are a number of techniques to do this that can be drawn from multi-objective optimization. Early work focused on combining the multiple objectives into a single scalar metric that embodies the user-defined relative importance of these objectives.⁴⁷ This relative importance can also be changed throughout a campaign, which has led to strategies such as adaptive weighted sum methods for multi-objective optimization.⁴⁸ More recently, it is common to base experiment selection on hypervolume optimization,^{25,46,49} which converts each potential $\tilde{y}(\vec{x})$ into the predicted amount of additional hypervolume in the output space that is expected to be enclosed from that experiment.

Even when only a single metric is important, there are important choices to make in how this variable is quantified. While linear transformations such as normalization should not affect SDL progress, provided that the numbers are not made so large or small such that they lead to numerical precision errors, non-linear transformations can also be used to compress or expand important domains in y . For example, when y has a finite range (say an efficiency that varies from 0 to 1), the model will often predict non-

physical values (i.e. $\tilde{y} > 1$ or $\tilde{y} < 0$). Thus, performing transformations that enforce these physical constraints can be helpful. For example, building a model to predict $y' = \text{arctanh}(2y - 1)$ both preserves the physical constraints on y while expanding the domains associated with high and low values.

2.3.3: Setting C: The surrogate model used to approximate the experimental response.

A core decision in the BO process is deciding how to model y . In principle, any type of mathematical model can be used here, provided it can be used to predict both the expectation value and uncertainty at any \vec{x} . The most commonly applied model is the Gaussian process regression (GPR). This model is highly applicable in the low- to medium-data regimes and is relatively bias-free, despite only being able to model surfaces that are infinitely differentiable.

To gain intuition about the function of a GPR, it can be thought of as a tool for predicting the expectation $\tilde{y}(\vec{x})$ at any \vec{x} as the weighted average of the prior belief $\mu_0(\vec{x})$ and the results of all prior experiments $y_i(\vec{x}_i)$. Commonly, the prior is assumed to be constant and therefore usually considered uninformative, although more advanced methods can be used such as setting this prior using low-fidelity measurements including simulation-based approximations of the experimental function.²⁸ The weighting used for this prior is based on a hyperparameter θ_0 , which represents the total range of experimentally realizable values. More interesting are the weights chosen for each experimental measurement. The underlying assumption is that points closer to the sampling point are more highly correlated, although each dimension of \vec{x} could feature correlations with different length scales. A common strategy is to define a kernel function

that models how correlations vary in space with common choices being squared exponential or Matérn functions. Such functions feature a set of hyperparameters $\vec{\theta}$ that represent the distance in each dimension over which the output is expected to be highly correlated. This kernel function, together with the physical measurement uncertainty σ_y , together determine the weighting of each previously measured experiment. A similar process can be used to predict the uncertainty in $\tilde{y}(\vec{x})$, termed $\tilde{\sigma}(\vec{x})$. The hyperparameters $\vec{\theta}$, θ_0 , σ_y are crucial to GPR performance and these hyperparameters are typically set using maximum likelihood estimation. It is important to note that σ_y is an estimation of the measurement uncertainty used by the GPR and it can either be independently estimated by performing control experiments or it can be chosen using statistical methods such as maximum likelihood estimation. Either way, the decision to use a single value of σ_y across the whole parameters space is an approximation that deserves careful consideration. For a more nuanced look at GPRs, their mathematical underpinnings have been usefully outlined in prior publications.^{45,50}

It should be noted that there are vastly more sophisticated approaches that can be taken than the simple GPR-based approach described above. For instance, rather than assuming that correlations between points are translationally invariant (that correlations depend solely on distance between points but not on their location in parameter space), one can define non-stationary kernel functions that allow correlations that capture nuances in the parameter space itself.⁵¹⁻⁵⁴ Additionally, there are hybrid methods such as deep kernel learning in which a deep neural network is trained to transform the input data

to a GPR.^{55,56} This concatenation of models makes the system able to accommodate variable length-scale correlations and discontinuities in the input space.

In addition to data-driven methods to modify the expected correlations in space, there are considerable opportunities for the experimenter to bring physical intuition into the modeling effort. In particular, it is possible to transform the input space based on compound variables that are expected to more closely connect to the physical output. Simple examples of this process could include dividing the masses of reagents by their volumes to learn on concentrations or taking the ratio of geometric parameters to learn on aspect ratios.⁴² While such anecdotes appeal to our intuition, the choice of how to represent the variables that define the parameter space is an underappreciated facet of SDL operation and deserves concerted study.⁵⁷

2.3.4: Setting D: The acquisition function used to evaluate the value of sampling at a given location

The second half of BO is choosing an acquisition function a_f . The goal of this function is to convert the predictions $\tilde{y}(\vec{x})$ and $\tilde{\sigma}(\vec{x})$ into a scalar quantity that relates to the predicted benefit of sampling at that location. A key consideration here is balancing the needs of exploration and exploitation.⁴⁵ Exploration prioritizes sampling in regions with little prior information while exploitation prioritizes sampling in regions believed to be high performing as a means of optimizing this performance. One acquisition function that displays this dichotomy well is the upper confidence bound (UCB) in which $a_{UCB} = \tilde{y}(\vec{x}) + \lambda \tilde{\sigma}(\vec{x})$, where λ reflects a weighting in which $\lambda = 0$ would be considered pure exploitation while $\lambda \rightarrow \infty$ would be considered pure exploration. Other commonly used

acquisition functions are expected improvement (EI), in which a_{EI} is proportional to the amount by which the sampling point is expected to be above the highest previously observed point. Like surrogate modeling, there have been innovations in the development of advanced acquisition policies, such as those that incorporate simulation data^{28,58} or expert knowledge.⁵⁹

Another consideration when choosing an acquisition function is whether experiments are performed individually or in batches.^{60–63} When experiments are performed in batches, it is necessary to take action to make sure that each point will target a different region in \vec{x} . One approach, for instance, is to select the first experiment in the batch, then recondition the GPR assuming that the point was sampled and either returned the GPR mean (Kriging believer) or a constant value (constant liar), and use this updated GPR to predict the next point iteratively until all batch points have been selected.⁶¹ Alternatively, a simple penalty can be applied after selecting each point to force subsequent points to be further away in \vec{x} .

2.3.5: Setting E: Finding the maximum of the acquisition function

Once the surrogate model is conditioned and a decision policy selected, the maximum of the acquisition function must be determined. Because GPR models do not provide an equation where the maximum can be located in a closed form, it is necessary to choose discrete locations in \vec{x} to evaluate as candidate experiments. Exhaustively evaluating every possible location is nearly always impractical, especially when dealing with continuous variables. Thus, it becomes necessary to choose a subset of points to sample. Grid and random sampling are unfavorable for opposite reasons: Grid based

search will fail to locate a maximum point if it is not on the grid, while random sampling fails to ensure evenly spaced sampling points (Figure 2.2a-b). In contrast, Latin hypercube sampling (LHS) presents a way to cover all of parameter space while introducing some randomness that prevents the same precise point from being reconsidered multiple times in subsequent steps (Figure 2.2c). In terms of how many points to select, the curse of dimensionality plays a major role. Specifically, if \vec{x} has d_m dimensions, adding 10 conditions per d_m leads to $\sim 10^{d_m}$ points, which becomes quickly intractable when $d_m > 7$ despite 10 conditions providing only a sparse sampling of space. Thus, in any reasonably high dimensional space, it becomes necessary to couple an initial sampling with a refinement or optimization process. For example, a second cluster of LHS points may be collected near the maximum observed during the first round (Figure 2.2d). Alternatively, the maximum or maxima of the first round can be used as input points to other optimization algorithms such as gradient ascent or genetic algorithms to find local maxima of a_f .⁴⁶

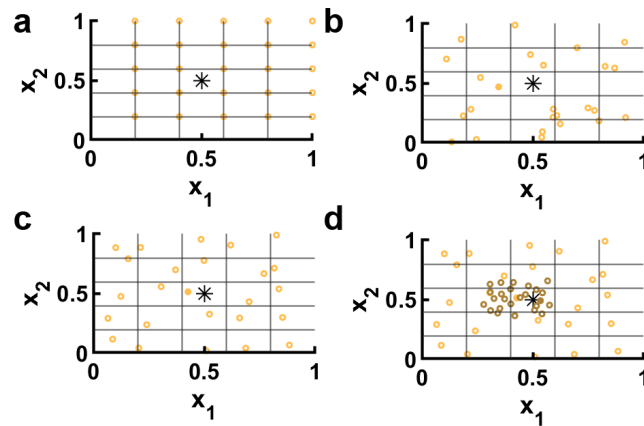


Figure 2.2: Selection methods for sampling points. a, Grid based sampling (orange open circles) will sample the same points each time. The target maximum (black asterisk) and the closest point (solid orange circle) are also shown. b, Random sampling is statistically likely to

cluster in local regions, leaving other regions with no points. c, Latin hypercube sampling (LHS) points are selected by dividing the space into a grid and selecting one point randomly in each domain of the resulting grid. d, A second round of LHS points (dark orange open circles) can be used to further search in a promising region, finding a better point (solid orange) than the originally selection points.

2.3.6: Setting F: Which experimental data points are used to condition the surrogate model

It may seem obvious that one should use all available data to improve the predictive power of the surrogate model, but as campaigns shift from exploration to exploitation, it is useful to focus on high performing regions. This counterintuitive approach has two rationales. First, as the number of experiments increases, the time to condition the model and predict the performance at the sampling points increases to the point where it can slow the progress of the SDL. By focusing on the experiments that are most likely to be relevant to the area of interest, the experiment selection process can avoid becoming a bottleneck. Second, focusing on the region of interest allows the hyperparameters to be tuned to the region of interest. This focusing may be necessary if the output space is inherently different on average than the region of interest, *i.e.* it is a needle in a haystack. This challenge can also manifest if the density of experiments in the region of interest is higher than it is in the rest of the parameter space, allowing shorter length scale hyperparameters to be employed in a manner that increases prediction accuracy. The opportunities inherent in pruning the available data have been recently noted in both the chemical and material spaces.^{64,65} There are several algorithms that have been developed to do this pruning including zooming memory-based initialization (ZoMBI)⁶⁶ and trust region Bayesian optimization (TuRBO).⁶⁷

An interesting related point is how to deal with outliers. Algorithmically rejecting outliers is a risky strategy as in a needle-in-a-haystack type search, the needle is likely to look like an outlier. For this reason, it is easier to justify including or excluding data based on its position in \vec{x} rather than its measured value of y . That said, all available data and characterization should be used to assess whether a given experiment was conducted correctly and therefore whether the measured value can be trusted.

2.4: Sources of Feedback to Adjust the Knobs – the Gauges

Having outlined the main variables that the experimenter can adjust, now we turn to the ways in which the experimenter can determine when to make adjustments. In other words, we seek to identify the avenues for obtaining real-time information during a campaign that provides actionable feedback that can be used to change settings **A-F** listed above. Crucially, given the complexity of SDL campaigns and the opacity of many ML models, we seek to draw direct connections between graphical observables that exist in spaces of any dimensionality and each setting **A-F**. In this way, the experimenter can build intuition for the connections between these items and exercise their agency in fruitfully guiding the SDL. In particular, we have identified four sets of graphs that the experimenter should be examining and listed how features of these graphs can provide guidance when adjusting SDL settings.

2.4.1: Response scatter plots

After each additional experiment, a set of d_m plots should be generated that each feature y vs. a single dimension of \vec{x} with the most recently added data highlighted. Older

data points can be colored to illustrate the order in which they were collected (Figure 2.3). Looking at one dimensional slices of a multidimensional space can only offer limited insight, but one thing that such plots can do quite well is illustrate whether top performing points are on a boundary. Such an observation would be an indication that the boundaries of \vec{x} should be enlarged if possible to capture this trend, thus providing feedback on parameter range (A). We note that if boundaries are not independent, rather than plotting y vs each component of \vec{x} , it may be more useful to plot y vs. the distance from the nearest boundary.




Figure 2.3: Response scatter points. a, Observed response y vs. x_1 . Points colored based on the order that they were tested (light gray to black), with the most recent test data shown as a blue star. In this panel, the peak of y is firmly within the boundaries of x_1 , indicating that the range is likely appropriate. b, Observed y vs. x_2 . In contrast with x_1 which is a design variable that can be chosen continuously, x_2 corresponds to a material property and therefore only discrete values

corresponding to known materials can be selected. In this panel, the peak of y is on the boundary that cannot be expanded due to material constraints, pointing to potential material development goals.

2.4.2: Parity plots

A crucial plot to consistently observe is the parity plot, or the plot of \tilde{y} vs. y (Figure 2.4). Naturally, data on this plot should fall on the line $\tilde{y} = y$, but imperfections in the model and experimental noise will prevent this from occurring. Of particular interest is the result of the most recent experiment and so the model that was used to select the last experiment should be the one used to compute \tilde{y} . In this way, it is possible to assess the goal of the experiment (how good the point was predicted to be) as well as its accuracy (how close the experiment was to the prediction). The general appearance of this parity plot is a crucial metric in determining the acquisition function (**D**), namely, to adjust the priority of the policy to more heavily lean towards exploration or exploitation. In brief, if the parity plot exhibits a very poor correlation, then any attempt at exploitation will likely fail and exploration should be prioritized. It should be noted that while some policies, namely UCB, feature a discrete knob that allows one to change the relative focus on exploration versus exploitation, not all policies have such a setting, so changing a_f entirely may be necessary.³⁹ It is also possible to use combinations of multiple acquisition policies.^{17,42,68}




Figure 2.4: Parity plots. a, Predicted response \tilde{y} vs. y built from 13,250 data points. Red line represents $\tilde{y} = y$ and pink lines represent estimated measurement error σ_y . Blue star represents the latest experimental prediction and the subsequent result. b, Same model as a, but zoomed in on the upper section of the data to highlight the flat plateau that occurs at $\tilde{y} \approx 0.7$, leading to significant underprediction of performance for the latest experiment. c, Model built from 4,060 data points near **max** (y) shows improved performance at high values, which are the most important in a maximization problem. Prediction of σ_y is also decreased. d, Model built from 410 data points illustrating deteriorated performance and a very small estimated σ_y due too little data.

The overall behavior of the parity plot is also crucial term in validating the surrogate model (C). Two aspects of fit quality should be immediately apparent when examining the parity plot. First, the spread of the data about the trend line provides an excellent proxy for prediction accuracy and this can be quantified using the square of the Pearson's correlation coefficient r^2 . Since these are in-training sample predictions, training conditions exist for a perfect match between model and data. However, it is important to remember that the model is Bayesian with the expectation that experimental data has uncertainty and thus the spread about the parity line should approach σ_y . As

such, one would expect $\sim 65\%$ of data points to fall between lines parallel to the parity line but spaced apart by $\pm\sigma_y$, shown as light red lines in Figure 2.4. If nearly all of the data falls between these lines, this is a signal that the model is overfitting. While generally r^2 will increase as the volume of data increases, it is also highly dependent upon the model hyperparameters θ and any transformations made to the input space. Thus, low r^2 reflects the need to pay careful attention to these terms. We recommend making decisions to set these terms based upon minimizing cross-validation error or through maximum likelihood estimation. In addition to spread about the parity line, if the residuals of the data are not uniformly distributed, it may reflect an incorrect bias in the model. As GPRs are nearly bias-free, they should not exhibit this phenomenon.

The parity plot can also clearly communicate what data should be included in the model (**F**). In particular, for maximization problems, the data that reflects the largest values of y are in many ways the most important. A common occurrence with GPRs is that they will effectively smooth out sharp peaks in parameter space due to their inability to model wide low-performing regions while also modeling narrow high-performing regions, or needles in a haystack. This problem is manifest in the parity plot through a flattening of \tilde{y} at the high end (Figure 2.4b). Under these circumstances, it will be very challenging to tease out accurate predictions in the high-performing region. One solution is to adjust the data used to train the model (**F**). In particular, one can retrain the model using only the data in the proximity of the high-performing points, thus ensuring that the model accurately captures this region (Figure 2.4c). Variations exist to this approach such as omitting data based on y (only including high performing experiments), the location in

\vec{x} relative to $\text{argmax}(y)$ (only including experiments near the best performing experiment), or the location in \vec{x} relative to $\text{argmax}(\tilde{y})$ (only including experiments near the best predicted sampling point of an initial GPR model). However, if too few data points are included, the model performance can deteriorate (Figure 2.4d).

2.4.3: Proximity plots

While hunting for extrema in complex and high dimensional parameter spaces, a crucial metric to examine is how close experiments are to one another. In order to readily visualize this metric as the campaign progresses, y should be plotted vs. $\rho =$

$\sqrt{\sum(x_j - x'_j)^2}$, where \vec{x}' is the location of the most recently sampled point (Figure 2.5).

and the subscript j corresponds to an iterator over each dimension in the parameter space. A virtue of plotting data in this way is that it makes it very clear the degree to which the SDL is prioritizing exploration vs. exploitation for the most recent experiment. Simply – the further the sampled point is from nearby points, the more exploration focused the choice of the policy. This makes such plots very useful for evaluating the behavior of the decision policy (**D**) and ensuring that it is behaving as expected. This is also very useful when using policies like EI that themselves balance exploration and exploitation. The distance of the nearest point is a clear indicator of how the policy is resolving this tradeoff. In addition to providing feedback into the acquisition function, proximity plots can be useful in helping select the data used to train the model (**F**). When training a model with some subsample of the available data, it is useful to indicate this data on the proximity plot to illustrate the complexity of the subspace under consideration.

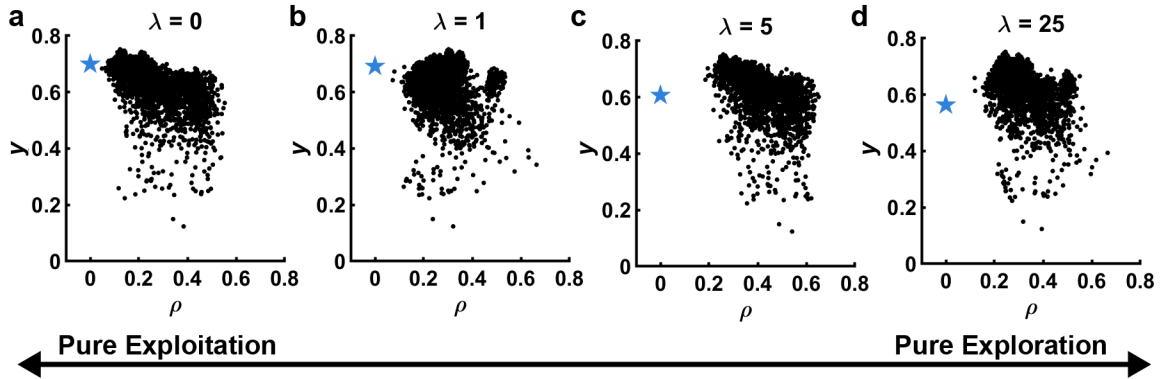


Figure 2.5: Proximity plots. Observed y vs. distance to selected point ρ . Black points represent previous experiments, while the blue star represents the selected experiment and its predicted performance \tilde{y} . The distance of the point closest to the star is an indicator of the exploration/exploitation tradeoff. In this panels, this tradeoff is illustrated by selecting a point based on upper confidence bound with varied hyperparameter λ . a, $\lambda = 0$, pure exploitation. b, $\lambda = 1$, balance of exploration/exploitation. c, $\lambda = 5$, focus on exploration. d, $\lambda = 25$, nearly pure exploration.

Despite the utility of proximity plots, a question arises in how to normalize ρ to make the data most meaningful. For instance, considering a normalized d_m -dimensional parameter space, the furthest two points can be from one another is $\sqrt{d_m}$, making Euclidian distance in a normalized space a very useful way of thinking about distances. That said, the relevance of distances depends on the topography of the response surface and not just the chosen bounds. One way to take this into consideration when using a GPR is to normalize each dimension by the length scale hyperparameter θ_j associated with that dimension. In this way, the magnitude ρ is closely connected to covariance with $\rho \sim 2$ indicating that only 5% of the sample point's variance is captured by the neighbor. Unfortunately, this approach can be challenging as $\vec{\theta}$ can change and in particular shrink as additional data becomes available that elucidates high-spatial-frequency variations. This makes it hard to track the performance of this plot over time. Interestingly, when

using proximity plots normalized by GPR hyperparameters, the behavior of points $\rho < 2$ illustrates how well the GPR is performing in the experimental region. If y is spread widely despite being considered statistically close by the GPR, it is a signal that the model is not performing well in the region of interest and the surrogate model **(C)** needs to be reevaluated.

2.4.4: Performance gamut

Since the ultimate goal of the SDL, in at least this example application, is the optimization of performance y , a crucial piece of information is evaluating the full spectrum of y that has been observed thus far (Figure 2.6). This plot is ultimately the main avenue for collecting feedback on the choice of the metric **(B)**. Even though it is ultimately necessary to make decisions based on a single metric, this is not always the most useful for conceptualizing the response surface. Humans visualize information easily in two dimensions and thus, it is very useful to conceive of ways to define two dimensions to aid in visualization, even if they are not directly related to the SDL decision-making process. For a high dimensional \vec{y} , unsupervised approaches such as principal component analysis or autoencoders may be useful for reducing the performance space to two variables. When optimizing a scalar metric, it is useful to introduce another value that provides a second dimension for visualization. Ideally, this second metric would be related to the first, but not directly connected. For example, in a campaign to optimize toughness, material stiffness could be a useful second property for visualization as it would help distinguish between tough but compliant materials and brittle but stiff materials, even if they had the same toughness.

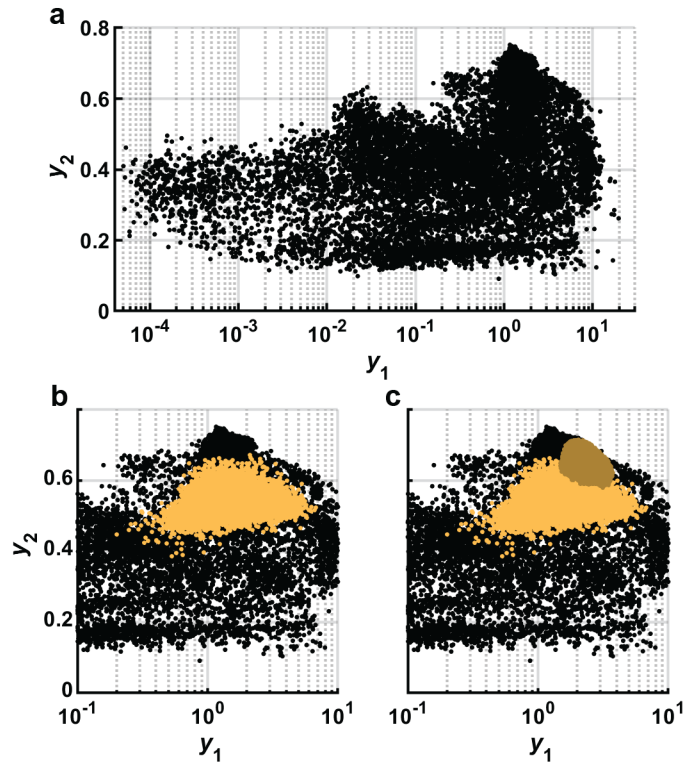


Figure 2.6: Performance gamut. a, Experimental data (black dots) are represented by two dimensions: y_1 and y_2 . b, To find the maximum of the acquisition function, LHS is used to evaluate sampling points (orange) using the zoomed-in GPR model from Figure 2.4c. c, A second pass of LHS (dark orange) is done with the same GPR model but in a smaller region to more closely locate the maximum of the model.

In addition to plotting the actual values of y , it is useful to plot on the same axis the \tilde{y} that were in consideration for the most recent experiment as this can be useful in determining the number of sampling points used when finding the maximum of the acquisition function (**D**). The predicted points should be dense enough that they fill a reasonable region in the performance gamut with predicted points that extend above the highest observed experiments. If the observed prediction space appears too sparse, one solution is to increase the number of sampling points. However, this is not the only consideration when deciding the number of points to sample. In particular, the process of

training a model and selecting sampling points can take non-trivial amounts of time. The easiest way to modulate this time is by adjusting the number of sampling points (**E**).

Given that the tempo of an SDL campaign is naturally set by the pace of collecting experimental data, one approach is to adjust the number of sampling points until the time required to perform experiments is commensurate with the time to choose subsequent samples. Further considerations are necessary when experiments are performed in batch or using asynchronous agents, but the availability of supercomputing clusters means that running prediction models in parallel can increase the prediction rate.

2.5: Concluding Remarks

Collectively, the principles and heuristics presented here provide a blueprint for sustained operation of an SDL in which SDL-driven discoveries and human-led interventions together collectively accelerate the research enterprise. The focus on two complementary aspects – the settings that govern SDL operation and data showing the SDL performance in real time – draws a close analogy with operating automobiles in which drivers have a clear idea of how to act on feedback from each gauge. While this work draws from our experience running SDLs, there are many other practitioners with valuable insight, and so it is expected that these methods will require updating as more sophisticated methods and tools become available. Indeed, we anticipate that this is a facet of human-machine interaction that is sure to become increasingly complex in the coming years. It should be noted that there are also opportunities for meta learning in which processes such as reinforcement learning are used to adjust the knobs discussed in this work, which could be further leveraged to increase the pace of learning.

CHAPTER 3: SUPERLATIVE MECHANICAL ENERGY ABSORBING EFFICIENCY DISCOVERED THROUGH SELF-DRIVING LAB-HUMAN PARTNERSHIP

This chapter is adapted from work published in *Nature Communications*⁶⁹.

3.1: Introduction

Structural motifs define the ways we efficiently use materials. For instance, the ubiquity of I-beams in architecture is due to the efficiency of this shape in resisting both shear and bending.^{70,71} Natural structures feature similar examples such as the hollow circular cross-section of bamboo providing high bending and torsional resistance.^{72–75} For the large class of structures designed to provide protection under a compressive load, the key property to consider is the total mechanical energy absorbed during compression.^{76–79} This desire to discover tough structures has motivated a focus on metrics like energy dissipated per unit volume or per unit weight. However, in compression, it is nearly always possible to increase the applied stress to absorb more energy, thus when considering specific energy absorbed, or another metric with a dimension, one must also define an operating stress. Along these lines, there are practical restrictions to absorbing energy in any engineering application, for example, that the stress must be held below a level that would damage the system to be protected.⁸⁰ Collectively, these restrictions mean that terms like specific energy absorption are not easily applied in comparing tough structures at different stresses. Therefore, it is useful to define energy-absorbing efficiency K_S , a non-dimensional measure of how much energy is absorbed without surpassing a given threshold stress.^{7,81} Unfortunately, K_S is difficult to optimize because

most of the energy absorbed by a structure designed for mechanical protection occurs beyond the elastic regime where deformations are highly non-linear, often feature dynamic self-self contacts, and are challenging to model.

