Diversity-Based Topology Optimization of Soft Robotic Grippers

Pinskier, Josh; Wang, Xing; Liow, Lois; Xie, Yue; Kumar, Prabhat; Langelaar, Matthijs; Howard, David

DOI
10.1002/aisy.202300505

Publication date
2024

Document Version
Final published version

Published in
Advanced Intelligent Systems

ADVANCED INTELLIGENT SYSTEMS
Open Access
www.advintellsyst.com

Diversity-Based Topology Optimization of Soft Robotic Grippers

Josh Pinskier, Xing Wang, Lois Liow, Yue Xie, Prabhat Kumar, Matthijs Langelaar, and David Howard*

Soft grippers are ideal for grasping delicate, deformable objects with complex geometries. Universal soft grippers have proven effective for grasping common objects, however complex objects or environments require bespoke gripper designs. Multi-material printing presents a vast design-space which, when coupled with an expressive computational design algorithm, can produce numerous, novel, high-performance soft grippers. Finding high-performing designs in challenging design spaces requires tools that combine rapid iteration, simulation accuracy, and fine-grained optimization across a range of gripper designs to maximize performance, no current tools meet all these criteria. Herein, a diversity-based soft gripper design framework combining generative design and topology optimization (TO) are presented. Compositional pattern-producing networks (CPPNs) seed a diverse set of initial material distributions for the fine-grained TO. Focusing on vacuum-driven multi-material soft grippers, several grasping modes (e.g. pinching, scooping) emerging without explicit prompting are demonstrated. Extensive automated experimentation with printed multi-material grippers confirms optimized candidates exceed the grasp strength of comparable commercial designs. Grip strength, durability, and robustness is evaluated across 15,170 grasps. The combination of fine-grained generative design, diversity-based design processes, high-fidelity simulation, and automated experimental evaluation represents a new paradigm for bespoke soft gripper design which is generalizable across numerous design domains, tasks, and environments.

1. Introduction

Soft gripping has rapidly become the *de facto* approach for grasping delicate, compressible, and geometrically diverse objects across a range of valuable target domains, including the robotic harvesting, manufacturing, and medical industries. Several classes of soft gripper including pneumatic soft fingers (Pneunets),[1] vacuum-driven universal jamming grippers,[2,3] and passive soft grippers such as Fin Rays[4] have been extensively investigated, with several designs now commercially available. Regardless of their actuation mechanism, the aforementioned approaches are united in leveraging flexible materials to conform to an object's surface without applying large forces,[5] and as such many robotics researchers aspire to create universal end-effectors based on this concept. Such "swiss army knife" grippers that offer high performance across a broad range of applications are an attractive proposition, however recent research points to limitations with this approach, for example, showing that "Universal" jamming grippers are in fact highly amenable to optimization of both membrane geometry[6] and constituent granular materials,[7] achieving significant performance gains over standard grippers of the same type. The same is true of Pneunets, which despite being widely used as universal soft grippers, exhibit significant performance gains when both the material and geometry are optimised for their task.[8]

Bespoke gripper morphology tuned to unique features of the task and environment allows for a level of "behavioural niching" and heightened performance. More specialized soft gripper designs have therefore emerged in recent years (e.g.,[9–11]). Despite promising results, a reliance on designer intuition and manual experimental evaluation means that these design

J. Pinskier, X. Wang, L. Liow, D. Howard
CSIRO
1 Technology Ct, Pullenvale 4069, QLD, Australia
E-mail: josh.pinskier@csiro.au

The ORCID identification number(s) for the author(s) of this article can be found under https://doi.org/10.1002/aisy.202300505.

DOI: 10.1002/aisy.202300505

Y. Xie
Bio-Inspired Robotics Lab
University of Cambridge
Trinity Lane, Cambridge CB2 1TA, UK

P. Kumar
Department of Mechanical and Aerospace Engineering
Indian Institute of Technology Hyderabad
Kandi 502284, Telangana, India

M. Langelaar
Department of Precision and Microsystems Engineering
Delft University of Technology
Delft 2628, The Netherlands

approaches have been unable to scale, both in terms of 1) the time and effort required to create a gripper for each new application, and 2) the ability to explore increasingly expansive design spaces offered by, for example, multi-material printing, where the opportunity for bespoke performance is high, if only the space could be efficiently sampled.[12] Computational design offers a pathway towards scalable, bespoke gripping for increasingly challenging and diverse domains, which cannot be met by existing universal and anthropomorphic grippers, nor through conventional design techniques.[13,14]

The cornerstones of computational soft robotics design are 1) the design algorithm, which governs how the design space is explored, and 2) the simulator, which allows for rapid *in silico* evaluation of successive iterations of potential designs to map the design space. Two broad approaches to computational soft robotics design can be delineated according to the pairing of simulator and design algorithm used in each case. Most popularly, topology optimization (TO), coupled to a finite element analysis (FEA) model, allows fine-grained, high-resolution optimisation that accurately captures system features and kinematics (materials, actuation, etc), and generalizes to myriad applications including compliant mechanisms[15,16] and self-sensing structures.[17] However, TO requires slow, expensive solves[12,18–20] and a very structured problem. It assumes the problem is well-defined, and that a suitable design domain, loads (forces, voltages, pressure, etc) and constraints can be specified a priori, from which it generates only a single solution. A second body of work focuses on fast simulation coupled with evolutionary or reinforcement learning, permitting a much wider design exploration, but with a loss in both simulation accuracy (typically using mass-spring methods rather than FEA) and in the resolution of the final design,[21–23] for example, focusing on macro-scale voxel-based abstractions[24] which do not approach the potential for fine-grained optimization offered by TO. In short, each approach has its own benefits and drawbacks, and no soft robotics simulators exist that are both accurate and fast.[25] Interestingly, these benefits and drawbacks are frequently mutually exclusive, hinting that combining the two approaches may be a fruitful avenue of research.

Motivated by the clear need for bespoke soft grippers, and the absence of sufficiently accurate and efficient simulators for automated soft robotic design, we adopt a reality-coupled computational design approach to generate diverse sets of soft grippers. To overcome the infeasibility of capturing every feature of interest (e.g., size, cost, manufacturability, performance) and simultaneously optimise all of them, we instead aim to generate numerous valid designs with distinct features. We hypothesise that from the diverse set, high-performing designs will emerge, that are specialized for specific tasks without explicit prompting. To generate diversity, we couple TO and evolutionary design, providing a "best of both worlds" approach that permits both fine-detailed optimization and broad design space exploration. Compositional pattern producing networks (CPPNs), a compact design encoding created for evolutionary algorithms, generate patterns of initial material distribution within a given design domain; each pattern then permits a unique solution when solved using TO and allows diverse designs to be produced by a gradient-based solver. The resulting designs are then 3D printed using a multi-material Polyjet printer and evaluated for their grasp strength, robustness and durability using an automated robotic platform, allowing us to capture extensive experimental data for evaluation totalling 15,170 grasps – far more than is seen in the state of the art.

Aside from the core novelty of combining the two main computational design methods for soft robotics, our automated experimental platform is also novel in that efficiently gathers vast amounts of data on grasp quality and robustness and the design's durability, enabling an efficient two-stage design process. Rather than simulating an entire grasp and optimizing in simulation, we abstract environmental contact away from the optimization, such that a high-fidelity optimization can be performed using the finite element method (FEM) and the true grasp performance verified in an automated experimental process. Two insights guide this approach: firstly, in the presence of a sufficiently good object approximation soft grasping quality can be estimated by the grippers structure without a contact model. Decomposing grasp quality into appropriate functions of local compliance and global stiffness can predict the grippers ability to conform to a surface to ensure stable attachment, and apply sufficient normal forces to hold the object. Secondly, given the options of fast/inaccurate physics simulators, slow/restrictive FEM solvers, a multi-fidelity solution is preferable to any single solver. We demonstrate the efficient design of sets of high-performing soft robotic grippers and verify their real-world performance. We investigate multi-material vacuum-based soft grasping, and observe the automated generation of diverse grasp methods, including pinch grips and the creation of enveloping basket-like structures, all without external prompting. Examples of grippers generated using this approach are presented side-by-side in **Figure 1**.

The main contributions of this research are: 1) A framework for generative topology optimization (TO) through CPPN design

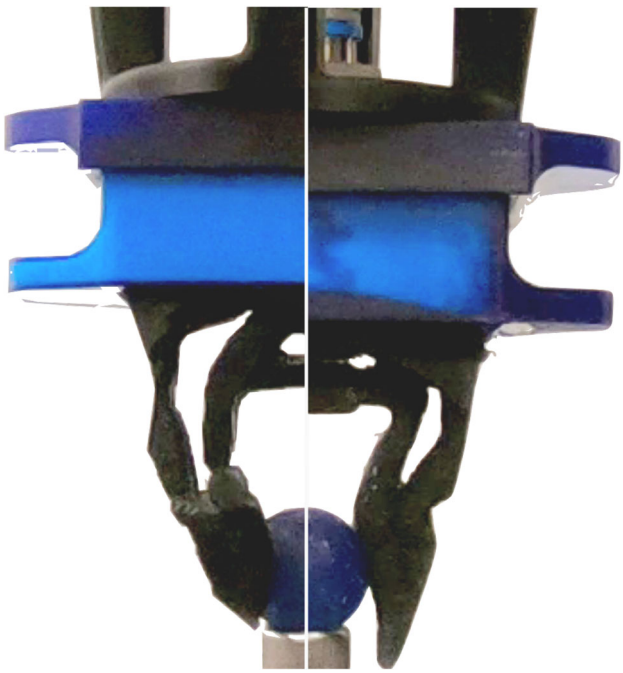

Figure 1. 50–50 split of two optimised soft gripper candidates grasping a spherical object. The two grippers found different grasp modes (envelop and pinch) which produce high-quality grasps.

seeding. 2) The establishment of a TO method for multi-material pneumetic soft robots, which is generalizable across the majority of pneumatic soft robots. 3) Demonstration of the algorithm being used to generate a range of morphologically and functionally diverse soft grippers that provide high performance across a range of test cases. 4) The development of a set of benchmarks and dataset which characterises soft grasping under ideal conditions and their robustness to environmental perturbations. This second element is rarely featured in robotic grasping literature, despite it being critical to future use.

