

Evolving variable stiffness fiber patterns for multi-shape robotic sheets

Atoosa Parsa^{1*}, Medha Goyal^{2*}, Maggy Lambo², Bilige Yang², Josh Bongard¹, and Rebecca Kramer-Bottiglio²

Abstract—Thin, planar sheets can be programmed to morph into complex shapes through stretching and out-of-plane bending, with applicability to shape-shifting soft robots. One way to make a morphing sheet is to use variable stiffness fibers that can modulate their tensile stiffness attached to the surface of a volumetrically expanding sheet. Adjusting local stiffnesses via tensile fiber jamming during sheet expansion allows control of the local shape tensor. However, finding the fiber placements and jamming policies to achieve a set of desired shapes is a non-trivial inverse design problem. We present an additive inverse design framework using an evolutionary algorithm to find optimal jamming fiber patterns to match multiple target shapes. We demonstrate the utility of our optimization pipeline with two input curvature pairs: 1) cylinder and sphere curvatures and 2) simple saddle and monkey saddle curvatures. Our method is able to find a diverse set of sufficient solutions in both cases. By incorporating hardware constraints into our optimization pipeline, we further explore the transfer of evolved solutions from simulation to reality.

I. INTRODUCTION

Biological systems are able to adjust their morphology and behavior to accommodate various tasks and environmental conditions. Inspired by such dynamic plasticity, soft robots are often designed to respond to changing tasks and environments through adaptation of their physical structure using an increasing diversity of approaches [1]. One approach is via shape-shifting stimuli-responsive sheets programmed to morph into a target curvature upon actuation. This shape change occurs when local strains arising from differential in-plane growth cause out-of-plane buckling, as seen in flowers that unfold [2] and seed pods that break open [3]. The inverse design problem of finding local differential growth required to morph any reference surface into any target curvature has been solved analytically for the unconstrained case [4].

Unlike organisms, however, synthetic shape-shifting sheets struggle to morph into more than a single pre-programmed target shape. Demonstrations of shape-shifting sheets comprised of hydrogels [5], dielectric elastomers [6], and pneumatic bilayers [7], [8] result in fixed shapes that cannot be reconfigured without fabricating a new device.

Yang et al. [9] tackled this problem with strain-limiting fibers (first described by Kim et al. [10]), which can be adhered, removed, and re-adhered on a pneumatic bilayer sheet. Fiber arrangements to produce curvatures without

analytical solutions, such as the monkey saddle, were found using a genetic algorithm [9]. Once adhered, the inextensible fibers created local strains for a bilayer inflating uniformly in-plane, biasing deformation to ensure that the minimal energy curvature is the target shape. New curvatures could only be accessed via time-consuming manual intervention—fibers must be removed, rearranged, and re-adhered—rendering the approach infeasible in the field. A new calculation is required to determine fiber placements for each new curvature.

To bypass manual reconfiguration, several approaches have been employed to create multi-shape sheets. Some include patterning actuating components on inert substrates [6], [11], or variable stiffness components on expanding substrates [12], [13], and selectively activating the components to produce new curvatures. In prior works, however, the dynamic components are placed in simple pre-determined positions, which limits their accessible curvatures.

To attain a multiplicity of accessible curvatures, we propose finding novel, non-intuitive placements of variable stiffness components that can turn their strain-limiting properties on/off on the surface of an inflating sheet. We use dual-mode “tensile jamming fibers” [12], which can rapidly modulate tensile stiffness while maintaining a low bending stiffness in all directions, switching from inert (up to 200% tensile strain) to strain-limiting (low tensile strain) via 1D stretchable vacuum layer jamming. When a new curvature is desired, fibers need not be moved. Rather, a different combination of variable stiffness fibers required for the differential growth pattern can be activated while the rest can be kept inert.

Yet, the inverse design problem compounds when more than one shape is desired due to the complex relationship between the sets of jammed fiber patterns and the sheet. To address this, we present a multi-objective evolutionary algorithm to optimize tensile jamming fiber placements on an inflatable sheet for multiple desired inflated curvatures. The evolutionary algorithm searches the space of all possible fiber patterns, where each pattern can vary in fiber count, positions, lengths, and orientations. The algorithm also incorporates fabrication and hardware constraints to facilitate the transfer of selected designs from simulation to reality.

Fig. 1 shows our optimization pipeline that takes two target curvatures as inputs and finds two independent sets of fiber placements that achieve the target curvatures when activated independently. We demonstrate the utility of our pipeline with two target curvature pairs: 1) sphere & cylinder and 2) simple saddle & monkey saddle. Our optimization framework is able to produce a diverse set of “good” solutions in both cases. We transfer one shape-pair solution

* Authors contributed equally.

¹Department of Computer Science, University of Vermont, 82 University Pl, Burlington, VT 05405, USA. (email: josh.bongard at uvm.edu)

²Department of Mechanical Engineering and Materials Science, Yale University, 10 Hillhouse Avenue, New Haven, CT 06520, USA.

