Discovering A Library of Rhythmic Gaits for Spherical Tensegrity Locomotion

Colin Rennie and Kostas E. Bekris

Abstract— **Tensegrity robots, which combine both rigid and soft elements, provide exciting new locomotion capabilities but introduce significant control challenges given their high-dimensionality and non-linear nature. This work first defines an effective parameterization of a spherical tensegrity for generating rhythmic gaits based on Central Pattern Generators (CPG). This allows the definition of periodic and rhythmic control signals, while exposing only five gait parameters. Then, this work proposes a framework for optimizing such gaits by exploring the parameter space through Bayesian Optimization on an underlying Gaussian Process regression model. The objective is to provide gaits that allow the platform to move along different directions with high velocity. Additionally, kNN binary classifiers are trained to estimate whether a parameter sample will result in an effective gait. The classification biases the sampling toward subspaces likely to yield effective gaits. An asynchronous communication layer is defined between the optimization and classification processes. The proposed gait discovery process is shown to efficiently optimize the parameters of gaits defined given the novel CPG architecture and outperforms less holistic approaches and Monte Carlo sampling.**

I. Introduction & Related Work

Tensegrities are made up of rigid elements interconnected with flexible, dynamic modules, such as cables. Tensegrity robots, an example of hybrid soft-rigid robots, can dynamically control the lengths and tensions of the dynamic elements. In this way, they can dynamically adapt their global shape and rigidity, as well as locomote [1], [2]. This allows them to adapt to different terrains [3], distributed forces throughout their structure and achieve high durability. They are ideally suited for planetary exploration [4], [5], which involves navigating unstructured and unknown terrains.

Tensegrity systems introduce their own set of challenges that complicate the development of robust locomotion strategies. They are high dimensional, both in terms of state and control variables, highly nonlinear and exhibit complex contact dynamics. Furthermore, they do not have direct control of their position but only of their shape.

This work focuses on a gait generation problem [6] for a prototypical spherical tensegrity platform, known as SUPERBall and shown in Fig. 1. This system and physics engine modeling the system were developed at NASA Ames [2]. This work proposes a parameterization for this platform that allows the generation of rhythmic gaits. Furthermore, a data-efficient optimization method for generating effective gaits is defined.
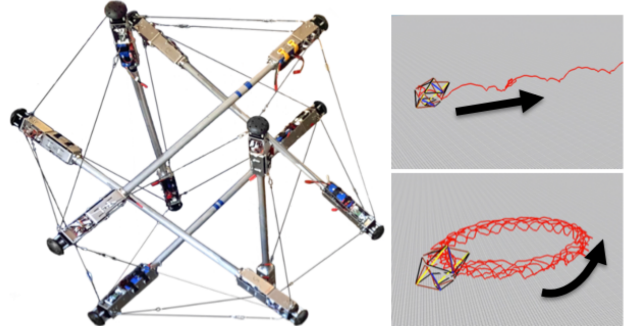
Fig. 1. (left) The SUPERBall prototype at NASA Ames is composed of rigid rods connected by tension cables. Cable lengths are actively adjusted to achieve deformations and varying rigidity. This work develops periodic, stable gaits with predefined motions, given a physics engine that models the system. (right) Two example gaits: one moving the center of mass along a straight line (top) and one resulting in a tight rotation (bottom).

A. Gait Generation for Complex Systems

Gait generation involves designing curves in a robot's base space, i.e., the space representing its shape [6]. This often results in significant computation cost. Discretizing the control space and exhaustively searching it or sampling random controls, does not scale for the 12-actuated tensegrity robot of Fig. 1. Computational requirements are further increased when expensive calls to a physics engine [7] are needed to express effects of contact forces.

Prior work has provided control policies via Monte Carlo search and evolutionary algorithms for the target robot focused at removing the platform from craters [8]. This work generates gaits that are more dynamic in nature and provides a data-efficient process for optimizing these gaits given an appropriate robot parameterization. To the best of the author's knowledge, the resulting gaits are the fastest generated for the system's simulated version.