2. Related Work

2.1. Computational Gripper Design

In the domain of rigid grasping, computational tools have been developed to optimize active and passive grippers for specific objects. Under the assumption that both the object and gripper are known rigid bodies, these find a geometry[26] or combination of geometry and path,[27,28] which allows a printed end-effector to perform grasping operations. This critical assumption is obviously not valid in soft robotics, where deformability is a key design feature.

The grasping performance of soft grippers arises from a strongly non-linear interplay between material properties (stiffness, Poisson ratio, coefficients of friction), geometry, and actuation. As a result computational soft gripper design research has focused on design exploration and optimisation within established gripper classes, such as jamming grippers[6,7] and pneumatic bending fingers.[29–31] However, several design toolkits have recently emerged which enable exploration of large design spaces using parameterised or implicit geometric descriptors.[32–34] Whilst vastly increasing tractable search space, these toolkits remain restricted to designs using hollow pneumatic chambers.

2.2. Computational Soft Robotics

A body of research exists into computational design of "soft agents" within the soft robotics, embodied intelligence and artificial life communities. These typically use inexpensive physics engines with soft material primatives to evolve or learn high performing designs for mobile "soft robots" and artificial life forms.[21,35–38] As the field has developed, evermore accurate simulators have been developed, increasing modelling fidelity, expanding environment realism and adding features. The current state of the art still lacks physical grounding, however; generated designs are either unsuited to physical manufacture or unable to cross the reality gap.[39] A recenly developed platfrom using reinforcement learning in FEM have begun to challenge this. But to date, it has only been demonstrated in learning controllers, rather than designs[40]

2.3. Soft Gripper Topology Optimisation

Compliant gripper optimization is a benchmark problem in TO and has been evaluated for more than 2 decades, producing numerous scissor-like designs.[41,42] More recently, these have extended into the soft robotics domain by with single-physics structural optimizations, which consider the design's geometry but not actuation. For example compliant (fin-ray like) structures have been optimized which deform around an object to perform grasping.[43,44] Because of their widespread use in soft robotics, pneumatic soft fingers have also been a popular optimization target. For convenience it is normally assumed that the inflating chamber is fixed,[45–47] but multi-physics formulations have also optimized both the pressure chamber shape and material layout in 2D[48,49] and 3D.[50,51]

3. Multi-Material Pressure-Driven Topology Optimisation Formulation

3.1. SIMP Formulation

In this work, we use the solid isotropic material with penalization (SIMP) TO method as the basis of our multiphysics optimization framework.[41] Topology optimization is built around FEM. As such, the problem is specified by a design domain (design geometry), loads, and boundary conditions. SIMP discretises the design domain into a set of finite elements, each which containts a continuous "pseudo-density", ρ between 0 (void space) and 1 (solid material). A penalty exponent is applied to each element to drive a final result towards a binary solution. In single-material optimisations this follows the interpolation law:

$$E_i = (1 - \overline{\rho}_i{}^p)E_{\min} + \overline{\rho}_i{}^p E_s \tag{1}$$

where E_i is elastic modulus of element i, E_s is the elastic modulus of the base material, $E_{\min}$ is the modulus of the void material and is non-zero to avoid singularity, $\overline{\rho}_i$ is the filtered (smoothed) form of the design variable ρ_i and p is the SIMP penalty exponent where $p \geq 3$. To remove non-physical checkerboard patterns and intermediate (i.e., non-binary) densities from the final design, we use standard convolutional density filtering as in.[52]

The SIMP approach generalises to multiple materials by assigning each material a unique design variable. Rather than interpolating between solid and void, the SIMP optimization interpolates between the different material options and void.[53]

For example, the two-material formulation is given by:

$$E_i = (1 - \overline{\rho}_{i1}^p)E_{\min} + \overline{\rho}_{i1}^p((1 - \overline{\rho}_{i2}^p)E_1 + \overline{\rho}_{i2}^p E_2) \tag{2}$$

where E_1 and E_2 are elastic moduli of materials 1 and 2, respectively, $\overline{\rho}_{in}$ and is the filtered density of design variable ρ_{in} for $n \in \{1, 2\}$. The design variables have a cascading effect, $\overline{\rho}_{i1}$ (the topology variable) determines the presence of a void or solid element and $\overline{\rho}_{i2}$ selects whether a solid element contains material 1 or 2. More formally, $\{\overline{\rho}_{i1} = 0, \overline{\rho}_{i2} = 0\}$ results in a void element, $\{\overline{\rho}_{i1} = 1, \overline{\rho}_{i2} = 0\}$ is filled with material 1, and $\{\overline{\rho}_{i1} = 1, \overline{\rho}_{i2} = 1\}$ with material 2.

Increasing the number of materials is achieved by adding new selector variables, each of which has the effect of switching an element between material $n - 1$ and n. For example in the 3 material case:

$$E_i = (1-\overline{\rho}_{i1}^p)E_{\min} + \overline{\rho}_{i1}^p[((1-\overline{\rho}_{i2}^p)E_1 + \overline{\rho}_{i2}^p((1-\overline{\rho}_{i3}^p)E_2 + \overline{\rho}_{i3}^p E_3))] \tag{3}$$

where E_3 is the modulus of material 3 and $\overline{\rho}_{i1} = 0$ selects between material 2 and 3, given the presence of material 2. That is: $\{\overline{\rho}_{i1} = 0, \overline{\rho}_{i2} = 0, \overline{\rho}_{i3} = 0\}$, $\{\overline{\rho}_{i1} = 1, \overline{\rho}_{i2} = 0, \overline{\rho}_{i3} = 0\}$, $\{\overline{\rho}_{i1} = 1, \overline{\rho}_{i2} = 1, \overline{\rho}_{i3} = 0\}$ and $\{\overline{\rho}_{i1} = 1, \overline{\rho}_{i2} = 1, \overline{\rho}_{i3} = 1\}$ give a void, material 1, material 2, and material 3, respectively.

3.2. Darcy Method

The Darcy method for pneumatic soft robot TO is then applied to the multi-material TO problem, it builds on our previous work into TO pressure-loaded structures.[51,54,55] The method is unique in capturing the design-dependency of pneumaticTO, in which the movement of the fluid–solid interface during the optimization also changes the loading applied to the structure. The Darcy method evaluates these coupled problems by solving two physical systems at each optimization iteration: first it uses a modified form of Darcy's law for flow of a fluid through a porous medium, to evaluate the pressure field in the design domain. Pressure diffusion is modelled as a function of each elements pseudo-density, such that it behaves like a porous media. The pressure field is then applied to the structure to drive deformation in the domain and evaluate the structures compliance. A detailed description of the Darcy formulation can be found in[54]

3.3. Cost Function

In this work we aim to automate the design of vacuum based soft grippers. Given the balance of compliance and stiffness required in soft grasping, a single optimization parameter cannot effectively capture the entire problem. Hence we investigate multi-objective formulations to identify the trade-off between parameters of interest.

The major features of the design relevant to the soft robotic grasping problem are as follows: 1) Tip displacement: the gripper must enable sufficient displacement to close around objects of different sizes. 2) Stiffness: the gripper must be able to exert sufficient normal force on the object to grasp and hold it. 3) Closure: The gripper should be airtight, preventing energy loss due to leakage. 4) Adaptability: The design should ideally enable multiple grasping points and adapt to multiple objects.

The first three features are considered directly in the cost function, whilst the fourth is implicitly addressed by enforcing the use of soft materials. That is, the cost function comprises three terms:

$$\phi_1(\rho_{1,2,3}) = -Lu \tag{4}$$

$$\phi_2(\rho_{1,2,3}) = \frac{1}{SE^{1/n}} \tag{5}$$

$$\phi_3(\rho_{1,2,3}) = (E_l - E_{ld}) \tag{6}$$

where u is the design's global displacement vector; L is a binary vector which selects relevant nodes in grasping edge; $SE = u^T K u$ is the strain energy of the mechanism under pneumatic loading, penalising SE restricts the generation of thin design components and encourages stiffer designs; n is a constant penalty exponent, which can be tuned to adjust the relative weighing of terms in the cost function; $K(\rho)$ is the global stiffness matrix; E_l is the energy lost due to air flowing across the boundary of the design domain, and E_{lt} is the target loss (calculated as in[51]). The Darcy formulation permits a small flow even in solid elements, hence $E_l > 0$ even in within a closed space.

To evaluate the significance of each of these components, we investigate three cost functions: Firstly, a formulation which considers only compliance and strain energy, without penalizing energy loss, enabling the scale of energy to be evaluated:

$$\begin{aligned}\phi_{\text{unpen}}(\rho_{1,2,3}) &= \frac{\phi_1(\rho_{1,2,3})}{\phi_2(\rho_{1,2,3})} \\ &= \frac{-Lu}{SE^{1/n}}\end{aligned} \tag{7}$$

Then two functions which penalize energy loss, one using a linear penalty:

$$\begin{aligned}\phi_{\text{lin}}(\rho_{1,2,3}) &= \frac{\phi_1(\rho_{1,2,3})}{\phi_2(\rho_{1,2,3})} + \frac{1}{B}\phi_3(\rho_{1,2,3}) \\ &= \frac{-Lu}{SE^{1/n}} + \frac{1}{B}(E_l - E_{ld})\end{aligned} \tag{8}$$

and a second using an exponential heat penalty:

$$\begin{aligned}\phi_{\text{exp}}(\rho_{1,2,3}) &= \frac{\phi_1(\rho_{1,2,3})}{\phi_2(\rho_{1,2,3})} + \frac{1}{B}\phi_3(\rho_{1,2,3}) \\ &= \frac{-Lu}{SE^{1/n}} + Ae^{\frac{1}{B}(E_l - E_{ld})}\end{aligned} \tag{9}$$

where A and B are scaling constants.