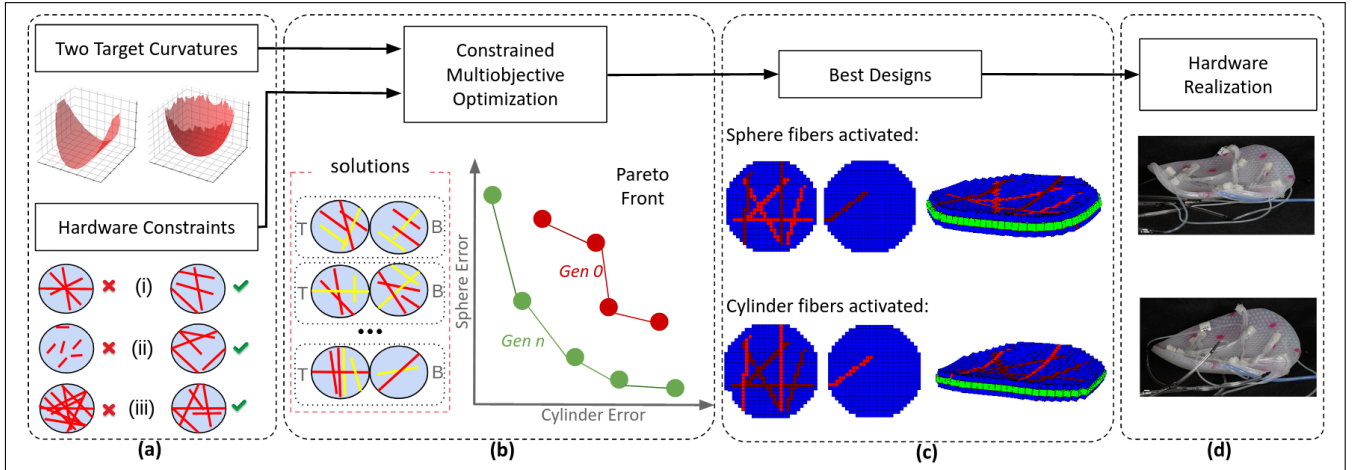


Fig. 1. **Pipeline overview.** (a): The pipeline begins with two target curvatures and constraints imposed by the hardware setup. Constraints are too many intersecting fibers (a-i), overly short fibers (a-ii), and too many fibers (a-iii). (b): A population of random fiber patterns is generated (*Gen 0*). Each pattern is composed of four fiber sets: two on the [T]op and two on the [B]ottom. Activation of the red fibers is an attempt to match the first target shape; Activation of the yellow fibers is an attempt to match the second target shape. An evolutionary multi-objective optimization algorithm finds a set of non-dominating solutions that satisfy the objectives to different degrees (*Gen n*). (c): Each solution is a pair of fiber patterns that, when activated independently, cause the inflated sheet to form the corresponding desired curvature. (d): We choose solutions from the Pareto front to fabricate and evaluate in hardware.

for the sphere/cylinder targets to reality and demonstrate that, because the hardware constraints have been incorporated into the optimization process, the designs can be transferred to reality with some initial success.

II. METHODS

A. Simulation Setup

To simulate the sheet, we used a GPU-accelerated voxel-based physics engine called *Voxcraft* [14]. This physics engine was developed based on a 3D simulator that quantitatively models the statics, dynamics, and nonlinear deformation of heterogeneous soft bodies using a Finite Element Method (FEM) [15]. The inflatable sheet is modeled as three layers of soft voxels placed in a circular grid. The top and bottom layers are made of passive voxels, while the middle layer is filled with active voxels that expand during the simulation time. The fiber patterns are simulated by replacing the corresponding passive voxels on the top or bottom layers of the bladder with passive voxels of a higher modulus that resist tensile deformation (if the fiber is jammed) or passive voxels of a lower modulus (if the fiber is unjammed). Table I includes the simulation parameters used in all experiments, as measured from the physical hardware. A number of additional simulator parameters control the expansion of the active voxels, which were set by calibrating the simulation to match a real cylinder curvature with three parallel fibers on the top face of the sheet.

B. Optimization

Evolutionary algorithms are a class of population-based gradient-free optimization methods that are inspired by natural evolution. They generally start with a random population of candidate solutions and at each step of the optimization, the fittest solutions based on problem objectives are selected to reproduce and survive to the next generation. In this paper,

TABLE I
SIMULATION PARAMETERS.

Parameter	Value
Grid Size	28×28 voxels
Diameter	14 (cm)
Density	938.33 (kg/m^3)
Moduli of passive layers	82700 (Pa)
Modulus of jammed fiber	1330688 (Pa)
Modulus of unjammed fiber	68900 (Pa)
Modulus of expanding layer	160000 (Pa)

we used a multi-objective evolutionary algorithm called Age-Fitness Pareto Optimization (AFPO) [16]. This approach was developed to address the premature convergence problem by allowing newly added solutions to survive in the population before being dominated by the more fitted ones, thereby increasing the overall diversity of the population. In the last generation of the optimization, the algorithm will converge to a population of non-dominated Pareto optimal solutions.

The evolutionary algorithm has multiple components to encode the solutions (genome representation), evaluate them based on the objectives (fitness definition), select the solutions to reproduce (selection operator), and produce the next generation of solutions (crossover and mutation operators). The following sections provide the details of these components in our optimization framework.