A popular idea for gait generation is dimensionality reduction. Early work has shown that by utilizing symmetries it is possible to project system dynamics onto the shape space [9], [10], [11]. Analytical expressions can be defined in certain cases expressing the relationship between shape and position velocities giving rise to a non-linear control system. Such analytical expressions, however, are difficult to define for systems as complex as tensegrities. Reducing problem dimensionality and achieving periodic controls, however, are important for generating gaits.

B. Central Pattern Generator Control Structures

A popular paradigm for gait generation comes from biomimicry [12], [13], [14]: biological snakes locomote by undulating their bodies, and many bacteria locomote by changing their shapes. Central Pattern Generators (CPG)
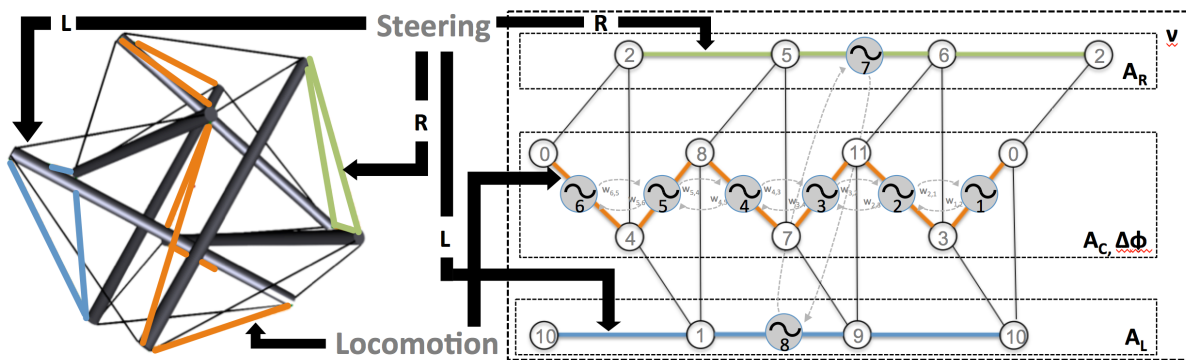
Fig. 2. The 12 actuated cables of the SUPERBall are color-coded and mapped to nodes of the CPG architecture. Locomotion is achieved by periodically contracting and expanding the lengths of the longest actuation chain (orange/center ring). This action shifts the system's CoM and causes it to "roll" onto the next surface. Steering is achieved by contracting the cables on the system's sides (green, blue), which causes a lateral motion. In light gray: the nodes and interconnections of the CPG. In black: The free parameters of the reduced-dimensionality CPG control, and the nodes they affect.

is as an oscillatory gait generation mechanism specifically explored for redundant, bio-inspired systems. The framework has been widely adopted in the control of snake robots [15], [16], [17]. CPGs define sinusoidal equations for variables, which define desired setpoints for the system's joints and allow the setpoints to evolve in constant phase offset with one another [18]. Coupled with a PD controller, this results is oscillatory control, which is parameterized by frequency, amplitude, and phase offset between the joints. A single parameter setting defines a gait, which provides oscillatory control input even at steady state. Grid search of the parameter space has been attempted to identify individual gaits corresponding to the most rapid displacement of a snake-like robot [19]. Evolutionary algorithms have also been used to search the parameter space of CPGs [20].

Contribution A. A CPG structure for SUPERBall, which utilizes the robot's symmetries and provides stable limit cycle behavior, while exposing only 5 parameters.

C. Sample Efficiency in Learning Parameters

The next step is to identify the parameters of the CPG architecture so as to generate a library of fast, dynamic gaits that can move the center of mass (CoM) of SUPERBall in different directions. The motivation comes from the path planning literature on path diversity [21], [22]. Once gaits are available that can move the CoM along different directions, they can then be used for planning purposes among obstacles.

Optimization of these objectives is achieved by choosing a small set of target variables, i.e., forward and angular velocity of the system's CoM, and modeling these variables as a function of the CPG parameters. The process begins with little knowledge of the parameter space and selects samples on-line so as to maximize velocities along different directions of motion. The space of angular velocities for the robot's CoM is discretized so that gaits within a predefined range define a single class. The objective is to maximize the system's velocity within each gait class, while minimizing the search effort required to identify the optimized gaits.