As a result of the challenge of designing tough structures, much work has focused on known, relatively simple motifs such as honeycomb lattices or cylindrical shells that have an analytical basis for performing well.^{7,82} Others have drawn inspiration from nature to identify more complex structural motifs.^{83–86} Computational approaches including finite element analysis (FEA) and machine learning-based approaches have also been widely used to design complex geometries.^{87–93} These computational approaches pair well with additive manufacturing, which allows the fabrication of complex designs.^{10,94–97} Nevertheless, the fabrication of candidate structures is often the limiting step in the design process and is commonly limited to validating designs. And furthermore, despite the speed and versatility of computational approaches such as FEA, it is very challenging and sometimes impossible to accurately capture K_S using computation because of the complex interplay of material plasticity, material non-linearities, structural non-linearities, and dynamic self-self contacts.^{98–100} Furthermore, models studied by FEA often deviate from reality due to unavoidable processing-dependent defects and variability of the physically realizable structures. Thus, improvements to the maximum energy absorbing efficiency of a part at any stress, known as K_S^* , remain slow: to date, additively manufactured structures designed for energy absorption typically feature $K_S^* < 50\%$ (Figure 3.1). There exist better synthetic materials, the best being a plastic foam reported to have reached $K_S^* = 68.1\%$ (Table 3.1).¹⁰¹

However, this record is surpassed by nature: Balsa wood has the highest previously achieved K_s^* , 71.8%, showing the value of millions of years of evolution.¹⁰² It is clear that new approaches are needed if the performance envelope of this important property is to be improved.




Figure 3.1: Common and superlative structures and materials. Values of critical stress threshold σ_t^* and critical energy absorbing efficiency K_s^* for synthetic structures (blue diamond) and natural materials (green triangle) gathered from literature, with superlative plastic and hyperelastic generalized cylindrical shells (GCS-this study) components (red circle).

σ_t^* (MPa)	K_s^* (%)	Material origin	Reference
8.50×10^{-3}	38.5	Synthetic	https://doi.org/10.1016/j.matdes.2017.11.037
1.00×10^{-2}	32.0	Synthetic	https://doi.org/10.1016/j.matdes.2017.11.037
2.85×10^{-2}	63.2	Synthetic	This work – ADTS ID 22335
5.55×10^{-2}	44.4	Synthetic	https://doi.org/10.1016/j.matdes.2017.11.037
9.17×10^{-2}	37.4	Synthetic	https://doi.org/10.1016/j.matdes.2017.11.037
6.47×10^{-1}	29.7	Synthetic	https://doi.org/10.1002/admt.201800419
9.01×10^{-1}	47.3	Synthetic	https://doi.org/10.1002/admt.201800419
1.01	68.1	Synthetic	https://doi.org/10.1177/0021955X06063519
1.04	36.9	Synthetic	https://doi.org/10.1016/j.actamat.2004.05.039
1.16	75.2	Synthetic	This work – ADTS ID 21285
1.47	33.5	Natural	https://doi.org/10.1016/j.jmbbm.2019.103603
2.42	32.3	Synthetic	https://doi.org/10.1177/0021955X06063519
2.45	70.9	Natural	https://doi.org/10.1016/S0167-6636(02)00268-5
3.08	58.2	Synthetic	https://doi.org/10.1016/j.msea.2004.03.051
3.89	59.8	Synthetic	https://doi.org/10.1177/0021955X06063519
4.07	32.4	Synthetic	https://doi.org/10.1177/0021955X06063519
4.30	52.9	Synthetic	https://doi.org/10.1016/j.msea.2004.03.051
4.76	46.8	Synthetic	https://doi.org/10.1016/j.actamat.2004.05.039
5.53	44.5	Synthetic	https://doi.org/10.1016/j.ijimpeng.2010.03.007
6.47	71.8	Natural	https://doi.org/10.1016/S0167-6636(02)00268-5
6.88	44.2	Synthetic	https://doi.org/10.1177/0021955X06063519
9.99	41.5	Synthetic	https://doi.org/10.2140/jomms.2013.8.65
1.01×10^1	45.0	Synthetic	https://doi.org/10.1016/j.actamat.2004.05.039
1.17×10^1	23.8	Synthetic	https://doi.org/10.1002/admt.201800419
1.21×10^1	39.4	Synthetic	https://doi.org/10.1177/0731684419868018
1.24×10^1	47.9	Synthetic	https://doi.org/10.1177/0731684419868018
1.34×10^1	45.6	Synthetic	https://doi.org/10.1177/0731684419868018
1.42×10^1	43.2	Synthetic	https://doi.org/10.1177/0731684419868018
1.46×10^1	52.8	Synthetic	https://doi.org/10.1177/0731684419868018
1.50×10^1	56.1	Natural	https://doi.org/10.1016/S0167-6636(02)00268-5
1.52×10^1	51.9	Synthetic	https://doi.org/10.1177/0731684419868018
1.74×10^1	53.3	Synthetic	https://doi.org/10.1177/0731684419868018
1.92×10^1	50.3	Synthetic	https://doi.org/10.1177/0731684419868018
1.98×10^1	47.6	Synthetic	https://doi.org/10.1177/0731684419868018
1.98×10^1	33.2	Synthetic	https://doi.org/10.1002/admt.201800419
2.22×10^1	47.7	Synthetic	https://doi.org/10.2140/jomms.2013.8.65
2.22×10^1	46.1	Synthetic	https://doi.org/10.1177/0731684419868018
2.25×10^1	47.5	Synthetic	https://doi.org/10.2140/jomms.2013.8.65
2.29×10^1	53.5	Natural	https://doi.org/10.1016/S0167-6636(02)00268-5
2.69×10^1	45.8	Synthetic	https://doi.org/10.2140/jomms.2013.8.65
2.86×10^1	32.1	Synthetic	https://doi.org/10.1177/0021955X06063519
3.88×10^1	45.8	Natural	https://doi.org/10.1016/S0167-6636(02)00268-5
5.84×10^1	34.3	Synthetic	https://doi.org/10.1016/j.ijstr.2015.02.020
8.86×10^1	53.6	Synthetic	https://doi.org/10.1016/j.msea.2004.03.051
1.40×10^2	51.5	Synthetic	https://doi.org/10.1016/j.msea.2004.03.051

Table 3.1. Common and superlative structures and materials: Values of stress threshold σ_t^* and energy absorbing efficiency K_s^* not directly reported were computed based on reported force-displacement or stress-strain plots.

Here, we utilize a self-driving lab (SDL) to test >25,000 additively manufactured structures in a large-scale data-driven campaign to discover tough structures with superlative K_S^* . SDLs are robotic research systems that select, perform, and analyze physical experiments without needing human intervention,^{31,103} and they have been productively employed in chemistry,^{17,104} materials science,¹⁵ mechanics,²⁶ and microscopy.^{23,105} Motivated by the observations that SDLs can progress toward user-chosen goals faster than either high-throughput experimentation²⁷ or tests chosen by subject matter experts,^{29,44,106,107} we hypothesize that an SDL allowed to explore seven polymers in an 11-dimensional parameter space over trillions of possible designs can discover new structural motifs that advance the frontier of K_S^* . The result of this sustained human-machine collaboration is that we realize a structure with $K_S^* = 75.2\%$. In addition to showing the opportunities for SDLs to overcome design barriers, this campaign results in a vast, labeled dataset that has implications for both mechanics and design more generally. For instance, we explore two high-performing structural motifs and find that they exhibit consistent performance within classes of materials, namely plastic or hyperelastic polymers. Finally, aggregate analysis of this data provides general design heuristics that allow for the efficient selection of materials and structures.

3.2: Results

3.2.1: Defining a campaign to study generalized cylindrical shells

As a motivating example to explore the considerations that define and limit K_S^* , we consider the compressive behavior of a cylindrical shell composed of a hyperelastic thermoplastic polyurethane (TPU). When tested in compression, the resulting force F -

displacement D curve shows an initial elastic region, a yield point, and then complex post-yield behavior that originates from combinations of plastic deformation, buckling and other large elastic deformations, and reentrant contact (Figure 3.2a). To compute K_S^* , F - D is first converted to stress σ vs. strain ε for the effective medium using the dimensions of the component, as explained in section 3.4.3. K_S is specified at a threshold stress σ_t that is typically associated with the strength of the system to be protected. Graphically, K_S represents the amount of energy absorbed by the component while $\sigma \leq \sigma_t$ (Figure 3.2a – blue region) relative to the maximum energy that could be absorbed during complete compression ($\varepsilon = 1$) without exceeding σ_t (Figure 3.2a – red rectangle). To compute K_S at a specific σ_t , we numerically evaluate,

$$K_S = \sigma_t^{-1} \int_0^{\varepsilon_t} \sigma(\varepsilon) d\varepsilon, \quad (3.1)$$

where ε_t is the greatest strain at which $\sigma \leq \sigma_t$ for all $0 < \varepsilon \leq \varepsilon_t$. Interestingly, for most structures, $K_S(\sigma_t)$ has a single well-defined maximum K_S^* at an optimum threshold stress σ_t^* . This σ_t^* often corresponds to the initial yield stress of a structure, although it can occur at higher stresses, particularly when the component densifies at low strains or does not significantly soften after yielding. As such, K_S^* and σ_t^* were found for this and all structures numerically as described in Section 3.4.3. In the example of Figure 3.2a, the cylindrical shell is limited to $K_S^* = 39.8\%$ due to significant post-yield softening. To maximize K_S^* , a flat post-yield region and a delay of densification until large ε are both desirable. Unfortunately, this knowledge alone does not provide a prescription for how to adjust the structure to obtain these desired behaviors.

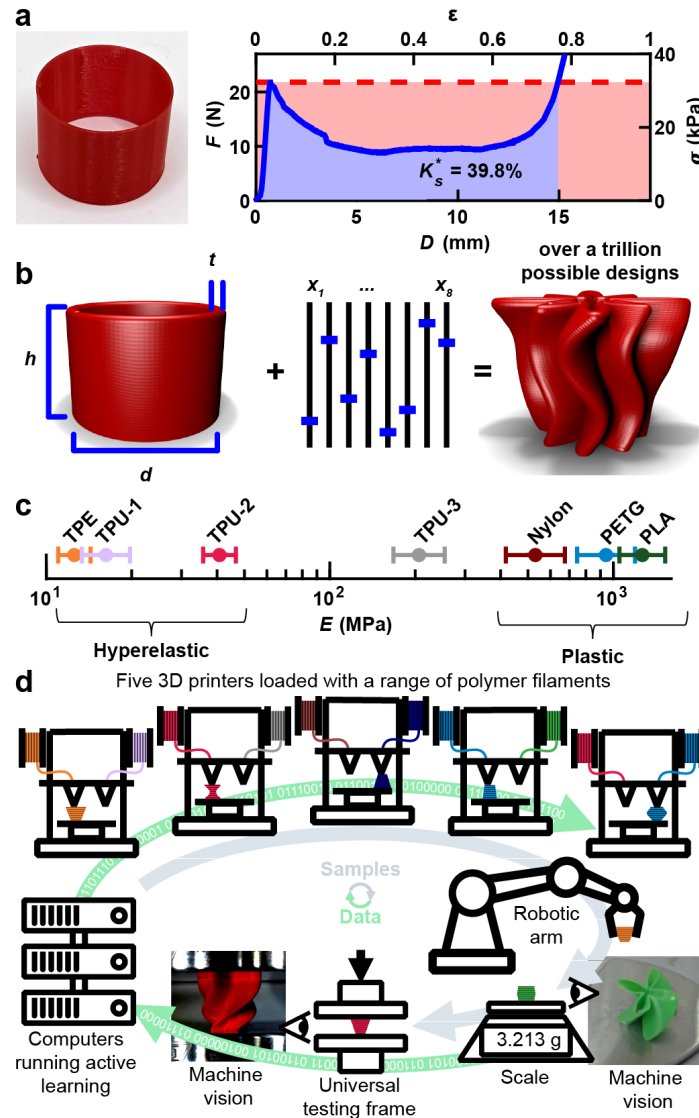


Figure 3.2: Challenge of designing energy absorbing structures. a, Force F vs. displacement D and effective medium stress σ vs. compressive effective medium strain ϵ for an additively manufactured cylindrical shell made of thermoplastic polyurethane (TPU). Maximum energy absorbing efficiency K_s^* is calculated at an optimum threshold stress σ_t^* (dashed line) by dividing the energy absorbed while $\sigma \leq \sigma_t^*$ (blue region) by the theoretical maximum amount absorbed (red rectangle). b, Eleven independent geometric parameters including diameter d , height h , wall thickness t , and eight other parameter x_{1-8} that together define a generalized cylindrical shell (GCS). When combined, at least trillions of unique designs are possible. c, Elastic modulus E of the seven polymers studied as determined by compression tests. Error bars represent one standard deviation. d, Schematic showing an autonomous research system in which five 3D printers are used to fabricate polymeric structures that are automatically weighed, imaged, and tested using quasistatic compression. The output of this testing is automatically interpreted and used to select subsequent designs to test.

We hypothesized that programmed perturbations to the geometry of a cylindrical shell could tailor the complex post-yielding behavior to drastically increase K_s^* . While cylindrical shells are typically defined by a small number of geometric parameters, namely their diameter d , height h , and thickness t , we augmented these to form an 11-parameter family of structures termed generalized cylindrical shells (GCS) (Figure 3.2b). In addition to t , h , and d , a GCS is defined by eight additional parameters x_{1-8} including four that adjust the cross-sections of the top and bottom of the shell, one to define the perimeter of the top relative to the perimeter of the bottom, and three to describe the rotation of the perturbations from top to bottom (Figure 3.12). Furthermore, because the GCS space does not have circular cross sections for most designs, we define $d = \frac{P_0}{2\pi}$, where P_0 is the average perimeter of the design. We note that each of the eleven parameters would only need twelve unique values for their combinations to surpass a trillion unique designs. Given that each parameter is continuous and can be assigned many more than twelve values, we consider trillions of unique designs to be a lower limit to the size of the parameter space. While we have previously used FEA to compute mechanical metrics like resilience, stiffness, and yield force,²⁸ our inability to rapidly or accurately calculate the shape of the post-yield region using FEA led us to not pursue this method for accelerating the study of K_s . Therefore, experiments were necessary for assessing K_s . Because all the resulting structures are topologically equivalent to cylindrical shells, they can be fabricated using extrusion-based additive manufacturing by continuously extruding material, thus making this parameter space intrinsically designed for additive manufacturing. In addition to the geometric parameters, we sought to explore

a variety of polymers. Therefore, we considered seven materials that included those that are hyperelastic such as a thermoplastic elastomer (TPE) and TPU, those that plastically deform such as polylactic acid (PLA), and materials that fall in between these two distinct classes (Figure 3.2c). These materials can be characterized based on their elastic modulus E , their plateau strength σ_p , and rebound strain.

To efficiently search the effectively infinite GCS parameter space, we employed the Bayesian experimental autonomous researcher (BEAR), a customized SDL developed to combine additive manufacturing of polymers and mechanical testing (Figure 3.2d). The BEAR is a closed-loop system in which samples are printed using one of five fused filament fabrication (FFF) printers, automatically retrieved using a six-axis robotic arm, and then characterized using a scale, machine vision, and uniaxial compression testing. After testing, the information was automatically analyzed to determine whether the test was of acceptable quality. This learning process featured a fault-tolerance that was a combination of intrinsic and explicit factors. Specially, if components had no structure due to large overhangs, small contact area with the print bed, non-manifold surfaces due to high twist values, or print-head collisions caused by excessive extrusion, they would be tested and exhibit low efficiency, naturally guiding the algorithm away from this region of design space. If the sample was so poorly defined that it could not even be tested, these experiments were automatically marked as unprintable to prevent the algorithm from repeatedly selecting this or similar designs. In both cases, these assignments were manually confirmed asynchronously by the human team. Subsequent experiments were selected using Bayesian optimization, which entailed conditioning a

surrogate model of the mechanical performance using all previously measured GCS components and then selecting combinations of designs and materials that maximized a specified acquisition function. As this experimental campaign began with a very small amount of data, Bayesian optimization using Gaussian process regression was selected to model this data because it is efficient in the low-data regime.¹⁰⁸ Now that a large database exists, more sophisticated models such as variation autoencoders could be employed to more accurately model the design space and facilitate in the selection of new experiments.¹⁰⁹ The SDL autonomously performed these tasks to choose, perform, and analyze experiments at a typical pace of ~50 experiments per day. Collectively, 25,387 experiments were performed using seven different materials, with 13,250 experiments resulting in valid data. It is worth emphasizing that even experiments not included in the final dataset provide value, for example in determine the subspace of GCS designs that are printable using FFF. From the beginning of the campaign to the end, experiments were running for ~60% of the total wall clock time, showing the robustness of this process. This system is an evolution of an SDL developed by our research group.²⁷ A picture of the system (Figure 3.13) as well as full details on the hardware (Figure 3.14) and software (Figure 3.15) used as part of the BEAR are provided in Section 4.2.

3.2.2: Discovering high-performing structures

Over the course of the multi-year campaign, the details of how experiments were chosen were altered based on the intuition of the experimenters and by evaluating the progress of the BEAR. Examples of these changes include, the introduction of sinusoidal twist, the switch to LHS sampling (from grid-based sampling), allowing components to

have both sinusoidal and linear twists combined, and switching to cooling plastic materials after printing. The timing of these changes is shown in Figure 3.3a. Researchers also controlled which filaments were loaded into which nozzles. New filaments were introduced during the campaign and the mix of filaments was changed to pursue different goals, as summarized in Figure 3.3b. Finally, 23 different decision policies were used throughout the campaign, as shown in Figure 3.3c and Table 3.4.

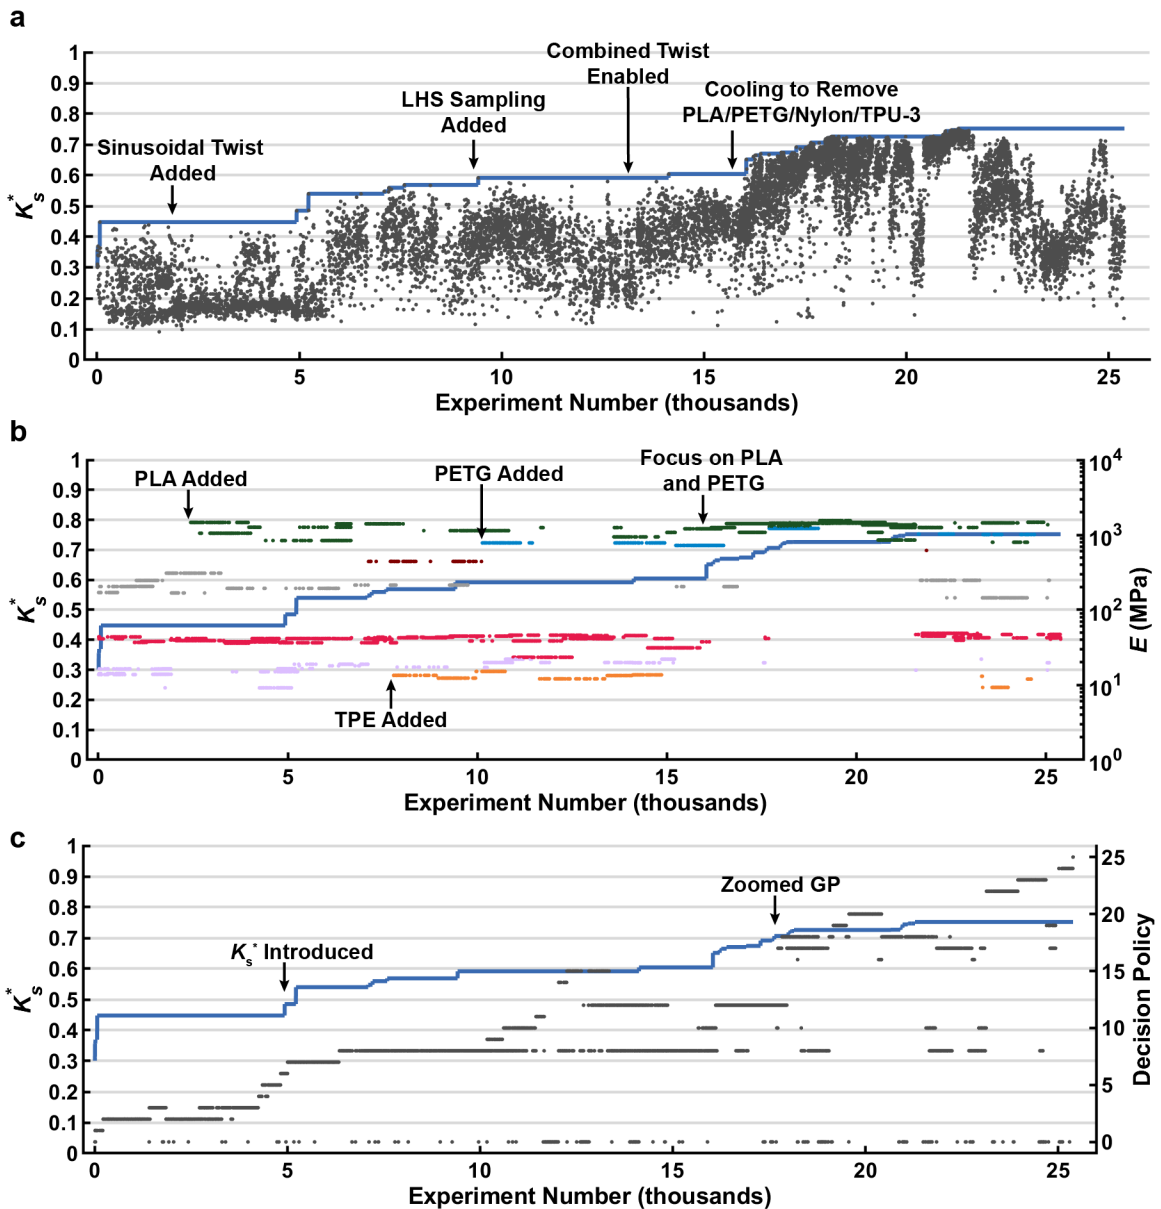


Figure 3.3 Details of the human/machine collaboration. a, K_s^* of each successful test (gray dots), along with the highest K_s^* to date (blue line). Key changes to the processing and sampling space are marked. b, Modulus of each experiment's filament roll plotted in semi-log (right axis) and colored according to the Figure 3.2c, along with the highest K_s^* to date (blue line – left axis). c, Decision policy of each experiment (right axis), along with the highest K_s^* to date (blue line – left axis). Decision policies are listed with descriptions in Table 3.4.

Progress in the campaign can be visualized by tracking K_s^* measured for each experiment along with a running maximum throughout the campaign (Figure 3.4a). The

continuous progression was a result of both persistent experimentation by the SDL and choices made by the human team based on the progress of the SDL. Interestingly, large jumps in performance were typically either due to serendipity (*i.e.* the SDL chose a fortuitous region) or a human-led intervention. For example, we highlight three human interventions based on observing the progress of the SDL. First, prior to experiment 4,829, the SDL was programmed to select experiments based on K_S at one specific σ_t . However, we noted that there were several different reasons why a specific sample would have a low K_S , so we needed to provide the SDL with more information. We hypothesized that tracking both K_S^* and σ_t^* from each experiment would allow for more meaningful information to be extracted by the SDL. After implementing this change, the SDL rapidly increased the maximum K_S^* from 45% to 55%. As a second example, at experiment 15,678, we noted that a large fraction of plastic components was failing the height quality control check but passing the mass check. We had been heating the print bed after printing to facilitate the automated removal of components but determined that the forces exerted during removal could deform plastic components. Upon changing the SDL to cool plastic components prior to removal, the system proceeded to make a series of jumps in maximum K_S^* from 60 to 68%. Finally, at experiment 17,730 we noted that the predictive model used by the SDL was systematically underpredicting K_S^* for high performing components, so we implemented a process where the proposed experiment was selected using a model built only on data close to the best observed experiment, a process similar to algorithms such as TURBO or ZOMBI.^{66,67} This intervention led the SDL to progress from 70.6% to 75.2% in maximum K_S^* .

The culmination of these adjustments and continued experimentation by the SDL resulted in the observation of $K_s^* = 75.2\%$, a value that was higher than had been previously reported. The performance of this superlative experiment is shown in Figure 3.4b, which shows the $\sigma - \varepsilon$ curve and photographs of the component at different stages of compression. It is clear from the flatness of the post-yield region together with the photographs that the SDL has discovered a way for buckling and other large elastic deformations, plasticity, and reentrant contact to work together to achieve a remarkably flat plateau until densification initiated at $\sim 80\%$ strain. Interestingly, this component was composed of PLA, which is not commonly regarded as a high-performance material. Upon repeated experiments, the design, which we termed Palm, had an average $K_s^* = 73.1 \pm 0.9$. Although Palm printed in PLA had the largest single value of K_s^* , 75.2% , observed in the entire campaign, we discovered other components that had higher average K_s^* values than Palm.

