Behavioral Repertoires for Soft Tensegrity Robots

Kyle Doney †, Aikaterini Petridou †, Jacob Karaul †, Ali Khan †, Geoffrey Liu † and John Rieffel *
* Computer Science Department
Union College
Schenectady, NY
Email: rieffelj@union.edu
† Undergraduate

***Abstract*—Mobile soft robots offer compelling applications in fields ranging from urban search and rescue to planetary exploration. A critical challenge of soft robotic control is that the nonlinear dynamics imposed by soft materials often result in complex behaviors that are counter-intuitive and hard to model or predict. As a consequence, most behaviors for mobile soft robots are discovered through empirical trial and error and hand-tuning. A second challenge is that soft materials are difficult to simulate with high fidelity - leading to a significant *reality gap* when trying to discover or optimize new behaviors. In this work we employ a Quality Diversity Algorithm running model-free on a physical soft tensegrity robot that autonomously generates a behavioral repertoire with no *a priori* knowledge of the robot's dynamics, and minimal human intervention. The resulting behavior repertoire displays a diversity of unique locomotive gaits useful for a variety of tasks. These results help provide a road map for increasing the behavioral capabilities of mobile soft robots through real-world automation.**

***Index Terms*—Soft Robotics, Quality Diversity Algorithms, Tensegrity, Evolutionary Robotics**

I. Introduction

One of the "Grand Challenges" of robotics is to "translate fundamental biological principles into engineering design rules...to create robots that perform like natural systems" [1]. In pursuit of that challenge, soft robotics seeks to use materials that more closely mimic the properties of natural systems [2]–[7]. Unfortunately, soft-materials introduce considerable elasticity and deformability, resulting in robots with nearly infinite degrees of freedom, and significant dynamical complexities. As a result, it can often be difficult to find effective controllers for soft robots, particularly mobile soft robots [2], [4], [6]. Consequently, most soft robotic locomotive behaviors are therefore developed by hand through empirical trial-and-error [6], and their actions tend to be limited to a single gait.

This paper describes methods to *autonomously* discover novel and effective soft robotic locomotive behaviors without the biases and limitations of these human-in-the-loop approaches. Moreover, we demonstrate that these methods can be run on a physical robot, without the need for simulation or modeling. Specifically we show how Quality Diversity Algorithms (QDAs) can autonomously discover a diverse *repertoire* of unique locomotive behaviors for a tether-free soft tensegrity robot. The discovered behaviors could, in turn, be incorporated into higher-level control strategies. These methods offer to advance the state of the art in soft robotics by allowing researchers to more fully exploit the dynamical complexities of soft robots – leading to more versatile and robust mobile soft robots. To the best of our knowledge, no other mobile soft robot has demonstrated as wide a variety of behaviors, much less autonomously generated behaviors.

978-1-7281-2547-3/20/$31.00 ©2020 IEEE

II. Background and Related Work

The inherent complexities of soft materials makes soft robots incredibly difficult to simulate with sufficient fidelity to address the reality gap [8]. Several soft robotic simulators exist, including VoxCraft [9] and SOFA [10], both of which have produced transferable simulations only in very limited contexts [11], [12]. The process of discovering locomotive behaviors for soft robots is often accomplished through empirical trial-and-error on physical robots [6], a method that can be both challenging and time consuming. As a consequence, the ensuing behaviors tend to either be slow, and quasi-static, or else largely uncontrolled [4]–[6], [13]. The use of *ad hoc* and hard-coded controllers also means that these soft robots are generally unable to autonomously adapt to internal or external changes, for instance when they are physically damaged, or when they encounter a unknown terrain. Moreover, many existing mobile soft robots tend to be designed only to perform one task, such as crawling or swimming, and exhibit little or no steering, turning, or any wide diversity of behaviors.

A. Tensegrity-Based Soft Robots

Tensegrities (Figure 1) are relatively simple mechanical systems, consisting of a number of rigid elements (struts) joined at their endpoints by tensile elements (cables or springs), and kept stable through a synergistic interplay of pre-stress forces. At every scale, tensegrity structures have an impressive strength-to-weight ratio, are structurally robust, and stable in the face of deformation [14], [15]. Moreover, by choosing appropriate spring constants they can be used to construct *soft* robots with a high dimensional configuration space [16], high compliance [17] and capable of rapid shape change and tunable stiffness [18]. Unlike many other soft robots, tensegrity structures are inherently modular and are therefore, in principle, relatively easy to construct.