Contribution B. A multi-objective optimization framework for CPG gait parameters so as to provide a library of efficient and diverse gaits.

D. Classification-aware Bayesian Optimization

Since appropriate CPG parameter bounds are not known a priori, the part of the parameter space that results in useful gaits is very small, and different objectives are maximized in different effective subspaces. For this reason, this work models the effect of the CPG parameters on the forward velocity of the system's CoM at steady state by a set of Gaussian process (GP) regression models: a class of predictive models in the family of Bayesian nonparametrics [23]. These models are often chosen when the goal is to approximate a smooth nonlinear function from sparse data [24], [25], [26]. The GP models are optimized through Bayesian Optimization (BO), which aims to optimize the output by selecting samples to evaluate in an on-line manner. The current work uses one BO process per gait direction. For n gait directions, n BO processes are executed in parallel and communicate with each other asynchronously.

In addition, a set of K-Nearest Neighbor (kNN) classifiers are chosen to model the effect of parameters on the angular velocity of the system's CoM. The kNN processes learn the effective subspaces residing within the bounds of the CPG parameter space and bias samples fed to the BO optimization processes toward the effective subsets. Training and communication of the kNN models is also performed on-line, in parallel and asynchronously.

Contribution C. An inherently parallel framework for both learning the effective subspaces of the CPG parameter space and using these estimates to bias sampling within a Bayesian Optimization framework. This is shown to outperform experimentally random sampling, only classification or only optimization in terms of sample-efficiency.

II. CPG DESIGN FOR SUPERBALL

A. SUPERball Design Choices

In designing the prototype for the SUPERball, size and weight requirements led to including motors on 12 of the system's cables. The other 12 cables are "passive". As shown in Fig. 2, the scheme consists of a single actuated ring around the robot's surface and two opposing actuated equilateral triangles on each side. Each of the system's faces contains at least one actuated cable. The benefit of this pattern is

that the system is never in a stable resting configuration without at least one actuated cable on the ground. In this way, the system never gets "stuck" by resting on an entirely passive face. Basic locomotion can be achieved by "flopping" from face to neighboring face by retracting cables of the resting surface such that the center of gravity of the system moves outside of this support surface. Steering the system is achieved by expanding/restricting the side triangles in order to shift the system's center of mass in an effective way to facilitate this flopping behavior.

B. Biomimetic Design for SUPERBall CPG

Previous work on CPG structures mimics the oscillatory center in the lamprey: thousands of neurons located in the animal's spinal cord can produce oscillatory motions [19]. This behavior has been replicated by designing a CPG structure for an 8-link salamander robot. By defining differential equations governing the time evolution of phase and amplitude state variables, the resulting double chain oscillatory structure produces traveling wave motions from the head to the tail similar to those in the biological counterpart.

This prior work has taken advantage of a variety of symmetries to constrain the free parameters to only 5 in the CPG structure: ν the oscillatory frequency, $\Delta\phi$ the phase offset between nodes, A_L and A_R the left and right oscillatory chain's amplitudes, and α governing an increasing or decreasing multiplicative coefficient on the amplitudes of individual signals from head to tail. Each actuated joint of the robot receives a time-varying oscillatory control calculated by combining the time-varying outputs of its corresponding nodes from each of the left, x_L, and right, x_R, oscillatory chains. The desired angular setpoint for the i-th actuated joint in the robot is then calculated as:

$$\varphi_i = x_L - x_R$$

and fed to a PD controller. This control structure results in asymptotically stable limit cycle behavior:

$$\varphi_i^\infty(t) = \alpha_i(A_L - A_R + (A_L + A_R) \cdot \cos(2\pi\nu t + i\Delta\phi + \phi_0))$$

where ϕ_0 depends on initial conditions. This gives rise to a 5-dim. parameter space for control and defining gaits, which converge to predictable, periodic behaviors. For salamander robots, the resulting gaits are largely composable, i.e., due to the underlying system dynamics, transitions between gaits happen in a smooth and continuous manner.