The exponential penalty will aggressively drive an initial design towards one which minimizes energy loss, hence the exponential and linear penalties are compared to assess the impact of the convergence rate on the final design. Finally, a constraint is placed on the total volume of each material, such that:

$$\sum_i \rho_{\text{in}}/n_{\text{ele}} < \rho_{n-\max} \tag{10}$$

$\rho_{n-\max}$ is the volume limit for material n and n_{ele} is the total number of elements in the design domain.

3.4. Design Domain

The design domain of the vacuum gripper is presented in **Figure 2**b. It represents one half of the final gripper, with symmetry applied to one the mid-plane to reduce computation time. A vacuum is applied at the upper face, which drives the deformation of the gripper, as measured at the bottom edge. For efficiency a geometrically linear finite element solver is used, with a small, 150 Pa pressure input.

The design domain is 100mm × 50mm × 50mm. 8 Node hexahedral elements with 2 mm sides are used, giving a total of of $n_{\text{ele}} = 31250$ elements for each material. Three materials are used in this work to, giving a total of 93750 elements. Three materials are used to spam the three orders of magnitude of elastic moduli which can be printed. The materials have moduli $E_1 = 0.46$ Mpa, $E_2 = 11.51$ Mpa, and $E_3 = 39.98$ Mpa,

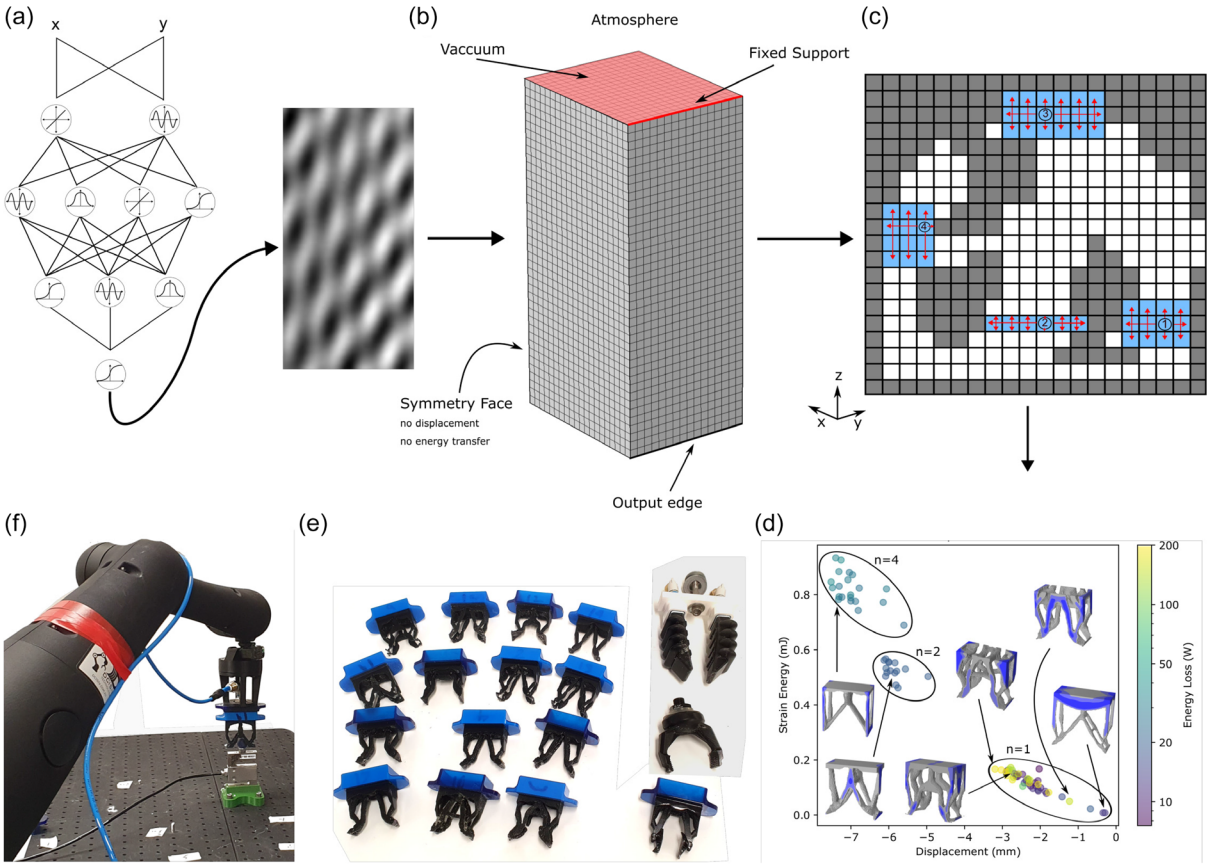

Figure 2. Method for generating diverse soft grippers and evaluating their performance. a) A CPPN with randomized weights generates an 2D pattern by evaluating the density of each pixel after a series of activation functions. The pattern sets the initial distribution for the topology optimization problem. b) The design domain for the 3 material, 2 physics (pressure/solid) topology optimization problem. c) 2D illustration of Cuboid floodfill method. A point is randomly selected in the space to be filled (white pixels), the floodfill algorithm then searches left, right, up and down until it encounters a boundary (gray pixel) in each direction. The process repeats until all pixels are filled. d) Combined results of 100 (5×20) simulations showing 3 features of interest, with material distribution shown on overlaid designs: Gray - Agilus30, light blue - Agilus85, dark blue - Agilus95 e) Complete set of experimentally evaluated optimised grippers, comprising 14 computational designs (left), the Soft Robotics 2 fingered gripper and 3D printed replica (top right), a 2 Fingered Piab gripper and 3D printed replica (middle right), and the optimised gripper with enhanced grasping surface (bottom right). f) Experimental Grasp Strength Testing Configuration, it comprises a 7DOF robot arm, load-cell and pneumatic subsystem (not shown) comprising of compressor, pressure regulator, solenoid valves, and vacuum generator.

corresponding the experimentally derived moduli of three 3D printable material blends (discussed in Section 10). In this work $\rho_{1-\max} = 0.2$, $\rho_{2-\max} = 0.1$, and $\rho_{3-\max} = 0.05$, corresponding to material budgets of 0.1 (10% of the total volume), 0.05 and 0.05, for materials 1, 2, and 3, respectively.

The derivatives of the cost function and constraints with respect to the design variables $\rho_{1,2,3}$ can be derived analytically, enabling its efficient solution through a large scale gradient-based solver. In this work we use the method of moving asymptotes.[56]

4. CPPN Initialization

TO emphasizes finding a *single* optimal designs, but in most cases finding diverse *sets* of high-quality designs is a more desirable outcome. Regardless of optimization method, some manual post-processing is required to integrate optimized components into a larger assembly. Having sets of designs rather than a single one increases the engineers freedom to select advantageous features which are not explicitly targeted in the optimization. Optimizing for diversity also forces a broader search of the design space than gradient optimizations, which overcomes local minima traps and frequently finds higher-performing candidates than a pure objective optimization.[57] To encourage diverse solutions we initialise candidates' density distribution using a compositional pattern producing network (CPPN),[58] rather than a constant initial density which is conventionally used. CPPNs are a class of neural-network which takes a pixel/voxel's coordinates as an input and passes it through several layers of weighted functions (e.g., triganometric functions, sawtooths, sigmoids, and other periodic and aperiodic functions) to generate patterns and images.

In this work we use a 2D CPPN generator to initialize the topology optimization with patterned pseudo-densities. To transform them into a 3D shape, the 2D patterns are copied along the third dimension. The network configuration is the same as in.[59] It uses four activation functions (sin, gaussian,

sigmoid, and identity) and 1–10 hidden layers to produce the patterns, where the number of hidden layers is defined by taking the floor of the number of design iterations divided by 10. That is: $n_1 = \lfloor ii/10 \rfloor + 1$, where n_1 is the number of layers and ii is the design number. Increasing the number of hidden layers typically results in more complex patterns. To create unique patterns, all weights and activation functions are randomized prior to each pattern being produced. In this work, 20 designs are evaluated for each cost function, giving 2 repeats of each CPPN configuration.

The generated patterns are exported as grey-scale images with each pixel having a value between 0 and 1, corresponding to the material pseudo-densities. Pixel values are linearly scaled such that the average density of the pattern is equal to the material limit. The pattern is then set as the initial density distribution for all three of the materials in the optimization. This process is illustrated in Figure 2a

5. CAD Geometry Generation

The resulting designs are in the form of a grid of elements with pseudo-densities in the range 0–1, which need to be processed into a final design. Applying a threshold of $\rho = 0.5$ enables the element list to be constructed into an occupancy grid for each material. That is, each material in the design domain is transformed into a 3D binary array, representing the presence or absence of material at each location.

The conversion of a 3D voxel grid to a surface mesh is a common procedure, which can be performed by meshing algorithms such as Marching cubes.[60] However the process of meshing changes the structure, as a surface of triangular elements is formed from solid voxels.

This creates a number of practical challenges when integrating the optimised component into an assembly, including: 1) the large number of triangles needed to create an accurate mesh, often orders of magnitude larger than the original number of faces. 2) flat surfaces can become uneven after meshing; parallel and orthogonal faces don't retain that relationship. 3) meshing artefacts often prevent watertight mesh formation.

For engineering purposes, a CAD file format is preferred. The voxel grid can be converted into CAD geometry either by forming a mesh from the square faces and enclosing it to make solid geometry, or by transforming the voxels into geometric primitives and performing boolean operations on them.

The latter approach is used here, as the number of number of filled voxels is typically fewer than the number of surface faces, and they are simpler to identify.

The CAD geometry is hence formed by creating cuboids from the binary voxel map and then taking the union of all the cuboids. The merging process used to take the union is computationally expensive, so to reduce the run time of the algorithm, it is desirable to find the minimum number of cuboids required to fill all the voxels without covering any unoccupied spaces. However, as overlapping voxels are permitted as these are removed during the merge operation, voxels need not be filled uniquely. That is, each voxel must be included in a cuboid, but can be in multiple.