As robots, tensegrity-based robots are typically controlled by changing the lengths of the tensile strings, causing large-scale deformations of the structure, which in turn cause the robot to step or tumble [19]–[23], in a method NASA calls

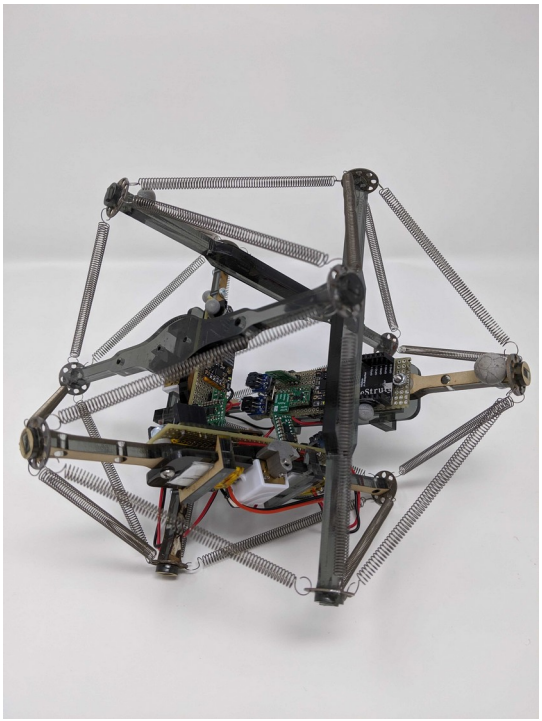

Fig. 1. Tensegrity Robots exhibit many of the complex behaviors of more fully soft robots. This robot, consists of 3 struts vibrational motors, 3 passive struts without motors, and 18 springs.

'flop and roll,'. More recently, researchers have begun to explore more dynamic methods of control. Bliss *et al.* for instance [24], [25] use central pattern generators (CPGs) to control tensegrity-based swimming robots. Compellingly, they were able to demonstrate a control technique based upon the resonant *entrainment* of the robot. Mirletz *et al.* have used CPGs to produce goal-directed behavior in simulated tensegrity-spine-based robots [26]. More recently, Böhm and Zimmermann developed a tensegrity-inspired robot actuated by an single oscillating electromagnet [27].

The appeal of using tensegrities as the basis for soft robotics research is that they are relatively inexpensive to build, they are inherently modular, and, like more fully soft robots they are capable of a wealth of dynamic behaviors including crawling, walking, tumbling and hopping. The challenge of tensegrity robots, like soft robots, lies in finding ways to explore and exploit these behaviors.

In prior work, we describe the design of a relatively inexpensive and tether-free soft tensegrity robot [28] (Fig. 1). In contrast to many other soft robotic platforms, which require tethers for pneumatic or electrical power, this robot is capable of autonomous and tether-free operation – allowing for a wide range of behaviors without the risk of tangling a tether. The low cost, modularity, and ease of assembly help lower the barrier to entry for soft robotics research.

Earlier work of ours also demonstrates how machine learning algorithms can be used to optimize gaits for the robot across level terrain [28]–[30]. During this prior work we observed that the robot is actually capable of a wide variety of translational and rotational behaviors and that the nonlinearity of the system means that the changes are consistent across trials and yet non-intuitive. The diversity of these behavioral modalities far exceeds those typically seen in soft robots. It is this last insight that guides our inquiry: how best to explore and exploit the entire range of a soft robot's behavioral diversity.

B. Quality Diversity Algorithms

Quality Diversity algorithms (QDAs) [31]–[34] search for *novelty* alongside *optimality* in a given space. When applied to robotics, this can mean generating mappings between a robot's *parameter space* (those features that can be directly affected by a controller), and its *behavior space* (a description of the outcomes of those actions), as illustrated by Figure 2. The outcome of a successful quality diversity search is a *behavioral repertoire* describing a diverse variety of reachable behaviors that the robot is capable of – essentially a mapping between parameter space and behavior space. QDAs have been employed with considerable success with both simulated virtual agents [32], [33] and physical robots [34]. Most notably, Cully *et al.* [34] used a hybrid simulation/physical approach to behavior build repertoires for a hexapod robot.

However, most soft robots, including our tensegrity robots, are in principle capable of a considerably wider range of behaviors than those exhibited by the systems in the above work, due to their soft material composition and wide range of degrees of freedom. Moreover, the complexity of soft systems, including tensegrities, means that repertoires cannot be developed in simulation and then transferred into physical system with any degree of expected fidelity [8]. In this work, we show how quality diversity algorithms can be effectively leveraged to generate behavioral repertoires for soft robotics operating in the real world.

III. Behavioral Repertoires for Tensegrity Robots

In the context of our tensegrity robots, we use the MAP-Elites quality diversity algorithm to generate a behavioral repertoire for the tensegrity robot [34]. In particular, the tensegrity robot has three vibrational motors each of which can be assigned a frequency, and so a set of parameters $p \in P$ (referred to as Map of parameters P in Algorithm 1) can be expressed as $p = (f_1, f_2, f_3)$ where $f_i \in Z$ (integers) and $255 \geq f_i \geq 0$. A locomotive behavior of the robot is described by the robot's displacement in the Cartesian (x/y) coordinate plane and its rotation around the z axis (yaw) over an evaluation period. Formally a behavior b in the behavior space B to equal $(\Delta X, \Delta Y, \Delta\Psi)$ where ΔX and ΔY are within the range -360mm:360mm, and $\Delta\Psi$ is within the range -180°:180°. Reasonable bounds on the limits of the behavior space were determined empirically based on the displacements achieved in prior experimentation and are equivalent to 3.5 body-lengths of our robot.