In designing a CPG system structure for controlling the 12-actuator SUPERball prototype, the guiding principle was to build upon this line of work, while also taking advantage of the robot's physical structure and symmetries [27], i.e., separately controlling the central actuation ring for forward locomotion and the two side triangles for steering. The CPG system design consists of a single chain controlling the main actuation ring of the prototype, and a second chain consisting of two nodes controlling the behavior of the left and right actuated triangles (as in Fig. 2). In each chain, nodes are coupled in series, with neighboring nodes exhibiting bi-directional influence on each other.

The present design sets all neighboring connections' weights to equal values similar to the salamander robot [18]. The free parameters in the end are: ν the oscillation frequency, $\Delta\phi$ the phase offset between nodes in the main actuation ring (nodes in the secondary chain are constrained to anti-phase), and A_L, A_C, and A_R the amplitudes of the left actuated triangle, main center ring, and right triangle, respectively. Collectively, these 5 parameters are denoted as Θ. Experimentally, this design exhibits stable limit cycle behavior in the SUPERBall, which is denoted as a function of the free parameters in the system:

$$g(\Theta) = \varphi_{1\ldots N}^\infty(t)$$

While the coupling of the second series of nodes renders it entirely separate from the first, parameters governing frequency are shared, with the frequency of the CPG chain governing the side triangles being 3x higher than that of the main actuation ring Experimentally, this gave the best results in terms of achieving flopping locomotion.

III. MULTI-OBJECTIVE BAYESIAN OPTIMIZATION

Given the above 5-dim. CPG control parameterization, the goal becomes to optimize the gait parameters in order to maximize the effectiveness of a gait "library" according to desirable objectives. As proof-of-concept objectives, this work presents a center of mass (CoM) task space abstraction. Each optimization objective is defined so as to allow the generation of gaits that move the platform along a different direction with high velocity.

A. Optimization Objective

The motivation is to define gaits that can be used by a motion planner, i.e., a discrete set of gaits that maximize the probability of finding a successful path from a wide variety of start location to a variety of end location under 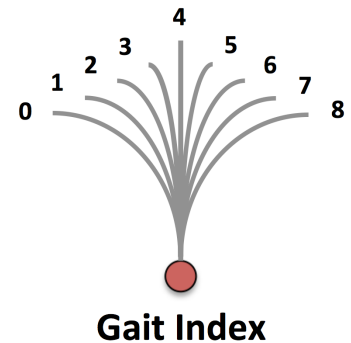various obstacle configurations. The robot also needs to reach the desired goal location fast. The problem of selecting the best sparse set of candidate paths has been studied in both the static and dynamic planning settings [22], [28], [29], [30].

Fig. 3. Illustration of the pre-selected directions used as optimization objectives for the CoM motion.

The objective is to define constant curvature arcs and achieve a low-dimensional parametrization of these gaits. A task space abstraction of the SUPERBall system is first defined as the forward velocity and angular velocity of the system's center of mass (CoM). The angular velocity dimension is discretized into the following 9 directions, illustrated visually in Fig. 3:

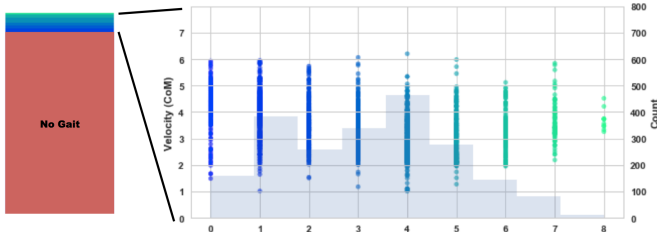| Gait Idx	0	1	2	3	4	5	6	7	8
Min	-10.0	-5.0	-3.0	-1.5	-0.5	0.5	1.5	3.0	5.0
Max	-5.0	-3.0	-1.5	-0.5	0.5	1.5	3.0	5.0	10.0

Fig. 4. Resulting velocities of 20,000 random samples of the CPG parameters. (Left) 89.4% of samples correspond to gaits without a stable limit cycle. (right) Among the stable gaits, a crude normal distribution can be seen over the angular velocity ranges (shaded columns). Individual data points correspond to forward velocities achieved per angular velocity range.