3.2.3: Material influence on design and performance

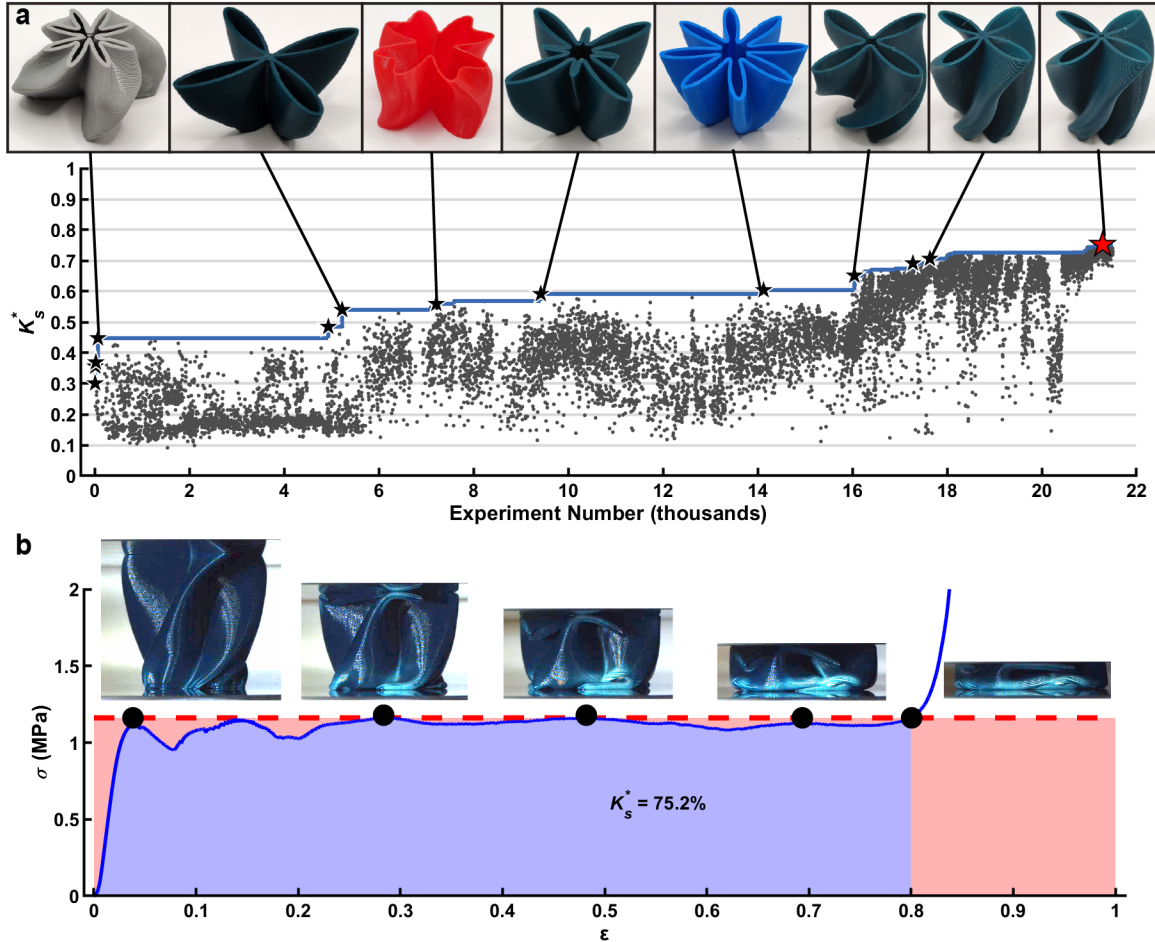


Figure 3.4: Research campaign to find highly efficient structures. a, Each maximum energy absorbing efficiency K_s^* measured over the first $\sim 21,500$ experiments performed. Pictures highlight noteworthy components (black stars) and the highest performing structure (red star). The color of the pictured components is indicative of the material used, with Green indicating PLA, Blue indicating PETG, and Red/Gray indicating different types of TPU. The solid blue line denotes the running best K_s^* observed. b, effective medium stress σ vs. effective medium strain ϵ for experiment 21,285, named Palm, which resulted in $K_s^* = 75.2\%$. Inset photographs show the state of the component at various points indicated on the curve (images enhanced to improve clarity). Shading denotes regions used to compute K_s^* as described in Figure 3.2a.

To explore variations in performance across different material classes, we selected two high-performing designs discovered in different materials. The design discovered in PLA with the highest average K_s^* , termed Willow, is tall and has a compact

center region (Figure 3.5a). Testing 15 identically prepared samples of the Willow design resulted in a tight distribution of yield forces with variations in the post-yield plateau. Nevertheless, PLA components made using this design exhibit $K_s^* = 73.3 \pm 0.9\%$, showing a consistent performance above previously reported maxima. That said, the repeatability of superlative designs on multiple FFF systems suggests that defects in the fabrication process do not limit performance. The highest performing design discovered for TPU-2 is termed Iroko and consists of a relatively short and open design (Figure 3.5b). We observed more substantial variations among the 15 Iroko $\sigma - \varepsilon$ profiles, and the average $K_s^* = 53 \pm 4\%$ was substantially lower than that of the best plastic components. The differences in attainable K_s^* between PLA and TPU-2 can be explained by considering that these are different material classes, with PLA being a glassy polymer that exhibits substantial plastic deformation while TPU-2 is a hyperelastic elastomer. This difference in properties is most evident in their post compression behavior in which the TPU component recovers $\sim 99\%$ of its height one minute after compression while the PLA component is permanently flattened to $\sim 23\%$ of its initial height.




Figure 3.5: Exploration of high-performing designs discovered in elastic and plastic materials. a, Rendering of Willow, a high-performing design discovered using the plastic polymer polylactic acid (PLA) together with effective medium stress σ vs. effective medium strain ϵ for 15 identically prepared PLA Willow components. b, Rendering of Iroko, a high-performing design discovered using the hyperelastic polymer TPU-2, together with σ vs. ϵ for 15 identically prepared TPU-2 Iroko components. c, Measured maximum energy absorbing efficiency K_s^* vs. polymer elastic modulus E for Iroko and Willow components made from one of five polymers. Dashed lines show a sigmoidal fit to guide the eye. Error bars represent one standard deviation. Marker colors denote component composition, as shown in Figure 3.2c.

While Willow and Iroko represent optimizations for their original materials, the question remains of whether the performance of these shapes can translate to other materials or if it is a highly bespoke optimization of this combination of material and design. To explore this, components based on the Willow and Iroko designs were fabricated using four additional materials and tested in triplicate. Studying K_s^* of these components showed the limitations of the transferability of these designs (Figure 3.5c). While each design performed with comparable K_s^* for materials in their respective classes (*i.e.* hyperelastic vs. plastic), a transition region was observed at intermediate E . This is an example of how the mechanical outcome can depend on the confluence of material properties and structure. By examining the individual $\sigma - \epsilon$ curves for the Willow design

(Figure 3.6a), the lower performance of the hyperelastic materials is caused by significant post yield softening. In contrast, for the Iroko design several of the TPU-3 and PETG components experienced buckling events during compression, leading to sharp declines in the force that recovered at higher values of displacement (Figure 3.6b). Because the Iroko design utilizes elastic wall bending to absorb energy, stiffer plastics can buckle abruptly, leading to decreased efficiency. Overall, the comparison of Willow and Iroko confirmed that designs perform well within specific classes of materials, but that these geometric motifs do not directly translate to different material classes.

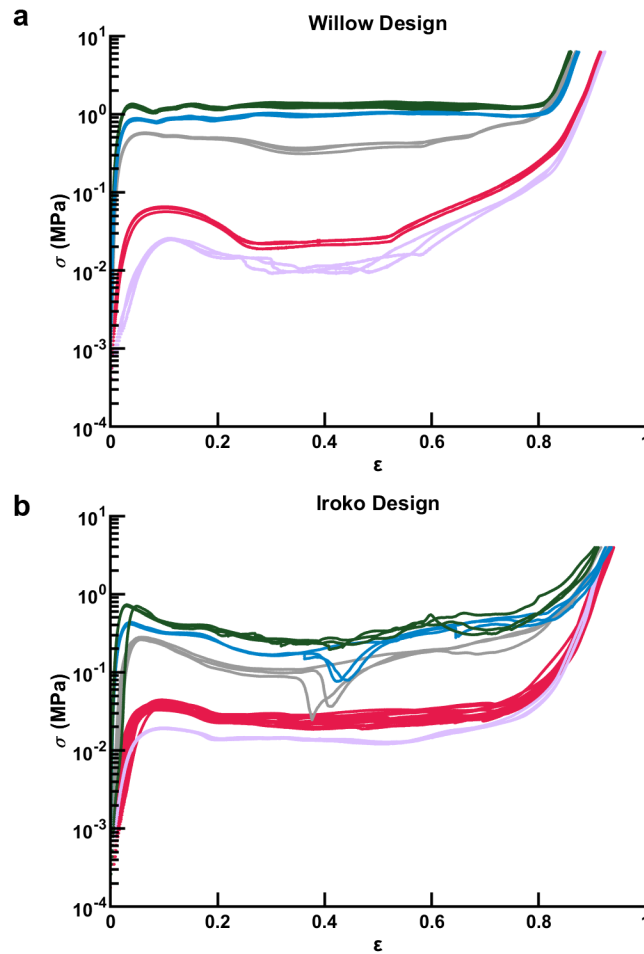


Figure 3.6: Willow/Iroko by material. a, Stress σ –strain ϵ curves for components made using the Willow design printed in TPU-1, TPU-2, TPU-3, PETG, and PLA. PLA, the original Willow material, has 15 tests, while the other materials have three each. b, σ – ϵ curves for components made using the Iroko design printed in the same five materials. TPU-2, the original Iroko material, has 15 tests while the other materials have three each. Colors depict the material as in Figure 3.2c.

3.2.4: Broader design considerations

While the SDL-based campaign was able to discover highly efficient designs, we hypothesized that the broader corpus of mechanical tests performed during this campaign could provide further mechanical insight. As an initial exploration of this idea, the results of all the experiments performed are grouped by material (Figure 3.7). The attainable

envelope of K_s^* and σ_t^* for each material was estimated by computing a convex hull around all experimentally measured points, and the best performance observed for each material occurs at a single σ_t^* , which we denote the material peak threshold stress σ_{tp} . To obtain a measure of the uncertainty in this term, we retroactively step through the campaign and determine each time the σ_{tp} would change and report the expected value as the median of these terms with the error being the standard deviation in their values (in logarithmic space). Interestingly, all materials studied exhibit a similarly shaped convex hull with a distinct peak, highlighting both the importance and the feasibility of tuning the material properties to the specific energy-absorbing application.




Figure 3.7: Convex hulls for seven materials. All tests for each of the seven materials studied, with their final σ_{tp} marked. The materials are TPE (a), TPU-1-3 (b-d), nylon (e), PETG (f), and PLA (g).

We found that over the seven tested materials, σ_{tp} was strongly correlated with the polymer plateau stress σ_p (Figure 3.8a), providing an algorithmic process for selecting a material to optimally match use cases across a wide range of threshold stresses determined from different systems to be protected. Then the material can be selected and structured to maximize K_s at that σ_t . Importantly, we do not believe that σ_{tp} is dependent upon 3D printing capabilities or choice of experiments as this value was not found to drift to higher or lower stresses throughout the campaign. As such, this relationship may have applicability beyond the GCS design family.




Figure 3.8: Design insights that emerge from mechanical dataset. a, σ_{tp} vs. polymer plateau stress σ_p for seven polymers tested during the campaign together with a power law fit shown as a dashed line. Marker color indicates the material type. Error bars indicate one standard deviation of σ_{tp} found throughout the campaign. b, K_s^* vs. relative density ρ_r for all components tested during the campaign with point color denoting K_s^* . c, Normalized height ω vs. d/t for all components tested during the campaign in which point color denotes K_s^* as in b.

In addition to tuning material properties, we hypothesized that unifying features of high-performing designs could be extracted to provide transferable guidance for realizing efficient structures. For example, it is reasonable to expect that the relative density ρ_r of the component would strongly influence K_s^* (Figure 3.8b). This hypothesis is motivated by the observation that the two factors that together bound K_s^* are the flatness of the plateau region and the strain ϵ_d where this plateau drastically rises due to

densification. It has been observed for foams that ε_d is bounded by relative density.^{7,101} Because $\varepsilon_d \geq K_s^*$, we hypothesized that low ρ_r is necessary for high K_s^* . Examining K_s^* vs. ρ_r , we found that K_s^* is peaked at $\rho_r \sim 0.1$ with all designs with $K_s^* \geq 65\%$ having $0.05 \leq \rho_r \leq 0.21$, providing guidance for structural design. Interestingly, because ρ_r can be calculated prior to fabrication, limiting physical testing to designs with ρ_r in this range can eliminate potential components that are unable to achieve high K_s^* . To leverage this fact, we implemented a decision-making policy that focused on designs in this low-density regime. This is both an example of using a metric that is quick to compute to accelerate learning and represents a facet of the human-machine teaming where an observation by the human team helped the SDL search more efficiently.

Beyond the aggregate details of the design, there is a great deal of work exploring the mechanical regimes present for cylindrical shells under uniaxial compression. For example, the ratio d/t of a cylindrical shell determines whether plastic cylinders fail through plastic deformation (thick wall limit) or fail elastically through the formation of local buckles (thin wall limit).¹¹⁰ This transition has been observed to occur at $d/t \sim 100$. Further, the height of cylindrical shells is often characterized by the dimensionless length parameter $\omega = h/\sqrt{dt}$.¹¹¹ Here, cylinders are considered to be short when $\omega < 1.7$. Plotting K_s^* vs. d/t and ω reveals that all of the highest performing structures (*i.e.* $K_s^* \geq 70\%$), which were made from plastic materials, had $16 < d/t < 24$ and $6.75 < \omega < 8.25$, which can be considered thick-walled medium-length cylindrical shells (Figure 3.8c). For simple cylinders in this region, one would expect elastic buckling and plastic deformation to both play major roles. Thus, one way to understand the data-driven optimization

process is that the other eight geometric parameters that define a GCS component have been tuned to guide these complex buckling and plastic interactions to interact constructively.

Interestingly, we may use the tools of machine learning to identify which geometric motifs are most responsible for this improvement. Specifically, we built a neural network to predict K_s^* for components made out of PLA. Figure 3.9a depicts the parity plot of this network. By applying Shapley additive explanations (SHAP),¹¹² we were able to separate the individual contributions made by each GCS parameter to the neural network's predictions of K_s^* . Inspired by Shapley values in Game Theory, SHAP assigns a value to each feature in a machine learning model, indicating its impact on the prediction. We seek to understand the difference in influence between a component and an ideal cylindrical shell (same diameter, height, and wall thickness). To achieve this, we subtract the SHAP values of Willow from the SHAP values for a pure cylindrical shell to obtain a "delta" in explanations. Our analysis of Willow revealed that the four-lobed profile of the cylinder (x_{2-3}) together with the linear and sinusoidal twist of this profile along the shell (x_{6-8}) are together responsible for 90% of the improvement over a simple cylindrical shell (Figure 3.9). Mechanically, this suggests that the key feature for improving the efficiency is producing local plastic deformation events that result in sufficient self-self contacts to prevent post-yield weakening. This feature allows the structure to maximize the material plasticity that occurs while maintaining a flat post-yield region.

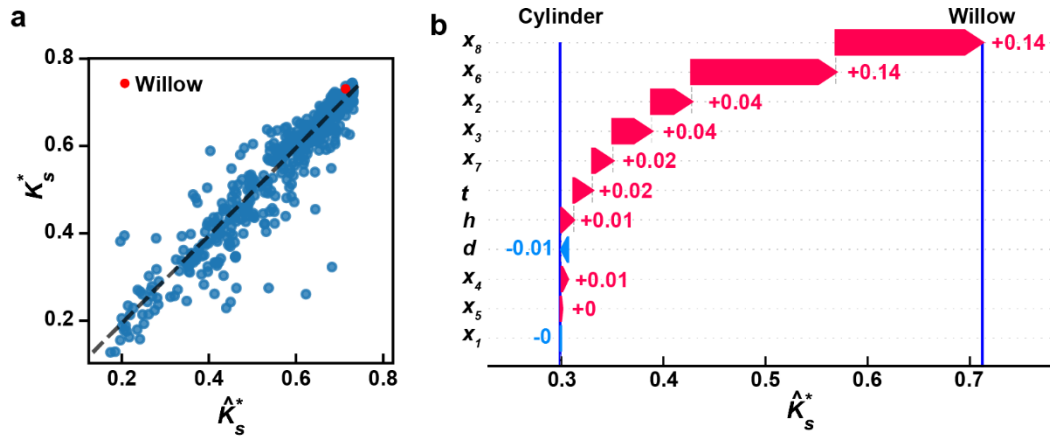


Figure 3.9: Shapley additive explanations (SHAP) analysis of Willow as the superlative design. a, Parity plot of the neural network built on all data taken using PLA with Willow highlighted. b, SHAP waterfall plot for the Willow design tested in PLA relative to a PLA cylindrical shell with the same height, diameter, and thickness. These values show the cumulative effect of positive (red) or negative (blue) contributions of individual feature values to model predictions.

While the main focus of this work has been the discovery of isolated structures that exhibit high mechanical efficiency, and therefore high toughness, it is interesting to consider whether these components could act in concert to be useful for larger systems. For instance, these structures could be tiled in two dimensions to form a regular lattice. To facilitate this, the minimum tileable area of each component was used as the reference area to compute stress. To explore whether the aggregate K_s^* of multiple high-performing components would retain their high performance, we computed the predicted K_s^* values for combinations of the Willow replicates reported as part of Figure 3.5a. Interestingly, not only was the composite predicted to retain its superlative K_s^* , but this value was predicted to increase as more components were added to a maximum of 74.2% when 10 components were combined (Figure 3.10). This improvement with averaging can be understood by considering that the largest variability observed for this component was in

the plateau region and so regression of this curve to its mean is expected to result in closer agreement between the initial yield stress and this plateau region. These results point to examining combinations of GCS components as a promising area of future research.

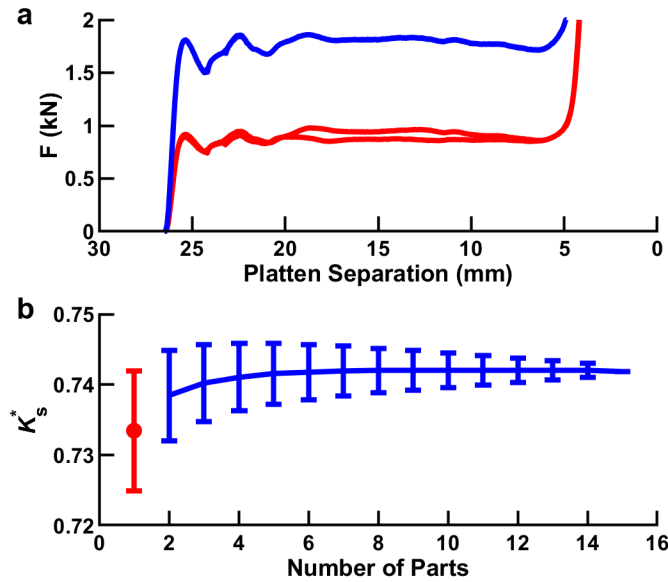


Figure 3.10: Simulation of efficiency for combinations of tested willow components. a, Two or more experimental F -separation curves (red) can be combined to create a simulated test of in-plane tiled components (blue). Curves are added in platten separation space to account for possible variations in component height. b, Mean and one standard deviation in K_s^* for original components (red) and combined prediction (blue) for varying number of components.

This study focused on optimizing the maximum energy absorbing efficiency, but it is also interesting to consider the performance of these components through the lens of other commonly used metrics such as energy absorption normalized by mass or volume (Figure 3.11). While this study focused on energy absorbing efficiency, both the toughness per unit mass U_m (Figure 3.11a) and toughness per unit volume U_v (Figure 3.8b) are interesting properties to examine. Despite this not being the optimization target, structures that absorbed as much as 22.8 J/g or 10 J/cm³ were

identified as part of this study. For both metrics, performance is correlated with both material strength and component effective medium density. Despite this, the toughness of a component is unbounded in compression due to the densification of most structures. To illustrate this issue, consider both the Willow and Iroko structures from Figure 3.5. By measuring the specific toughness below a given σ_t , it demonstrates that most of the energy is absorbed before the critical stress, but that U_v continues to increase as σ_t increases (Figure 3.11c–d). This makes it difficult to compare toughness tested on instruments that have different force limits unless the performance is bounded by some force or stress threshold. For this reason, employing K_s allows not only comparison across different materials, but also mitigates different testing limitations based on machine capabilities or user choices.




Figure 3.11: Examining toughness metrics beyond energy absorbing efficiency. a, Toughness per mass U_m vs. component effective medium density for all successful experiments. The color of each dot corresponds to the material used, as designated in Figure 3.2c. b, Toughness per volume U_v vs. component effective medium density for all successful experiments. c, U_v vs σ_t for the original Willow component with the point representing K_s^* marked by a black dot and the parity line showing the maximum possible U_v . d, U_v vs σ_t for the original Iroko component.

3.3: Discussion

This work reports a series of mechanical insights that resulted from performing an extensive experimental campaign using an SDL. Through the exploration of a vast parameter space, we were able to identify components with superlative K_s^* , advancing the frontier of energy absorption and finally overcoming the record previously held by nature. These designs were found to be general within material classes when additively manufactured, showing the opportunities and limitations of transferability of the designs.

A longstanding goal in the mechanics community is to identify simple shapes that are predicted to obtain high performance, such as the analytically predicted 75% efficiency of honeycomb lattices.⁷ However, a noteworthy feature of this work is that the final components contain repeatable geometric details defined by the FFF process that are not captured in the simple GCS digital model, such as the scalloping of the walls. This acknowledges a key reality that, in practice, all structures have to be physically realized and any manufacturing technique contains characteristic microstructures and idiosyncrasies that are unique to that method. It may be that obtaining superlative efficiency requires optimizing the structure within the context of how it would be physically realized by the fabrication method – as we have done here through our extensive experimental campaign. We expect future work will focus on studying the transferability of these high performing GCS geometries to other fabrication methods.

In addition to identifying superlative designs, this SDL campaign also illustrated how optimization and learning can be complementary goals in that the generated corpus of data allowed for the extraction of design insights for optimal use of polymers. These insights include matching polymer materials to target use cases, highlighting the use of ρ_r as an aggregate descriptor, and gaining connections to the broader literature on mechanics of cylindrical shells. While here we only present limited mechanical trends that capture the breadth of the dataset, we anticipate that the availability of this data will motivate others to dive deeper into specific regions and provide insight into phenomena such as failure mechanisms and what geometric aspects would prevent them. The assembled database of GCS properties under uniaxial compression could also be very

valuable in accelerating the discovery of other categories of mechanical response through the principles of transfer learning. For example, we have previously shown that quasistatic performance can be used to predict impact performance in data-driven manner.¹¹³ Further, transfer learning using feature transformation can be used to develop output spaces in which correlations are easier to learn, as we have shown using FEA calculations.²⁸ As such, the GCS dataset is likely to be useful to accelerate learning other categories of mechanical response. From a mechanics standpoint, we expect that this work will provide geometric motifs that lead to more efficient and safer protective equipment.

From a broader learning perspective, this work shows how the iterative and collaborative interaction between SDLs and human teams can lead to sustained progress. We should note that this study was not dedicated to benchmarking the acceleration inherent to using SDLs. Prior work, including our own,^{27,28} has focused on such benchmarking and ongoing research is focused on developing algorithms and processes to efficiently select experiments.^{66,114} In this work, our main focus was discovering new mechanical structures and we believe that this type of sustained campaign is an example of how SDLs can fruitfully exist in the materials discovery pipeline. An important point to consider is that the process used here allowed the human-SDL partnership to leverage the strengths of each member. In particular, it has been widely seen that algorithmic decision making is more efficient at navigating high-dimensional spaces than even expert selection of experiments.^{29,44,106,107} The approach used here allowed the SDL to autonomously and continuously explore this space while the humans provided periodic

guidance in the form of making adjustments to the algorithm or accessible space. Having the human out of the loop allowed the autonomous system to continuously progress while requiring human decision-making at each point would have made the pace of experimentation impossible. While it would have been useful from a benchmarking perspective to construct a predetermined process for the humans to interact with the SDL within a narrow scope, this rigid framework would not allow for the effective use of human creativity to adapt to unexpected problems. The fruitful partnership of human and SDL raises the question of what general lessons or approaches can be applied in these cases. Some studies have emerged that begin to address the concept,^{30,115} and it is a promising area for further study. Finally, an important avenue for future research is learning how to effectively combine simulations such as FEA with experimental campaigns such as this to increase the rate of learning while reducing the burden of physical experimentation.

3.4: Methods

3.4.1: Design of generalized cylindrical shells

To provide a vast space of potential designs that are topologically identical to cylindrical shells, we developed a generalized cylindrical shell (GCS) family of structures, which is an 11-dimensional parameter space (Figure 3.12). A cylindrical shell is often defined in terms of its height h , wall thickness t , and diameter d . Here, we design generalized cylindrical shells (GSC) that are topologically consistent with cylindrical shells and have a consistent wall thickness and height, but vary in their cross-sectional profile along the axial direction. As a diameter is not an appropriate measure for such a

complex shape, we parameterize these using their average perimeter P_0 . A GCS design is realized by deforming cylindrical shells using three distinct transformations: variable perimeter, variable cross section, and twist (Figure 3.12a).

Briefly, the variable perimeter is realized by linearly varying the perimeter from the top of the GCS to the bottom of the GCS (Figure 3.12b). In this campaign, the perimeter of the top was constrained to be larger than the perimeter of the bottom to ease with component removal. The cross sections of the GCS were transformed using a summed cosine function (Figure 3.12c). The top cross section and bottom cross section are specified, and each intermediate layer is calculated as a linear interpolation of these two faces, ensuring a manifold surface. Finally, both sinusoidal and linear twist can be applied to these cross sections in a height-dependent manner (Figure 3.12d). Collectively, these transforms allow for more than trillions of unique designs.