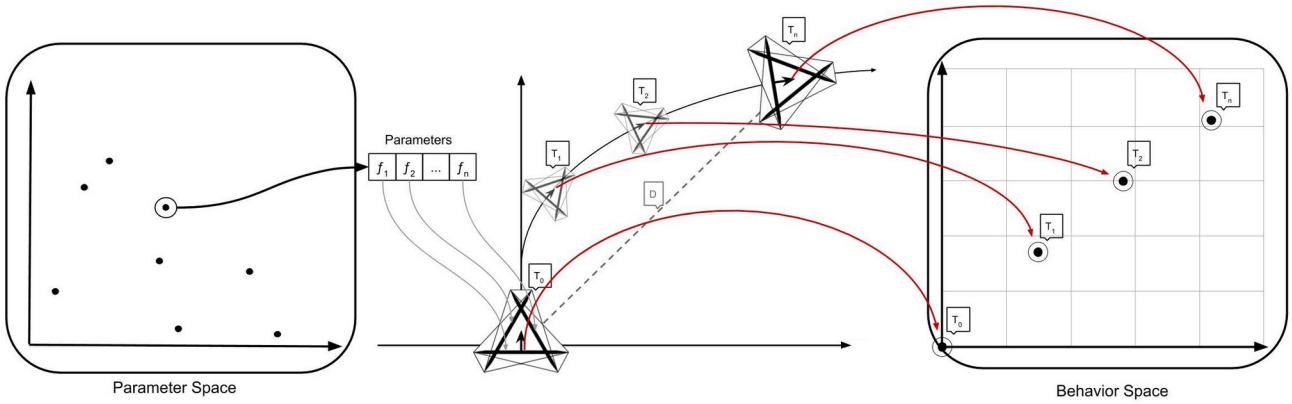

Fig. 2. Quality Diversity algorithms generate a mapping between parameter space P and behavior space B. Typically, the behavior space is discretized into a set of bins, and when necessary a *quality metric* is used to compare behaviors in said bins.

As an example of this behavior parameterization, a set of motor parameters $p = (189, 30, 251)$ that caused the robot to displace 50mm along the x axis and 25mm along the y axis, without rotating, would result in $b = (50, 25, 0)$. And a robot that spun 90° without displacement would result in $b = (0, 0, 90^{\circ})$.

When discretizing the behavior space into intervals (bins), we address the stochasticity of our robot by choosing a interval size based on standard deviation testing. When considering the usefulness of a behavior repertoire, it is important to maximize the probability that if set p generated behavior b in bin x, then any further evaluating of p also generate $b^0 \in x$. This repeatability testing consisted of re-running the same set $p \in P$ on the robot. For each p, we ran 30 tests: 10 five second tests, 10 ten second tests, and 10 fifteen second tests. We tested 15 parameter sets that represented a diverse selection of the behavior space. Based on the results of this testing, we discretized B into 864 bins. The x displacement and y displacement axes of B were discretized into 12 bins at 6cm intervals, and the ψ displacement axis B was discretized into 6 bins at 60 degree intervals. We also determined that 10 second trials were optimal to maximize repeatability of behaviors, while also minimizing trial length.

At a high-level the MAP-Elites generates the behavior repertoire P in two phases. The first phase consists of an initial random sampling of the parameter space, where sets of parameters $p \in P$ are randomly selected and evaluated on the robot. The resulting behaviors (the elites) are stored as mappings in the behavior repertoire P. The second phase of the algorithm selects an elite previously stored $\in P$, slightly mutates the parameter set associated with the elite, then evaluates the new parameter set on the robot.

If in either phase a set of parameters is evaluated where multiple behaviors occupy the same bin, a fitness function (quality metric) is employed to select the "better" behavior. Our *fitness metric* $f(b) = |\Delta x|+|\Delta y|+|\Delta\varphi|$ selects competing behaviors within a niche with higher absolute displacement, indicating a preference for behaviors at the more extreme end of each

Algorithm 1 MAP-Elites, Modified from [35]

function MAP-ELITES ($f_{itness}()$, $variation()$, $X_{initial}$)
  $P \leftarrow \varnothing$, $F \leftarrow \varnothing$ · *Map of params P, and fitnesses F*
  $B_{initial} \leftarrow behavior\_descriptor(P_{initial})$
  $P(B_{initial}) \leftarrow P_{initial}$ · *Place initial solutions in map*
  $F(B_{initial}) \leftarrow f_{itness}(B_{initial}))$
  for iter = $1 \rightarrow I$ **do**
    $\mathbf{p}^0 \leftarrow variation(P)$. *Create new solns from elites*
    $\mathbf{b}^0 \leftarrow behavior\_descriptor(\mathbf{p}^0)$
    $\mathbf{f}^0 \leftarrow f_{itness}(\mathbf{p}^0)$
      · *Replace if better*
    if $F(\mathbf{b}^0) = \varnothing$ or $F(\mathbf{b}^0) < \mathbf{f}^0$ **then**
      $F(\mathbf{b}^0) \leftarrow \mathbf{f}^0$
      $P(\mathbf{b}^0) \leftarrow \mathbf{p}^0$
  return (P, F) · *Return illuminated map*

behavioral niche. In other words, f optimizes each behavior bin $\in \mathsf{B}$ so that the most active behaviors of the robot are not discarded. All three components of b have be normalized to the range 0-1, so they factor equally into the *fitness metric*.