Maximizing velocity is defined as a separate objective for each of the discretized directions. By maximizing the system's velocity along several discrete directions of angular velocity, the resulting gaits follow the spirit of a constant curvature arc strategy, while also maximizing the velocity achieved along each direction for the spherical SUPERBall. Thus, each objective is formally defined as:

$$\text{Objective i:} \quad \arg\max_{\Theta_i} Vel(g(\Theta_i))$$

where Θ_i are the CPG parameters of a gait in the i-th angular velocity range, and $Vel(g(\Theta_i))$ returns the velocity of the SUPERBall's CoM for the corresponding parameters.

B. Bayesian Optimization

Each of the objective functions to maximize takes as input a 5-dim. CPG parameter set, evaluates the resulting trajectory of the SUPERBall using a physics engine until a steady state is reached, and returns the resulting forward and angular velocities of the robot's CoM at steady state. Optimizing these functions lends itself to Bayesian optimization (BO) since the evaluation is expensive, as it requires physics-based simulation over a long time horizon. The BO framework searches for a global function maximum with a minimal amount of samples by solving a sequential decision making problem: at each step the framework aims to find the sample location, which maximizes the surrogate function, and is the current best estimate of the objective function maximum [31]. The target variables are modeled by a Gaussian process (GP) regression model using a radial basis function kernel.

While optimizing a single gait has been addressed using BO before [32], the current work deviates from standard practice. Instead of searching for a single global maximum, developing a library of gaits dictates that several maxima must be found, each corresponding to an individual gait direction. In the current work, these objectives are achieved by specifying the acquisition function for the ith BO node for a given control input Θ:

$$f_{acq}^i(\Theta) = \widehat{Vel}(g(\Theta)) + \sigma(\widehat{Vel}(g(\Theta)))$$

where $\widehat{Vel}(g(\Theta))$ represents the BO node's estimated velocity resulting from input Θ and the second term (σ) represents the uncertainty about the estimate. While individual objectives will have maxima in distinct locations in the parameter space, the objectives operate in the same space – sharing both domain (input CPG controls) and range (CoM velocity).

IV. Parallelized Optimization Architecture

A parallelized architecture, which facilitates communication between simultaneous processes, is proposed to achieve maximization of multiple objectives within a single space. This also addresses the sparseness of good quality solutions within the input space.

A. Challenge: Sparsity of the Solution Space

The solution space for individual gaits can be very sparse (Fig. 4). Among 20k randomly sampled points in the CPG control space, only roughly 10% of these meet a minimum criteria for stability, i.e., after converging to the limit cycle behavior, the angular velocity of the robot varies by no more than $\pm5°/s$, and the linear velocity is greater than 1.0m/s. If the goal was to find a stable gait with steady state angular velocity in the range $[2.5, 5.0]°/s$ - gait index 7 in the presented analysis - only 0.4% of samples would meet this criterion.

B. CPG Control Space Classification

The creation of a gait library can be broken down into two subproblems: (a) space exploration in search of the valid regions, which yield usable and desirable motions, and, (b) within these regions, searching for individual gaits that optimize the current objective. The presented framework tailors predictive models for each task individually. A set of N binary kNN classifiers, one per CoM direction, are trained on samples - both positive and negative - in a streaming fashion as the space is being explored.

The classifiers' predictions bias the sampling of the BO framework toward subspaces predicted to be likely to yield good results. This is achieved through rejection sampling: Candidate points are sampled until a target number of points have been found predicted to belong to a gait class. Once found, these promising points are supplied to the corresponding BO process. Thus, any optimization, and ultimately the next evaluated sample location, occurs at a point predicted by the classifier to be promising.