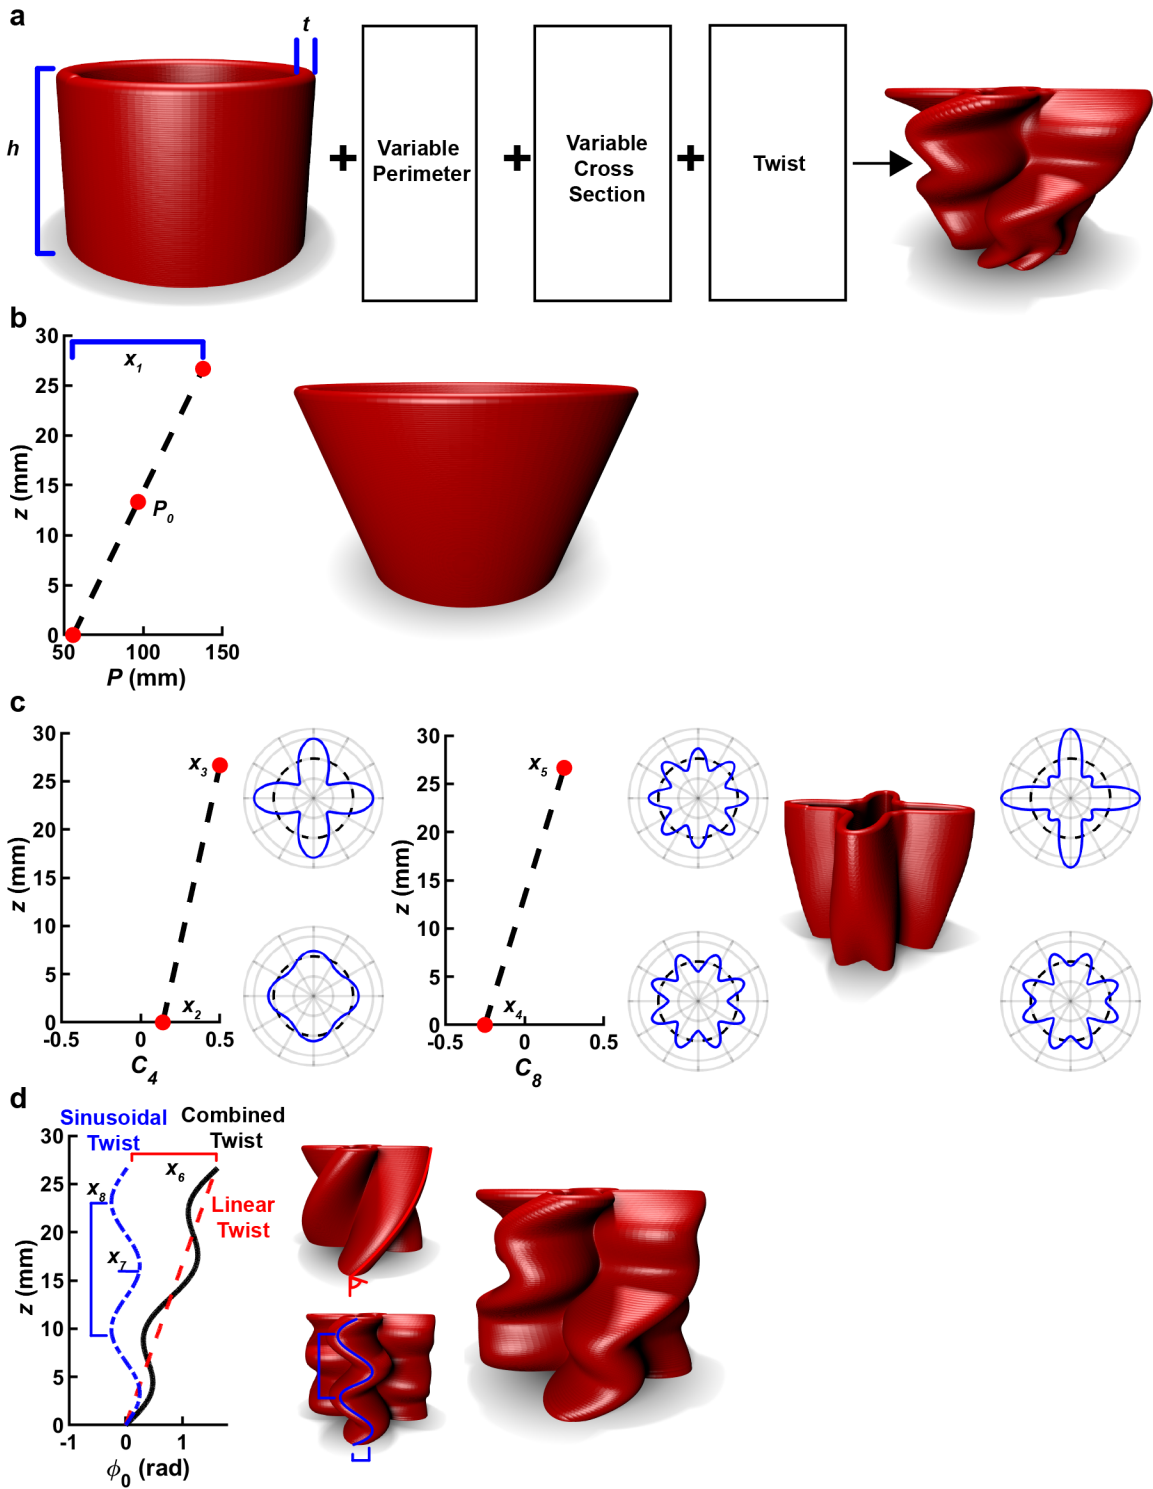


Figure 3.12: Generalized cylindrical shells. a, Generalized cylindrical shells (GCS) are realized by transforming a cylindrical shell of height h and wall thickness t to create interesting shapes that preserve the topology of the shell. b, The perimeter P varies linearly along the height z of the

shell based on an average perimeter P_0 and a perimeter difference x_1 . c, The cross sections of each layer are deformed in a z -dependent manner using a summed cosine function with 4-period amplitude C_4 and 8-period amplitude C_8 . These are defined at the top and bottom by four variables x_2 , x_3 , x_4 , and x_5 , and linearly interpolated to determine the cross section at any z . d, The cross sections of the design are rotated about the cylinder axis in a z -dependent manner by rotation angle ϕ_0 using both linear and sinusoidal twists as defined by linear twist x_6 , sinusoidal twist amplitude x_7 , and sinusoidal twist period x_8 .

The azimuthally-dependent radius $r(z, \phi)$ is shifted by adding two cosine functions with set periodicities as inspired by the summed cosine design of mechanical structures.¹¹⁶ In particular, r at any given height z and azimuthal angle ϕ is given by,

$$r(z, \phi) = r_0(z) \left[1 + C_4(z) \cos(4(\phi + \phi_0(z))) + C_8(z) \cos(8(\phi + \phi_0(z))) \right], \quad (3.2)$$

where $C_4(z)$ and $C_8(z)$ are amplitude prefactors to the summed cosines, $\phi_0(z)$ is a rotational offset, and $r_0(z)$ is a prefactor adjusted to set the height-dependent perimeter $P(z)$ of the shell. Each of these functions is defined by terms of \mathbf{x} . Specifically, we define,

$$P(z) = P_0 \left[1 + x_1 \left(\frac{z}{h} - \frac{1}{2} \right) \right], \quad (3.3)$$

such that x_1 is the difference between the perimeter at the top of the component and the perimeter at the bottom. The Python function `scipy.optimize.minimize` was used to minimize the error between the $P(z)$ and the actual perimeter of Equation (3.2), estimated using Simpson's rule, by adjusting $r_0(z)$ at each layer.

Each summed cosine term is defined by specifying its value at the top and bottom of the shell and linearly interpolating between these points, specifically,

$$C_4(z) = x_2 \frac{h-z}{h} + x_3 \frac{z}{h}, \quad (3.4)$$

and

$$C_8(z) = x_4 \frac{h-z}{h} + x_5 \frac{z}{h}. \quad (3.5)$$

To determine the azimuthal offset of each layer, we include two ways in which this can vary with height, namely a linear and sinusoidal shift. Specifically, we define,

$$\phi_0(z) = x_6 \frac{z}{h} + x_7 \sin(2\pi x_8 z). \quad (3.6)$$

Code to generate standard triangle language (STL) models based on the GCS family of shapes is available <https://github.com/bu-shapelab/gcs>.

3.4.2: Development of the Bayesian experimental autonomous researcher

In order to study the mechanical energy absorbed by additively manufactured components under uniaxial compression, we developed and utilized a self-driving lab (SDL). This system incorporated several distinct instruments, computers, and algorithms that worked in concert to select experiments, construct physical samples, and test them with minimal human intervention. From a hardware perspective, this system consisted of five fused filament fabrication (FFF) printers (MakerGear M3-ID) arrayed in an arc. In the center of this arc was a 6-axis robot arm (Universal Robotics UR5e). Also in the working radius of this arm was a scale (Sartorius CP225D) and a universal testing machine (UTM) (Instron 5965). The arm had a webcam (Logitech C930e) to allow for monitoring the flow of experiments and there was a video camera (PixelLINK PL-D722) with lens (Infinity InfiniMite Alpha) mounted facing the bottom platen of the UTM to record videos of the tests. A picture of the SDL is shown as Figure 3.13. The hardware and software organization of this system is shown in Figure 3.14.




Figure 3.13: Picture of the Bayesian experimental autonomous researcher (BEAR), consisting of five fused filament fabrication 3D printers, a six-axis robot arm, a scale, and a universal testing machine.

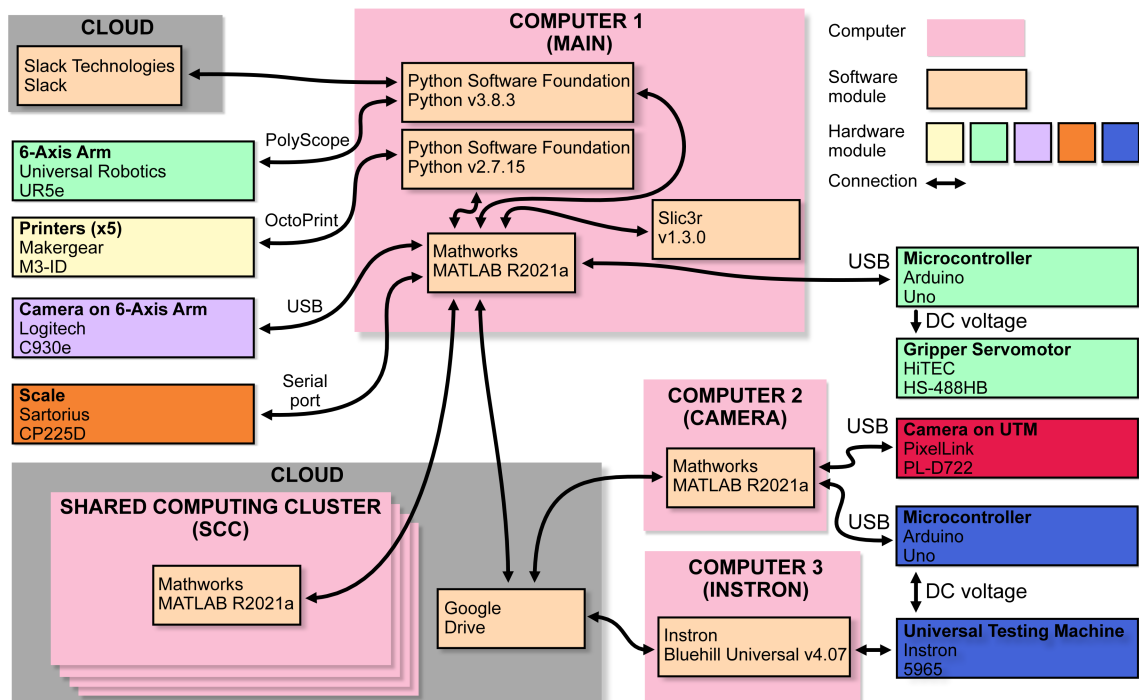


Figure 3.14: Hardware and software Organization of the Bayesian experimental autonomous researcher (BEAR).

A flow chart describing the core actions of this system is shown in Figure 3.15. At the most basic level, the system comprised a loop implemented in MATLAB 2021a (MathWorks Inc) in which the system repeatedly iterated through six potential actions, namely: (1) Use the accumulated data to select the design and material to be tested next given the available printer and materials. (2) Generate the digital design files needed to run the available printer. (3) Send the G-code file to the printer and begin printing the component. (4) Retrieve the completed component from the printer and weigh it using the scale. (5) Retrieve the component from the scale, place it on the platen of the universal testing machine, run the mechanical testing program, and then clear the component from the platen. (6) Read the results of the mechanical testing and update the accumulated data. The order of priority was tuned to maximize the throughput of the system by giving priority to actions that were likely to become bottlenecks.

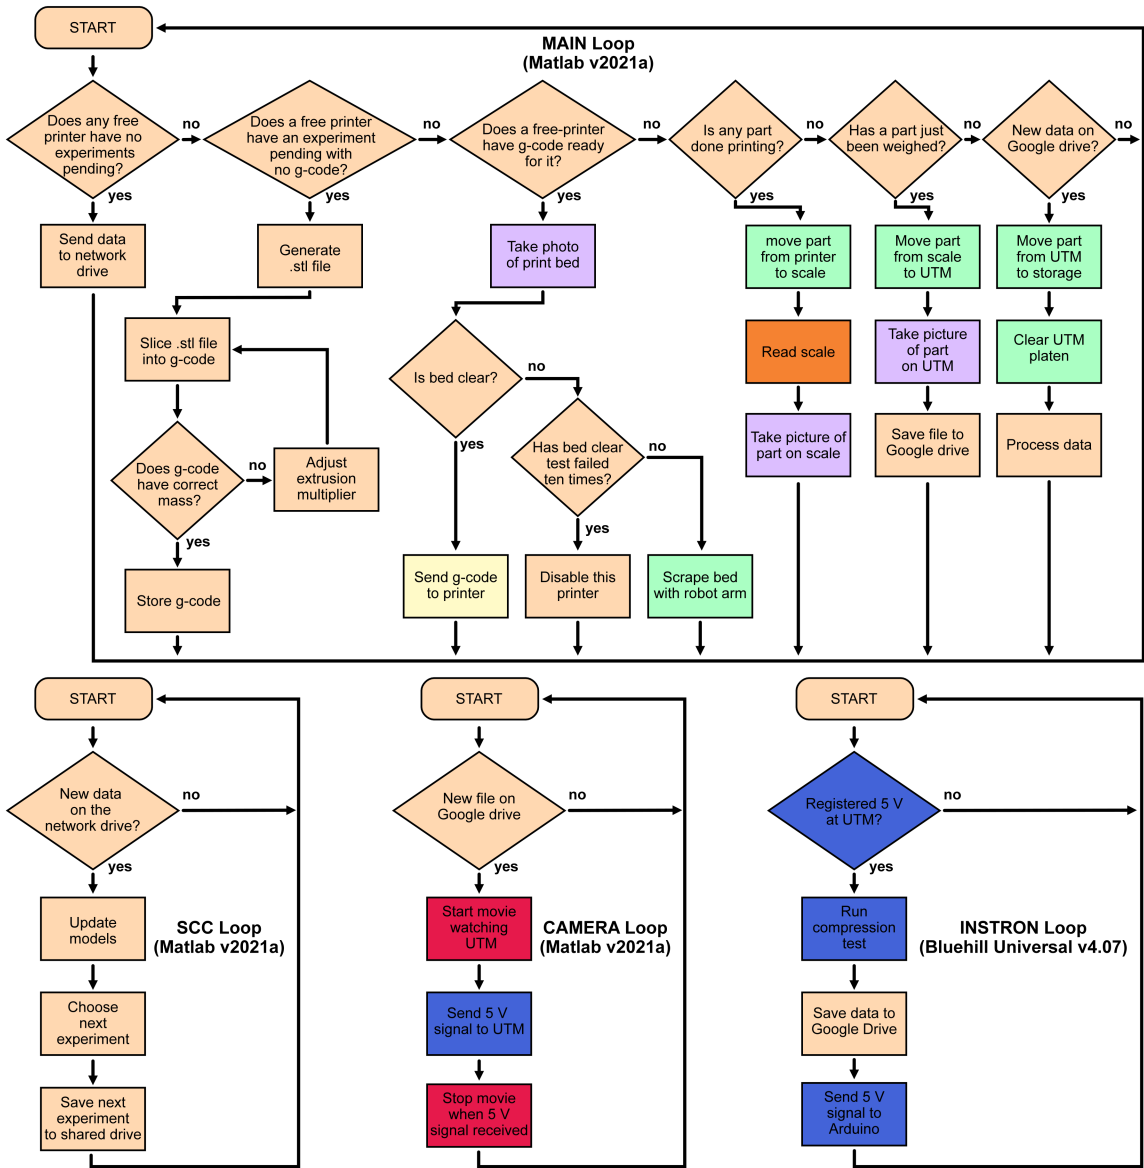


Figure 3.15: Flowchart of the software loops executed in the four computers running as part of the BEAR. Colors on the panels correspond to systems in Figure 3.14. Order of Main Loop actions can be adjusted by researchers to maximize throughput by prioritizing potential bottlenecks.

3.4.3: The research campaign

Over the course of the campaign, 25,387 experiments were performed (Figure 3.16). Although individual experiments were rarely selected by hand, a variety of decisions were made by the researchers along the way. Changes were made to the parameter space under consideration, such as adding sinusoidal twist or switching to Latin hypercube sampling (LHS). Additionally, new materials were added to the campaign, and the mix of filaments loaded into the printers was adjusted to focus on specific goals.




Figure 3.16: Experiments carried out by the BEAR. Experiments performed during the campaign, which are defined by eleven GCS parameter values. The color of each dot corresponds to the material used, as designated in Figure 3.2c. Black and red stars correspond to breakthrough experiments, as designated in Figure 3.4a.

Select Experiment

Bayesian optimization was used by the BEAR to algorithmically select additional experiments. This process includes the conditioning of a surrogate model to approximate the connection between input space and output space and then the use of an acquisition function to evaluate this model to find experiments that are believed to be most useful to perform. Since the goal of this work was to identify structures with high K_S^* , we treated

this as a maximization problem. The input space for this maximization was both the design of the GCS and the material used to realize a component out of this design. As such, we required a 13-dimensional input (11 geometric parameters and two material properties). For the output space, we were not just interested in K_s^* , but we also found it necessary to predict σ_t^* . Gaussian process regressions (GPRs) were used to predict K_s^* and σ_t^* . A neural network with one hidden layer equal to the input size was used to predict component printability. Specifically, predicting printability is done purely through empirical means based on the results of past experiments. Specifically, each experiment is assigned a tag of “printable” or “unprintable” based on a combination of automated processes (*i.e.* was a component delivered to the testing frame) and out-of-the-loop manual inspection. These tags were used to train a classifying neural network that can rapidly predict the printability of any point in the GCS parameter space. To evaluate potential experiments, the output of this neural network was multiplied by their predicted K_s^* to bias the search into printable regions of parameter space.

In developing surrogate models, transformations were done to the GCS design parameters and material parameters to improve the accuracy of the models. The overall motivation of these transformations was to improve correlations in the input space and thus improve predictions in the output space. The full list of the transformed input and output spaces are given in Tables 3.2 and 3.3. For example, the logarithms of σ_p , σ_t^* , and E were taken because their values varied over several orders of magnitude and the points were more evenly spaced when considered logarithmically rather than linearly. Additionally, rather than specify P_0 , we preferred to specify the target mass m

normalized by h , or the mass per height m/h . Additionally, a wall angle θ was used as it was hypothesized that the angle of the wall was more important than the absolute value of the change in P . This wall angle was estimated using the formula $\theta = \text{atan}\left(\frac{x_1}{2\pi h}\right)$.

Finally, rather than conditioning the GPR to directly predict K_s^* , we found that it was useful to transform K_s^* to emphasize differences at the high end while minimizing differences at the low end and to explicitly prevent the model from predicting physically impossible values (*i.e.* $K_s^* > 1$ or $K_s^* < 0$). Thus, we instead predicted $\text{atanh}(2K_s^* - 1)$. This function was chosen because it monotonically transforms inputs from 0–1 to outputs from negative infinity to positive infinity. Mechanical intuition suggests that a given design produced using different materials may have a similar efficiency, but its strength should vary based upon the strength of the material. Given sufficient time, this trend would naturally be learned by an active learning process. However, to accelerate this process, we hypothesized that learning the strength of a design relative to its material properties may be more efficient. Indeed, early in the campaign we plotted $\log \sigma_t^*$ vs. $\log E$ and found that the best fit trend line had a slope of 0.408. Thus, we began normalizing σ_t^* by $E^{0.408}$ in an attempt to capture this variation of component strength with material properties. It is important to note that variations about this trend were still allowed for and learned.

Model input variable	Description
h	Height (mm)
m/h	Mass per height (g/mm)
t	Wall thickness (mm)
$\text{atan}\left(\frac{x_1}{2\pi h}\right)$	Wall angle (degrees)
x_2	4-period amplitude of bottom cross section (dimensionless)
x_3	4-period amplitude of top cross section (dimensionless)
x_4	8-period amplitude of bottom cross section (dimensionless)
x_5	8-period amplitude of top cross section (dimensionless)
x_6 / h	Linear rotation per height (rad/mm)
x_7	Sinusoidal rotation amplitude (rad)
x_8	Sinusoidal rotation wavelength (mm)
$\ln(E)$	Natural log of the polymer elastic modulus $\ln(\text{MPa})$
$\ln(\sigma_p)$	Natural log of the polymer plateau stress $\ln(\text{MPa})$

Table 3.2. Inputs to the machine learning models used for Bayesian optimization.

Model output variable	Description	Model type
$\text{atanh}(2K_s^* - 1)$	Transformed peak energy absorbing efficiency (dimensionless)	Gaussian process regression
$\log\left(\frac{\sigma_t^*}{E^{0.408}}\right)$	Log 10 of the ideal threshold stress normalized by the modulus and raised to an empirically determined power $\log(\text{MPa}^{0.592})$	Gaussian process regression
p	Printability (dimensionless)	Artificial neural network

Table 3.3. Outputs of the machine learning models used for most of the experimental campaign.

When selecting a subsequent experiment for a given printer, not all combinations of designs and materials were available. Specifically, each printer had two independent extruders, which allowed two different filaments to be loaded at once. Further, each extruder had either a 0.5 mm diameter nozzle or a 0.75 mm diameter nozzle. Values of t different from these diameters could be achieved by over or under extruding. We restricted $t < 0.7$ mm for the 0.5 mm diameter nozzle and $t \geq 0.7$ mm for the 0.75 mm nozzle.

To select an experiment, we define an acquisition function a_f that takes as its input positions in parameter space along with the current surrogate models and select the experiment that maximizes a_f . Throughout the campaign, three types of acquisition functions were used: maximum variance (a_f is equal to the variance in predicting K_s^*), expected improvement (a_f is the predicted amount of improvement beyond the previous best K_s^*), and upper confidence bound (a_f is the weighted sum of the predicted of K_s^* and the predicted uncertainty in predicting K_s^*). The combination of a_f and the strategy for finding its maximum is considered a decision policy. However, this process was not treated as a simple single-objective maximization. For instance, in all cases, a_f is multiplied by the predicted printability p to ensure that we are only considering components that are expected to be realizable in practice. Additionally, many of the decision policies are multi-objective, trying to find high values of K_s^* across a range in σ_t^* . When this was the case, multiple GPR model predictions were combined to select a component by penalizing the K_s^* prediction by the distance of its predicted σ_t^* from the target σ_t or by comparing the predicted K_s^* to the performance of other tests at that σ_t^* . A full list of considered decision policies is given in Table 3.4. These policies were added sequentially during the progression of the campaign, so their order reflects the evolution of our thought process during the campaign. Additionally, the GPR models can be retrained using only data from the region of interest, which was begun with decision policy 19. This allowed the GPR to capture finer correlations in the parameter space around the region of interest. All models were trained using MATLAB's built in functions and the code is available at https://github.com/KelseyEng/BEAR_ADTS.

Model Training was performed on Boston University's Shared Computing Cluster, where multiple compute nodes could work in parallel. GPR processing time scales with the number of experiments cubed.¹¹⁷ Therefore, the longer the campaign ran, the more computationally expensive model building and component selection became.

Decision Policy Number	Acquisition function	Metric	Number of Valid Experiments
0	Manually Selected	Researcher intuition or performance validation	730
1	Upper confidence bound	Full integral of force-displacement curve	24
2	Maximum variance	Full integral of force-displacement curve normalized by component mass	916
3	Expected improvement	Full integral of force-displacement curve normalized by component mass	775
4	Expected improvement	Expected acceleration of a simulated impact test	93
5	Expected improvement	K_S at a target σ_t	249
6	Expected improvement	K_S^* penalized by an amount proportional to the distance between σ_t^* and a target σ_t	97
7	Expected improvement	K_S^* penalized by an amount proportional to the distance between σ_t^* and a target σ_t with uncertainty in σ_t^* considered	383
8	Expected improvement	K_S^* minus the best K_S previously observed at the predicted σ_t^*	3,219
9	Expected improvement	K_S^* minus the best K_S previously observed at the predicted σ_t^* , but with limits imposed on the largest and smallest stresses considered	31
10	Expected improvement	K_S^* minus the best K_S previously observed at the predicted σ_t^* , but only considered components that could have been printed using the specific printer under consideration	501
11	Maximum variance	K_S^* , but only considering cylindrical shells	34
12	Expected improvement	K_S^*	1,608
13		Not Used	
14	Expected improvement	A weighted sum of the acceleration from a simulated impact test and the plateau stress of the component	41
15	Expected improvement	A weighted sum of the acceleration from a simulated impact test and the plateau stress of the component (different simulation model from DP 14)	212
16	Expected improvement	K_S^* times the ideal threshold force for that component	22
17	Expected improvement	K_S^* minus the best K_S^* that could have been printed using the specific printer under consideration	1,041
18	Expected improvement	K_S^* , but only considering components near the best previously found component	1,569

19	Upper confidence bound	K_S^* , but only considering components near the best previously found component	224
20	Upper confidence bound	K_S^* , but only considering components near the best previously found component that have effective densities ρ_d below 10%	542
21	Not Used		
22	Expected improvement	K_S^* minus the best K_S previously observed at the predicted σ_t^* , but only considering components that could be continuously extruded without a linear twist	523
23	Expected improvement	K_S^* minus the best K_S previously observed at the predicted σ_t^* , but only considering components that could be continuously extruded with a linear twist	286
24	Expected improvement	K_S^* of a two-component system	129
25	Expected improvement	K_S^* of a two-component system, but only considering components near the best previously found pair of components	1

Table 3.4. Descriptions of decision policies used during campaign.

Initially, sampling points were selected on a grid. Starting at ID 9,261, potential sampling points were selected using Latin hypercube sampling (LHS) to facilitate exploring space more finely. Starting with ID 11,763, after the proposed experiment had been selected, a second round of sampling points were added that were zoomed in a hypercube around the selected point to more closely find the maximum of a_f .

Generate G-code

Once a component has been selected for testing, the STL was generated using a custom Python script. This Python script (Python version 3.8.3) was run on the main computer and called from MATLAB using the command line function. The resulting STL was created as a solid object. In order to convert this STL file into the G-code needed for the printer, Slic3r (version 1.3.0) was run from the command line of MATLAB. Prior to sending the STL file, the Slic3er configuration file was edited using

string manipulation directly from MATLAB to set the nozzle temperature, bed removal temperature, and extrusion multiplier. The extrusion multiplier was set as part of a feedback system to maintain component weight (see below). Slic3er was configured to use vase mode (spiral mode), which removes the tops and bottoms of solid objects and turns the STL solid into a shell. The output of this process is G-code for the print and predicted amount of filament that is needed to print this component, which is read into MATLAB. Using an initial set of calibration prints and subsequent use of integral feedback, we predicted the mass of the component from the amount of filament predicted to be used by the slicer. Adjusting the extrusion multiplier and reslicing the component provided a reliable method of controlling the mass of the final component and standardizing performance across different printers and filament rolls (Figure 3.17). This gives effective control over t by over or under expanding the material leaving the nozzle. It also allows the computer to automatically compensate for variations in the thickness of the filament diameter or variations between the stepper motor of different printers. Relatively slow print speeds of 15 mm/min were employed to prevent clogging, which was especially important for the softer filaments.