1) *Genome Representation*: In our optimization problem, each solution is composed of two independent pairs of fiber patterns, each associated with one of the two target curvatures. Fiber patterns are modeled as straight line segments with four numbers indicating their starting and ending coordinates on the sheet. We used a variable-length direct encoding scheme for the genome representation. To map the genotype to fiber patterns, we used a variation of the Bresenham algorithm [17] as the rasterization process that draws the fiber segments on the simulated discrete grid.

2) *Mutation Operator*: To produce the next generation in the evolutionary algorithm, we implemented two mutation operators. The first operator (applied with a probability of 25%) changes the number of fibers by deleting one randomly chosen fiber or adding a randomly generated one. The second mutation operator (applied with a probability of 75%) causes a smoother variation in the configuration space by changing one endpoint of a randomly chosen fiber to a new position drawn from a uniform random distribution within the grid.

3) *Optimization Constraints*: We include fabrication constraints in the optimization process to ensure the selected fiber patterns can be transferred to reality. Fibers shorter than six cm are challenging to fabricate, thus providing the first constraint. Because the fibers are bulky, we cannot overlay more than two without losing their coupling to the underlying sheet (*i.e.*, a fiber laid on top of others would not make full contact with the sheet in the vicinity of the overlap). Thus, a second constraint limits the number of fibers at any given intersection to two. The last constraint is on the total number of fibers in a pattern: we limit this to a maximum of six fibers on each face of the sheet for the sphere/cylinder curvatures. However, this number is increased to 10 for the saddle/monkey-saddle curvatures due to their complexity. All constraints were imposed in the initialization step and checked during the optimization process, to ensure the optimal fiber patterns found by the algorithm were realizable in hardware.

4) *Fitness Definition*: With evolved fiber patterns, we simulated the sheet with both fiber patterns (*i.e.*, sphere-fibers, cylinder-fibers) on the top and bottom faces. We used the heights of specific and distributed points to determine how much the inflated sheet matches target curvatures when the corresponding set of fibers is jammed. The error between the target and the simulated solution (e_s) is defined as:

$$e_s = \frac{\sum_{i=1}^{n_v} |h_i - t_i|}{n_v} \quad (1)$$

where h_i is the height of voxel i at the end of the simulation time, n_v is the number of voxels on the top layer of the inflated sheet, and t_i is the height of the corresponding point from the target curvature.

The e_s metric represents the average error per voxel, measured in mm. In each of the case studies presented in this paper, we have two target curvatures. For every candidate solution, we calculated two errors associated with matching each of the two targets. These two errors are used as the fitness values that will be minimized by our multi-objective optimization algorithm. In the results section, we also report the error after normalizing the values by the minimum and maximum attainable heights in the simulated targets, to give an unbiased indication of matching the target shapes.

5) *Parameters and Setup*: We ran the evolutionary algorithm with a population size of $p = 100$ for $g = 300$ generations for the sphere/cylinder and $g = 350$ for the saddle/monkey-saddle. One experimental run takes about 15 hours to finish on a *Penguin Relion XE4118GTS* GPU

node. We conducted three independent runs for each of the two case studies. All optimizations¹ were run on the GPU cluster at Vermont’s advanced computing core.

C. Hardware Setup

The pneumatic bilayer sheet was fabricated following the method detailed in Yang et al. [9]. The sheet contains pillars connecting the top and bottom layers, thus constraining expansion to be in-plane. In prior work, strain-limiting fibers were adhered to a sheet’s surfaces to induce curvatures upon inflation, but the output was immutable. Here, we introduce variable tensile stiffness fibers to this platform to adjust local stiffness during volumetric expansion of the sheet, allowing control of the local shape tensor. The fibers were fabricated following the procedure described in Yang et al. [12]. This prior work included a demonstration of the tensile jamming fibers arranged in a grid on a pneumatic sheet. In this paper, we open the design space to allow the fibers to vary in length and placement. We adhere the fibers to the sheet using Silpoxy™ (Smooth-On).

Fiber patterns for specific target sheet curvatures obtained via evolutionary algorithm were transferred to hardware. Variable stiffness fibers were fabricated with lengths and placements on the sheet provided by the optimization. Select fibers were jammed with vacuum to activate a target shape. The sheet was inflated and held at 3.9 ± 0.1 psi (3.6 ± 0.1 psi for the neat sphere due to occlusions at higher psi) and heights along the inflated curvature were measured at 16 representative points using a commercial motion capture system (Phasespace, Inc.) and averaged over five trials.

D. Design of Experiments

In our prior work, we searched for a set of inert fiber placements, each of which enabled a sheet to obtain a desired curvature [9]. Here we find patterns that leverage the fibers’ stiffness-switching ability to achieve shape change. The simplest approach is to run two evolutionary algorithms—The first evolves a specialized pattern P_i that produces shape S_i ; The second evolves a specialized pattern P_j that produces shape S_j . The best pattern from each algorithm can then be extracted and layered on to the same sheet. Shape S_i can be achieved with the merged pattern by stiffening P_i and slackening P_j , and shape S_j can be achieved by stiffening P_j and slackening P_i . However, we found that this approach of simply layering optimal fiber patterns could not translate well to the physical hardware, as it resulted in many fiber intersections (3+ overlapping fibers) and otherwise ineffective combined pattern signatures.