C. Parallel BO Architecture

The design of the architecture initializes N BO nodes each focusing on optimizing only a single gait objective, which allows each node to search for a single global maximum. Moreover, it allows processing in a maximally parallel fashion. Nodes don't

Fig. 5. Parallel architecture. BO_{1-N}: BO nodes. $Pub_{1-N/Neg}$: communication nodes. kNN_{1-N}: binary classifiers.

have to wait for all other gait objectives to select and simulate their next sample location as this processing is done in parallel rather than sequentially. Each node is able to leverage information gained from other nodes through asynchronous communication channels.
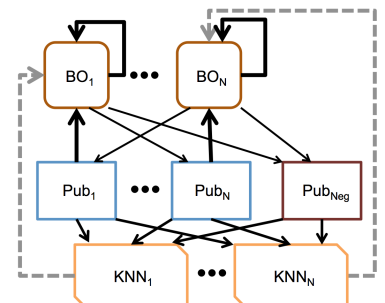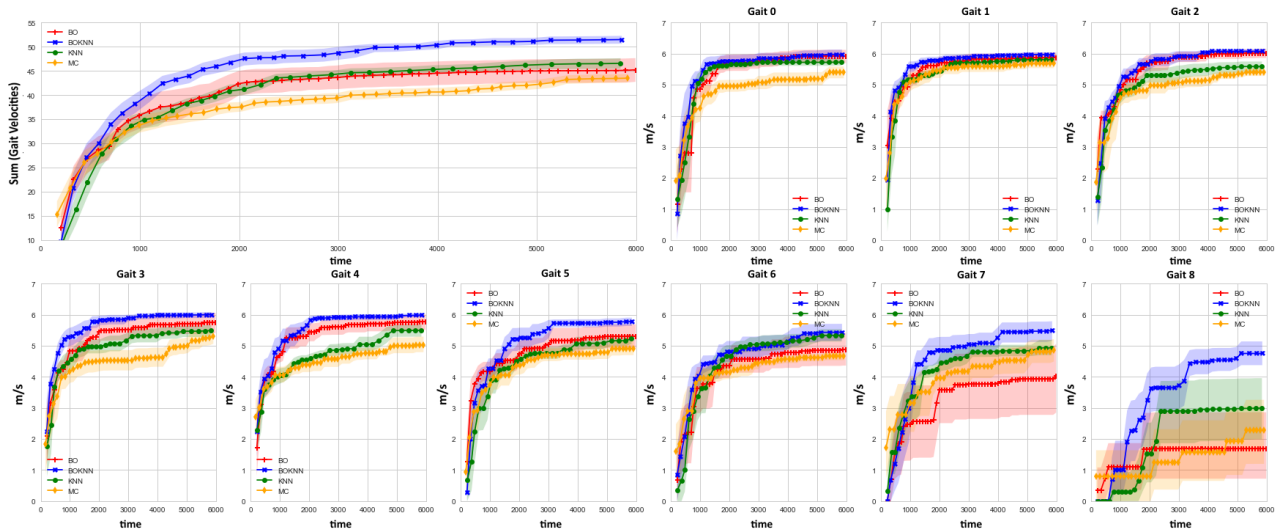
Fig. 6. The effectiveness of the combined (BOkNN) approach against Monte Carlo (MC), Bayesian Optimization alone (BO), and kNN alone (kNN). Resulting velocities of the best gaits in each direction over 6,000 seconds were recorded and averaged over 10 trials per method. (top left) aggregate velocities of all optimized gaits (rest) individual gait directions. 90% confidence intervals are shown as opaque shading.

Fig. 5 illustrates the proposed architecture. To optimize N gaits, N BO nodes are initialized and run in paralle. Simulation results in terms of linear and angular velocity of the system's CoM from any of these nodes are communicated to one of $N+1$ communication channels corresponding to the N gait objectives, e.g., if node BO_3 samples a gait, which ends up belonging to gait objective 4, that control and result is published to channel Pub_4. There is an additional negative sample class for the points, which are deemed unstable. BO nodes only take in and train on data from their corresponding channels (BO_k subscribes to channel Pub_k), effectively focusing each node on only the positive samples for the corresponding angular velocity range. The binary kNN classifiers, conversely, are continuously trained on data from all channels to make their predictions.