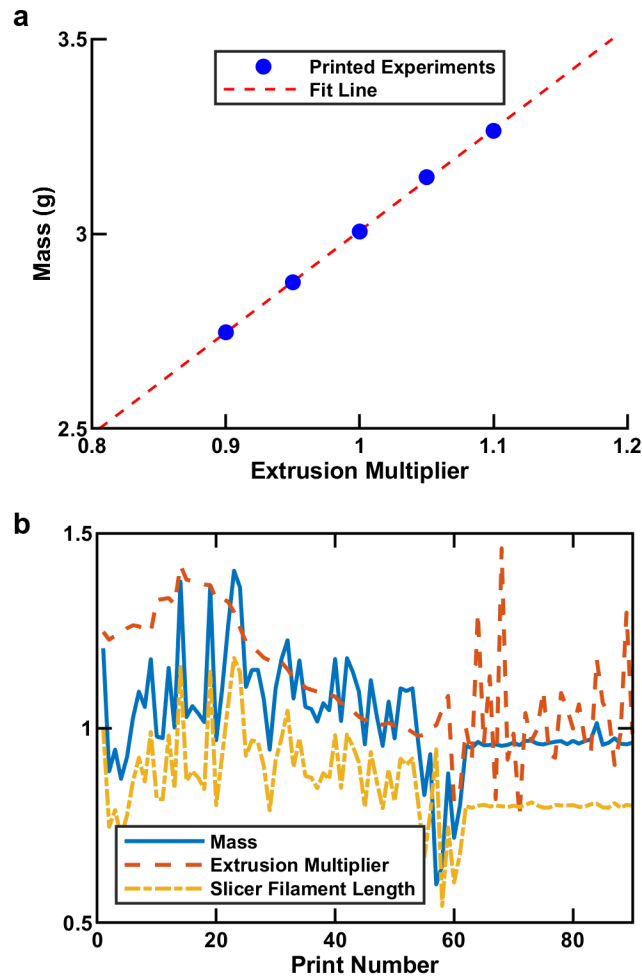


Figure 3.17: Mass calibration through extrusion multiplier adjustments. a, A single component printed five times with different extrusion multiplier shows that component mass is linearly correlated with extrusion multiplier. b, Applying integral tuning to the extrusion multiplier was ineffective because of variations in slicing complex curved structures, as seen in print number < 60 . However, when integral tuning was applied to the slicer filament length by adjusting the extrusion multiplier rapidly, consistent mass was obtained (print number > 60). Mass is normalized by target mass and slicer filament length is normalized by the initial slicer filament length (print number = 1).

Begin Experiment

Once the G-code had been prepared for a given printer, the physical experiment was ready to begin. First, the arm moved into position over the chosen printer and the arm-mounted camera took a photograph of the print bed to ensure that it was free from

debris and ready for the next print. To accomplish this, the picture was run through a neural net based on GoogLeNet¹¹⁸ and classified as ‘clear’ or ‘needs cleaning’. If the bed needed cleaning, the robot arm picked up a scraper and scraped the print bed. A new picture was then taken to verify that the print bed was clean. If the bed was still unclean, the BEAR would attempt to clean the bed up to ten times with the scraper. If the bed was still unclean, the printer would be deactivated and the operator was notified to clean the bed before further experiments. Once the bed was determined to be clear, the system sent G-code to the printer using OctoPrint.

Samples were fabricated using MakerGear M3 printers with either a 0.5 or a 0.75 mm nozzle at 80 °C bed temperature, 250 °C nozzle temperature (except for PLA, which was printed using 220 °C), and 15 mm/s print speed using vase mode. The cylindrical shell sample in Figure 3.2A was fabricated out of TPU-2 (Cheetah – Ninjatek) to be 19.5 mm tall, 27.9 mm wide, and have 0.5 mm thick walls.

Weigh Component

When a print was complete, as determined by querying the state of the printer through OctoPrint, the bed was heated (TPE, TPU-1-2) or allowed to cool (PLA, PETG, Nylon, and TPU-3) to facilitate removal of the component.¹¹⁹ Once the desired temperature had been reached, the robot arm removed the component from the printer and moved it to the scale, which determined its mass. This mass reading was read through a serial port by MATLAB. If no mass was registered on the scale, the system attempted to re-grab the component from the print bed up to three times. At this stage, the arm-mounted camera took a photograph of the component on the scale, which was used

to verify that the component was fully on the scale. Components that were misoriented, as determined by machine vision, were discarded before testing.

Test Component

If a component was on the scale and ready to be tested while the universal testing machine (UTM) was not performing any experiments, the component was moved to the UTM for compression testing. Once the component was in position, the main computer sent a command to the Instron computer to begin the test through a .mat file transferred by the cloud. The Video computer then told the Instron to start the test while it recorded a video of the compression testing. The test began with the top platen ~200 mm over the component. After zeroing the force sensor, the top platen moved at a rate of 50 mm/min toward the component until the force sensor registered 1 N. The platen then moved away from the component 1 mm so that it no longer was in contact with the component. At this point, the UTM started recording the force measurement while it lowered the top platen at 2 mm/min. A given test ended when either 1) the force exceeded the 4.5 kN force limit or 2) the top platen position fell below the safe height of 0.4 mm separation between the two platens. After a one-minute relaxation period, the platen was lowered again at a rate of 100 mm/min until the force exceeded 1 N to find the rebound height. After testing, the component was removed from the UTM and stored. The platen was then cleaned with the robot arm to ensure that the platens were clear and ready for the next test. Each mechanical test took approximately ten minutes. A third computer recorded the Instron data and saved it to the cloud. When the test was finished, the Video computer informed the main computer that the UTM was now free for another experiment.

Process Results

When converting from force-displacement curves to engineering stress-engineering strain curves (simply called stress and strain hereafter), it is necessary to define the area of the component and its height. For traditional materials, this process is straightforward as it amounts to defining the cross-sectional area of the component under study. However, for more complex structures, the area of the component is less clear. The as-printed height of the component was calculated by finding the platen separation when the moving median of twenty force measurements surpassed 0.3 N. Here, we define the cross-sectional area to be the amount of area that is required per component to tile the component infinitely on a plane. To calculate this algorithmically, we used the following steps (illustrated graphically in Figure 3.18):

1. Find the maximum radius r_{max} of the component by finding the maximum of the radius r at all heights z and azimuthal angles ϕ , as defined in Equation (3.2) in the Methods.
2. Enclose the component with a cylinder with radius r_{max} (Figure 3.18b).
3. Enclose the component in a hexagonal prism that circumscribes the cylinder (Figure 3.18c).
4. The area of the hexagonal prism is used as an estimate of the cross-sectional area needed to tile the component on a plane.

Although some designs may be more closely packed in a square lattice, most are more closely packed using the hexagonal approach (Figure 3.18d) due to the applied linear and sinusoidal twists, and this approach is invariant of the rotational orientation of the design.

Therefore, this hexagonal packing approach was used to estimate the cross-sectional area of all designs.

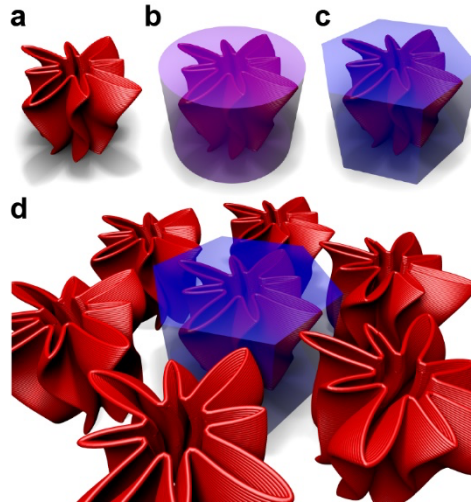


Figure 3.18: Calculating the effective area of GCS designs. a-b, To calculate the effective cross-sectional area of a design, it is fit into a cylinder based on its maximum radius. c, This cylinder is then enclosed in a hexagonal prism. d, The hexagonal prism can be tiled infinitely in a plane. Thus, the effective cross-sectional area of a design is estimated as the area necessary per design to tile it in a plane without collisions.

From the σ - ε curve, a variety of useful metrics were calculated. To find K_S^* and σ_t^* , K_S was calculated at 1,400 σ_t values that were logarithmically spaced between 10 Pa and 100 MPa. Ten additional σ_t sampling points were selected by dividing the σ - ε curve into ten equal sections in ε and finding the maximum σ in each section. Because σ_t^* is often a peak early in the σ - ε curve, these ten additional sampling points can often determine σ_t^* precisely. For each of these sampling points (1,400 evenly spaced points and ten extra points taken from the σ - ε curve), we compute $K_S(\sigma_t)$ and assign K_S^* and σ_t^* to the maximum and argmax of this calculation, respectively.

For most structures, K_S reaches its maximum energy absorbing efficiency K_S^* at an optimum threshold stress σ_t^* . This single point, (σ_t^*, K_S^*) can be used to describe the ideal

operating performance of a structure or material.

The densification strain ε_d is the ε value at which σ first exceeds σ_t^* . The relative density of the component ρ_r was calculated by dividing the mass of the component by the mass of solid material equal to the volume of the enclosing hexagon (Figure 3.18). Finally, quality control checks were performed to determine if the sample should be included in the complete dataset. Components that were not within 5% of their mass target or within 5% of their target height were excluded from the results. Additionally, components that hit the force threshold of the UTM when $\varepsilon < 0.3$ were excluded due to the high probability that σ_t^* was greater than the UTM's force threshold.

Filament rolls for 3D printing were purchased and used as received. They include three different types of thermoplastic polyurethane (TPU): TPU-1 (NinjaFlex-Ninjatek), TPU-2 (Cheetah-Ninjatek), and TPU-3 (Armadillo-Ninjatek). Additionally, four more filaments were used: thermoplastic elastomer (TPE) (Chinchilla-Ninjatek), nylon (Matterhackers Pro Series), polyethylene terephthalate glycol (PETG) (Matterhackers Pro Series), and polylactic acid (PLA) (eSun PLA+ and MakerGear). To determine the mechanical properties of each roll of filament, solid cylinders (100% infill) were printed measuring 16 mm tall and 8 mm in diameter. These cylinders were then tested in uniaxial compression at 2 mm/min. Force-displacement curves were converted to stress-strain curves by dividing the force by the component cross-sectional area and by dividing the displacement by the height, respectively. From the resulting stress-strain curves, three material properties were calculated: the modulus of the material, plateau stress σ_p , and the rebound fraction. The modulus E was calculated by fitting a series of lines in

windows of 0.05 to 0.25 strain and an initial strain location of 0 to 0.25 strain (to avoid toe regions), both in increments of 0.05. The largest slope observed was taken as the modulus for the sample. The σ_p was the stress value at 25% strain. The rebound fraction was the height of the component after a one-minute relaxation period divided by the initial height before testing, both measured by the Instron. One cylinder was tested for each roll of filament used.

The polymer materials studied in this work are provided in Table 3.5 along with the temperature at which they were printed, the temperature at which the print bed was held during removal, and the material class. As shown in Figure 3.19, the result of this testing are estimates of the elastic modulus E and plateau stress σ_p of each material. In addition, the degree to which the cylinder rebounded after a one-minute relaxation period was also recorded, although this is an imprecise measure of elasticity as a consistent force threshold was used for all tests, indicating that different materials experienced different total strains. Nevertheless, the plastic materials rebounded less than the hyperelastic materials, despite their total strain being lower.

Material	Manufacturer	Nozzle Temperature (°C)	Bed Removal Temperature (°C)	Class	Spools Used
TPE (Chinchilla)	NinjaTek	250	100	Hyperelastic	9
TPU-1 (NinjaFlex)	NinjaTek	250	100	Hyperelastic	16
TPU-2 (Cheetah)	NinjaTek	250	100	Hyperelastic	37
TPU-3 (Armadillo)	NinjaTek	250	30	Intermediate	11
Nylon	MatterHackers	250	30	Plastic	2
PETG	MatterHackers	250	30	Plastic	5
PLA	eSun/MakerGear	220	30	Plastic	29

Table 3.5. Filaments studied in this work along with their processing settings.




Figure 3.19: Material characterization of polymers studied. a, Plateau stress σ_p vs. elastic modulus E for seven materials used in this campaign. b, Rebound fraction vs. E . c, Rebound fraction vs. σ_p . Error bars represent one standard deviation. Here, σ_p is the stress at 25% strain. Rebound fraction is the height after 1 minute relaxation divided by the initial height.

Maintenance

At the beginning of the campaign and periodically thereafter, new filament rolls were loaded into the printers. After performing material characterization, a series of calibration components were printed to tune the extrusion multiplier of the printer to the density and diameter of the filament. The target mass for the calibration component was

3.3 g. If the component was too heavy, the extrusion multiplier was decreased. If it was too light, the extrusion multiplier was increased. This continued until the mass was within 5% of the target mass. In this way, it was possible to estimate the ratio of the filament length computed by Slic3er to the mass of the resulting component. As components were subsequently printed during the campaign, this ratio was slowly adjusted using integral tuning to remain accurate. Additionally, components were printed on polyimide tape that was applied to the glass bed of the printers. Whenever the tape showed signs of wear or became damaged, it was manually replaced.

3.4.3: SHAP analysis

The neural network used for SHAP analysis comprised six layers: a 64-dimension linear layer followed by a ReLU activation¹²⁰, repeated three times. A data split of 80% for training, 10% for validation, and 10% for testing was employed. The GCS parameters were normalized and no preprocessing was applied to K_s^* . The network was trained using the mean squared error loss function. The training process uses the Adam optimizer¹²¹ with a learning rate of 0.001, weight decay of 1×10^{-5} , and a batch size of 16. Training was performed for 500 epochs with early stopping. The PLA network achieved a test loss of 0.0032 and a coefficient of determination $R^2 = 0.88$. For interpreting the predictions generated by the neural networks, we used the SHAP DeepExplainer which is initialized using the training split data. To provide explanations for individual components, we use the default SHAP waterfall visualization.

CHAPTER 4: A PHYSICS-INFORMED IMPACT MODEL REFINED BY MULTI-FIDELITY TRANSFER LEARNING

1. Introduction

Structures that absorb energy in impact events are critical for everyday life. They are used in bicycle helmets^{80,122}, the crumple zones of cars¹, packaging for shipping^{123,124}, and even ballistic and blast protection¹⁰⁰. Despite these critical needs, development of protective impact structures is challenging because performance depends on the properties of the object to be protected, the energy of the impact, and the height and contact area of the protective structure, requiring thousands of time-consuming physical tests to optimize^{8,125–128}. Further complicating the design process is that it is extremely challenging to obtain precise agreement between experiment and traditional physics-based approaches, namely finite element analysis.^{10,89} Disagreement often arises due to the difficulty associated with capturing the complex interactions that occur at high strain rate, at high strain, and with self-self-contacts of the structures.

To overcome these challenges, many impact structures make use of simple bulk materials such as foams^{125,129}. Such materials are simple to work with because their density can be tuned in a manner that predictably changes their performance⁷. This scaling leads to the common use of empirically measured cushion curves, which compare the peak acceleration of an impactor versus the static pressure exerted by that impactor at rest. Such cushion curves are helpful in selecting the ideal foam height and density for a given application, but assume that the height can be changed without affecting performance, limiting their applicability when applied to non-homogenous materials.

Despite the widespread use of cushion curves for foams, significant scientific research focuses on developing materials and structures that outperform foams in weight, volume, and safety metrics.

The increasing availability and decreasing cost of additive manufacturing (AM) has opened the opportunity to manufacture structures that are more complex than traditional foams.¹²⁴ One common approach is to use AM to produce uniform lattice structures, once again depending on density scaling or repeated unit cells to predict performance^{130–132}. Some researchers have investigated how structure affects performance for designs with consistent density, such as functionally grading strut thickness along an axis^{133,134}, auxetic structures^{78,87}, or using tapered beams⁹². Our previous work showed that Gaussian process regression (GPR) models could be used to predict peak acceleration of impact tests for lattice structures where the mass of the lattice is shifted between the joints and different types of struts¹¹³. Despite their promise, optimizing novel lattice structures for impact performance can quickly become impractical due to the immense number of design choices possible and the burdensome testing requirements.

Many efforts have been made to decrease the burden of impact testing. Early attempts focused on predicting the amount of energy that foam samples could efficiently absorb^{135,136}. To account for strain-rate effects, a common strategy is to treat the material as viscoelastic or scale quasistatic curves by multiplying them by empirical constants or dynamic factors that include both strain and strain rate^{127,137–140}. Recently, more numerical models have appeared to model the dynamics of impact^{126,141}. Despite these

efforts, little work has been done to extend these predictive models beyond homogenous foams and facile models that accurately predict impact performance remain an open challenge.

To speed up development of new impact structures, we develop a physics-based transfer learning model to predict the ideal impact velocity V_0^* for polymer structures. To explore this concept, we experimentally study generalized cylindrical shells (GCS), which have superlative energy absorbing capacity in quasistatic (QS) compression tests and are inherently easy to print using fused filament fabrication (FFF)⁶⁹. The development of this model involves three key steps, (1) the extraction of key metrics from QS testing, (2) the use of intermediate strain rate (ISR) tests to determine strain-rate dependent strengthening, and (3) the refinement of the model using impact tests. After these initial training steps, the model is able to predict impact performance on previously unseen samples with a single QS test and even extrapolates to test data using different impactor mass, designs, and materials with an RMSE of 0.23 m/s.

2. Theory

An impact test consists of an impactor of mass m dropping on a test component and hitting it with initial impact velocity V_0 . The initial kinetic energy KE of the impact is given by,

$$KE = \frac{1}{2} m V_0^2. \quad (4.1)$$

During an impact experiment, the acceleration a of the impactor is tracked and used to assess the performance of the component. Typically, the impact performance of a component is defined by the maximum acceleration a_m observed during the test, which

should be minimized to prevent damage. However, a_m is not always the clearest way to compare the performance of components as a_m generally increases monotonically with V_0 and m . For ease of comparison, impact performance is often described using the non-dimensional Janssen factor J , which is defined as $J = a_m/a_t$ where a_t is the theoretical minimum acceleration that could stop the impactor given by,

$$a_t = \frac{v_0^2}{2h}, \quad (4.2)$$

where h is the height of the test component. While J depends on the details of the impact test, namely V_0 and m , there exists a critical initial velocity V_0^* for which the minimum Janssen factor J^* is observed for a given m .

In contrast with impact tests in which the strain rate varies during the experiment, fixed-velocity tests feature a constant strain rate. Under these conditions, the useful non-dimensional metric to consider is the energy absorbing efficiency K_s of a component, which describes how efficiently it absorbs energy below a force threshold F_t and is given by,

$$K_s(F_t) = \frac{\int_0^{D_t} F(D) dD}{F_t h}, \quad (4.3)$$

with the displacement threshold D_t taken as the largest D for which $F(D) \leq F_t$ for all $D \leq D_t$. Practically speaking, this equation quantifies the amount of energy absorbed before exceeding F_t and normalizes it by $F_t h$, which is the total amount of energy that could be absorbed if $F = F_t$ for $0 \leq D \leq h$. The maximum energy absorbing efficiency K_s^* occurs at the critical force threshold F_t^* . While there are similarities between K_s^* and

J^* , it is worth emphasizing that K_s^* can be found using a single QS test while finding J^* requires empirically searching for the minimum value in V_0 at the relevant m .

3. Results and Discussion

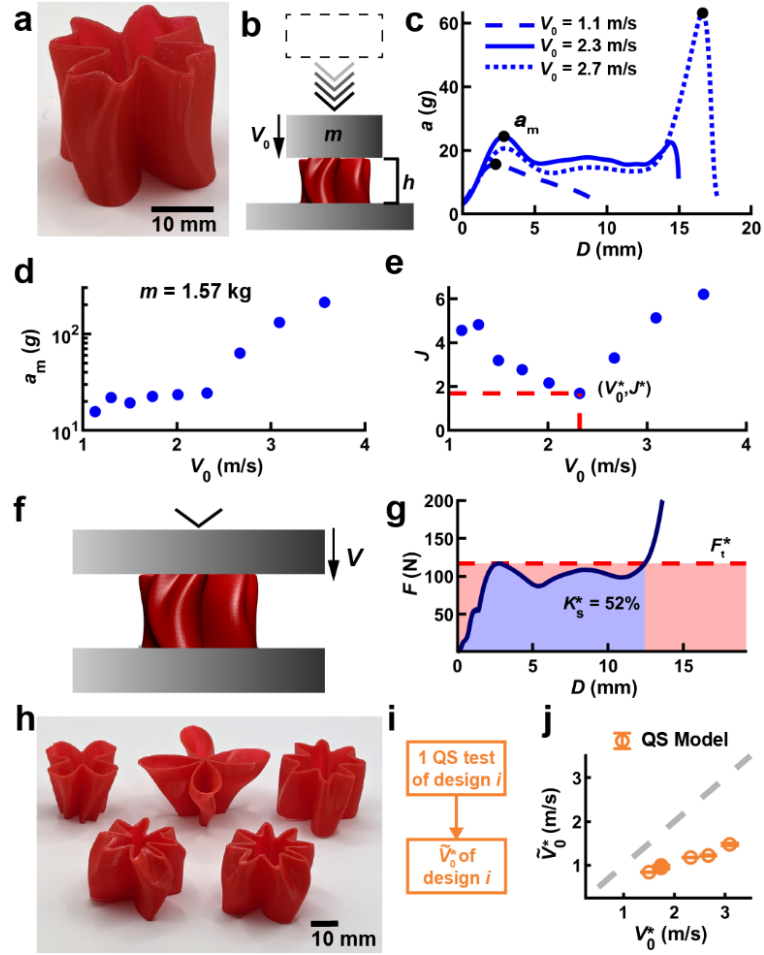


Figure 4.1: Predicting impact performance from quasistatic tests. a, Photograph showing a generalized cylindrical shell (GCS) component made from thermoplastic polyurethane (TPU). B, Schematic of an impact test with impactor of mass m hitting the component of height h at an initial velocity V_0 . c, Acceleration a vs. displacement D for impact tests of the same component at three different V_0 and $m = 1.57$ kg. The max acceleration a_m for each V_0 is denoted as a black dot. The units of a are the gravitational acceleration g . d, The a_m at various V_0 with $m = 1.57$ kg. e, Janssen factor J vs. V_0 . The test with the lowest J is designated as the point (V_0^*, J^*) . f, Schematic of a fixed speed test in which a platen is lowered at constant velocity V . g, Force F vs. D for quasistatic (QS) test ($V = 2$ mm/min). This data is used to compute the most efficient operating conditions in terms of the critical force threshold F_t^* and maximum mechanical energy absorbing efficiency K_s^* , which is defined by dividing the amount of energy absorbed (light blue area) by the maximum amount of energy that could be absorbed below F_t^* (light red area). h, Photograph showing five additional GCS designs made from TPU. i, QS model uses a single QS test to compute the predicted optimal impactor velocity \tilde{V}_0^* . j, Parity plot showing \tilde{V}_0^* vs. V_0^* for the six shown designs using the QS model for $m = 1.57$ kg. Error bars denote one standard deviation.

In order to examine the behavior of a component when impacted under various conditions, a specific GCS design was selected and fabricated out of thermoplastic polyurethane (TPU, Ninjatek – Cheetah) using FFF (Figure 4.1a) and subsequently tested in impact (Figure 4.1b). Representative a -displacement D curves illustrate the outcome of testing this component with too little KE , an ideal amount of KE , and too much KE , all relative to the ideal conditions for this component (Figure 4.1c). At low V_0 , a_m increases slowly with increasing V_0 due to strain-rate dependent material properties increasing the yielding force (Figure 4.1d). However, once V_0 increases to the point where densification is reached, a_m increases at a much faster rate with increasing V_0 as the component is unable to absorb KE without densifying further. This transition point coincides with (V_0^*, J^*) , which is the condition under which the design is most efficient (Figure 4.1e). Because of this, (V_0^*, J^*) is a highly useful way to quantify impact performance. However, determining this value required a large number of impact tests on independently prepared samples and is only valid for one value of m .

We hypothesize that fixed-velocity tests (Figure 4.1f) could be used to predict V_0^* . While there are substantial differences between impact and fixed-velocity testing, the move to fixed-velocity would be a tremendous reduction in the experimental burden as fixed-velocity tests use relatively common universal testing machines (UTMs), are amenable to automation, and feature fewer variables than impact tests. The result of a fixed-velocity test is a plot of force F vs. D (Figure 4.1g). One key question in exploring our hypothesis is determining a way to extract information from this QS curve that can be useful for predicting impact performance. In analogy with the non-dimensional J , we

consider K_s (defined in Equation (4.3)) as an important non-dimensional metric of a component's ability to absorb mechanical energy as previously explored by Gruenbaum, G. & Miltz¹³⁶. Graphically, K_s can be calculated by dividing the amount of energy absorbed (light blue area in Figure 4.1g) while $F \leq F_t$ by the maximum amount that could be absorbed at that F_t (light red area in Figure 4.1g). While F_t is an independent parameter that is most often chosen with consideration of the operational conditions of the component, there often exists an optimal force threshold F_t^* where the component exhibits its maximum energy absorption efficiency K_s^* . Based on these metrics, a simple prediction of the ideal impact conditions can be found by equating impact KE with the amount of energy most efficiently absorbed during QS testing,

$$\frac{1}{2}m\tilde{V}_0^{*2} = F_t^*K_s^*h, \quad (4.4)$$

and then solving for predicted optimal impact velocity \tilde{V}_0^*

$$\tilde{V}_0^* = \sqrt{\frac{2F_t^*K_s^*h}{m}}. \quad (4.5)$$

While the QS model is both simple and physics based, it neglects strain-rate dependent effects, which is expected to limit its predictive power.

To test the QS model, five additional designs were selected and fabricated from TPU (Figure 4.1h). To apply the model, a single QS test is needed to determine F_t^* , K_s^* , and h for each design. With this data, \tilde{V}_0^* for a given m can be predicted using Equation (4.5) (Figure 4.1i). For each of the six designs, components were fabricated and tested at 2 mm/min in triplicate (Figure 4.1j). While the predictions preserve rank order, they consistently underestimate V_0^* because the QS model does not account for the strain-

rate dependent strengthening. Despite this, because the prediction always underestimates V_0^* , it provides two key pieces of information. First, it provides a safe V_0 that will not reach densification. Second, if subsequent impact testing will be done to identify (V_0^*, J^*) more accurately, it provides guidance of the lower end of potential V_0^* values, potentially decreasing the number of experiments needed to identify the point of maximum performance. Nevertheless, it is clear that more information is needed to further improve this model.