We sought to resolve this problem by evolving patterns P_i and P_j in parallel and with imposed constraints that would increase the likelihood of successful transferal to reality. With the evolved fiber pattern pairs, we asked two questions:

- 1) How well do the evolved simulated shapes match the target shapes (simulation error e_s), and how well do the

¹<https://github.com/AtoosaParsa/robosoft-2023> contains the source code necessary for reproducing our results.

hardware-realized shapes match the evolved simulated shapes (transfer error e_t)?

- 2) When the fiber patterns are combined on the sheet, does the order of the layering matter? That is, does it matter if P_i is layered over P_j , or vice versa?

We designed our experiments to address these questions. First, we focused on simple curvatures with known fiber placement solutions: sphere/cylinder. We translated evolved fiber patterns for the sphere shape S_s and the cylindrical shape S_c to reality, and quantified differences between the simulated and hardware curvatures. Then, we layered the two patterns (cylinder on sphere; sphere on cylinder), and measured the shape fidelity when the patterns are combined and in different orders. We then extended the evolutionary algorithms to the more complex saddle/monkey-saddle shapes.

E. Calculation of Transfer Error

We quantified the error between the simulation output curvatures from evolved fiber patterns and the real hardware curvatures, e_t , as follows: Heights (measured via motion capture) for several points along the inflated sheets were normalized by the greatest height difference between the sheet when inflated and when uninflated. Normalization was also performed in simulation to ensure a fair comparison. The transfer error e_t is defined as:

$$e_t = \frac{\sum_{i=1}^{n_p} |r_i - s_i|}{n_p} \quad (2)$$

where r_i is the normalized height of the sheet at point i in hardware, s_i is the normalized height in simulation, and n_p is the number of points measured. Normalization was to ensure a direct comparison between curvature in simulation and hardware without considering scaling differences, since the heights achieved in hardware were highly dependent on the inflation pressure.

III. RESULTS

We demonstrate the performance of our optimization pipeline with two target shape-pairs: 1) sphere and cylinder curvatures, and 2) saddle and monkey-saddle curvatures. Fig. 2a,c depicts the optimization progress in terms of the average population fitnesses. As expected, the average error between the target shapes and evolved simulated shapes, e_s , decreases during the optimization. Fluctuations in the average error are due to the multi-objective nature of the problem and the underlying principles of the optimization algorithm. Fig. 2b,d shows the Pareto-optimal solutions found by one multi-objective optimization run. These solutions each satisfy the problem objectives to different extents, and since we are interested in minimizing error values for two curvature inputs in parallel, we break the trade-off by devising a selection metric based on distance from the average error of randomly produced solutions.

A. Spherical and Cylindrical Curvatures

Spherical and cylindrical curvatures are two Gaussian curvatures (zero, cylinder; positive, sphere) that can be made using only six jamming fibers placed in a perpendicular grid pattern on the top face of the sheet [10]. We used this case study as a “sanity check” to verify our sim2real pipeline.

1) *Simulation*: Target curvatures are made by simulating the known fiber pattern solutions shown in Fig. 3a(i-ii). The best evolved designs—limited to six fibers per top/bottom face—from the three independent trials are shown in Fig. 3b-d. The low simulation errors (e_s) indicate that the evolutionary algorithm discovers new solutions that are different from the known optimal solution but use the same number of fibers. Normalizing e_s according to the maximum and minimum heights of the target curvatures, we find the percent errors range between 10% and 30%, as shown in Table II.

TABLE II
ERROR BETWEEN TARGET SHAPES & EVOLVED SIMULATED SHAPES.

Design	Cylinder Curvature		Sphere Curvature	
	e_s [mm]	% error	e_s [mm]	% error
1	0.93	16.94%	1.19	19.28%
2	0.74	9.97%	1.29	29.72%
3	0.90	11.21%	1.00	14.00%

2) *Transfer to Reality*: Design-pairs produced in simulation with the lowest e_s and most pronounced curvatures were chosen as possible solutions, and of these, Design 1 was transferred to hardware (Fig. 4). We then assessed the error between the evolved simulated shapes and the realized hardware shapes, e_t , as given in Table III. The results indicate that sheets patterned only with fibers for a single target curvature have lower errors than sheets with both sets of fibers (*i.e.* the transfer errors e_t are lowest when only the sphere pattern or only the cylinder pattern exist for sphere and cylinder curvatures, respectively). We surmise that in sheets patterned with both sets of fibers, the presence of unjammed fibers (*e.g.*, fibers for the cylinder design that remain unjammed when only the sphere design is activated) impacts the overall stiffness of the sheet, producing some general curvature not associated with the activated design. Interestingly, the error for the neat cylinder is much lower than for the neat sphere. Yet, the errors for each curvature are similar for the combined sheets.