As opposed to a method such as random selection, by employing this fitness function the algorithm ensures that the more time-efficient behaviors are kept. In other words, behaviors which displace the robot more over the same time-period are inherently more interesting from a use-ability standpoint. While the primary goal of the QDA in this instance is to build a repertoire of diverse behaviors (fill many unique bins on the map), the quality of each kept behavior is still important. During automatic behavior generation it is much preferred that each of locomotive behaviors in the constructed behavior set is the optimal behavior (the one with the greatest displacement) in that grid section, as this behavior maximizes robot displacement over time.

The displacement of the robot was tracked using a 20-camera Qualisys Oqus 700+ system and the QTM Tracking Software at a frame rate of 300 frames per second. Four tracking markers placed on each of three of the tensegrity

struts to provide 6DOF position and rotation data for each of the marked struts. This tracking setup can been seen in Figure 3. Global reference frame tracking data was then converted into the robot's local reference frame in order to produce a behavior descriptor as described above. The "front" of the robot, corresponding to a global reference frame yaw of 0 is essentially arbitrary (because the robot is rotationally symmetric), and so was determined by using the reported Euler angle of one of the tracked struts.

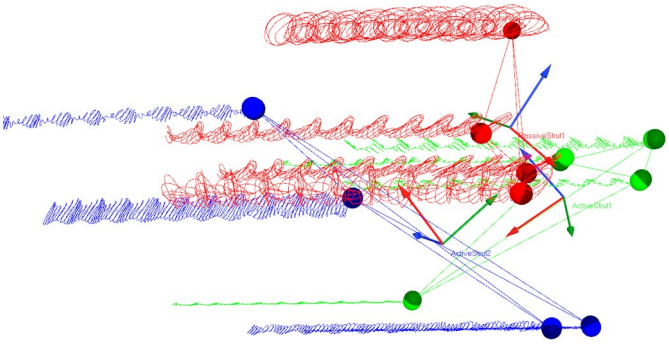

Fig. 3. A 3D rigid body model close-up view of the dynamics of the tensegrity robot's locomotion. Each dot represents a marker being tracked with sub-millimeter accuracy on the tensegrity robot. The resonances imposed by the three vibrating motors can be clearly seen.

IV. Experiment Design

This experiment seeks to quantitatively determine if the QDA MAP-Elites is a data-efficient method for illuminating the behavior space of a tensegrity robot. This is done by comparing MAP-Elites illumination against a pseudo-random control. A control computer running the MAP-Elites algorithm was used to generate frequencies that were sent to the robot's struts' Arduino controllers via Bluetooth. Details of the construction of the robot, and its wireless control are provided in [28]. An initial 100 random parameter trials were run (MAP-Elites phase 1). These trials were shared between both the control experiment and MAP-Elites experiment to ensure an equal baseline. If this was not done, the randomness of these initial seed trials could more greatly sway how well MAP-Elites performs in comparison with the control. Each time a unique behavior is found, it becomes slightly more difficult to find the next novel behavior, so if one experiment continued with 25 discovered behaviors, where the other continued with 60 behaviors, this would heavily bias the results.

The behavior repertoire generated from the shared 100 trials is used as the initial behavior repertoire for the second half of the experiment. The random control evaluates another 400 pseudo-randomly generated parameters on the robot. The non-control experiment continues by running the mutation phase of MAP-Elites (phase 2). This also evaluates another 400 sets of parameters on the tensegrity robot. A visual depiction of this is displayed in Figure 4.

Overall, the whole experiment required 900 robot trials, each of which consisted of 10 seconds of motor run-time, and an additional amount of cooldown time. This equates to approximately 3.5 hours of operating time. It is important to note that the number of trials for both sections of this experiment was empirically chosen. From some prior testing, it appears that increasing the number of trials during the initial random phase does not significantly change the outcome of the experiment. More work is needed to determine a lower-bound for the number of required initial random trials. Such a lower-bound number could be based on the percentage of bin's occupied in the behavior space.

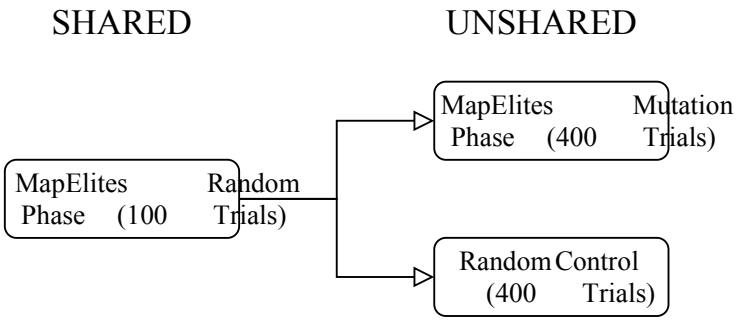

Fig. 4. This is a visual flowchart of the MAP-Elites vs. Control experiment. An initial amount of random trials are shared as to generate the initial behavior repertoire for MAP-Elites. This allows the mutation phase of the MAP-Elites algorithm to be directly compared with the random search.