This minimum cubes problem is NP-hard, hence a heuristic is developed which floodfills the space using only cuboids. The heuristic is illustrated in Figure 2c. It works as follows: 1) Generate a list containing all voxels to be filled, *voxels_to_fill*. 2) Initialize an empty list of identified cuboids, *min_cuboids*. 3) While *voxels_to_fill* is not empty: a) Randomly select a voxel from the list. b) Advance in the **positive *x*-direction** until a boundary voxel is reached. Boundaries are either a voxel which is not needed to be filled, or the edge of the design domain. c) Advance in the **negative *x*-direction** until a boundary voxel is reached. d) Advance in the **positive *y*-direction** until a boundary voxel is reached. The 1D line of voxels now forms a 2D rectangle, hence each voxel along the upper edge must be checked to ensure it is within bounds. e) Advance in the **negative *y*-direction** until a boundary voxel is reached. f) Advance in the **positive *z*-direction** until a boundary voxel is reached. The 2D rectangle now forms a 3D cuboid of voxels, hence each voxel along the advancing plane must be checked to ensure it is within bounds. g) Finally, advance in the **negative *z*-direction** until a boundary voxel is reached. h) Generate a new cuboid primitive from the identified vertices and append to list, *min_cuboids*. i) Remove newly filled voxels from the *voxels_to_fill* list. 4) Merge all cuboids in *min_cuboids* into a single design using Boolean "or" operation.

It is implemented using opencascade's python API, which enables the geometry to be saved as a multibody STEP file for integration into a larger assembly.

6. Optimisation Results

5 sets of experiments were undertaken in total, with 20 design optimisations in each. The 5 sets of experiments consist of one without an energy penalty; one with a linear energy penalty; and three with an exponential energy penalty and strain energy exponents $n \in \{1, 2, 4\}$ The combined results of the 5 are presented in Figure 2d, showing the performance of the optimized designs across the three features of interest: strain energy, output displacement, and energy loss. The results approximate a pareto front as displacement and strain energy are conflicting objectives. The ideal design would occupy the lower left corner of the plot (low strain energy, high displacement), but are clustered at the top-left (high strain energy, high displacement) and the bottom-right (low strain energy, low displacement). As is discussed below, there is significant (and desirable) variability in the morphology of designs within each set. Notwithstanding this, a few key features are noticeable by clustering the data: 1) The cost function parameter n sets the relative contribution of objectives; if n were allowed to occupy a continuous distribution, rather than integer values, we would likely generate a smooth front, enabling a design to be found anywhere along the curve. 2) Morphological similar designs emerge throughout the curve with thinner features in higher displacement designs and thicker ones in lower strain energy designs. 3) The energy penalty constraint does not materially impact the SE or displacement objectives. Within the $n = 1$ points, there are low and high energy loss designs clustered closely together. 4) Very low SE is undesirable, as it leads to impractically thin elements with almost no displacement. The lowest SE designs, have a semi-circular top section which is optimised to minimise compliance.

The results of each of the 5 sets of experiments are outlined in the remainder of this section.

6.1. Linear Pressure Penalty

A set of 20 designs are evaluated using a linear penalty on energy loss (Equation (8)) and the strain energy penalty $n = 1$. The results are presented in **Table 1**, for each design it shows the initial CPPN pattern, optimized topology, and reconstructed design,

Table 1. Reconstructed linear heat results. Cost = ϕ, Output displacement=D, strain energy = SE, pressure Loss = PL. Material colours: gray - Agilus30, light blue - Agilus85, dark blue - Agilus 95. Best results indicated in bold.

Iter.	Cost	Initial pattern	Optimized topology	Reconstructed geometry
0	Failed to converge		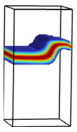	
1	ϕ: −90.32 SE: 0.099 mJ D: −1.99 mm EL: 10.15 W			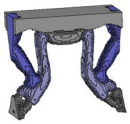
2	ϕ: −80.09 SE: 0.151 mJ D: −2.69 mm EL: 8.86 W		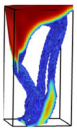	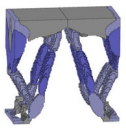
3	ϕ: −83.89 SE: 0.125 mJ D: −2.38 mm EL: 11.26 W		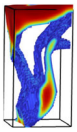	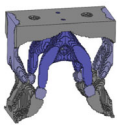
4	ϕ: −73.37 SE: 0.124 mJ D: −2.06 mm EL: 9.57 W		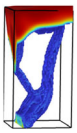	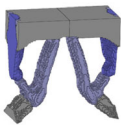
5	ϕ: −106.45 **SE:** 0.087 mJ D: −2.05 mm EL: 10.98 W		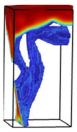	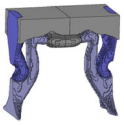
6	Failed to Converge		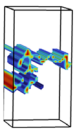	
7	Failed to Converge		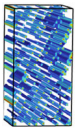	
8	ϕ: −74.33 SE: 0.148 mJ D: −2.47 mm EL: 9.09 W		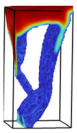	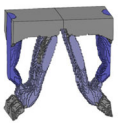

Table 1. Continued.

Iter.	Cost	Initial pattern	Optimized topology	Reconstructed geometry
9	ϕ: −89.30 SE: 0.106 mJ D: −2.12 mm EL: 10.72 W		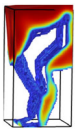	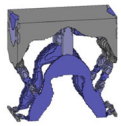
10	ϕ: −83.91 SE: 0.140 mJ D: −2.60 mm EL: 9.08 W		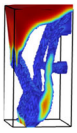	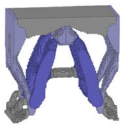
11	Failed to converge		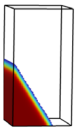	
12	ϕ: −97.50 SE: 0.105 mJ D: −2.25 mm EL: 9.40 W		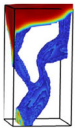	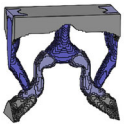
13	ϕ: −79.02 SE: 0.147 mJ D: −2.63 mm EL: 10.34 W		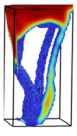	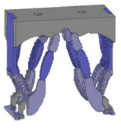
14	ϕ: −80.83 SE: 0.161 mJ **D:** −2.87 mm EL: 8.24 W		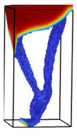	
15	Failed to converge			
16	Failed to converge			
17	ϕ: −82.25 SE: 0.117 mJ D: −2.14 mm EL: 9.13 W			
18	ϕ: −84.98 SE: 0.140 mJ D: −2.63 mm EL: 8.99 W			
19	ϕ: −52.96 SE: 0.168 mJ D: −2.03 mm **EL:** 7.57 W			

along with each component of the cost function. The optimized topology is the density field of ρ_1 where red gives $\rho_1 = 1$, blue is $\rho_1 = 0.5$, and $\rho_1 < 0.5$ is removed. The reconstructed geometry shows the three materials after thresholding. Gray is the the softest material, Agilus30 ($E = 0.46$ Mpa), light blue is the intermediate Agilus85 ($E = 11.51$ Mpa) and dark blue is the stiffest Agilus95 ($E = 39.98$ Mpa)

Of the 20 optimization runs, 14 succeeded in generating a final design, whilst 6 failed to find to converge and returned the initial topology. This occurred where the starting pattern had a large concentration of black pixels in the CPPN pattern.

The 14 designs produces final costs in the range of −52.96 to −106.45; Design #5 gave the lowest cost function overall. Further, the designs had displacements ranging from −1.99 to −2.87mm, strain energies 0.087 to 0.168 mJ and energy losses of 7.57 to 11.26 W. However, in all cases, an airtight design formed in the final reconstructed design, as the soft Agilus30 material formed a contiguous membrane at the vacuum interface.

Whist it was expected that high-performing designs may have hollow fingers reminiscent of pneunets, in all cases the optimizer converged to designs with a flat membrane. The 14 designs all present a variant on two-fingered pincers, however they exhibit a significant degree of variability. For example in designs #1 and #5, all of the Agilus95 (the stiffest material) is located on the extreme left of the optimized topology. The Agilus95 forms a relatively rigid support, which the softer Agilus85 can rotate around. Apart from sealing the vacuum chamber, the soft Agilus30 also formed a cross-bar linkage at the top of the design. This restricts the designs' strain energy under vacuum, giving these designs very lowest SE, with #5 best overall in this feature. In practice, this should increases grasp force by supporting the link in contact with the object (See Section 6.3 for experimental images of grasp behaviour).

Designs #4, #8, and #17 have broadly similar morphologies as #1 and #5 but without this upper cross-bar. Instead they have large sections of Agilus95 on their left side, which is joined directly to the input (top) face by an Agilus85 section. In #17 this Agilus85 section is a single, solid bar. However in #4 it is two smaller, parallel bars and in #8 it is 3 parallel bars.

Design #12 replaces the soft crossbar with a rigid Agilus95 linkage to connect the grasping arms to the input face, giving its opening a unique bell shape. Having rigid sections connecting the grasping arm to both the fixed side and centre of the vacuum face gives a small SE in the design without limiting displacement. However it narrows the gripper opening and reduces the size of graspable objects.

Most of the remaining designs invert the "rigid-on-the-outside" structure discussed above and instead use soft material to join a relatively rigid grasping arm to the fixed face. This facilitates large displacements at the expense of stiffness. Design #14 exemplifies this, it has a narrow, compliant hinge made from Agilus85 which joins the stiff Agilus95 gripping arms to the fixed side, and central Agilus85 arms connected the gripping arms to the fixed face. As a result it has the largest displacement of the optimised designs. #3 Uses a rigid 4-armed central core to force rotation of a soft "jaw-like" section of Agilus85, rather than the whole arm.

Aside from the connection to fixed face, we see considerable variation in the number and orientation of linkages in the designs. The simplest designs have just a few beams, whilst the most complex have several beams along the front and rear faces, with numerous cross-linkages (#2, #3, #9, #10, #13, #18, #19).