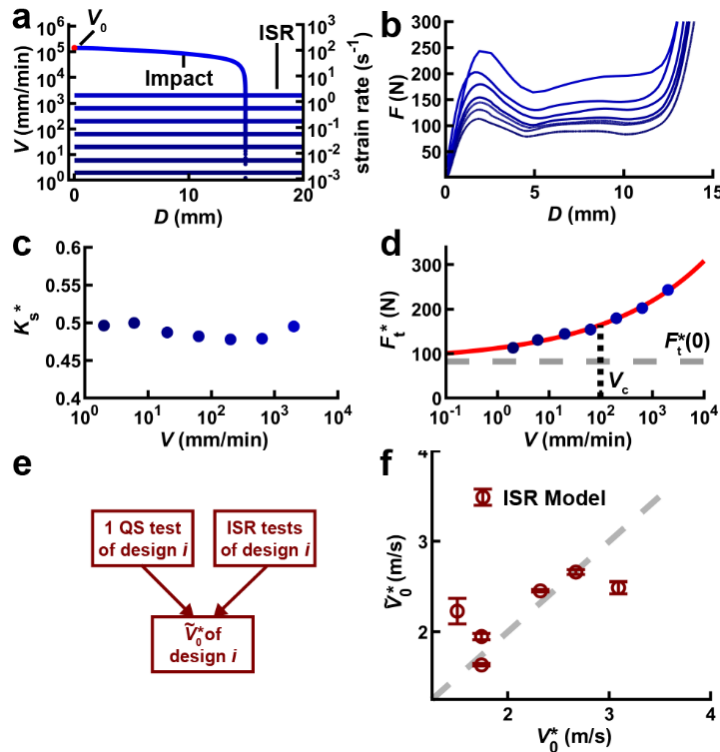


Figure 4.2: Predicting impact performance from intermediate strain rate tests. a, Velocity V vs. D during impact and intermediate strain rate (ISR) tests. While impact tests start at $V = V_0$ and decrease as energy is absorbed, ISR tests have a fixed V for the entire test. b, F vs. D for ISR tests of copies of the component shown in Figure 4.1a. c, K_s^* vs. V for the ISR tests shown in b. d, F_t^* vs. V for the ISR tests shown in b. The red line shows a fit to Equation (4.6) with key fitting constants $F_t^*(0)$ and V_c marked by gray and black dashed lines respectively. e, ISR model uses both QS and ISR tests of the target design to predict \tilde{V}_0^* . f, Parity plot showing \tilde{V}_0^* vs. V_0^* for the six original designs using the ISR model for $m = 1.57$ kg. Error bars denote one standard deviation. In panels a, b, c, and d, shades of blue indicate V .

In order to improve predictions of V_0^* , it is necessary to model the strain-rate dependent effects on the components. Using our UTM, it is possible to test at V of up to 2000 mm/min (Figure 4.2a). Although this is still more than an order of magnitude lower than V_0 during impact testing, it is several orders of magnitude higher than the 2 mm/min speed that was used in QS testing. Furthermore, when the F - D curves are shown for these ISR tests, strain-rate dependent strengthening is readily apparent (Figure 4.2b). Specifically, the curves have similar shapes, but F increases with increasing V . Interestingly, higher V has a modest effect on K_s^* (Figure 4.2c), which bears the advantage that QS tests can be used to estimate this value. In contrast, F_t^* monotonically increases with increasing V (Figure 4.2d), making this parameter a clear quantification of strain-rate dependent strengthening. Empirically, we find that these experiments can be fit using,

$$F_t^*(V) = F_t^*(0) \left(1 + \left(\frac{V}{V_c} \right)^\alpha \right), \quad (4.6)$$

with fitting constants $F_t^*(0)$, V_c , and α . These fitting constants have a physical basis: $F_t^*(0)$ represents the asymptotic limit at low speeds, V_c represents the characteristic velocity where $F_t^*(V)$ doubles, and α represents how quickly $F_t^*(V)$ increases as V increases. The ISR model is formed by combining Equation (4.6) and Equation (4.5) to obtain a transcendental equation that can be solved numerically to predict \tilde{V}_0^* of a given design based upon QS and ISR tests for the design (Figure 4.2e). The predictive performance of the ISR model is significantly improved, although the rank order is no longer correct. Despite this improvement, the ISR model only uses data from a given design and requires a large number of fixed velocity tests for each prediction.

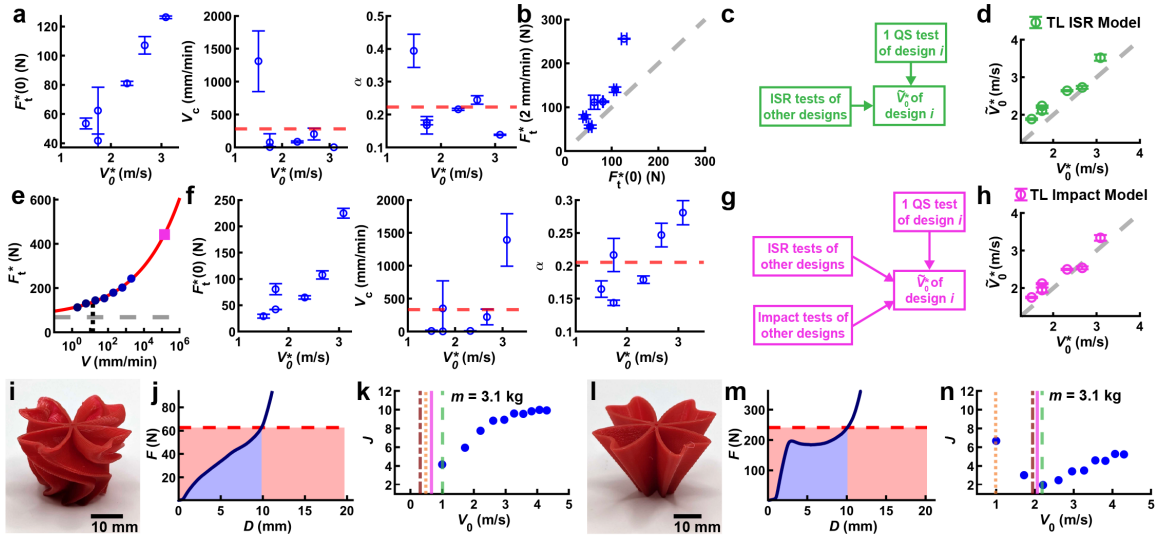


Figure 4.3: Predicting impact performance using transfer learning. a, Values from ISR fitting each of the six designs to Equation (4.6) with the mean value shown as a dashed red line. b, Plot of QS $F_t^*(2 \text{ mm/min})$ vs. fitting constant $F_t^*(0)$. c, Transfer learning (TL) ISR model that uses ISR tests of other designs to calculate V_c and α , allowing the prediction of \tilde{V}_0^* with a single QS test for each target design. d, Parity plot showing \tilde{V}_0^* vs. V_0^* for the six designs using the TL ISR model for impactor mass $m = 1.57 \text{ kg}$ using leave one out cross validation (LOOCV) to calculate V_c and α for each prediction. e, F_t^* vs. V for ISR tests with theoretical impact point (V_0^* , F_t^*) (pink square) calculated using Equation (4.7). f, Values from ISR and impact fitting each of the six designs to Equation (4.6) with the mean value shown as dashed red line. g, TL impact model uses ISR tests and impact F_t^* prediction from other designs to calculate parameters V_c and α , allowing the prediction of \tilde{V}_0^* with just a single QS test for each new target design. h, Parity plot of \tilde{V}_0^* vs. V_0^* for the six original designs using the TL impact model for $m = 1.57 \text{ kg}$ using LOOCV. i, New design printed from TPU. j, QS test for new design showing F vs. D with red dashed line indicating F_t^* . k, J vs. V_0 for the design shown in i with $m = 3.1 \text{ kg}$. Dashed lines represent \tilde{V}_0^* using the four models. l, Second new design printed from TPU. m, QS test for second new design showing F vs. D with red dashed line indicating F_t^* . n, Values of J vs. V_0 for design shown in l for $m = 3.1 \text{ kg}$. Dashed lines represent \tilde{V}_0^* using the four models. Error bars throughout represent one standard deviation.

Despite the improvements with the ISR model, taking a full sweep of ISR test data is much slower than taking a single QS test and high speeds tests have the potential to damage the load cells of UTMs. Therefore, we hypothesized that transfer learning could be employed to eliminate the need to take ISR tests for each new target design. By plotting the fitting constant values for $F_t^*(0)$, V_c , and α , some trends become clear (Figure

4.3a). While $F_t^*(0)$ clearly varies for each design, both V_c and α have consistent values for all tested designs, suggesting that their mean values could be used. Fortunately, $F_t^*(0)$ is correlated with $F_t^*(2 \text{ mm/min})$ determined by QS experiments of the target design (Figure 4.3b). Therefore, a transfer learning model can be created that uses the mean value of V_c and α from other designs while establishing $F_t^*(0)$, K_s^* , and h from a single QS test (Figure 4.3c) for each new target design. In this way, the time-consuming ISR tests can be done once for a representative set of designs to establish baseline values, and then subsequent designs can be predicted using a single QS test. Interestingly, when predicting impact performance using leave one out cross validation (LOOCV), the new TL ISR model performs significantly better than the ISR model trained on ISR data from each target point (Figure 4.3d). This suggests that V_c and α are not design dependent, as averaging the values for several designs provides better predictive performance than using the value from the specific design.

Motivated by the success of incorporating ISR data from other designs, we hypothesized that using limited impact testing could further improve the prediction of $F_t^*(V)$. Unfortunately, impact tests do not produce an estimate of F_t^* that can be directly used to estimate strain-rate dependent strengthening at these higher strain rates. Despite this, the value of V_0^* determined by impact experiments for each design can be used to compute what the required $F_t^*(V)$ would need to have been to arrive at the correct prediction, which can be computed by rearranging Equation (4.4) as,

$$F_t^*(V_0^*) = \frac{mV_0^{*2}}{2K_s^*h}. \quad (4.7)$$

This $(V_0^*, F_t^*(V_0^*))$ point can then be included in the $F_t^*(V)$ fit model to further improve its accuracy at high velocities (Figure 4.3e). Fitting the ISR and impact data together to Equation (4.6), we find that the additional point provided by impact testing barely changes $F_t^*(0)$, but slightly adjusts V_c and α (Figure 4.3f). Importantly, the scaling parameter α decreases from 0.223 to 0.205, suggesting that the TL ISR model overestimates the strain-rate dependent strengthening at higher velocities. With this addition, the TL impact model can predict \tilde{V}_0^* using a model initialized on a representative set of ISR and impact tests and a single QS test for each new design (Figure 4.3g). The TL impact model shows slight improvements to predictive accuracy at the cost of including some impact data to train the initial model (Figure 4.3h).

To further explore the transferability of this model, we tested two new designs at $m = 3.1$ kg, nearly twice the mass of the impactor used to generate all training data. The first design (Figure 4.3i) had low stiffness and no clearly defined yield point (Figure 4.3j). All models predict the very low \tilde{V}_0^* for this soft design, consistent with the test data obtained from impact testing. In contrast, another design (Figure 4.3l) was chosen that had a higher F_t^* than previous designs and a clear yielding region (Figure 4.3m). Here, the models predict higher \tilde{V}_0^* consistent with experimental results, with the QS model being the notable outlier (Figure 4.3n).

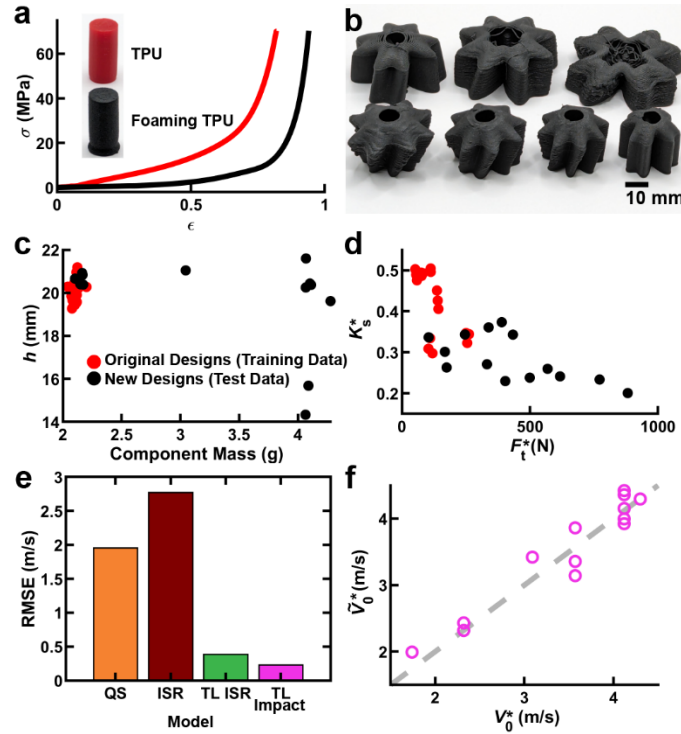


Figure 4.4: Extending transfer learning models to an additional material. a, Stress σ vs. strain ϵ for QS tests of solid cylinders 8 mm in diameter with $h = 16$ mm used in material testing. b, Picture of seven of the 16 additional designs fabricated with a foaming TPU filament. c, The h vs. component mass for original six designs and the new 16 foaming TPU designs. d, The K_s^* vs. F_t^* for original six designs and new 16 foaming TPU designs. e, Root mean squared error (RMSE) for predictions of V_0^* of new foaming TPU designs. Both TL models use fitting constants V_c and α trained solely on data from original six TPU designs. f, Parity plot showing \tilde{V}_0^* vs. V_0^* for the 16 foaming TPU designs using the TL Impact model for $m = 1.57$ kg.

As a final test of generalizability, we extended these models to a new material, a TPU filament mixed with a blowing agent (ColorFabb – VarioShore) such that when printed at 250 °C, it forms a foam with a density half of the previously studied TPU. When testing solid cylinder samples fabricated with both the TPU and foaming TPU (Figure 4.4a), the foaming TPU sample shows a significantly lower modulus and densifies at a much higher strain. Furthermore, we switched to a more complex GCS design space (Figure 4.4b), which adds a tapered region to provide an initial region with low stiffness. Some of the new foaming TPU designs had heights that included those

much lower than the previously studied designs and some were up to twice the mass of any components in the training set (Figure 4.4c). The new designs generally had much lower K_s^* due to the inefficiency of the tapered region (Figure 4.4d), but they also had much higher F_t^* , which were designed to target higher V_0^* . All four models were applied to the sixteen new foaming TPU designs, with both TL models' V_c and α parameter values being trained exclusively on data from the original six TPU designs. The TL impact model had an RMSE of 0.23 m/s, with the TL ISR model close behind and the ISR and QS models' errors significantly higher (Figure 4.4e). The TL impact model trained exclusively on the original six TPU data shows excellent accuracy at predicting impact performance using only a single QS test for each new foaming TPU design (Figure 4.4f), especially considering that the design space, h , component mass, and component material were all broadened from those contained in the training data. This suggests that the physics-informed transfer learning model may be able to extrapolate far beyond the limited training data used to train it.

4. Conclusion

This work reports a systematic exploration of how the impact performance of structures can be predicted using a set of models with varying complexity. In all cases, models are physics informed to meaningfully connect metrics from QS testing to impact performance. Our key finding is that the incorporation of related training data that includes testing in the ISR and impact regimes allows these models to powerfully extrapolate to new designs, materials, and impact conditions. From a design perspective, this work implies that once a representative set of data has been acquired, new designs

can be rapidly evaluated based on a single QS test. Furthermore, the success of this model in extrapolating beyond the training materials suggests that models trained on large databases of polymeric structures may be useful with new materials and architectures. To that end, the TL impact model based on all data presented in this study is given by $V_c = 331$ mm/min and $\alpha = 0.205$, which, when combined with a single QS test of a new sample, can be used to estimate \tilde{V}_0^* .

The ability to predict performance with a single QS test allows researchers to screen candidate structures for impact performance quickly and easily. Additionally, it allows researchers to leverage databases of QS tests already published or that can be generated using high throughput experimentation or self-driving labs. When impact testing is required to validate a chosen design, knowing the \tilde{V}_0^* before impact testing can significantly decrease the number of tests needed to pinpoint (V_0^*, J^*) . Furthermore, by using a physics-informed model instead of density or other approaches that depend on the specific structure of the individual design or material, it is more likely that the model will extrapolate to more disparate designs and structures.

5. Methods

GCS designs consist of 11 parameters that transform a standard cylindrical shell⁶⁹. The STLs were generated using a previously published python script¹⁴². For foaming TPU designs, tapered regions were added to the GCS components by adding two distinct regions. First, there is a transition zone where the top shape of the GCS component increases in diameter while the lobes are decreased size. The second region transforms the intermediate shape to a circular cross section with a small diameter, creating a

reentrant region that can collapse unto itself while offering low stiffness. These modifications retain the topological nature of the GCS space and can be considered a subset of GCS.

The STLs were converted into gcode using Slic3r and sent to the printers using Octoprint. Components were printed on one of five MakerGear M3-ID FFF printers on a glass bed covered with polyimide tape. Components were printed using either TPU (Ninjatek - Cheetah) or foaming TPU (ColorFabb - VarioShore) filament. Components were printed at 250 °C nozzle temperature and 80 °C bed temperature. The bed was heated to 100 °C before removal.

Components were removed from the printers using a UR5e (Universal Robotics) 6-axis robot arm and transferred to a scale (Sartorius CP225D) to record the mass. Afterwards, components for impact testing were placed on a table by the robot arm for temporary storage. Components for QS or ISR testing were moved by the robot to a UTM (Instron 5965). Three QS tests for each TPU design were performed at 2 mm/min until the force reached 4.5 kN. A single ISR test for each TPU design was performed at each speed (6, 20, 63, 200, 632, or 2000 mm/min). For higher testing speeds, the 4.5 kN threshold was lowered for some designs that densify quickly to protect the UTM from damage, but this did not prevent the proper calculation of K_s^* or F_t^* . A single QS test was performed for each foaming TPU design, and no ISR data was taken for foaming TPU designs. Components tested in impact were transferred to another location for testing on a drop tower impact system (Ceast 9350) with a flat-ended steel impactor with a weight of either 1.57 or 3.1 kg. Because impact testing took place at an alternate site, there was an

interval of several days between fabrication and testing of impact components.

QS and ISR tests record data as F - D data. The h of the component is calculated by measuring when the F first reaches 1 N. The K_s^* and F_t^* are then found by calculating the max and argmax of Equation (4.3). For impact testing, each design is tested once at each V_0 . Impact data is recorded as a vs. time. Because of noise in the signal, a is smoothed using a Gaussian filter by using the MATLAB function “smoothdata” with a smoothing factor of 0.13. The a_m is then taken as the max of this smoothed curve. Subsequently V_0^* and J^* are calculated by finding the minimum J for each design and its corresponding V_0 .

When calculating \tilde{V}_0^* , only a single QS test was used at a time. Therefore, to get standard deviation for the TPU models, the models were rerun with each of the three QS experiments. Likewise, when fitting Equation (4.6), a single QS test was paired with all ISR tests for each of the three predictions using the MATLAB function “fit” with the restriction that of $F_t^*(0) \geq 0$, $V_c \geq 2$, and $\alpha \geq 0$. For LOOCV predictions (Figure 4.3), ISR data from the target design was excluded when calculating the mean of fitting constant V_c and α . However, fitting constant $F_t^*(0)$ was calculated from the F_t^* of the single QS test used for each of the three predictions for that design. When predicting \tilde{V}_0^* for the new designs (Figure 4.3k,n and Figure 4.4), ISR data from all six original TPU designs were used when calculating the mean of fitting constants V_c and α . Only a single QS experiment was performed for each foaming TPU design, so no standard deviation was reported.

CHAPTER 5: EXPLORING SDLs AS A COMMUNITY RESOURCE

5.1: Introduction

SDLs are an efficient way to accelerate research, but they often have large start-up costs, both in money and in researcher time. Because their marginal costs to run tend to be more modest, it is important to utilize them as much as possible. One way to increase the utilization of existing systems is to collaborate with other researchers on topics on which they are already working. In this way, the SDL can serve as a community resource. Here, we partnered with researchers from Northeastern University to study viscous thread printing (VTP), an area that they had been working on previously for several years.

Viscous thread instability (VTI) is the phenomenon where a viscous threads will coil or oscillate in predictable patterns when falling upon a moving surface that moves slower than the exit velocity of the thread^{143–146}. This phenomenon has been exploited to create a wide variety of 3D foams on fused filament fabrication (FFF) printers^{147,148}.

Viscous thread printing (VTP) is characterized by four parameters: V^* , H^* , ΔL , and ΔZ .

V^* and H^* are both dimensionless variables defined by the following equations,

$$V^* = \frac{V_n}{V_t}, \quad (5.1)$$

$$H^* = \frac{H}{D_t}, \quad (5.2)$$

where V_n is the nozzle velocity, V_t is the thread velocity, H is the height of the nozzle above the part, and D_t is the thread diameter after exiting the nozzle (Figure 5.1a). The two remaining parameters, the distance that the print head moves between adjacent

toolpath lines ΔL and when starting a new layer ΔZ are necessary for 2D and 3D structures respectively (Figure 5.1b).




Figure 5.1: Parameters for viscous thread printing (VTP). a, In VTP, the nozzle speed V_n travels slower than the thread exit velocity V_t , while at a height above the print H significantly higher than the thickness of the extruded thread D_t , causing the thread to coil in patterns that can be predictable. b, To create 3D objects, rectilinear toolpath (blue) can be used while varying the spacing between lines in the x-y plane by ΔL and spacing layers by ΔZ . A cross hatch is used to ensure continuity from layer to layer.

Here, we utilize our SDL to map the VTP parameter space by running two individual campaigns studying PLA and TPU filament. We then develop GPR models for each material that predict both the layer height H_L and the effective modulus E . By combining the outputs of these two models, we are able to identify the homogenous subspace within the four-parameter space and target specific E . Additionally, we built a model to predict the stress-strain curve for VTP foams using PCA and GPR techniques, allowing us to replicate the stress-strain response of common commercial foams.

5.2: Results and Discussion

5.2.1: Modeling and homogeneous subspace

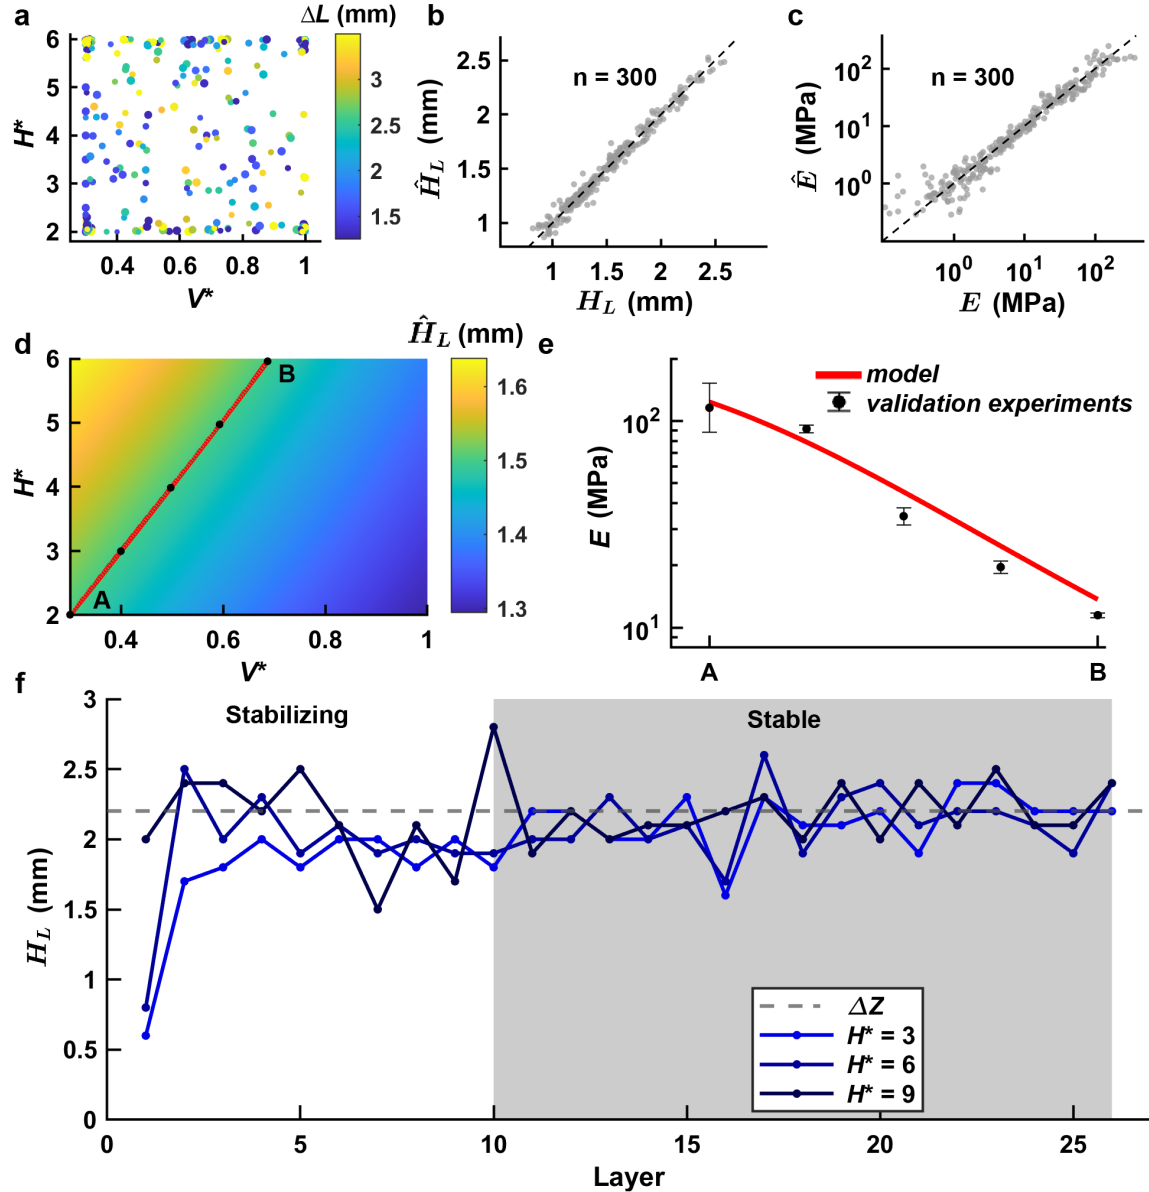


Figure 5.2: Mapping homogenous subspace for PLA filament. a, H^* vs. V^* for all 300 experiments performed in PLA with color indicating ΔL and size indicating ΔZ . b, Parity plot of predicted layer height \hat{H}_L vs. H_L for Gaussian process regression (GPR) model using leave one out cross validation (LOOCV). c, Log-log parity plot for predicted effective modulus \hat{E} vs. effective modulus E for the GPR model using LOOCV. d, Slice in 2D of H^* and V^* where ΔL and ΔZ are equal to 1.5 mm. Color indicates \hat{H}_L and the red line indicates predicted homogenous

subspace where $\hat{\mathbf{H}}_L = \Delta\mathbf{Z}$. Black dots indicate five equally spaced samples selected for validation testing. e, GPR model's \hat{E} (red line) vs. results of five validation experiments (black). Error bars represent one standard deviation in semi-log space from each condition being tested in triplicate. f, Layer height H_L vs. layer number for three PLA cubes with different H^* . After several layers, H_L for each cube stabilizes to 2.2 mm, which is equal to $\Delta\mathbf{Z}$.