V. Results

At the conclusion of our behavioral repertoire exploration our algorithm discovered 110 distinct behaviors during the non-control section of the experiment. This represents approximately 13% of the discritized behavior space. The same amount of control trials discovered only 71 distinct behaviors or 8% of the behavior space. To fully understand to what extent the MAP-Elites algorithm out performed the random control, we must compare specifically the unshared sections of the experiment. In these trials MAP-Elites discovered 83 unique behaviors, compared to the control discovering 44. In other words, the mutation phase of MAP-Elites found unique behaviors twice as fast as the control. The average quality of the distinct behaviors found by MAP-Elites is statistically significantly higher than the average quality of control behaviors. What this means is that the behaviors generated by the MAP-Elites mutation have on average larger displacement values. Table I summarizes the numeric data generated by this experiment.

Figure 5 illustrates the diversity of the behaviors produced. Not surprisingly, many of the behaviors discovered within the repertoire produce minimal motion: motor frequency parameter sets p can contain very low motor speeds that produce little or no displacement or rotation of the robot. This is somewhat mitigated by only running the experiment if a threshold value is reached. If the three motor speeds added together do not exceed the threshold value, then it is assumed the robot will remain stationary. Performing more complex thresholding to remove more parameters which produce minimal motion

Trial Sets:	Initial 100 Random (Phase 1)	Mutation 400 (Phase 2)	Random 400 (Phase 2)	Mutation 500 (Phase 1 + 2)
Number of Unique Behaviors	27	83	44	110
Avg. Fitness of Unique Behaviors	.67	.97	.87	0.90
Avg. Fitness of all Trials	.32	.75	.36	0.66

TABLE I

A SUBSET OF THE DATA COLLECTED FROM THE MAP-ELITES VS. RANDOM SEARCH EXPERIMENT. IT INCLUDES THE NUMBER OF UNIQUE BEHAVIORS (NUMBER OF BINS OCCUPIED IN THE BEHAVIOR SPACE) FOUND DURING EACH SECTION OF THE EXPERIMENT, AND THE AVERAGE QUALITY OF THOSE TRIALS, BASED ON THE FITNESS METRIC.

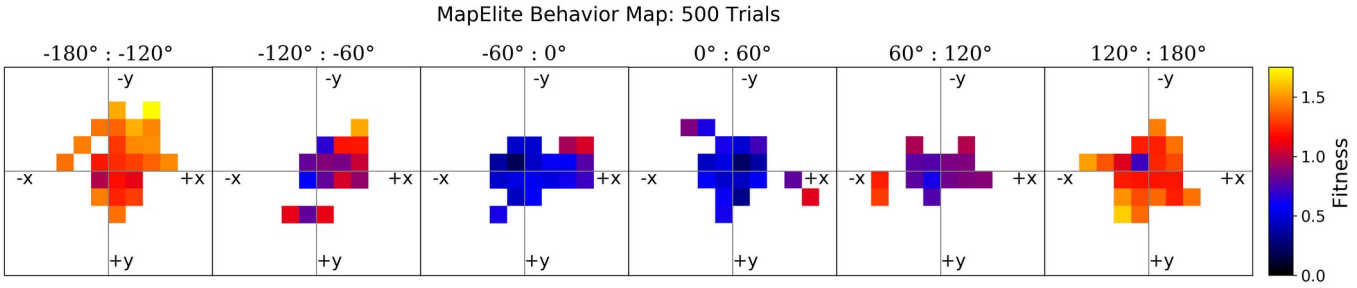

Fig. 5. A behavior map showing the diversity of behaviors discovered by the non-control MAP-Elites algorithm. Each graph displays the Δx and Δy displacements for all the behaviors within a specific rotational φ bin of the behavior map. The furthest left graph, for instance, illustrates all the behaviors discovered for which the robot rotated between -120° and -180° (negative rotations are anti-clockwise from initial position), and the adjacent graph illustrates behaviors for which the robot rotated between -60° and -120°. The color of a square maps to the fitness value of the behavior in that bin.

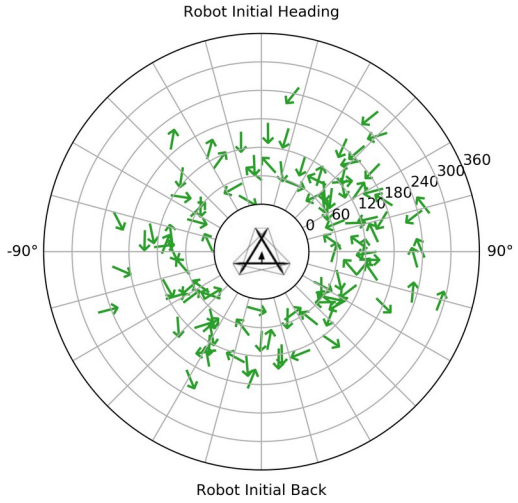

Fig. 6. A top-down view of the behavior spaced overlaid in the robot's coordinate frame. Each green arrow represents the location and orientation of the robot after a behavior $b \in \mathsf{B}$ is run for 10 seconds. This figure shows the behaviors discovered by the non-control MAP-Elites algorithm.

would be exceedingly difficult due to the counter-intuitive nature of tensegrity dynamics.