TABLE III
ERROR BETWEEN EVOLVED SIMULATED SHAPES & HARDWARE SHAPES

Fiber Placements	Cylinder e_t	Sphere e_t
Cylinder Neat	17.1% ± 0.5%	—
Sphere Neat	—	28.6% ± 0.9%
Cylinder on Sphere	25.3% ± 0.4%	31.1% ± 1.1%
Sphere on Cylinder	27.7% ± 0.6%	33.9% ± 1.0%

We predicted that the order in which fibers are attached to the sheet would impact the fidelity of each curvature. Fibers placed on top are further from the surface of the sheet at points of overlap, which reduces their ability to translate tensile strain during jamming into constraining the

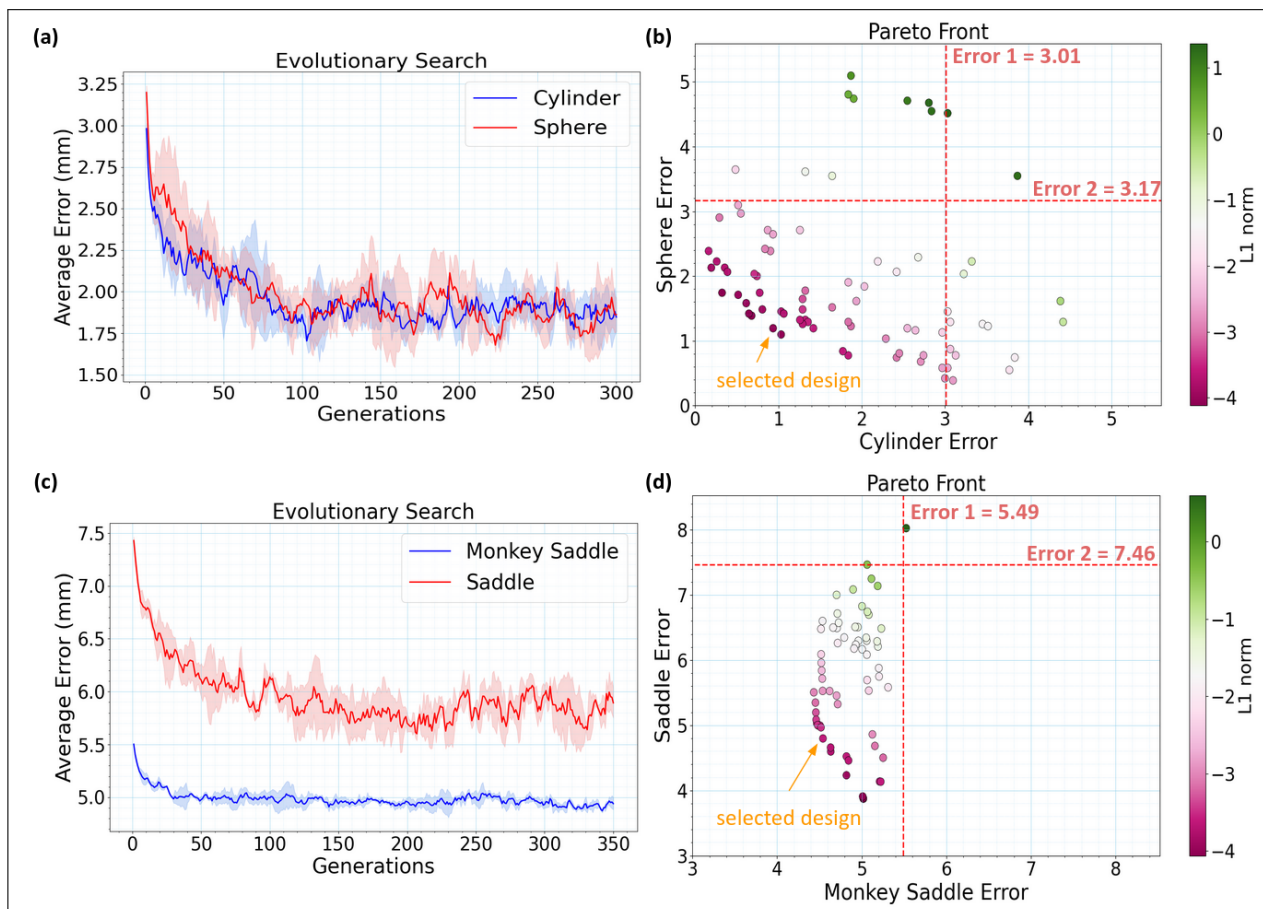


Fig. 2. **Optimization summary.** Optimization progress and the non-dominated Pareto front of the last generation for (a, b) cylinder/sphere curvatures and (c, d) saddle/monkey saddle curvatures. The evolutionary search plots (a, c) show the average errors of the population, averaged over 3 independent trials. The shaded area depicts the standard deviation of the error in the population. Each point in the scatter plots (b, d) represents one solution-pair (two independent fiber patterns) from the last generation of our optimization process. Colormap indicates the distance of the solution points from the average errors of 100 randomly generated configurations (L1 norm). As we move away from the vertical and horizontal red dotted lines (average random errors), the solutions become better at matching both targets at the same time. The selected solution from the Pareto front set is marked in orange.

bladder during inflation. Therefore, we suspected that the pattern closest to the bladder would produce a curvature with less error compared to the curvature produced by the pattern on top. However, the data show that there are negligible differences between the shape fidelity of each curvature, regardless of the pattern overlay order, suggesting that there is no need to prioritize one accessible state over another.