While the flexibility of VTP is a virtue from a materials design perspective, tuning the processing parameters to realize homogenous materials presents a challenge.

Specifically, when the first layer of a sample is printed, the processing variables V^* , H^* , and ΔL collectively determine the layer height H_L . For a structure to be homogeneous, ΔZ must equal H_L so that the distance between the print nozzle and substrate does not change when printing the next layer. Therefore, for each V^* , H^* , and ΔL there is a unique ΔZ that will lead to a homogeneous part.

In order to predict the correct ΔZ for each V^* , H^* , and ΔL triplet, we used a SDL to run an experimental campaign that iteratively selected and performed 300 experiments using Bayesian optimization to minimize the uncertainty across the full four-dimensional space, on VTP foams formed out of polylactic acid (PLA) (Figure 5.2a). After these experiments, the performance of the final Gaussian process regression (GPR) model to predict H_L was evaluated using leave one out cross validation (LOOCV) and found to have a R^2 value of 0.980 (Figure 5.2b). In addition, we also conditioned a GPR to predict an effective modulus E of each part, which was found to have a LOOCV R^2 value of 0.946 (Figure 5.2c). Importantly, these two models can be combined to allow precise control of the properties of a homogenous cube by enabling inverse design to target a desired E .

We hypothesized that the SDL-derived models for predicted layer height \hat{H}_L and predicted effective modulus \hat{E} could allow predictive control over the VTP outcome and enable allow us to fully exploit the capabilities of VTP printing. In particular, a major goal is to spatially vary the mechanical properties of a material in a manner that is easy and reliable to fabricate. Ideally, this could be achieved by maintaining a constant ΔL and ΔZ throughout the part, thus enabling a consistent rectilinear toolpath as shown in Figure 5.1b, but with varying V^* and H^* to change the E of different areas of the part. To explore this possibility, a 2D slice of H^* and V^* at a constant ΔL and ΔZ can be visualized (Figure 5.2d). The \hat{H}_L is shown, with the red line indicating the homogenous subspace where $\hat{H}_L = \Delta Z$. Note that while this red line appears to be linear, the underlying GPR model makes no assumption of linearity. Five points that are equally spaced from A to B are selected along this homogenous line. The E of the physical experiments is then compared to \hat{E} from the GPR model, showing that it can be modulated between ~ 10 and 100 MPa while remaining in the homogenous subspace with constant ΔL and ΔZ . (Figure 5.2e).

While the presence of a smooth homogenous subspace that facilitates the manufacture of graded materials is a powerful outcome of this study, the shape of the inhomogeneous region also revealed a fortuitous relationship in the underlying VTP process. Specifically, Figure 5.2d shows that there is a positive correlation between H^* and H_L . This positive correlation implies the existence of a stabilizing basin of attraction to correct errors in the selection of ΔZ . If ΔZ is too large, then $H_L < \Delta Z$ and the effective H^* will increase each layer until $H_L = \Delta Z$. In contrast, if ΔZ is too small, then $H_L > \Delta Z$

and the effective H^* will decrease until $H_L = \Delta Z$. Consequently, the component will homogenize at a new effective H^* if the selected parameters are not already in the homogenous subspace, although this process will change the properties of the foam and produce a non-homogenous zone for several layers during stabilization. For this reason, it is important to determine the correct ΔZ before printing. Nevertheless, the basin of attraction means that small errors in the selection of ΔZ will not be catastrophic to the printing process and that any anomalies during printing will self-correct. To demonstrate this basin of attraction, we printed three samples with the same V^* , ΔL , and ΔZ , that varied in their H^* . After printing these samples, we measured the thickness of each layer. As expected, H_L for each of the experiments stabilized at the selected ΔZ , regardless of H^* (Figure 5.2f).

5.2.2: Comparison to common foams

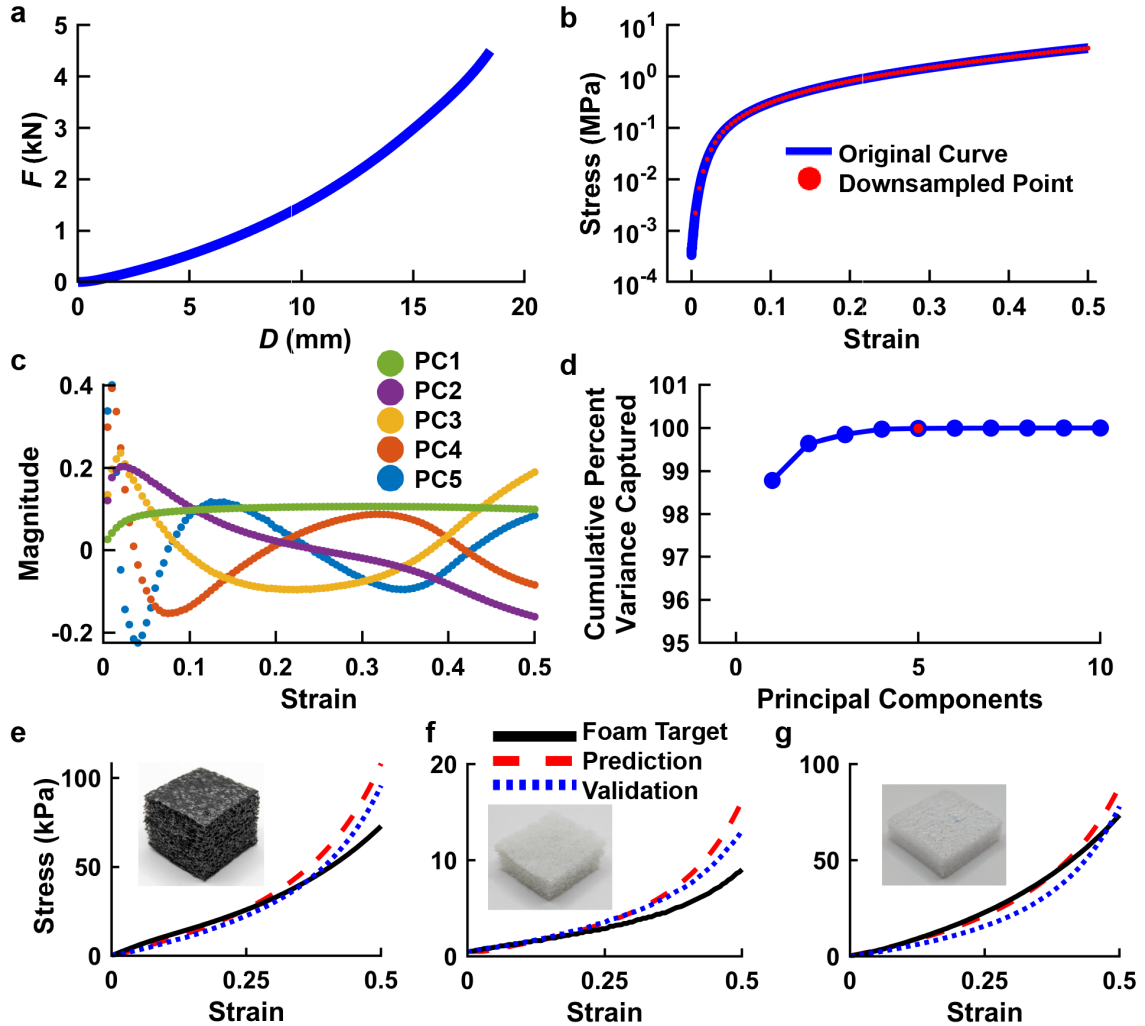


Figure 5.3: Predicting full stress-strain curves to replicate commercial foams. a, Force (F)-displacement (D) of a TPU cube to the 4.5 kN force limit. b, Log stress-strain curve (blue) converted from F - D curve is down sampled to 100 points (red) equally spaced to 50% strain. c, Principal component analysis breaks down sampled curves into 100 components (only top five shown for clarity). d, Cumulative variance captured by number of components. The vast majority of variance is captured by the first five components. e-g, Stress-Strain curves for three foam samples (black curves), the sample predicted to most closely match their performance (red dashed), and the tested performance of that predicted sample (blue dotted).

We hypothesized that the processing freedom afforded by VTP would allow VTP foams to replicate the material performance of common foams. To explore this, we ran an additional campaign with the SDL to explore VTP print parameters for TPU filament. As before, the campaign used active learning to fully explore the parameter space with this process requiring 155 experiments. Next, *F-D* curves (Figure 5.3a) were converted to stress-strain curves based on the sample geometry. The stress-strain curves were then down sampled to 100 stress points that were equally spaced in strain from $\varepsilon = 0.5$ to 50% (Figure 5.3b). We then took the logarithm of these stress values to preclude the possibility of predicting negative stress. Principal component analysis was then applied on all 155 TPU samples tested (Figure 5.3c–d). A separate GPR model was conditioned on the scores for each of the first five principal components as a function of their processing conditions. These GPR models could then be used to predict the performance of VTP foams that have not yet been tested.

In order to evaluate the VTP foam selection pipeline we selected three reference foam samples from everyday objects, consisting of a polystyrene foam used for packing/shipping (Figure 5.3e), a cellulose foam used for packing/shipping (Figure 5.3f), and a polystyrene foam used as part of a children’s car seat (Figure 5.3g). They were tested in quasistatic compression using the same protocol as the VTP foams. The performance of each foam was then compared to the predictions of 125,000 potential experiments selected using Latin hypercube sampling. An error function minimized the mean squared error of the difference between the log stress prediction and the foam stress, as described in section 5.4.3. The VTP cube with the lowest error was selected and

tested, showing the ability to match the performance of common foam samples.

5.3: Conclusion

This chapter develops GPR models that accurately predict the H_L and E of VTP foams, allowing the fabrication of homogenous foams with targeted mechanical properties. Furthermore, by using PCA and GPR, we demonstrated the ability to predict stress-strain responses for VTP foams and used these predictions to replicate the stress-strain response of common commercial foams.

Furthermore, this work pioneers the use of an SDL as a community resource by opening up the BEAR for use with outside researchers. This provides a template for how researchers can collaborate to increase the utilization of preexisting SDLs and further accelerate scientific discovery and innovation.

5.4: Methods

5.4.1: Layer height measurements

Layer height measurements (Figure 5.2f) were made through a manual process. First, a single cube was printed from PLA for $H^* = 3$, $V^* = 0.3$, $\Delta L = 1.5$ and $\Delta Z = 2.2$. The height of the cube was measured using a universal testing machine (UTM) (Instron 5965) by lowering the platen manually until the force reached 1 N, ensuring that stray strands of filament did not artificially increase the height measurement. The height of the cube was taken as the platen separation of the UTM. Then, the bottom layer of the cube was removed with wire cutters. Because the filament is more tightly bonded within the layer than between layers, it was possible to cut away a single layer without disturbing the remaining layers. However, because of the thin nature of the removed layer, it was

destroyed in the process. The height of the remaining cube was then measured in the UTM with the 1 N threshold. By comparing the difference in height between the two measurements, the height of the removed layer was inferred. This process was repeated until only two layers remained. Unfortunately, because the layers were so thin, measurement of the final two layers was not possible. This process was repeated for the other two cubes with $H^* = 6$ and 9.

5.4.2: Automated testing

Automated testing leveraged the Bayesian autonomous experimental researcher (BEAR) previously developed^{27,69}. The system consists of a UTM (Instron 5965), a scale (Sartorius CP225D), a fused filament fabrication (FFF) printer (Prusa Mk3S+), and a six-axis robot arm (Universal Robotics UR5e).

The learning loop consisted of several steps performed in sequence. First, an experiment is selected using active learning. GPR models were conditioned in MATLAB using the built-in function `fitrgp` using a squared exponential kernel with automatic relevance detection. Potential experiments were calculated using Latin hypercube sampling and an experiment was selected using a maximum variance decision-making policy. Initially, E was used as the target metric for the GPR, but it was later changed to H_L to ensure that both metrics were mapped fully.

G-code for the selected experiment was then generated using a custom python script. The g-code was sent to the printer using OctoPrint through the python package OctoRest. For PLA (eSun PLA+ gray), printing temperature was set to 215 °C and the bed temperature was set to 60°C. For TPU (Ninjatek Ninjaflex blue), printing

temperature was set to 230 °C and the bed temperature was set to 50 °C. For both materials, the part was removed by the robot after the bed had cooled below 32 °C.

The part was then transferred to the scale where it was weighed. If the measured mass was more than 5% from the expected value, the part was discarded due to the likelihood of a print error. After recording the mass, the part was transferred to the UTM for testing. During testing, the UTM lowered its platens at 2 mm/min until the force reached 4.5 kN or the platens were separated by less than 0.4 mm.

The E was calculated by converting the force-displacement curve into a stress-strain curve. The height for this conversion was measured by the UTM as when the force exceeds 0.3 N. The cross-sectional area was assumed to match the designed dimensions for the cube of 30 x 30 mm². The E was then calculated using a linear fit of the stress between 5% and 15% strain to avoid any toe regions.

The H_L of the part was calculated by taking the height of the final part and dividing by the number of layers. For a part to be considered homogenous, H_L must be within 5% of ΔZ .

5.4.3: Principal component analysis

The stress-strain curve was converted to log10 and truncated at 50% strain. It was then down sampled to 100 equally spaced points. PCA was performed using MATLAB's built in `pca` function. Because 99.99% of cumulative variance was captured by the first five components, only the first five components were used for stress-strain curve prediction. For each of these components, a GPR model was conditioned to predict the score of the component based on the four-dimensional input space (V^* , H^* , ΔL , and ΔZ).

Sampling points (125,000) were then selected using Latin hypercube sampling for V^* , H^* , and ΔL . ΔZ was selected using the H_L GPR model to ensure $\Delta Z = H_L$, therefore ensuring that the resulting sample was homogenous. The predicted score components were then transformed back into the log10 stress/strain space. The predicted curves were compared to the foam samples using the error function shown in Equation (5.3). The closest curve was selected for each foam type and tested.

$$Error = \frac{\sum(\log(\hat{\sigma}_n) - \log(\sigma_n))^2}{n} \quad (5.3)$$

Foam samples were cut to 30 x 30 mm² rectangular prisms to match the cross-section of the sample VTP foams, but the heights of the foam samples were not modified. The foams were weighed by hand and manually transferred to the UTM. The UTM tested them using the same 2 mm/min speed and the same 4.5 kN stop threshold protocol.

CHAPTER 6: CONCLUSION

In addition to reporting on the experimental work presented in the previous chapters, it is important to reflect on broader lessons learned and opportunities for the SDL community. In this chapter, we consider the lessons that we have learned while operating an SDL for many years and offer advice for the researcher and the community at large on how to develop SDL systems and maximize their utility.

6.1 Lowering barriers to develop and use SDLs

In Chapter 2, we shared insights on how to run an SDL that we gleaned after running our own campaign for more than two years. Our guidance focused on six key settings that can be tuned to optimize an active learning campaign and four plots that should be monitored to evaluate performance and inform changes to the settings. These insights will provide guidance for future researchers that are looking to run an SDL campaign for the first time or looking to run longer campaigns that would benefit from periodic adjustments. Despite this, setting up and running an SDL is still daunting for researchers that are new to this research strategy. More work remains to be done by the community to decrease the developmental costs of setting up a new SDL and make it easier to run the SDL on a day-to-day basis.

When developing a new SDL, one main challenge is integrating new hardware. Often, equipment needed to run experiments is designed for human use or is controlled by proprietary software. Ideally, equipment will be manufactured to have a common interface plug, such as a USB port, and a software development kit (SDK) with full documentation that supports all hardware features. Unfortunately, hardware is generally

manufactured to somewhere in between these two extremes, and it is often difficult to know before purchase how well it is supported. To accelerate the development of SDLs, researchers should encourage manufacturers to develop full-featured SDKs for all major operating systems and be transparent about the level of support before purchase.

An additional barrier to SDL adoption is the time it takes for a researcher to learn the machine learning needed to implement active learning. Many researchers come from disciplines that do not generally require a background in programming, which makes the task of developing code for machine learning very challenging. Furthermore, all effort focused on optimizing the machine learning code is time diverted from the main research focus of the researcher. To lower this burden, standard active learning software should be developed that allows users to implement active learning with a simple graphical user interface. ChemOS is a Python-based software package that is intended on becoming a universal SDL management system, similar to how the Robot Operating System (ROS) has become a common choice for controlling and coordinating robots¹⁴⁹. Several other attempts have been made to unify decision making and equipment control into a single package, but none of the competing platforms have fully matured. Hopefully the community will consolidate around a single platform that is easier to use for non-programmers.

6.2 Leveraging dataset for additional mechanical insights

Chapter 3 presents our two-year SDL campaign that resulted in the discovery of a structure with the highest mechanical energy absorbing efficiency reported to date. This campaign also led to structures with optimized energy absorbing efficiency across a

broad stress range from 100 Pa to 10 MPa, explored how material and design are interconnected, and offered insights on general design motifs. Finally, we provided a template for how to report the progress and adjustments made during a long-term SDL campaign. Despite these accomplishments, there are still many opportunities to expand upon this work.

First, more can be done to study the interesting GCS design family. While our study focused on the multi-objective optimization of K_s^* and σ_t^* , countless other interesting metrics could be studied, such as U_m , U_v , hysteresis, durability, or recovery after compression. Additionally, although we showed that tiling GCS components in 2D space could in principle improve the performance by averaging out printing defects and variance in the post yield softening, more work can be done to evaluate the mechanical performance of GCS components tiled in 2D and 3D. Finally, GCS components can be printed using other materials, with metals being a particularly exciting potential option. It may be interesting to see how much performance could be improved in broader range of σ_t^* . Of particular interest, it may be possible to use transfer learning to guide design of metal GCS parts by leveraging the large polymer GCS dataset.

In addition, utilizing an SDL to study interesting and complex families of structures offers the opportunity to expand upon general design principles by discovering commonalities across disparate design families. Already by studying the GCS family we have discovered principles that we believe will transcend the design family itself. As more distinct and complex design families are studied thoroughly, more trends may become apparent. However, for this to happen, robust data sharing principles must be

employed, such as the findable, accessible, interoperable, and reusable (FAIR) principles for scientific data management¹⁵⁰. In addition, permanent data storage must be considered carefully, especially because the amount of data produced by SDLs can quickly become burdensome. Including up front permanent data storage costs in grant application budgets should be considered for researchers, and research institutions should consider formalizing data management policies and solutions for their researchers.

6.3 Working towards fundamental impact performance metric

In Chapter 4, we developed a physics-informed model to predict V_0^* for GCS designs based on QS tests and refined that model using limited ISR and impact data to account for strain-rate strengthening. The model was then validated by applying it to a new impactor m , new designs, and a new material. The RMSE for these new designs was 0.23 m/s, showing the ability for the model to extrapolate far beyond the limited training data. Because the model can make predictions using only a single QS test, it opens the opportunity to screen performance from large datasets like the one we produced in Chapter 3 or to do active learning using predicted impact performance as the metric of interest.

Furthermore, many opportunities exist to expand upon our study of impact performance for GCS components. Developing a model to predict the strain-rate dependent strengthening for the entire σ - ε curve could be useful for a variety of applications. With our SDL, it is possible to create sweeps of ISR data for hundreds of designs, rather than the few designs we have tested to date. This process could unlock additional insight into how structures respond to increasing strain rates in a general sense.

Perhaps the most interesting prospect is building a model that can predict a global J^* for all m and V_0 combinations. Just as our quasistatic work sought to represent the performance of a part as a single point in 2D space (σ_t^*, K_s^*) , the impact performance of a GCS part should have a globally optimal J^* at a single combination of m and V_0 . Building a model that could predict this (m^*, V_0^*, J^*) triplet without thousands of tedious impact experiments for each new design could allow even faster optimization of performance in impact and provide additional insights into how design and material affects impact performance.

6.4 Increasing productivity and utilization of SDLs

Finally, Chapter 5 demonstrated the ability to collaborate with outside researchers to use an existing SDL for new research opportunities. Using our SDL, we ran two independent campaigns to study how VTP parameters affect H_L , E , and σ - ϵ curves. After hundreds of experiments, we built GPR models that predicted H_L and E with great accuracy. The combination of these models enables the inverse design of complex 3D mechanisms by programming the local stiffness while maintaining the continuity and integrity of the global structure. Furthermore, we combined PCA and GPR techniques to predict full σ - ϵ curves for homogenous VTP samples and found VTP parameters that match the stress response of common commercial foams.

Not only was this collaboration fruitful in understanding VTP mechanics, but it also demonstrates how opening SDLs up as a community resource can increase their productivity and utilization. Broadly speaking, SDLs have large upfront costs, not only in equipment but also in development time. However, the marginal cost to run each

additional experiment tends to be much more modest. Therefore, once an SDL is built, a high utilization rate is desirable. Similar to user facilities at national labs or the shared use of particle accelerators and astronomical telescopes, grant agencies can fund the development of shared SDLs to serve as a resource to the whole community.

In addition to helping researchers in a single domain, the development of complementary SDLs can provide synergistic benefits for researchers. For example, our SDL focused on compression testing, but a complementary SDL could conduct 3-point bending tests, incorporate an oven for post-print annealing, or include a photogrammetry or 3D laser scanning station to check for print defects. By combining the capabilities of multiple SDLs, researchers could quickly gather a variety of data about their topic of interest. In this way, each individual SDL can become more useful as they proliferate.

6.5 Practical advice for new SDL users

Developing and running an SDL is a challenging endeavor, but it can unlock opportunities that would be nearly impossible to pursue with traditional experimental approaches. When embarking on the development of a new SDL, there are some practical tips to follow to make it easier to succeed.

First, it is important to gather a multi-disciplinary team that contains the broad range of skills necessary to successfully develop and run an SDL, which include knowledge in robotics/automation, machine learning, data science, and subject matter expertise in the research domain. While it is possible for a single person to develop all these skills, it is very challenging to truly be an expert in all the necessary domains. Building a team where people share responsibility according to their skills allows each

individual to develop greater expertise in their area of core competency while also allowing limited cross-training to help manage the transitions when personnel leave and join the project. Furthermore, SDLs can often generate a large quantity of data, overwhelming a lone researcher. Having a team of researchers that can collaborate on the running of the SDL and analysis of the resulting data can permit the SDL to run at full capacity without having to halt experiments to analyze and publish data.

Second, it is important to identify a good scientific problem that is amenable to study with an SDL. While in theory an SDL can be used for any experiment, SDLs have a competitive advantage when certain criteria are met. For example, SDLs are particularly good when studying large combinatorial problems where high number of experiments are necessary. The metric of interest should be difficult or impossible to simulate, thus requiring physical experiments that are enabled by the SDL. The response space should depend on the interaction of the input parameters and should be complex, otherwise small numbers of experiments may be sufficient to map even large, multidimensional input spaces. Additionally, the experiments must be amenable to automation. Tedious, time consuming, or dangerous experiments are prime candidates for SDL, as they free the researcher from needing to physically perform the experiments. Finally, experiments with low marginal costs are preferable, as expensive experiments may make it difficult to fund an SDL for full-time operation.

Third, researchers should consider using GitHub or another version control system to document any changes made to the code base. During development and operation of an SDL, it is likely that changes will need to be made to the control code.

Having a robust documentation system in place can provide transparency to the changes made and can help identify the onset of inevitable bugs. It may also be useful to mark key code changes down in a lab notebook to easily locate them later. GitHub can also allow researchers to easily share or collaborate on code development. However, if the code repository is intended for public sharing, care should be taken to protect passwords, API keys, or other sensitive information.

Finally, researchers should consider how they will record and publish the data that they generate. One of the key advantages of SDLs is that they enable the generation of large, high-quality datasets with rich metadata. As much metadata should be documented as possible, including data that may not be traditionally recorded but is easy to automate, such as temperature and humidity in the laboratory and automated photographs or video of key moments in the experiment. However, this brings additional data management challenges. Perhaps the most fundamental choice is how to identify experiments. It can be tempting for the identifier to have key parameter information built into the filename so that the researcher can quickly understand the experiment by the filename alone. However, this method quickly becomes unwieldy when the input space has many dimensions, and it is not robust to repeat experiments or changes to the parameter space. Instead, we advocate for a simple sequential numbering system (exp1, exp2, exp3...). This makes the identifier short and robust to changes in parameter space. Then, all metadata associated with the experiment can be contained in a table or csv file that links the experiment number to the experiment's metadata.

BIBLIOGRAPHY

1. Lukoševičius, V., Juodvalkis, D., Keršys, A. & Makaras, R. Investigation of Functionality of Vehicle Crumple Zones Recovered after a Traffic Accident. *Applied Sciences* **13**, (2023).
2. Hála, P., Frydrýn, M., Máca, P. & Sovják, R. Comparative study on the crashworthiness of brittle blocks with thin-walled cellular structures. *Thin-Walled Structures* **148**, 106578 (2020).
3. Neice, R. J. *et al.* An Experimental Platform Generating Simulated Blunt Impacts to the Head Due to Rearward Falls. *Annals of Biomedical Engineering* (2021) doi:10.1007/s10439-021-02809-9.
4. Clough, E. C. *et al.* Elastomeric Microlattice Impact Attenuators. *Matter* **1**, 1519–1531 (2019).
5. Rowson, S. & Duma, S. M. Development of the STAR evaluation system for football helmets: Integrating player head impact exposure and risk of concussion. *Annals of Biomedical Engineering* **39**, 2130–2140 (2011).
6. Perkins, P., Jamieson, A., Ferguson, G., Spratford, W. & Hahn, A. Iterative design of impact-damping gloves for safer boxing. *Open Journal of Safety Science and Technology* **8**, 49–97 (2018).
7. Zhang, J. & Ashby, M. F. Mechanical selection of foams and honeycombs used for packaging and energy absorption. *Journal of Materials Science* **29**, 157–163 (1994).
8. Ge, C. & Huang, H. Corner foam versus flat foam: An experimental comparison on cushion performance. *Packaging Technology and Science* **28** 217–225 (2015).
9. Kang, D. *et al.* Multi-lattice inner structures for high-strength and light-weight in metal selective laser melting process. *Materials and Design* **175**, 107786 (2019).
10. Smith, M., Guan, Z. & Cantwell, W. J. Finite element modelling of the compressive response of lattice structures manufactured using the selective laser melting technique. *International Journal of Mechanical Sciences* **67**, 28–41 (2013).
11. Panesar, A., Abdi, M., Hickman, D. & Ashcroft, I. Strategies for functionally graded lattice structures derived using topology optimisation for Additive Manufacturing. *Additive Manufacturing* **19**, 81–94 (2018).