6.2. No Energy Penalty

For comparison, the 18 coverged designs resulting from 20 runs of the cost function ϕ_{unpen} (Equation (10)) are presented in **Figure 3**. Note the grippers here show simulated material distribution, rather than the final reconstructed design, giving a smoother appearance than the reconstructed designs in Table 1. Without an explicit energy penalty the simulation places no limits on airflow through the vacuum and the solver is able to create leaky (non-airtight) designs. These create a pressure gradient throughout the design domain, rather than a discrete boundary, which applies forces directly onto the grasping limbs. However, its effect on performance is relatively small as the loss of pressure prevents large forces being transferred to the design. Whilst the resulting designs are have similar features to the previous results, they contain noticeably thicker linkages. This is a result of the fixed material budget used in the optimisations; instead of Agilus30 forming a membrane at the vacuum interface, it is distributed throughout the design domain to reinforce linkages. However, the relatively soft material (an order of magnitude lower elastic modulus than Agilus95) only marginally increases the stiffness of these sections.

6.3. Exponential Energy Penalty

To assess the impact of convergence rate on the final designs, an exponential penalty on energy loss is also evaluated. As the penalty term is initially very large, the optimizer is driven to seal the design domain within a few iterations. In most cases this convergence rate does not significantly influence gripper morphology. Most of the designs produced with $n = 1$ have visibly similar structures to those produced with a linear heat penalty. However, in a few cases the exponential penalty generated designs with semicircular upper sections, rather than flat membranes. These designs have strong similarities to those used to minimize compliance in pressure loaded structures. The semicircular section is very stiff and results in very low SE in the actuated design at the expense of displacement. Clearly these extreme restriction of SE is undesirable and do not result in a usable soft gripper.

Reducing the penalty on SE (increasing n) removes this effect and results in more compliant, higher displacement designs. At $n = 2$, a set of relatively homogeneous designs emerge with a triangular opening and 3 or 4 arms on each side of the gripper. However, at $n = 4$ the penalty is insufficient to generate contiguous, sealed designs. In several cases, regions of low pseudo-density increase simulated displacement, but are removed during discretisation, leaving disconnected linkages and unsealed designs. However, a number of high-displacement designs emerge which would be beneficial for grasping lightweight objects with complex geometries.

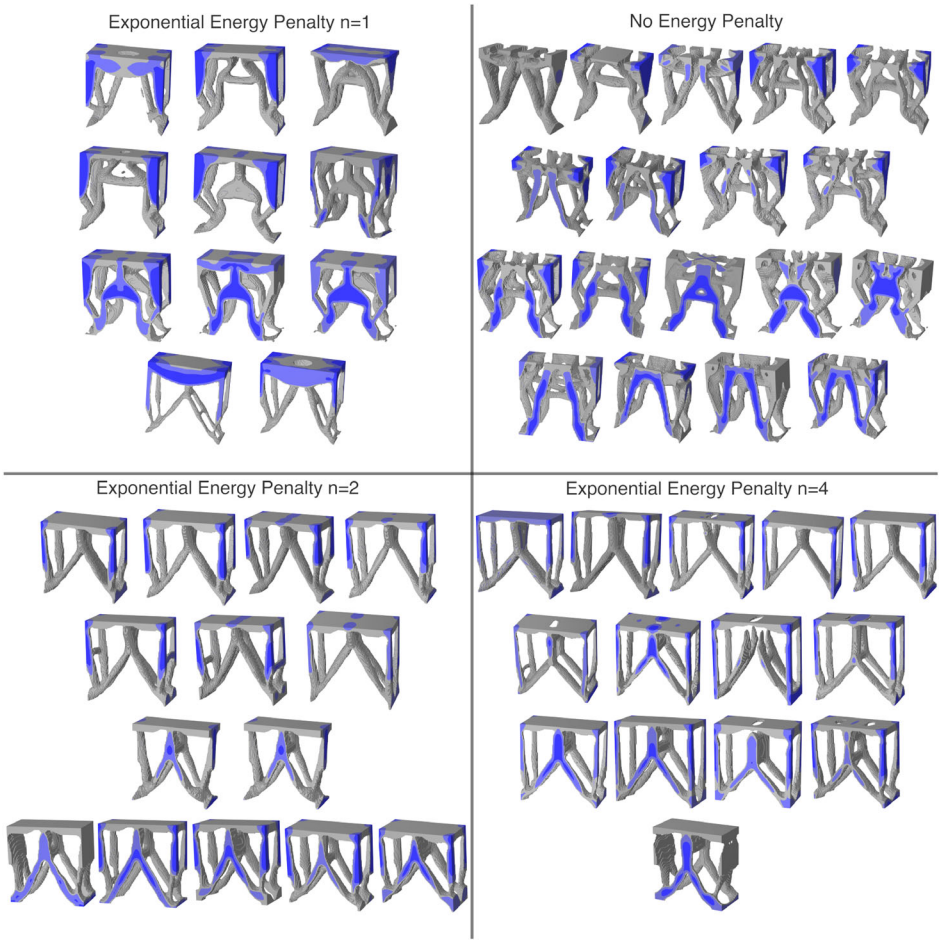

Figure 3. Converged results for the remaining 4 cost functions: No energy penalty; and strain energy penalty with $n = 1$, $n = 2$, and $n = 3$. Material colours: Gray - Agilus30, light blue - Agilus85, dark blue - Agilus 95.

7. Experimental Grasp Performance

The TO process uses highly accurate finite element modelling, but still requires linear assumptions to be made to tractability produce a design. Geometric non-linearities and contact mechanics are essential features of soft grasping, which are only coarsely approximated in the optimisation.

Further, salient features of the gripper such as contact surface and grasping mode are not explicitly optimized and instead emerge indirectly. Optimizing an end-effector which performs well across an array of objects and poses is an intractable task. However, by leveraging compliance in each design and diversity in the set of designs, these features emerge naturally.

To assess the correlation between design features and grasp performance, we print and test the 14 linear energy penalty designs (Table 1) in an automated robotic grasping facility (Figure 2f). Details of the manufacturing process and automated experimental facility can be found in Section 10.

We assess the grasping performance of the 14 linear energy penalty optimized gripper designs and 4 reference designs (Figure 2e). The 4 reference designs comprise 3D printed replicas of commercially available soft grippers: a Piab piSOFTGRIP 2 fingered vacuum gripper (P850), and a Soft Robotics MGrip 2 fingered inflatable finger (S850). Printing the designs geometry in-house allows us to isolate the effects of geometry from material, both were printed in shore 70 material. A final reference design manually post-processes the highest performing optimized gripper to further improve grasp quality (17–30 and 17–85). This extends gripper 17 by increasing the contact surface with flat gripping pads made from Agilus30 and Agilus85, respectively.

Each of the 18 grippers is tested across 410 grasps on each object, comprising 41 grasp positions with 10 repeats of each. The grippers were each tested on the three objects: coin, cube and sphere, sequentially.

This allows 4 features of the grippers to be evaluated: 1) Grasp Strength: The maximum vertical force the gripper can apply to an object to support its mass. 2) Durability: The number of cycles to failure. 3) Robustness: The variation in grasp strength due to uncertainty in the position of the object. 4) Generality: The change in grasp strength across objects.

For each set of tests the process is: 1) Centre the gripper on the object. 2) With the end-effector oriented vertically, grasp 10 times

from centre point. 3) Test translation sensitivity: For $x, y \in \{0\text{mm}, 2.5\text{mm}, 5\text{mm}, 7.5\text{mm}, 10\text{mm}\}$: a) Move end-effector by (x,y) relative to grasp-centre. b) Perform 10 grasps. 4) Test rotation sensitivity: For $\theta_x, \theta_y \in \{0°, 15°, 30°, 45°\}$: a) Rotate end-effector by $(\theta_x, \theta_y$ relative to grasp-centre, and align end-effector with object. b) Perform 10 grasps.

Where each grasp consists of moving the gripper from its pre-grasp pose to its grasp point, applying a vacuum, moving 50mm vertically whilst holding vacuum, releasing vacuum and returning to the starting position and orientation.

7.1. Grasp Modes and Strength

The grasping behaviour of elected grippers are shown in **Figure 4**a

From the results three distinct methods of grasping were identified, pinch, envelop, and wrap. The majority of grippers applied a pinch-grasp, where two points pushed against the gripping objects with a large normal force (#1, #4, #5, #8, #12, #14, #17). Others formed an enveloping grasp, where the contact points wrapped around the object, increases the contacting surface