As noted above, the transfer error e_t for the neat cylinder is lower than for the neat sphere. The neat spherical curvature possesses areas of unwanted curvature, as indicated by the $\sim 29\%$ transfer error (compared the $\sim 17\%$ transfer error for the neat cylinder). We speculate that these areas of unwanted curvature could be redirected by the presence of additional fibers, which could be incorporated by relaxing our hardware constraints on the number, lengths, and proximity of fibers in the constrained multi-objective optimization.

We further note that simulated jammed fibers are represented as stiffer voxels embedded within the top or bottom layers of the three-layer simulated sheet (as described in §IIA), while in hardware the variable stiffness fibers are attached to the sheet's surface. This discrepancy between

simulated and hardware instantiations is another source of transfer error, which could be mitigated by simulating the fibers as separate voxels that are unique from the top and bottom faces of the simulated sheet.

Finally, we found that the sheet curvature was dependent on inflation pressure. The transfer error for the neat sphere at 2.6 psi was $30.5\% \pm 1.5\%$, at 3.0 psi was $34.5\% \pm 1.4\%$, and at 3.5 psi it was $28.6\% \pm 0.9\%$. Additionally, we observed time-dependent relaxation effects in the sheets, and we needed to wait a sufficient time between the inflation of unique shapes. The relaxation time was taken into consideration when switching between shapes during data collection and could be quantified more formally in future work.

B. Simple Saddle and Monkey Saddle Curvatures

We next investigated a more complex target shape-pair: saddle and monkey saddle. While there is a known fiber pattern solution for the simple saddle, the monkey saddle represents an unknown solution space. The results of three independent optimization runs are given in Fig. 5 and Table IV. We see that optimization discovers a diverse set of solutions with each random initialization. Inspecting Table IV, we see

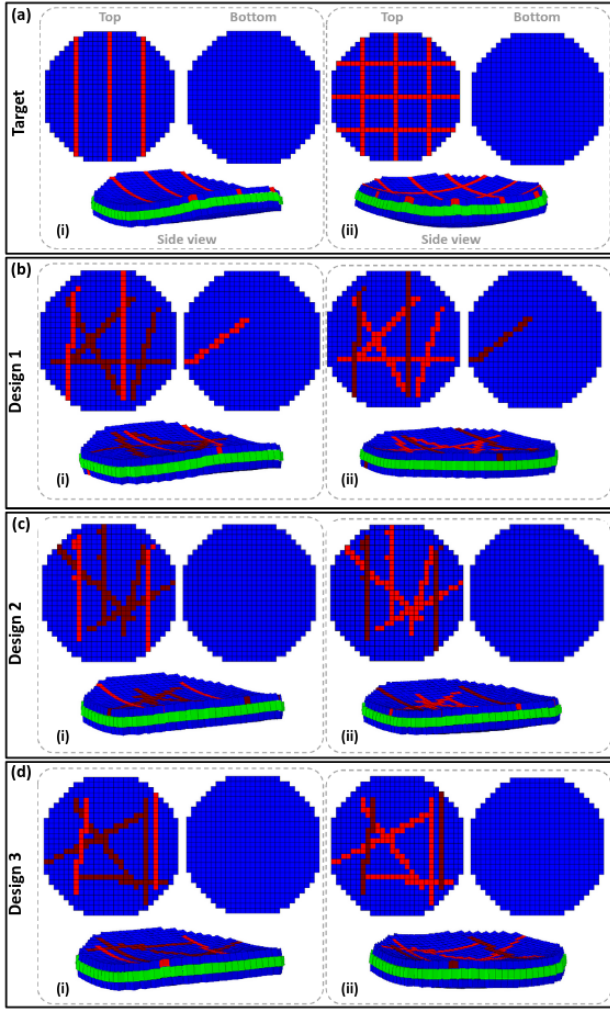


Fig. 3. **Target and evolved simulated solutions for sphere/cylinder curvatures.** Target sphere/cylinder curvatures in (a) and the three best designs for fiber placements from three independent evolutionary runs in (b-d). Fiber patterns are shown for the top (left) and bottom (right) faces of the sheet. Each line on the circular sheet represents a fiber that is jammed (light red) or unjammed (dark red). The blue voxels represent passive voxels on the top and bottom faces of the sheet. Green voxels in the middle layer are active voxels expanding during the simulation. The side view in each panel shows the inflated sheet at the end of the simulation time.

that the simple saddle generally has a lower simulation error ($e_s \sim 8\text{-}10\%$) than the monkey saddle ($e_s \sim 20\text{-}36\%$).