The most qualitatively interesting behaviors are less common: those behaviors that result in large $\Delta x \Delta y$ displacements are found somewhat more often at the outer edges of the high (positive or negative) rotation angle graphs (those being the leftmost and rightmost graphs of Figure 5). It might be expected that those behaviors that translate the robot the furthest are less likely to rotate the robot, and vice versa, but this is not the case: many of the behaviors that produce a large translation also produce a correspondingly large rotation.

Another way to visualize this data is displayed in Figure 6. The tail of each arrow indicates the displacement of a specific behavior within its coordinate frame, and the orientation of the arrow indicates the behavior's rotation $\Delta\varphi$. The upper right most arrow, for instance, illustrates a behavior that caused the robot to displace forwards just around 240mm, and rotate approximately $+40^\circ$ (clockwise). What this visualization demonstrates is that distinct behaviors were discovered by our algorithm which move the robot in all 360 degrees. This demonstrates how behavior repertoire can be very easy converted into a robotic controller, as we already have known parameter sets which move the robot in essentially all directions.

Figure 7 provides examples of some of the behaviors discovered during our search. Each behavior is both quantitatively and qualitatively distinct. The colored tracks on each image show the rigid model overlay for the three tracked struts. The four markers per strut all share a common color. The two behaviors on the left hand side are both high rotation behaviors, with relatively low translation on the plane. The behaviors on the right, by contrast, produce linear translation (although along different vectors), with minimal rotation.

It is worth leveraging the 300fps frame rate and sub-millimeter accuracy of the Qualisys motion capture system in order to scrutinize the details of the robot's motion in high resolution. Figure 3 provides this. The three motors in this gait are spinning at different frequencies, and we can see the periodic trace of the how those tracked points on the tensegrity robot vibrate.

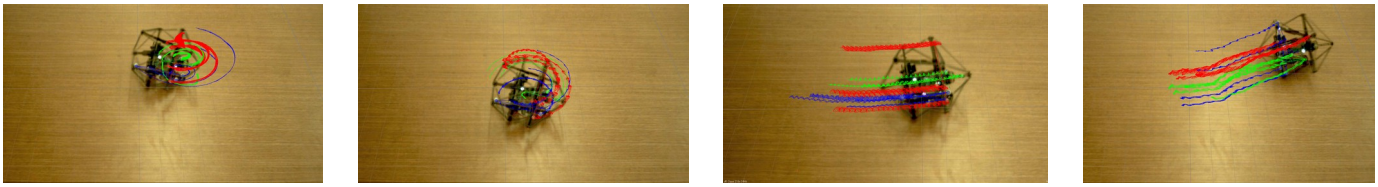

Fig. 7. Overhead camera views and rigid body model overlays of several of the diverse behaviors discovered by our QDA. In the two left-hand figures, the robot rotates with minimal displacement in the x/y plane. In the two right-hand images, the robot translates with relatively little rotation.

VI. Discussion

Our ambition in this paper is to demonstrate that Quality Diversity Algorithms (QDAs) running on real-world soft robots can, with little *a priori* knowledge, autonomously illuminate and exploit the tremendous diversity exhibited by soft robots. The range and variety of behaviors discovered by these methods are significantly more dynamic and diverse than those demonstrated by other mobile soft robots.

Beyond the dynamics and algorithm, a design detail of our robot that facilities this diversity is the fact that unlike almost all other soft robots it is wireless and completely tether-free, unlike many mobile soft robots that are pneumatically actuated and tethered to externap pressure and electrical sources, and thereby limited in their ability to roll, twist, and turn in interesting ways. Because our robot is entirely tether free and controlled via Bluetooth, we are able to explore a wider range of behaviors (indeed, earlier tethered versions of our robot became quite tangled in their tethers when rotating [29]).

An added value of this process is that the automated nature of our algorithm is not confined by human intuition when exploring the behavior space, and is therefore significantly less fettered by human biases - for instance a preference for ”forwards” behaviors (indeed our robot has no pre-defined front). Small changes to the motor speeds of our robot can produce profound and unexpected changes in behavior – for instance a switch from forward locomotion mode to a purely rotational mode. Recent work of ours optimized a single behavior such as forward locomotion [28], [36], however, these methods by nature discard interesting behaviors that are non-optimal in regards to a single fitness function but that might otherwise be useful in some other context – for instance rotation, or a behavior that mixes rotation and translation. Indeed, in our earlier approaches to optimizing locomotion we discarded any behavior that produced any gaits that happened to also rotate the robot outside of a narrow window of angles. This QDA approach, by contrast, allows for a balance between a diversity of behaviors (across different discretized bins of the behavior space) and optimization of behaviors (within bins).

This raises the question of how exactly to make use of this repertoire. Fortunately recent work on QDAs on more conventional robots in both simulation and the real world show how behavioral repertoires can be incorporated into high-level goal oriented tasks through methods such as policy networks [32] or Monte Carlo Tree Search [37]. Our interest going forward, therefore, is to work toward applying these hybrid QDA-planning methods toward the signicantly more complex behavior space of soft robots.