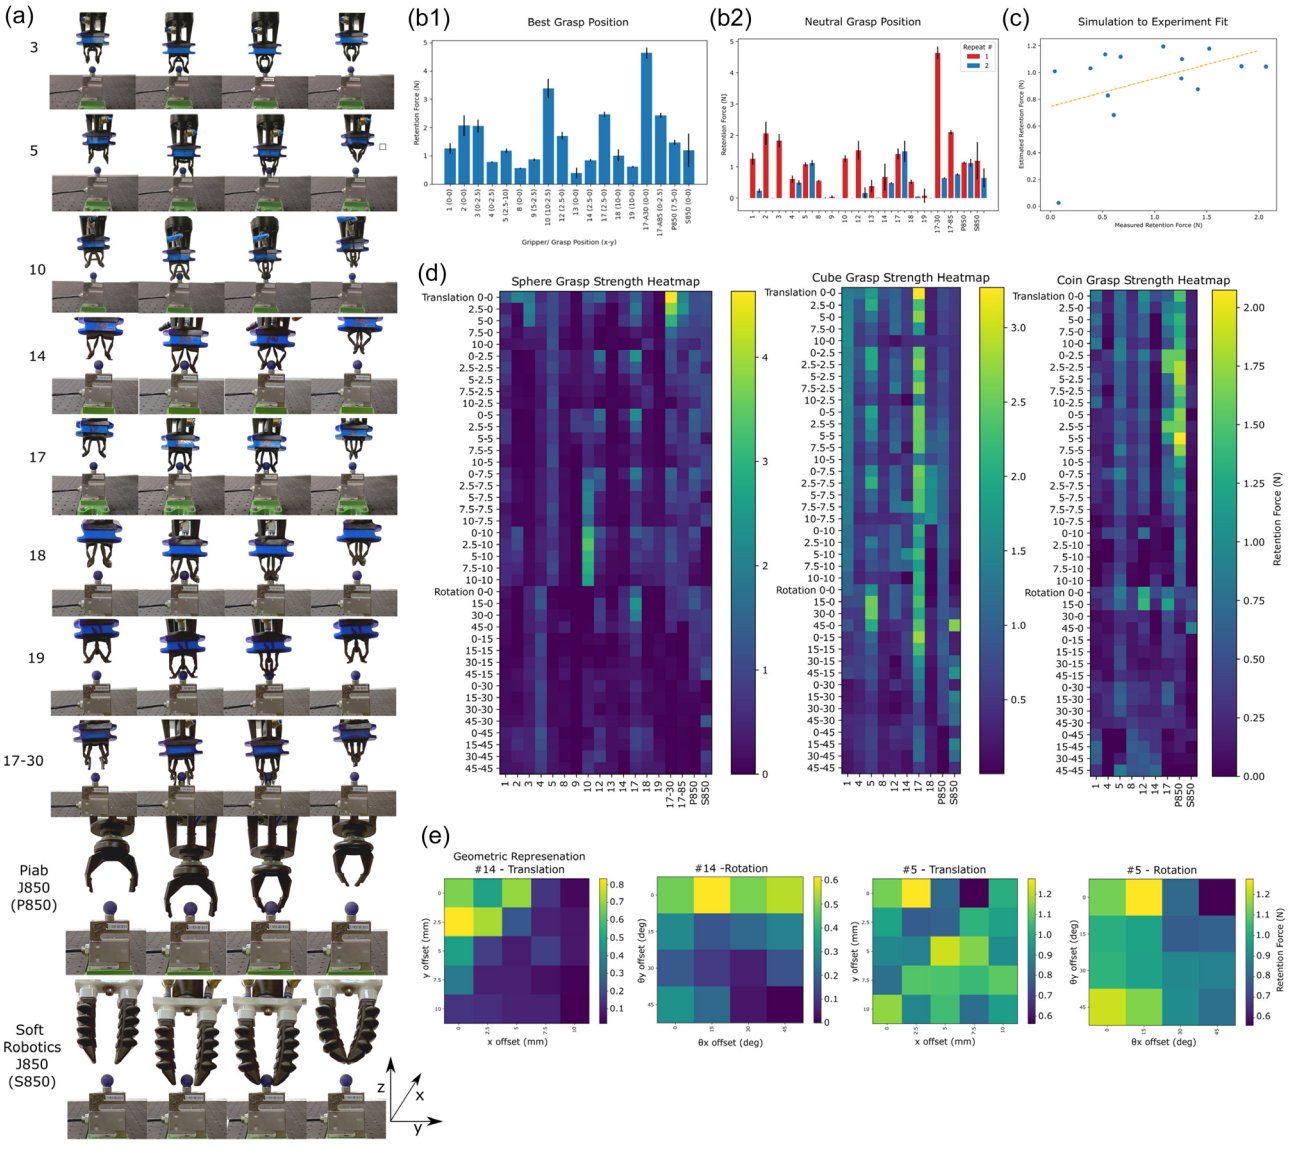

Figure 4. Experimental grasp test results a) grasping sequence of representative set of grippers, showing cycle of: initial position, pregrasp, grasping, and release. Spherical object shown has radius 18.5mm, b1) Grasp strength of 20 grippers on sphere at best position, indicating highest average grasp strength across 10 grasps. b2) Grasp strength of 20 grippers on sphere at initial position, showing response of freshly printed grippers and when retested after 250 grasp cycles. c) Comparison of estimated retention force to measured results, where estimated force is the linear fit of simulated SE and displacement to the measured results. d) Heatmaps of grasping retention force of 20 grippers tested on sphere, cube, and coin object across 41 poses. e) Geometric view of design #5 and #14 heatmaps, highlighting sensitivity to pose and positioning uncertainty. #14 is relatively sensitive to pose, whilst #5 is not.

(#3, #10, #18). Finally a few grippers used a thin member to wrap around the bottom surface of the object and lift directly (e.g., #2, #9, #13, #19). In comparison, both reference grippers used a pinch grasp.

Of the optimized designs, the largest grasp force (the highest average force across 10 samples at any pose) on the spherical object was an envelop grasp (#10) which pulled 3.38 N, followed by a pinch grasp (#17) at 2.45 N (Figure 4b1). Whilst there was significant variability between designs, pinch grasps performed best on average. Most grippers achieved maximum grasp strength when centred and grasping vertically, however gripper #10 was more than twice as effective when offset by 10mm in the x direction (towards rear of experimental table) as without offset (1.26 N vs. 3.38 N). Using the large offset, the gripper better conformed to the object as the front arms aligned with the spheres mid-plane. Gripper #19 was able to use its webbing to passively grasp objects with some success even after failure.

A comparison between the simulated and measured results is given in Figure 4c. The grasping force is estimated by linearly fitting the simulated SE and displacement to the measured force at its neutral (zero offset) position. In a perfect fit, the points would sit along a 45 degree line passing through the origin. Because of 1) the unmodelled contact interactions. 2) material degradation, and 3) other unomodelled nonentities, the estimate captures only 19.2% of the variation in the experimental data. Nevertheless, the fitted curve gives an initial estimate of retention force, which can be used to screen designs for experimental validation.

7.2. Durability

The centred grasp strength is presented in Figure 4b2, showing both the initial strength, and the strength when retested after 250 grasp cycles. Many designs failed within the first 250 grasps, and hence recorded no force when retested (#2, #3, #8, #9, #10, #13, #19). Designs with thin members or large deformations were most prone to failure through linkage fracture or membrane tears, respectively. For example, the membrane in #19 ruptured during the first test cycle.

The stresses caused during grasping and viscoelasticity of the agilus material resulted in a significant performance deterioration. This occurred not just in optimized designs, but also in the S850 printed replica. The P850 undergoes relatively little strain during grasping and hence does not exhibit the same performance deterioration.

7.3. Robustness and Generality

The complete set of grasp results are presented in Figure 4d. Showing the 41 sets of grasps for each gripper across the 3 objects. All grippers were tested top-to-bottom, then left-to-right in the order shown by the heatmaps. Grippers that failed were not tested on subsequent objects, hence there are more columns in the sphere grasp than the cube or coin.

Through these tests we investigate the ability of a high-performing design to generalize across poses and object. Ideally a design which grasps well from the vertical would also grasp well at a 30 degree offset, for example. Its assumed that the optimal grasp point is the neutral point, i.e., with the grippers oriented vertically and grasping the centre of the object, and that performance will degrade with an offset in translation or rotation. However the sensitivity varies significantly between designs. For example design #14 exhibits a high sensitivity to pose uncertainty, it shows a gradient of grasp strength with the strongest grasps in the top left corner and the weakest in the bottom right. In contrast #5 is robust to uncertainty, with little difference in performance between poses. This is illustrated in Figure 4e, where the line of data from the heatmap is reshaped as a grid to highlight the effect of pose offset. Although both gripper #5 and #14 use the same grasping mode (pinch), the higher stiffness and larger contact surface of #5 produces a more robust solution. Designs #1, #4, #12, and #17 (all pinch-grasps) are similarly robust, whilst all others perform poorly. Enveloping and Lifting grasps proved to be both sensitive to positioning, due to their requirement to completely enclose the object, and prone to failure because of their relatively thin members. The commercial replicas, P850 and J850 are both highly robust and generalizable. They display relatively little difference in grasp performance between objects and poses. The designs are able to achieve this through their large deformation and contact surfaces - both designs are significantly larger than optimised ones.

8. Discussion

In this work, the sensitivity of the gradient-based topology optimization solver to initial conditions was exploited to generate diverse designs, however many other features of the problem were assumed to be fixed. Key features of the problem and resulting design are explicitly or implicitly captured by the design domain, expanding the search space to incorporate features of the design domain would allow entirely new classes of grippers to emerge without designer input. Whilst the curse of dimensionality makes such a search intractable using current methods, an efficient method for identifying a 'good' design space or initial condition would be a large step forward for the field.

Combining generative design with high fidelity simulation, our diversity based topology optimization method is capable of generating a large set of design features. However, because of the complex interactions between actuator, material, geometry, and object, mapping these features onto the desired behaviours (grasp strength, robustness and generality, and gripper durability) remains a challenging tasks; even state of the art simulators (FEM or otherwise) cannot reliably calculate retention force in soft grippers. To address this challenge we present a two step design process, in which a diverse set of designs is first generated using generative topology optimisation, and then evaluated using an automated experimental platform. Thus filtering the large number of generated designs into a small set of high-quality ones. In this work the process is purely open-loop (i.e., the designs are generated then evaluated), requiring each generated design to be tested experimentally, however, closing the loop is of major interest going forward. Using a simple, linear regression (Figure 4c), we show that there is a positive correlation between optimization variables and grasp performance. However, there is also large degree of variability in the experimental results, which arise from unmodelled physical phenomena. Improving the quality of this fit, and hence improving optimization quality

and reducing the simulation to reality gap, is of significant interest going forward. This could be achieved by better understanding the relationship between optimisation variables and grasp quality, improving simulation fidelity, or directly learning accurate surrogate models from data. Continuously learning from the vast amounts of generated data could generate and reduce the need for experiments over time, hence allowing a much larger number of designs to be generated and evaluated. Further, evolving the CPPN in the closed loop, rather than randomizing it, would enable a broader search of the design domain and generate greater diversity.