TABLE IV
ERROR BETWEEN TARGET SHAPES & EVOLVED SIMULATED SHAPES

Design	Monkey Saddle Curvature		Saddle Curvature	
	e_s [mm]	%error	e_s [mm]	%error
1	4.54	25.28%	4.80	10.68%
2	4.31	19.31%	4.46	8.72%
3	4.58	36.51%	4.74	8.40%

It is worth noting that although the errors seem to be high ($e_s \sim 5$ mm), the optimized fiber patterns are substantially better than a randomly generated fiber pattern (see Fig. 2b). Plotting the simulated output curvatures from evolved fiber patterns (Fig. 6a,c) and target (Fig. 6b,d) curvatures shows

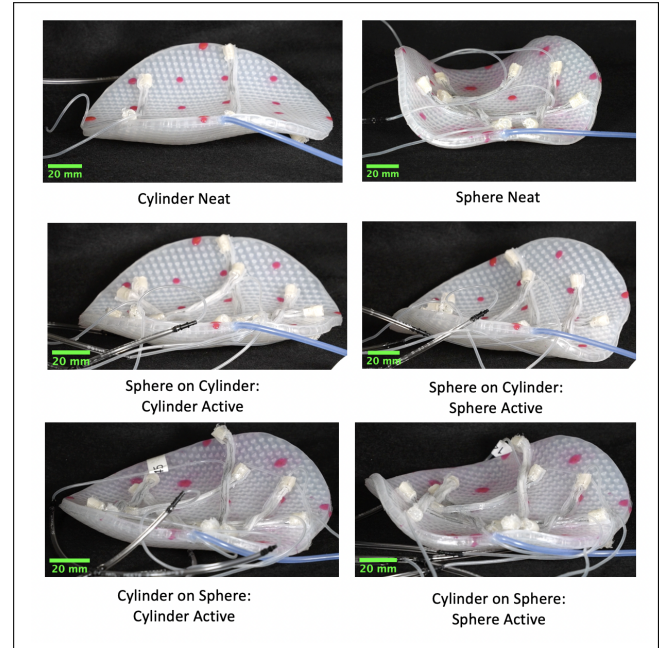


Fig. 4. **Sphere/cylinder curvatures in hardware: Design 1.** Inflated sheet in sphere/cylinder shapes, with neat and combined evolved fiber placements.

their visual similarities.

IV. CONCLUSION AND FUTURE WORK

We presented a multi-objective evolutionary algorithm to optimize variable stiffness fiber patterns on the surface of an inflatable sheet for multi-shape sheets. Leveraging the fibers' stiffness-switching ability, we evolved fiber patterns that enabled the sheet to access two unique shapes by stiffening/slackening the subsets of fibers corresponding to each shape. We demonstrated the utility of the optimization for two target shape-pairs: 1) sphere/cylinder curvatures, and 2) saddle/monkey-saddle curvatures. We showed that the evolutionary algorithm successfully produced a diverse set of designs for each case. We also transferred one design from the sphere/cylinder case study to reality and concluded that, because the optimization algorithm takes the hardware constraints into account, the transfer of evolved solutions from the Pareto front to hardware is feasible. Furthermore, we found that when combining fiber patterns on the sheet's surface to attain multi-shape access, the order of pattern overlay is not an important consideration (*i.e.*, combined fiber patterns display no dominant shape, regardless of attachment order). Future work involving the transfer of additional designs from target to simulation to hardware would deepen our understanding-of and confidence-in our design pipeline.

Because our optimization pipeline is additive—it produces independent sets of jamming fiber patterns, which are then overlaid on top of each other on the inflatable sheet—we believe it is scalable to an increasing number of target shapes and could produce patterns matching highly complex target curvatures. In the future, a post-processing step could further simplify the patterns by merging the most similar fibers (via thresholding distance, length, and orientation). A merged

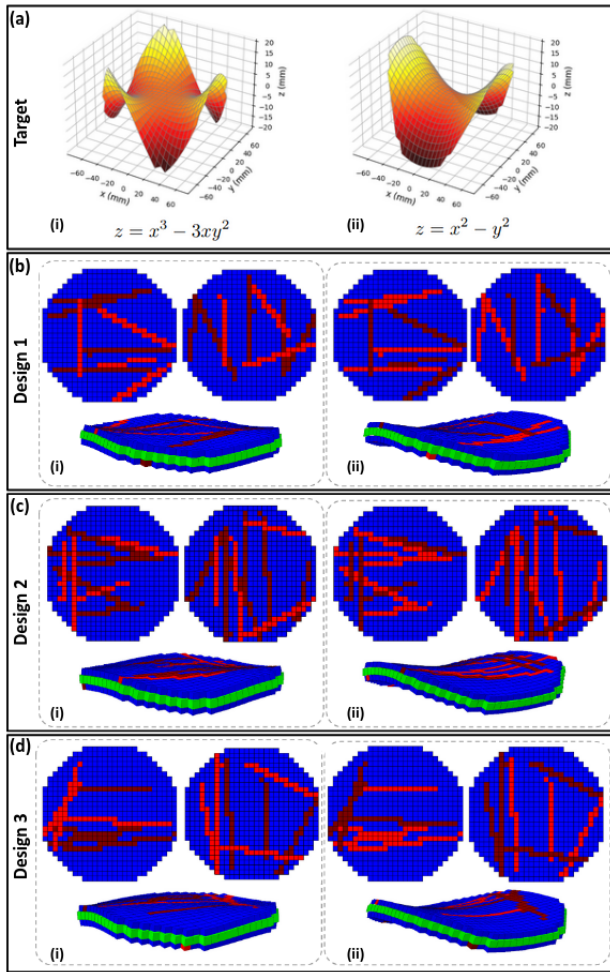


Fig. 5. **Target and evolved simulated solutions for saddle/monkey-saddle curvatures.** Target saddle/monkey-saddle curvatures in (a) and the three best designs for fiber placements from three independent evolutionary runs in panels (b-d). The targets are generated using the equations and normalized in the desired range. Fiber patterns are shown for the top and bottom faces of the sheet. Each line on the circular sheet represents a fiber that is jammed (light red) or unjammed (dark red). The side view in each panel shows the inflated sheet at the end of the simulation time.

fiber could belong to two or more fiber patterns, rather than keeping the fiber patterns mutually exclusive. Such a merging approach could reduce the total number of fibers, stay within our hardware constraints, and open access to a multiplicity of sheet curvatures that transfer from target to simulation to reality with low simulation and transfer errors.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under the Emerging Frontiers in Research and Innovation (EFRI) program (EFMA-1830870).