VII. Conclusions

In this work we have illustrated how Quality Diversity Algorithms (QDAs) can be used to replace hand-tuning with automated locomotive behavior discovery. Moreover, we demonstrated their utility on physical robots, without the need for low-fidelity soft robotic simulators and with minimal human intervention. The QDA can relatively efficiently explore the large and complex search space of possible locomotive behaviors for the given soft robot. The behavioral repertoire developed through this process shows a diversity rarely seen in other mobile soft robots, which tend to focus on one particular task domain such as forward locomotion. The long-term outcome of this thread of inquiry will lead to new methods by which increasingly complex soft robots can not only outperform conventional robots across a range of environments, but also learn and resiliently adapt to changing conditions and damage.

Many soft organisms in nature quite effectively leverage their material properties: *Manduca sexta* caterpillars, for instance, have a mid-gut which acts like a “visceral-locomotory piston” to assist in locomotion [38]. Jellyfish use their elasticity in order to recover energy during swimming [39]. Understanding how to exploit soft material properties, and developing tools to explore their dynamics are both critical tasks for advancing the state of the art. Our results are therefore an important step toward the soft robotic grand challenges of developing engineered systems that are are capable of the type of behavioral complexity seen in the natural world.

Ackowledgements

This material is based upon work supported by the National Science Foundation under Grant Nos. NSF-1939930 and NSF-1827495.

References

[1] G.-Z. Yang, J. Bellingham, P. E. Dupont, P. Fischer, L. Floridi, R. Full, N. Jacobstein, V. Kumar, M. McNutt, R. Merrifield, B. J. Nelson, B. Scassellati, M. Taddeo, R. Taylor, M. Veloso, Z. L. Wang, and R. Wood, “The grand challenges of science robotics,” *Science Robotics*, vol. 3, no. 14, 2018. [Online]. Available: http://robotics.sciencemag.org/content/3/14/eaar7650

[2] H. Lipson, “Challenges and opportunities for design, simulation, and fabrication of soft robots,” *Soft Robotics*, vol. 1, no. 1, pp. 21–27, 2014.

[3] M. Wehner, R. L. Truby, D. J. Fitzgerald, B. Mosadegh, G. M. Whitesides, J. A. Lewis, and R. J. Wood, “An integrated design and fabrication strategy for entirely soft, autonomous robots,” *Nature*, vol. 536, no. 7617, pp. 451–455, 2016.

[4] B. A. Trimmer, "New challenges in biorobotics: incorporating soft tissue into control systems," *Applied Bionics and Biomechanics*, vol. 5, no. 3, pp. 119–126, 2008.

[5] H.-T. Lin, G. Leisk, and B. Trimmer, "Goqbot: A caterpillar-inspired soft-bodied rolling robot," *Bioinspiration and Biomimetics*, vol. 6, p. 026007, 2011.

[6] R. F. Shepherd, F. Ilievski, W. Choi, S. A. Morin, A. A. Stokes, A. D. Mazzeo, X. Chen, M. Wang, and G. M. Whitesides, "Multigait soft robot," *Proceedings of the National Academy of Sciences*, vol. 108, no. 51, pp. 20 400–20 403, 2011.

[7] R. V. Martinez, J. L. Branch, C. R. Fish, L. Jin, R. F. Shepherd, R. Nunes, Z. Suo, and G. M. Whitesides, "Robotic tentacles with three-dimensional mobility based on flexible elastomers," *Advanced Materials*, vol. 25, no. 2, pp. 205–212, 2013.

[8] N. Jakobi, P. Husbands, and I. Harvey, "Noise and the reality gap: The use of simulation in evolutionary robotics," in *European Conference on Artificial Life*. Springer, 1995, pp. 704–720.

[9] 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.

[10] C. Duriez, "Control of elastic soft robots based on real-time finite element method," in *2013 IEEE international conference on robotics and automation*. IEEE, 2013, pp. 3982–3987.

[11] S. Kriegman, A. M. Nasab, D. Shah, H. Steele, G. Branin, M. Levin, J. Bongard, and R. Kramer-Bottiglio, "Scalable sim-to-real transfer of soft robot designs," in *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2020, pp. 359–366.

[12] S. E. Navarro, S. Nagels, H. Alagi, L.-M. Faller, O. Goury, T. Morales-Bieze, H. Zangl, B. Hein, R. Ramakers, W. Deferme *et al.*, "A model-based sensor fusion approach for force and shape estimation in soft robotics," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5621–5628, 2020.

[13] N. W. Bartlett, M. T. Tolley, J. T. Overvelde, J. C. Weaver, B. Mosadegh, K. Bertoldi, G. M. Whitesides, and R. J. Wood, "A 3d-printed, functionally graded soft robot powered by combustion," *Science*, vol. 349, no. 6244, pp. 161–165, 2015.

[14] K. Snelson, "The art of tensegrity," *International Journal of Space Structures*, vol. 27, no. 2-3, pp. 71–80, 2012.

[15] R. E. Skelton and M. C. de Oliveira, *Tensegrity systems*. Springer, 2009, vol. 1.

[16] A. Tibert and S. Pellegrino, "Review of form-finding methods for tensegrity structures," *International Journal of Space Structures*, vol. 26, no. 3, pp. 241–255, 2011.

[17] V. Böhm, S. Sumi, T. Kaufhold, and K. Zimmermann, "Compliant multistable tensegrity structures with simple topologies," in *New Trends in Mechanism and Machine Science*. Springer, 2017, pp. 153–161.