From the analysis of different designs in this work, it is clear that a high-quality design must perform well across numerous features. In practice, performing well on a single benchmark is not sufficient for a usable design. However, the number of relevant features is strongly coupled to the design's use case, several salient features may reduce to just a few in narrowly specified tasks, for example a bespoke gripper design for a single object with a known and repeatable grasp pose. Given a broader set of conditions, a unique solution cannot be found to the multi-objective problem. Here, several possible candidates emerge which are worthy of further development. Design #10 gave the highest retention force on the sphere object but lacked durability, #17 gave a large retention and generalized well, whilst #5 gave a modest retention but was insensitive to pose errors. Through an ongoing process of large-scale, closed-loop data collection a multi-dimensional pareto front could be generated capturing each salient feature. However, it is preferable to explicitly optimize only one or two features explicitly and let the remaining features emerge indirectly. With a sufficient pool of candidate, low quality designs can be filtered out, leaving only the high performing ones.

9. Conclusion

Soft grippers are uniquely suited to grasping complex and flexible objects. Leveraging multi-material 3D printing and expressive design algorithms, bespoke soft gripping designs can be produced which are tailored to the needs of specific applications and environments. In this work we present a topology optimisation method which uses a multi-physics optimisation and CPPN seeding to generate diverse soft gripper designs. An automated experimentation platform was used to collect a vast dataset (15 170 grasps) and evaluate generated designs. Segmenting the design task into design generation and evaluation stage transforms a fuzzy, ill-posed problem with numerous goals into a tractable one with a quantifiable cost. Using the method, we generated 71 designs, spanning a pareto front of 2 major design goals (minimizing strain energy and maximising displacement). Three unique grasping modes and numerous design morphologies emerged from the 71 designs without explicit prompting, highlighting the method's ability to explore the complex space and find multiple high-performing solution. The optimized designs' simulated performance was mapped to actual task performance across several key factors using multimaterial polyjet 3D printing and an automated experimental facility, enabling the identification grippers which outperform commercial benchmarks in grasping strength. The method presents a new paradigm for bespoke soft gripper design which generalizes across tasks and environments simply by adjusting the FEM and experiments.

10. Experimental Section

Material Characterisation: The 3D printed material blends are characterised through a set of ASTM D412 standard tensile tests on an Instron 34SC-5 universal testing machine. We evaluated 7 preset 'digital material' blends, with Shore-A values of 30, 40, 50, 60, 70, 85, and 95. The elastic modulus was then identified by fitting a curve to the experimental stress-strain curve of each material. See[61] for a complete test procedure and results. Although the materials display a slight hyperelasticity, we approximate their behaviour as linear elastic for simulation efficiency. The fitted moduli, tensile strength and maximum elongation are presented in **Table 2**. To maximise the range of stiffness, 3 materials were chosen for the optimizations: A-30, A-85, and A-95. Their stress–strain curves are shown in **Figure 5**.

Experimental Facility: The grasping abilities of each gripper were evaluated using a robotic grasp-testing facility (Figure 2f. It consists of a 7 DOF Haddington HDI Robotic Arm, Load cell, and Pneumatic infrastructure (compressor, vacuum generator, and solenoid valves). A vacuum is generated using a Venturi generator which converts positive pressure into vacuum. The Festo generator is capable of producing −93 kPa vacuum pressure, however because of frictional losses in the lines and fittings, the pressure at the gripper was measured to be −74 kPa. The two reference designs using positive pressure are inflated to 50 kPa

Manufacturing Process: The optimised designs are 3D printed using a 3D printing on J850 using printed seals a Stratasys J850 multimaterial polyjet printer. It allows continuous blending of multiple base resins to generate custom material properties and distribute them to individual voxels.

To allow easy removal of support materials, each gripper is printed in two parts: 1) The main gripper, comprising gripping fingers, membrane, and housing. 2) A lid with printed gasket seal.

A central hole in the lid is tapped to screw in an air fitting with. The gripper is assembled by pressing the two parts together, screwing in the air fitting and attaching the pneumtic tube.

A set of printed grippers and lids immediately after printing is shown in **Figure 6**. They are printed vertically to enable easy cleaning of the (clear) support material.

Table 2. Results (mean and standard deviation) for Elastic's modulus, tensile strength, and elongation at break for all digital material blends.

Material		Elastic modulus [MPa]	Tensile strength [MPa]	Elongation at break [%]
Shore A-30	Mean	0.46	1.42	304
(Agilus30)	S.D.	0.05	0.06	18.8
Shore A-40	Mean	0.50	1.30	253
	S.D.	0.03	0.11	6.5
Shore A-50	Mean	0.60	1.44	230
	S.D.	0.05	0.11	3.6
Shore A-60	Mean	0.90	2.37	247
	S.D.	0.05	0.05	10.4
Shore A-70	Mean	1.39	2.88	200
	S.D.	0.08	0.13	12.7
Shore A-85	Mean	11.51	4.78	104
	S.D.	1.75	0.29	9.4
Shore A-95	Mean	39.98	9.76	86
	S.D.	2.70	0.59	5.3

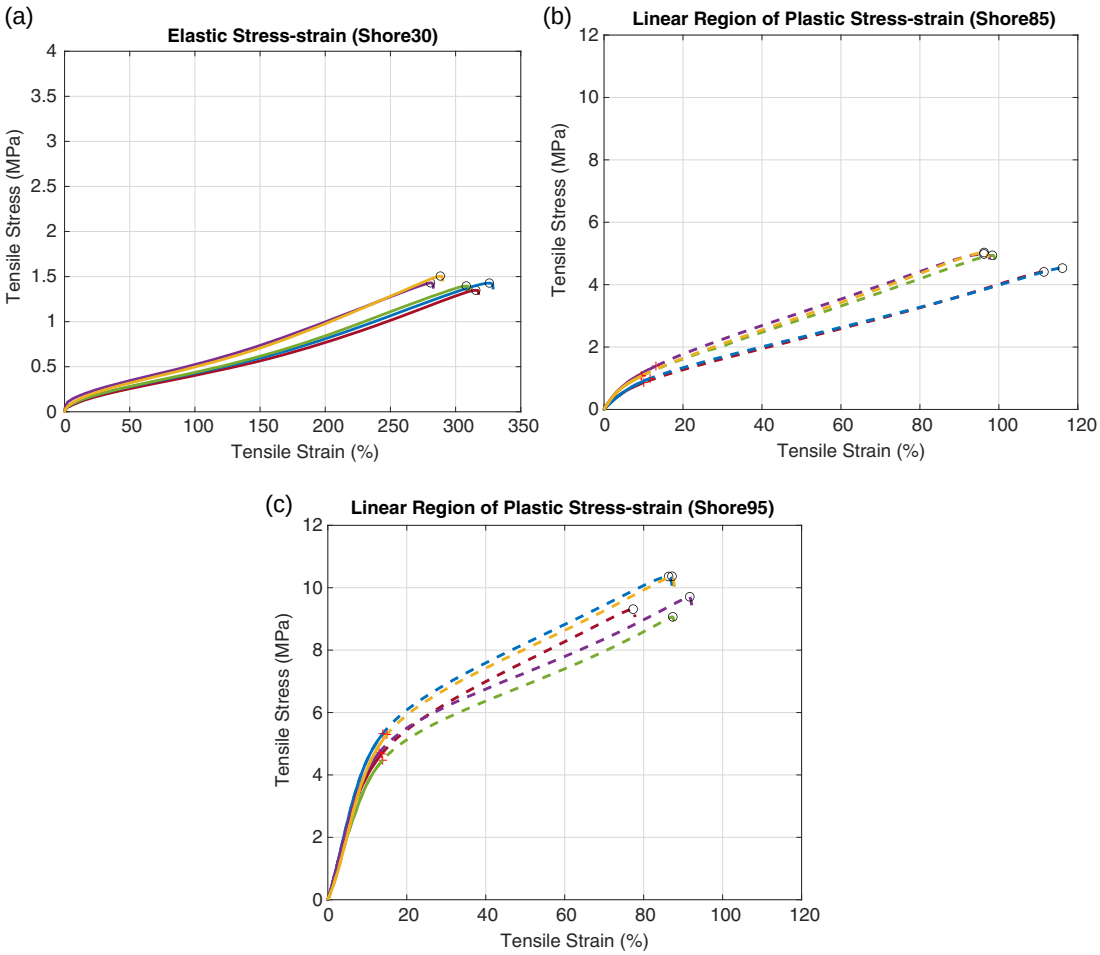

Figure 5. Experimental stress–strain Curve for a) Agilus30, b) Agilus85, and c) Agilus95.

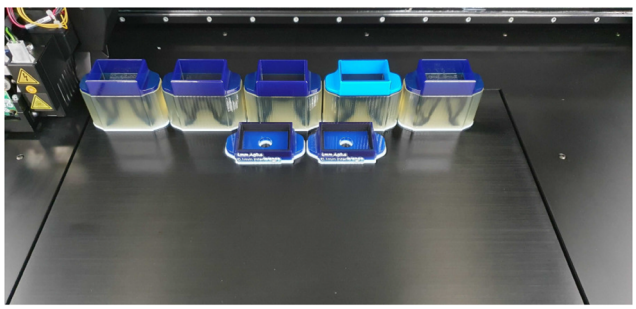

Figure 6. Set of 5 grippers and 2 lids on Stratasys J850 print bed.

Three test objects are also printed, a 18.5mm radius sphere, a cube with 18.5mm sides, and a coin with radius 28.65mm (equal to an Australian 20 cent coin).

Acknowledgements

This research was supported by the Science Industry and Endowment Fund.

Conflict of Interest

The authors declare no conflict of interest.

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Keywords

computational design, soft robotics, topology optimization

Received: August 23, 2023
Revised: November 15, 2023
Published online: January 18, 2024

[1] F. Ilievski, A. D. Mazzeo, R. F. Shepherd, X. Chen, G. M. Whitesides, *Angew. Chem.* **2011**, *123*, 1930.
[2] E. Brown, N. Rodenberg, J. Amend, A. Mozeika, E. Steltz, M. R. Zakin, H. Lipson, H. M. Jaeger, *Proc. Natl. Acad. Sci. U. S. A.* **2010**, *107*, 18809.