REFERENCES

- [1] D. Shah, B. Yang, S. Kriegman, M. Levin, J. Bongard, and R. Kramer-Bottiglio, "Shape changing robots: bioinspiration, simulation, and physical realization," *Advanced Materials*, vol. 33, no. 19, p. 2002882, 2021.
- [2] E. Reyssat and L. Mahadevan, "Hygromorphs: from pine cones to biomimetic bilayers," *Journal of the Royal Society Interface*, vol. 6, no. 39, pp. 951–957, 2009.




Fig. 6. **Optimized and target complex curvatures.** Optimized (red) and target (blue) monkey saddle shape (a-b) and simple saddle shape (c-d) for the first design from the first evolutionary trial.

- [3] S. Armon, E. Efrati, R. Kupferman, and E. Sharon, "Geometry and mechanics in the opening of chiral seed pods," *Science*, vol. 333, no. 6050, pp. 1726–1730, 2011.
- [4] W. M. Van Rees, E. Vouga, and L. Mahadevan, "Growth patterns for shape-shifting elastic bilayers," *Proceedings of the National Academy of Sciences*, vol. 114, no. 44, pp. 11 597–11 602, 2017.
- [5] A. Sydney Gladman, E. A. Matsumoto, R. G. Nuzzo, L. Mahadevan, and J. A. Lewis, "Biomimetic 4d printing," *Nature Materials*, vol. 15, no. 4, pp. 413–418, 2016.
- [6] E. Hajiesmaili, N. M. Larson, J. A. Lewis, and D. R. Clarke, "Programmed shape-morphing into complex target shapes using architected dielectric elastomer actuators," *Science Advances*, vol. 8, no. 28, p. eabn9198, 2022.
- [7] J. Pikul, S. Li, H. Bai, R. Hanlon, I. Cohen, and R. F. Shepherd, "Stretchable surfaces with programmable 3d texture morphing for synthetic camouflaging skins," *Science*, vol. 358, no. 6360, pp. 210–214, 2017.
- [8] E. Siéfert, E. Reyssat, J. Bico, and B. Roman, "Bio-inspired pneumatic shape-morphing elastomers," *Nature Materials*, vol. 18, no. 1, pp. 24–28, 2019.
- [9] B. Yang, J. Powers, A. Parsa, J. Bongard, and R. Kramer-Bottiglio, "Shape matching: Evolving fiber constraints on a pneumatic bilayer," in *2021 IEEE 4th International Conference on Soft Robotics (RoboSoft)*. IEEE, 2021, pp. 630–635.
- [10] S. Y. Kim, R. Baines, J. Booth, N. Vasios, K. Bertoldi, and R. Kramer-Bottiglio, "Reconfigurable soft body trajectories using unidirectionally stretchable composite laminae," *Nature Communications*, vol. 10, no. 1, p. 3464, 2019.
- [11] Y. Bai, H. Wang, Y. Xue, Y. Pan, J.-T. Kim, X. Ni, T.-L. Liu, Y. Yang, M. Han, Y. Huang *et al.*, "A dynamically reprogrammable surface with self-evolving shape morphing," *Nature*, vol. 609, no. 7928, pp. 701–708, 2022.
- [12] B. Yang, R. Baines, D. Shah, S. Patiballa, E. Thomas, M. Venkadesan, and R. Kramer-Bottiglio, "Reprogrammable soft actuation and shape-shifting via tensile jamming," *Science Advances*, vol. 7, no. 40, p. eabh2073, 2021.
- [13] R. Chellattoan and G. Lubineau, "A stretchable fiber with tunable stiffness for programmable shape change of soft robots," *Soft Robotics*, vol. 9, no. 6, pp. 1052–1061, 2022.
- [14] S. Liu, D. Matthews, S. Kriegman, and J. Bongard, "Voxcraft-sim, a gpu-accelerated voxel-based physics engine," <https://github.com/voxcraft/voxcraft-sim>, 2020.
- [15] J. Hiller and H. Lipson, "Dynamic simulation of soft multimaterial 3d-printed objects," *Soft Robotics*, vol. 1, no. 1, pp. 88–101, 2014.
- [16] M. D. Schmidt and H. Lipson, "Age-fitness pareto optimization," in *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, 2010, pp. 543–544.
- [17] J. E. Bresenham, "Algorithm for computer control of a digital plotter," *IBM Systems Journal*, vol. 4, no. 1, pp. 25–30, 1965.