[18] A. Tibert and S. Pellegrino, "Deployable tensegrity reflectors for small satellites," *Journal of Spacecraft and Rockets*, vol. 39, no. 5, pp. 701–709, 2002.

[19] C. Paul, F. J. Valero-Cuevas, and H. Lipson, "Design and control of tensegrity robots for locomotion," *IEEE Transactions on Robotics*, vol. 22, no. 5, 2006.

[20] A. Iscen, A. Agogino, V. SunSpiral, and K. Tumer, "Controlling tensegrity robots through evolution," in *Proceedings of the 15th annual conference on Genetic and evolutionary computation*. ACM, 2013, pp. 1293–1300. [Online]. Available: https://dl.acm.org/citation.cfm?id=2463525

[21] ——, "Flop and roll: Learning robust goal-directed locomotion for a tensegrity robot," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*. IEEE, 2014, pp. 2236–2243. [Online]. Available: http://ieeexplore.ieee.org/abstract/document/6942864/

[22] A. Agogino, V. SunSpiral, and D. Atkinson, "Super ball bot-structures for planetary landing and exploration," *NASA Innovative Advanced Concepts (NIAC) Program, Final Report*, 2013.

[23] D. Zappetti, S. Mintchev, J. Shintake, and D. Floreano, "Bio-inspired tensegrity soft modular robots," in *Biomimetic and Biohybrid Systems: 6th International Conference, Living Machines*, 2017, pp. 497–508.

[24] T. K. Bliss, T. Iwasaki, and H. Bart-Smith, "Resonance entrainment of tensegrity structures via cpg control," *Automatica*, vol. 48, no. 11, pp. 2791–2800, 2012.

[25] T. Bliss, T. Iwasaki, and H. Bart-Smith, "Central pattern generator control of a tensegrity swimmer," *IEEE/ASME Transactions on Mechatronics*, vol. 18, no. 2, pp. 586–597, 2013.

[26] B. T. Mirletz, P. Bhandal, R. D. Adams, A. K. Agogino, R. D. Quinn, and V. SunSpiral, "Goal-directed cpg-based control for tensegrity spines with many degrees of freedom traversing irregular terrain," *Soft Robotics*, vol. 2, no. 4, pp. 165–176, 2015.

[27] V. Böhm and K. Zimmermann, "Vibration-driven mobile robots based on single actuated tensegrity structures," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 5475–5480.

[28] J. Kimber, Z. Ji, A. Petridou, T. Sipple, K. Barhydt, J. Boggs, L. Dosiek, and J. Rieffel, "Low-cost wireless modular soft tensegrity robots," in *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2019, pp. 88–93.

[29] M. Khazanov, J. Jocque, and J. Rieffel, "Evolution of locomotion on a physical tensegrity robot," in *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, 2014, pp. 232–238.

[30] M. Khazanov, B. Humphreys, W. Keat, and J. Rieffel, "Exploiting dynamical complexity in a physical tensegrity robot to achieve locomotion," *Advances in Artificial Life, ECAL*, vol. 12, pp. 965–972, 2013.

[31] J. K. Pugh, L. B. Soros, and K. O. Stanley, "Quality diversity: A new frontier for evolutionary computation," *Frontiers in Robotics and AI*, vol. 3, p. 40, 2016. [Online]. Available: https://www.frontiersin.org/article/10.3389/frobt.2016.00040

[32] M. Duarte, J. Gomes, S. M. Oliveira, and A. L. Christensen, "Evolution of repertoire-based control for robots with complex locomotor systems," *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 2, pp. 314–328, 2018.

[33] J. Lehman and K. O. Stanley, "Evolving a diversity of virtual creatures through novelty search and local competition," in *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. ACM, 2011, pp. 211–218.

[34] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, "Robots that can adapt like animals," *Nature*, vol. 521, no. 7553, pp. 503–507, 2015.

[35] J.-B. Mouret and J. Clune, "Illuminating search spaces by mapping elites," *arXiv:1504.04909 [cs, q-bio]*, Apr. 2015, arXiv: 1504.04909. [Online]. Available: http://arxiv.org/abs/1504.04909

[36] J. Rieffel and J.-B. Mouret, "Adaptive and resilient soft tensegrity robots," *Soft robotics*, vol. 5, no. 3, pp. 318–329, 2018.

[37] K. Chatzilygeroudis, V. Vassiliades, and J.-B. Mouret, "Reset-free trial-and-error learning for robot damage recovery," *Robotics and Autonomous Systems*, vol. 100, pp. 236–250, 2018.

[38] M. A. Simon, W. A. Woods Jr, Y. V. Serebrenik, S. M. Simon, L. I. van Griethuijsen, J. J. Socha, W.-K. Lee, and B. A. Trimmer, "Visceral-locomotory pistoning in crawling caterpillars," *Current biology*, vol. 20, no. 16, pp. 1458–1463, 2010.

[39] B. J. Gemmell, J. H. Costello, S. P. Colin, C. J. Stewart, J. O. Dabiri, D. Tafti, and S. Priya, "Passive energy recapture in jellyfish contributes to propulsive advantage over other metazoans," *Proceedings of the National Academy of Sciences*, vol. 110, no. 44, pp. 17 904–17 909, 2013.