[3] S. G. Fitzgerald, G. W. Delaney, D. Howard, *Actuators* **2020**, *9*, 1.
[4] O. Pfaff, S. Simeonov, I. Cirovic, P. Stano, *Ann. DAAAM Proc.* **2011**, *22*, 1247.
[5] J. Hughes, U. Culha, F. Giardina, F. Guenther, A. Rosendo, F. Iida, *Front. Robot. AI* **2016**, *3*.
[6] D. Howard, J. O'Connor, J. Letchford, J. Brett, T. Joseph, S. Lin, D. Furby, G. W. Delaney, in *2022 IEEE 5th Inter. Conf. on Soft Robotics, RoboSoft 2022*, Edinburgh, Scotland **2022**, pp. 531–538.
[7] S. G. Fitzgerald, G. W. Delaney, D. Howard, F. Maire, in *GECCO2021*, **2021**, pp. 102–110.
[8] Z. Wang, S. Hirai, A. S. F. Design, in *2018 IEEE Inter. Conf. on Robotics and Biomimetics (ROBIO)*, Kuala Lumpur, Malaysia **2018**, pp. 612–617.
[9] T. Joseph, S. Baldwin, L. Guan, J. Brett, D. Howard, in *2023 IEEE Inter. Conf. on Soft Robotics, RoboSoft 2023*, Singapore **2023**, pp. 1–6.
[10] K. Becker, C. Teeple, N. Charles, Y. Jung, D. Baum, J. C. Weaver, L. Mahadevan, R. Wood, *Proc. Natl. Acad. Sci. U. S. A.* **2022**, *119*, 1.
[11] M. N. Zadeh, M. Garrad, C. Romero, A. Conn, F. Scarpa, J. Rossiter, *IEEE Robot. Autom. Lett.* **2022**, *7*, 10352.
[12] J. Pinskier, D. Howard, *Adv. Intell. Syst.* **2022**, *4*, 2100086.
[13] J. K. A. Langowski, P. Sharma, A. L. Shoushtari, *Sci. Robot.* **2020**, *5*, eabd9120.
[14] Y. Hao, Y. Visell, *Front. Robot. AI* **2021**, *8*.
[15] J. Pinskier, B. Shirinzadeh, *Precis. Eng.* **2019**, *55*, 397.
[16] J. Pinskier, B. Shirinzadeh, M. Ghafarian, T. K. Das, A. Al-Jodah, R. Nowell, *Mech. Mach. Theory* **2020**, *150*, 103874.
[17] Z. J. Wegert, A. P. Roberts, T. Bandyopadhyay, V. J. Challis, *Materials* **2023**, *16*, 1.
[18] F. Chen, M. Y. Wang, *IEEE Robot. Autom. Mag.* **2020**, *27*, 155.
[19] M. S. Xavier, C. D. Tawk, A. Zolfagharian, J. Pinskier, D. Howard, T. Young, J. Lai, S. M. Harrison, Y. K. Yong, M. Bodaghi, A. J. Fleming, *IEEE Access* **2022**, *10*, 59442.
[20] C. Tawk, G. Alici, *Robotics* **2020**, *9*, 9.
[21] T.-H. Wang, P. Ma, A. Spielberg, Z. Xian, H. Zhang, J. B. Tenenbaum, D. Rus, C. Gan, in *ICLR 2023*, Kigali, Rwanda **2023**.
[22] Y. Hu, L. Anderson, T. M. Li, Q. Sun, N. Carr, J. Ragan-Kelley, F. Durand, in *ICLR 2020*, **2020**, pp. 1–20, ISSN 23318422.
[23] M. A. Graule, C. B. Teeple, T. P. McCarthy, G. R. Kim, R. C. St. Louis, R. J. Wood, in *IEEE Inter. Conf. on Intelligent Robots and Systems*, Prague, Czech Republic **2021**, pp. 3934–3941.
[24] J. D. Hiller, H. Lipson, in *Proc. of the 11th Annual Genetic and Evolutionary Computation Conf., GECCO-2009*, Shanghai, China **2009**, pp. 1521–1528.
[25] J. Collins, S. Chand, A. Vanderkop, D. Howard, *IEEE Access* **2021**, in press.
[26] H. Ha, S. Agrawal, S. Song, in *4th Conf. on Robot Learning (CoRL 2020), CoRL*, Cambridge, MA, USA **2020**, https://fit2form.cs.columbia.edu/
[27] M. Kodnongbua, I. Good, Y. Lou, J. Lipton, A. Schulz, *ACM Trans. Graph.* **2022**, *41*, 1.
[28] J. Xu, T. Chen, L. Zlokapa, M. Foshey, W. Matusik, S. Sueda, P. Agrawal, in *Proc. of Robotics: Science and Systems*, July **2021**.
[29] W. Hu, R. Mutlu, W. Li, G. Alici, *Robotics* **2018**, *7*, 1.
[30] H. K. Yap, H. Y. Ng, C. H. Yeow, *Soft Robot.* **2016**, *3*, 144.
[31] R. Su, Y. Tian, M. Du, C. C. L. Wang, *IEEE Robot. Autom. Lett.* **2022**, *7*, 10430.
[32] Y. Yao, L. He, P. Maiolino, in *2022 IEEE 5th Inter. Conf. on Soft Robotics, RoboSoft 2022*, Edinburgh, Scotland **2022**, pp. 29–34.
[33] L. Smith, R. MacCurdy, *IEEE Trans. Autom. Sci. Eng.* **2023**, *20*, 1475.
[34] S. E. Navarro, T. Navez, O. Goury, L. Molina, C. Duriez, *IEEE Robot. Autom. Lett.* **2023**, *8*, 6044.
[35] J. Rieffel, D. Knox, S. Smith, B. Trimmer, *Artif. Life* **2013**, *19*, 119.
[36] N. Cheney, R. Maccurdy, J. Clune, H. Lipson, in *GECCO 2013*, Amsterdam, Netherlands, **2013**, ISBN 9781450319638.
[37] J. S. Bhatia, H. Jackson, Y. Tian, J. Xu, W. Matusik, *Advances in Neural Information Processing Systems*, NeurIPS **2021**, Vol. *3*, NeurIPS 2201.
[38] S. Chand, D. Howard, *Front. Robot. AI* **2021**, *8*.
[39] S. Kriegman, A. M. Nasab, D. Shah, H. Steele, G. Branin, M. Levin, J. Bongard, R. Kramer-Bottiglio, in *3rd IEEE Inter. Conf. on Soft Robotics (RoboSoft) Yale University, USA Scalable*, Yale University **2020**, Yale University, USA pp. 359–366, ISBN 9781728165707, ISSN 23318422.
[40] P. Schegg, E. Ménager, E. Khairallah, D. Marchal, J. Dequidt, P. Preux, C. Duriez, *Soft Robot.* **2023**, *10*, 410.
[41] M. P. Bendsøe, O. Sigmund, *Topology Optimization: Theory, Methods, and Applications*, Springer, Berlin Heidelberg **2003**.
[42] X. Huang, Y. Li, S. W. Zhou, Y. M. Xie, *Eng. Struct.* **2014**, *79*, 13.
[43] C. H. Liu, T. L. Chen, C. H. Chiu, M. C. Hsu, Y. Chen, T. Y. Pai, W. G. Peng, Y. P. Chiang, *Soft Robot.* **2018**, *5*, 452.
[44] F. Chen, W. Xu, H. Zhang, Y. Wang, J. Peng, W. G. Cao, M. Y. Wang, H. Ren, J. Zhu, Y. Zhang, *IEEE Robot. Autom. Lett.* **2018**, *3*, 2463.
[45] C. H. Liu, L. J. Chen, J. C. Chi, J. Y. Wu, *IEEE Robot. Autom. Lett.*, April **2022**, *7* 1.
[46] H. Zhang, A. S. Kumar, F. Chen, J. Y. Fuh, M. Y. Wang, *IEEE/ASME Trans. Mechatron.* **2019**, *24*, 120.
[47] S. Chen, F. Chen, Z. Cao, Y. Wang, Y. Miao, G. Gu, X. Zhu, *IEEE/ASME Trans. Mechatron.* **2021**, *26*, 1745.
[48] B. Caasenbrood, Sorotoki - a soft robotics toolkit for matlab, https://github.com/BJCaasenbrood/SorotokiCode, **2020**.
[49] E. M. de Souza, E. C. N. Silva, *Struct. Multidiscipl. Optim.* **2020**, *61*, 1763.
[50] Y. Chen, Z. Xia, Q. Zhao, *IEEE/ASME Trans. Mechatron.* **2019**, *24*, 2873.
[51] J. Pinskier, P. Kumar, M. Langelaar, D. Howard, in *Robosoft 2023*, Singapore **2023**.
[52] O. Sigmund, *Struct. Multidiscipl. Optim.* **2007**, *33*, 401.
[53] O. Sigmund, K. Maute, *Struct. Multidiscipl. Optim.* **2013**, *48*, 1031.
[54] P. Kumar, J. Frouws, M. Langelaar, *Struct. Multidiscipl. Optim.* **2020**, *61*, 1637.
[55] P. Kumar, M. Langelaar, *Int. J. Numer. Methods Eng.* **2021**, *122*, 2205.
[56] K. Svanberg, *Int. J. Numer. Methods Eng.* **1987**, *24*, 359.
[57] J.-B. Mouret, J. Clune, arxiv:cs.AI **2015**, 1–15.
[58] K. O. Stanley, *Genet. Program. Evolvable Mach.* **2007**, *8*, 131.
[59] D. Howard, J. Kannemeyer, D. Dolcetti, H. Munn, N. Robinson, *Assessing Evolutionary Terrain Generation Methods for Curriculum Reinforcement Learning*, Vol. *1*, Association for Computing Machinery, New York **2022**.
[60] W. E. Lorensen, H. E. Cline, in *Proc. of the 14th Annual Conf. on Computer Graphics and Interactive Techniques, SIGGRAPH 1987*, Anaheim, CA, US **1987**, Vol. *21*, p. 163.
[61] L. Liow, J. Pinskier, D. Howard, *PolyJet Digital Materials Tensile Characterisation based on ASTM D412. v2. CSIRO*, *CSIRO Data Collection*.