12. Gongora, A. E., Saygin, V., Snapp, K. L. & Brown, K. A. Chapter 12 - Autonomous experimentation in nanotechnology. In *Intelligent Nanotechnology* (eds. Zheng, Y. & Wu, Z.) 331–360 (Elsevier, 2023). doi:<https://doi.org/10.1016/B978-0-323-85796-3.00012-3>.
13. Flores-Leonar, M. M. *et al.* Materials Acceleration Platforms: On the way to autonomous experimentation. *Current Opinion in Green and Sustainable Chemistry* **25**, 100370 (2020).
14. King, R. D. *et al.* Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature* **427**, 247–252 (2004).
15. Nikolaev, P. *et al.* Autonomy in materials research: A case study in carbon nanotube growth. *npj Computational Materials* **2**, 1–6 (2016).
16. Sparkes, A. *et al.* Towards Robot Scientists for autonomous scientific discovery. *Automated Experimentation* **2**, 1–11 (2010).
17. Burger, B. *et al.* A mobile robotic chemist. *Nature* **583**, 237–241 (2020).
18. Volk, A. A. *et al.* AlphaFlow: autonomous discovery and optimization of multi-step chemistry using a self-driven fluidic lab guided by reinforcement learning. *Nature Communications* **14**, 1403 (2023).
19. Jiang, Y. *et al.* An artificial intelligence enabled chemical synthesis robot for exploration and optimization of nanomaterials. *Science Advances* **8**, eabo2626 (2022).
20. Volk, A. A., Epps, R. W. & Abolhasani, M. Accelerated Development of Colloidal Nanomaterials Enabled by Modular Microfluidic Reactors: Toward Autonomous Robotic Experimentation. *Advanced Materials* (2020) doi:[10.1002/adma.202004495](https://doi.org/10.1002/adma.202004495).
21. Zhao, H. *et al.* A robotic platform for the synthesis of colloidal nanocrystals. *Nature Synthesis* **2**, 505–514 (2023). <https://doi.org/10.1038/s44160-023-00250-5>
22. MacLeod, B. P. *et al.* Self-driving laboratory for accelerated discovery of thin-film materials. *Science Advances* **6**, (2020).
23. Liu, Y. *et al.* Autonomous scanning probe microscopy with hypothesis learning: Exploring the physics of domain switching in ferroelectric materials. *Patterns* **4**, (2023).
24. Rooney, M. B. *et al.* A self-driving laboratory designed to accelerate the discovery of adhesive materials. *Digital Discovery* **1**, 382–389 (2022).

25. MacLeod, B. P. *et al.* A self-driving laboratory advances the Pareto front for material properties. *Nature Communications* **13**, 995 (2022).
26. Fan, D. *et al.* A robotic intelligent towing tank for learning complex fluid-structure dynamics. *Science Robotics* **4**, (2019).
27. Gongora, A. E. *et al.* A Bayesian experimental autonomous researcher for mechanical design. *Science Advances* **6**, (2020).
28. Gongora, A. E. *et al.* Using Simulation to Accelerate Autonomous Experimentation (AE): A Case Study Using Mechanics. *SSRN Electronic Journal* 1–15 (2021) doi:10.2139/ssrn.3751791.
29. Duros, V. *et al.* Human versus robots in the discovery and crystallization of gigantic polyoxometalates. *Angewandte Chemie - International Edition* **56**, 10815–10820 (2017).
30. Snapp, K. L. & Brown, K. A. Driving school for self-driving labs. *Digital Discovery* **2**, 1620–1629 (2023).
31. Stach, E. *et al.* Autonomous experimentation systems for materials development: A community perspective. *Matter* **4**, 2702–2726 (2021).
32. Winicov, H. *et al.* Chemical process optimization by computer — a self-directed chemical synthesis system. *Analytica Chimica Acta* **103**, 469–476 (1978).
33. Hickman, R. J., Bannigan, P., Bao, Z., Aspuru-Guzik, A. & Allen, C. Self-driving laboratories: A paradigm shift in nanomedicine development. *Matter* **6**, 1071–1081 (2023).
34. Abolhasani, M. & Kumacheva, E. The rise of self-driving labs in chemical and materials sciences. *Nature Synthesis* **2**, 483–492 (2023).
<https://doi.org/10.1038/s44160-022-00231-0>
35. Green, M. L., Maruyama, B. & Schrier, J. Autonomous (AI-driven) materials science. *Applied Physics Reviews* **9**, (2022).
36. Bennett, J. A. & Abolhasani, M. Autonomous chemical science and engineering enabled by self-driving laboratories. *Current Opinion in Chemical Engineering* **36**, 100831 (2022).
37. Kavalsky, L. *et al.* By how much can closed-loop frameworks accelerate computational materials discovery? *Digital Discovery* **2**, 1112–1125 (2023).
38. Annevelink, E. *et al.* AutoMat: Automated materials discovery for electrochemical systems. *MRS Bulletin* **47**, 1036–1044 (2022).

39. Liang, Q. *et al.* Benchmarking the performance of Bayesian optimization across multiple experimental materials science domains. *npj Computational Materials* **7**, 188 (2021).
40. Matsuda, S., Lambard, G. & Sodeyama, K. Data-driven automated robotic experiments accelerate discovery of multi-component electrolyte for rechargeable Li--O₂ batteries. *Cell Reports Physical Science* **3**, (2022).
41. Koscher, B. *et al.* Autonomous, multi-property-driven molecular discovery: from predictions to measurements and back. (2023).
<https://chemrxiv.org/engage/chemrxiv/article-details/6435f8c5a41dec1a56e64577>
42. Snapp, K. L. *et al.* Autonomous Discovery of Tough Structures. (2023). *arXiv*
<https://arxiv.org/abs/2308.02315>
43. Hippalgaonkar, K. *et al.* Knowledge-integrated machine learning for materials: lessons from gameplaying and robotics. *Nature Reviews. Materials* **8**, 241–260 (2023).
44. Kanarik, K. J. *et al.* Human--machine collaboration for improving semiconductor process development. *Nature* **616**, 707–711 (2023).
45. Frazier, P. I. A tutorial on bayesian optimization. *arXiv* 1–22 (2018).
46. Erps, T. *et al.* Accelerated discovery of 3D printing materials using data-driven multiobjective optimization. *Science Advances* **7**, eabf7435 (2021).
47. Krishnadasan, S., Brown, R. J. C., deMello, A. J. & deMello, J. C. Intelligent routes to the controlled synthesis of nanoparticles. *Lab on a Chip*. **7**, (2007).
48. Kim, I. Y. & de Weck, O. L. Adaptive weighted sum method for multiobjective optimization: a new method for Pareto front generation. *Structural and Multidisciplinary Optimization* **31**, 105–116 (2006).
49. Iwasaki, Y., Jaekyun, H., Sakuraba, Y., Kotsugi, M. & Igarashi, Y. Efficient autonomous material search method combining ab initio calculations, autoencoder, and multi-objective Bayesian optimization. *Science and Technology of Advanced Materials: Methods* **2**, 365–371 (2022).
50. Williams, C. K. I. & Rasmussen, C. E. *Gaussian processes for machine learning*. vol. 2 (MIT press Cambridge, MA, 2006).
51. Noack, M. M. *et al.* Autonomous materials discovery driven by Gaussian process regression with inhomogeneous measurement noise and anisotropic kernels. *Scientific Reports* **10**, 17663 (2020).

52. Noack, M. M. & Reyes, K. G. Mathematical nuances of Gaussian process-driven autonomous experimentation. *MRS Bulletin* **48**, 153–163 (2023).
53. Noack, M. M., Krishnan, H., Risser, M. D. & Reyes, K. G. Exact Gaussian processes for massive datasets via non-stationary sparsity-discovering kernels. *Scientific Reports* **13**, 3155 (2023).
54. Xu, Y., Farris, C. W., Anderson, S. W., Zhang, X. & Brown, K. A. Bayesian Reconstruction of Magnetic Resonance Images using Gaussian Processes. *arXiv preprint arXiv:2303.13700* (2023).
55. Liu, Y., Ziatdinov, M. & Kalinin, S. V. Exploring causal physical mechanisms via non-gaussian linear models and deep kernel learning: applications for ferroelectric domain structures. *ACS Nano* **16**, 1250–1259 (2021).
56. Valletti, M., Ziatdinov, M. A., Kalinin, S. V & others. Deep Kernel Methods Learn Better: From Cards to Process Optimization. *arXiv preprint arXiv:2303.14554* (2023).
57. Markovitch, S. & Rosenstein, D. Feature Generation Using General Constructor Functions. *Machine Learning* **49**, 59–98 (2002).
58. Liu, Z. *et al.* Machine learning with knowledge constraints for process optimization of open-air perovskite solar cell manufacturing. *Joule* **6**, 834–849 (2022).
59. Häse, F., Aldeghi, M., Hickman, R. J., Roch, L. M. & Aspuru-Guzik, A. Gryffin: An algorithm for Bayesian optimization of categorical variables informed by expert knowledge. *Applied Physics Reviews* **8**, (2021).
60. Honarmandi, P., Attari, V. & Arroyave, R. Accelerated materials design using batch Bayesian optimization: A case study for solving the inverse problem from materials microstructure to process specification. *Computational Materials Science* **210**, 111417 (2022).
61. Chevalier, C. & Ginsbourger, D. Fast computation of the multi-points expected improvement with applications in batch selection. In *International conference on learning and intelligent optimization* 59–69 (2013).
62. Tamura, R. *et al.* Machine-Learning-Based phase diagram construction for high-throughput batch experiments. *Science and Technology of Advanced Materials: Methods* **2**, 153–161 (2022).
63. González, L. D. & Zavala, V. M. New paradigms for exploiting parallel experiments in Bayesian optimization. *Computers & Chemical Engineering* **170**,

- 108110 (2023).
64. Graff, D. E. *et al.* Self-focusing virtual screening with active design space pruning. *Journal of Chemical Information and Modeling* **62**, 3854–3862 (2022).
 65. Li, K. *et al.* On the redundancy in large material datasets: efficient and robust learning with less data. (2023). *arXiv preprint*, <https://arxiv.org/abs/2304.13076>
 66. Siemenn, A. E., Ren, Z., Li, Q. & Buonassisi, T. Fast Bayesian Optimization of Needle-in-a-Haystack Problems using Zooming Memory-Based Initialization. *npj Computational Materials* **9**, 1–17 (2022).
 67. Eriksson, D., Pearce, M., Gardner, J. R., Turner, R. & Poloczek, M. Scalable global optimization via local Bayesian optimization. *Advances in Neural Information Processing Systems* **32**, (2019).
 68. Rupnow, C. C. *et al.* A self-driving laboratory optimizes a scalable process for making functional coatings. *Cell Reports Physical Science* **4**, 101411 (2023).
 69. Snapp, K. L. *et al.* Superlative mechanical energy absorbing efficiency discovered through self-driving lab-human partnership. *Nature Communications* **15**, 4290 (2024).
 70. Mela, K. & Heinisuo, M. Weight and cost optimization of welded high strength steel beams. *Engineering Structures* **79**, 354–364 (2014).
 71. Kravanja, S. & Šilih, S. Optimization based comparison between composite I beams and composite trusses. *Journal of Constructional Steel Research* **59**, 609–625 (2003).
 72. Zhang, X., Yu, Z., Yu, Y., Wang, H. & Li, J. Axial compressive behavior of Moso Bamboo and its components with respect to fiber-reinforced composite structure. *Journal of Forestry Research* **30**, 2371–2377 (2019).
 73. Ahmad, M. & Kamke, F. A. Analysis of Calcutta bamboo for structural composite materials: Physical and mechanical properties. *Wood Science and Technology* **39**, 448–459 (2005).
 74. Ba, Z. *et al.* Robust flexural performance and fracture behavior of TiO₂ decorated densified bamboo as sustainable structural materials. *Nature Communications* **14**, 1–12 (2023).
 75. Shao, Z. & Wang, F. *The fracture mechanics of plant materials*. (Springer Singapore, 2018).

76. Kooistra, G. W., Deshpande, V. S. & Wadley, H. N. G. Compressive behavior of age hardenable tetrahedral lattice truss structures made from aluminium. *Acta Materialia* **52**, 4229–4237 (2004).
77. Ashby, M. F. *et al.* *Metal foams: a design guide*. (Elsevier, 2000).
78. Yuan, S., Chua, C. K. & Zhou, K. 3D-printed mechanical metamaterials with high energy absorption. *Advanced Materials Technologies* vol. 4 (2019).
79. Evans, A. G. *et al.* Concepts for enhanced energy absorption using hollow micro-lattices. *International Journal of Impact Engineering* **37**, 947–959 (2010).
80. Caswell, S. V., Gould, T. E. & Wiggins, J. S. Protective helmets in sports. *Materials in Sports Equipment* 87–126 (2007) doi:10.1533/9781845693664.1.87.
81. Cheng, H. *et al.* Mechanical metamaterials made of freestanding quasi-BCC nanolattices of gold and copper with ultra-high energy absorption capacity. *Nature Communications* **14**, 1–7 (2023).
82. Quanjin, M. A. *et al.* The energy-absorbing characteristics of filament wound hybrid carbon fiber-reinforced plastic/polylactic acid tubes with different infill pattern structures. *Journal of Reinforced Plastics and Composites* **38**, 1067–1088 (2019).
83. Gu, G. X., Takaffoli, M. & Buehler, M. J. Hierarchically enhanced impact resistance of bioinspired composites. *Advanced Materials* **29**, 1–7 (2017).
84. Wegst, U. G. K. & Ashby, M. F. The mechanical efficiency of natural materials. *Philosophical Magazine* **84**, 2167–2186 (2004).
85. Ingrole, A., Aguirre, T. G., Fuller, L. & Donahue, S. W. Bioinspired energy absorbing material designs using additive manufacturing. *Journal of the Mechanical Behavior of Biomedical Materials* **119**, 104518 (2021).
86. Gibson, L. J. Biomechanics of cellular solids. *Journal of Biomechanics* **38**, 377–399 (2005).
87. Bohara, R. P., Linforth, S., Nguyen, T., Ghazlan, A. & Ngo, T. Novel lightweight high-energy absorbing auxetic structures guided by topology optimisation. *International Journal of Mechanical Sciences* **211**, 106793 (2021).
88. Ma, C. *et al.* Accelerated design and characterization of non-uniform cellular materials via a machine-learning based framework. *npj Computational Materials* **6**, (2020).

89. Helou, M., Vongbunyong, S. & Kara, S. Finite element analysis and validation of cellular structures. *Procedia CIRP* **50**, 94–99 (2016).
90. Habib, F. N., Iovenitti, P., Masood, S. H. & Nikzad, M. Cell geometry effect on in-plane energy absorption of periodic honeycomb structures. *International Journal of Advanced Manufacturing Technology* **94**, 2369–2380 (2018).
91. Gu, G. X., Chen, C. T. & Buehler, M. J. De novo composite design based on machine learning algorithm. *Extreme Mechanics Letters* **18**, 19–28 (2018).
92. Tancogne-Dejean, T. & Mohr, D. Stiffness and specific energy absorption of additively-manufactured metallic BCC metamaterials composed of tapered beams. *International Journal of Mechanical Sciences* **141**, 101–116 (2018).
93. Mo, C., Perdikaris, P. & Raney, J. R. Accelerated design of architected materials with multifidelity Bayesian optimization. *Journal of Engineering Mechanics* **149**, 1–10 (2023).
94. Ge, C., Priyadarshini, L., Cormier, D., Pan, L. & Tuber, J. A preliminary study of cushion properties of a 3D printed thermoplastic polyurethane Kelvin foam. *Packaging Technology and Science* **31**, 361–368 (2018).
95. Abueidda, D. W. *et al.* Mechanical properties of 3D printed polymeric Gyroid cellular structures: Experimental and finite element study. *Materials and Design* **165**, 107597 (2019).
96. Ling, C., Cernicchi, A., Gilchrist, M. D. & Cardiff, P. Mechanical behaviour of additively-manufactured polymeric octet-truss lattice structures under quasi-static and dynamic compressive loading. *Materials and Design* **162**, 106–118 (2019).
97. Zhu, Z., Ng, D. W. H., Park, H. S. & McAlpine, M. C. 3D-printed multifunctional materials enabled by artificial-intelligence-assisted fabrication technologies. *Nature Reviews. Materials* **6**, 27–47 (2021).
98. Li, L., Zhang, G. & Khandelwal, K. Topology optimization of energy absorbing structures with maximum damage constraint. *International Journal for Numerical Methods in Engineering* **112**, 737–775 (2017).
99. Pizhong, Q., Mijia, Y. & Florin, B. Impact Mechanics and High-Energy Absorbing Materials: Review. *Journal of Aerospace Engineering* **21**, 235–248 (2008).
100. Zhu, F., Lu, G., Ruan, D. & Wang, Z. Plastic Deformation, Failure and Energy Absorption of Sandwich Structures with Metallic Cellular Cores. *International Journal of Protective Structures* **1**, 507–541 (2010).

101. Li, Q. M., Magkiriadis, I. & Harrigan, J. J. Compressive strain at the onset of densification of cellular solids. *Journal of Cellular Plastics* **42**, 371–392 (2006).
102. Vural, M. & Ravichandran, G. Microstructural aspects and modeling of failure in naturally occurring porous composites. *Mechanics of Materials* **35**, 523–536 (2003).
103. Häse, F., Roch, L. M. & Aspuru-Guzik, A. Next-generation experimentation with self-driving laboratories. *Trends in Chemistry* **1**, 282–291 (2019).
104. Coley, C. W., Eyke, N. S. & Jensen, K. F. Autonomous discovery in the chemical sciences part I: Progress. *Angewandte Chemie - International Edition* **59**, 22858–22893 (2020).
105. Noack, M. M. *et al.* A kriging-based approach to autonomous experimentation with applications to x-ray scattering. *Scientific Reports* **9**, 1–19 (2019).
106. Shields, B. J. *et al.* Bayesian reaction optimization as a tool for chemical synthesis. *Nature* **590**, 89–96 (2021).
107. Nakayama, R. *et al.* Tuning of Bayesian optimization for materials synthesis: simulation of the one-dimensional case. *Science and Technology of Advanced Materials: Methods* **2**, 119–128 (2022).
108. Wang, X., Jin, Y., Schmitt, S. & Olhofer, M. Recent Advances in Bayesian Optimization. *ACM Computing Surveys* **55**, (2023).
109. Jin, H., Zhang, E. & Espinosa, H. D. Recent Advances and Applications of Machine Learning in Experimental Solid Mechanics: A Review. *Applied Mechanics Reviews* **75**(6), 061001(2023). <https://doi.org/10.1115/1.4062966>
110. Allan, T. Experimental and analytical investigation of the behaviour of cylindrical tubes subject to axial compressive forces. *Journal of Mechanical Engineering Science* **10**, 182–197 (1968).
111. Wang, J. & Sadowski, A. J. Elastic imperfect cylindrical shells of varying length under combined axial compression and bending. *Journal of Structural Engineering* **146**, 04020014 (2020).
112. Lundberg, S. M. & Lee, S. I. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems* **30** (2017). https://papers.nips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf

113. Gongora, A. E. *et al.* Designing lattices for impact protection using transfer learning. *Matter* **5**, 2829–2846 (2022).
114. Astudillo, R. & Frazier, P. I. Bayesian Optimization of Composite Functions. in *Proceedings of the 36th International Conference on Machine Learning* (eds. Chaudhuri, K. & Salakhutdinov, R.) vol. 97 354–363 (PMLR, 2019).
115. Liu, Y., Ziatdinov, M., Vasudevan, R. & Kalinin, S. V. Post-Experiment Forensics and Human-in-the-Loop Interventions in Explainable Autonomous Scanning Probe Microscopy. arXiv:2302.06577 (2023).
116. Overvelde, J. T. B., Shan, S. & Bertoldi, K. Compaction through buckling in 2D periodic, soft and porous structures: Effect of pore shape. *Advanced Materials* **24**, 2337–2342 (2012).
117. Noack, M. & Zwart, P. Computational Strategies to Increase Efficiency of Gaussian-Process-Driven Autonomous Experiments. in *2019 IEEE/ACM 1st Annual Workshop on Large-scale Experiment-in-the-Loop Computing (XLOOP)* 1–7 (IEEE, 2019).
118. Szegedy, C. *et al.* Going deeper with convolutions. in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 1–9 (2015). https://www.cv-foundation.org/openaccess/content_cvpr_2015/ext/1A_001_ext.pdf
119. Snapp, K. L., Gongora, A. E. & Brown, K. A. Increasing throughput in fused deposition modeling by modulating bed temperature. *Journal of Manufacturing Science and Engineering, Transactions of the ASME* **143**, 094502 (2021).
120. Nair, V. & Hinton, G. E. Rectified linear units improve restricted boltzmann machines. in *Proceedings of the 27th International Conference on International Conference on Machine Learning* 807–814 (Omnipress, 2010).
121. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. in *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings* (arXiv, 2015).
122. Clough, E. C. *et al.* Elastomeric Microlattice Impact Attenuators. *Matter* **1**, 1519–1531 (2019).
123. Ge, C., Cormier, D. & Rice, B. Damping and cushioning characteristics of Polyjet 3D printed photopolymer with Kelvin model. *Journal of Cellular Plastics* **57**, 517–534 (2021).

124. Thompson, M. K. *et al.* Design for Additive Manufacturing: Trends, opportunities, considerations, and constraints. *CIRP Annals - Manufacturing Technology* **65**, 737–760 (2016).
125. Tomin, M., Lengyel, M. Á., Párizs, R. D. & Kmetty, Á. Measuring and mathematical modeling of cushion curves for polymeric foams. *Polymer Testing* **117**, (2023).
126. Gibert, J. M. & Batt, G. S. Impact oscillator model for the prediction of dynamic cushion curves of open cell foams. *Packaging Technology and Science* **28**, 227–239 (2015).
127. Li, G., Rouillard, V. & Sek, M. A. Evaluation of static and dynamic cushioning properties of polyethylene foam for determining its cushion curves. *Packaging Technology and Science* **28**, 47–57 (2015).
128. Ge, C. Theory and practice of cushion curve: A supplementary discussion. *Packaging Technology and Science* **32**, 185–197 (2019).
129. Deshpande, V. S. & Fleck, N. A. High strain rate compressive behaviour of aluminum alloy foams. *International Journal of Impact Engineering* **24**, 277–298 (2000).
130. Ling, C., Cernicchi, A., Gilchrist, M. D. & Cardiff, P. Mechanical behaviour of additively-manufactured polymeric octet-truss lattice structures under quasi-static and dynamic compressive loading. *Materials and Design* **162**, 106–118 (2019).
131. Messner, M. C. Optimal lattice-structured materials. *Journal of the Mechanics and Physics of Solids* **96**, 162–183 (2016).
132. Hawreliak, J. A. *et al.* Dynamic Behavior of Engineered Lattice Materials. *Scientific Reports* **6**, 1–7 (2016).
133. Andrew, J. J., Ubaid, J., Hafeez, F., Schiffer, A. & Kumar, S. Impact performance enhancement of honeycombs through additive manufacturing-enabled geometrical tailoring. *International Journal of Impact Engineering* **134**, 103360 (2019).
134. Al-Saedi, D. S. J., Masood, S. H., Faizan-Ur-Rab, M., Alomarah, A. & Ponnusamy, P. Mechanical properties and energy absorption capability of functionally graded F2BCC lattice fabricated by SLM. *Materials and Design* **144**, 32–44 (2018).
135. Ansorge, T. & Nendel, K. Calculation of cushion diagrams using a physical model. *Packaging Technology and Science* **11**, 1–9 (1998).

136. Gruenbaum, G. & Miltz, J. Static versus dynamic evaluation of cushioning properties of plastic foams. *Journal of Applied Polymer Science* **28**, 135–143 (1983).
137. Qi, H. J. & Boyce, M. C. Stress-strain behavior of thermoplastic polyurethanes. *Mechanics of Materials* **37**, 817–839 (2005).
138. Miao, Y., He, H. & Li, Z. Strain hardening behaviors and mechanisms of polyurethane under various strain rate loading. *Polymer Engineering & Science* **60**, 1083–1092 (2020).
139. Sek, M. A., Minett, M., Rouillard, V. & Bruscella, B. A new method for the determination of cushion curves. *Packaging Technology and Science* **13**, 249–255 (2000).
140. Ramon, O. & Miltz, J. Prediction of dynamic properties of plastic foams from constant-strain rate measurements. *Journal of Applied Polymer Science* **40**, 1683–1692 (1990).
141. Joodaky, A., Batt, G. S. & Gibert, J. M. Prediction of cushion curves of polymer foams using a nonlinear distributed parameter model. *Packaging Technology and Science* **33**, 3–14 (2020).
142. Silverman, S. bu-shapelab/gcs: v1.2.3. <https://doi.org/10.5281/zenodo.10933597> (2024).
143. Chiu-Webster, S. & Lister, J. R. The fall of a viscous thread onto a moving surface: A ‘fluid-mechanical sewing machine’. *Journal of Fluid Mechanics* **569**, 89–111 (2006).
144. Morris, S. W., Dawes, J. H. P., Ribe, N. M. & Lister, J. R. Meandering instability of a viscous thread. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* **77**, 1–11 (2008).
145. Brun, P. T., Ribe, N. M. & Audoly, B. A numerical investigation of the fluid mechanical sewing machine. *Physics of Fluids* **24**, (2012).
146. Brun, P. T., Audoly, B., Ribe, N. M., Eaves, T. S. & Lister, J. R. Liquid ropes: A geometrical model for thin viscous jet instabilities. *Physical Review Letters* **114**, 1–5 (2015).
147. Lipton, J. I. & Lipson, H. 3D Printing Variable Stiffness Foams Using Viscous Thread Instability. *Scientific Reports* **6**, 2–7 (2016).

148. Emery, B. & Revier, D. Applied Viscous Thread Instability for Manufacturing 3D Printed Foams. *Proceedings - SCF 2022 - 7th Annual ACM Symposium on Computational Fabrication* 3–4 (2022) doi:10.1145/3559400.3565596.
149. Roch, L. M. *et al.* ChemOS: An orchestration software to democratize autonomous discovery. *PLoS ONE* **15**, 1–18 (2020).
150. Wilkinson, M. D. *et al.* The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data* **3**, 160018 (2016).

CURRICULUM VITAE