V. SIMULATION RESULTS

The efficiency of the presented framework (BOkNN) is compared against: Monte Carlo sampling of the CPG parameters (MC), kNN classification-based region biasing (kNN) and multi-objective Bayesian optimization (BO). All algorithms were executed for 10 repeated trials with different random seeds for 6,000 seconds per trial. The simulation was performed using NTRT, an open source tensegrity simulation software package developed at NASA, which has been experimentally validated using the SUPERBall robot [4]. Initial experiments were executed for even longer durations but it was observed that results largely leveled-off prior to 6,000 seconds. Each method was given the objective of maximizing the velocities of 9 individual gaits with angular velocities consisting of a non-overlapping discretization of the range $[-10°/s, 10°/s]$, and was correspondingly allowed 9 parallel processes for each trial. Every 25 iterations the best gaits for each objective in terms of the velocity of SUPERBall's CoM found up to that point were recorded. To aggregate and compare results from the repeated trials, a running average was computed from the 10 runs for each of the methods.

Fig. 6 shows results from these trials. Overall (top left), the velocities of all gaits discovered increase over time for each of the individual component biasing and optimization methods. The proposed combined framework, however, achieves a more significant improvement, which begins early on during a trial. Across the plots for individual angular velocity ranges, it can be seen that while there is some variation in the relative performances of the BO and kNN methods, the combined BOkNN approach outperforms the other methods across the board, discovering higher velocity gaits in faster time for every objective. Notably, the relative improvement of the combined method is the greatest in the areas of greatest sparsity (gaits 7 and 8), which are also the areas where optimization alone fails to outperform even the baseline Monte Carlo approach. This result supports the effectiveness of the biasing component in the framework, even in very sparse spaces.

Overall, each individual component of the combined framework outperforms the random sampling strategy frequently, which validates their sample-efficiency properties. Nevertheless, in situations when one component fails, e.g., when BO alone struggles in particularly sparse spaces, the other component is able to counterbalance by successfully focusing the optimization process on better quality samples. This allows the combined method to achieve superior results across the board.

VI. CONCLUSION & FUTURE WORK

This work describes the first successful CPG design for a complex spherical tensegrity robot, which yields periodic locomotive behaviors with varying steady-state effects on the linear and angular velocity of the robot's center of mass. Furthermore, this work also represents the first parallel optimization strategy for simultaneously developing an entire gait library. In particular, a parallel approach is proposed for biasing the sampling of promising gait parameters and optimizing the velocities of motions in several different

directions. The optimization framework is particularly useful for new systems with broad control spaces. It is sample efficient and can also be applied when optimizing gaits for a new environment. The gait velocities achieved in the simulated experiments outperform anything achieved via random sampling, even in trials run for significantly longer periods of time. Moreover, the velocities achieved by the optimized gaits are greater than those achieved by prior efforts in kinodynamic planning [33], [34].

A near-term goal of this work is to reproduce the gaits discovered on the real robot. While the simulation used in this work was experimentally validated on the real system, there is bound to be divergence. The parameters of the simulation may need to be re-calibrated and new gaits may need to be optimized which conform better to the limitations of the physical platform. Exciting future work includes examining the integration of the gaits with a search-based planner to leverage and evaluate the gaits' composability. This could also lead to more complex optimization objectives relative to those considered here, such as evaluating the transient affects when composing gaits. An interesting line of work would be to use the current framework to optimize for low velocity gaits of different nature, e.g., lateral movement gaits that can be used in the presence of an unknown obstacle. Finally, allowing a higher-level autonomous process to define the set of gait objectives on the fly depending on the environment it is exploring would be useful as part of a planning-optimization feedback loop.

REFERENCES

[1] Y. Koizumi, M. Shibata, and S. Hirai, "Rolling tensegrity driven by pneumatic soft actuators," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012, pp. 1988–1993.

[2] A. Iscen, A. Agogino, V. SunSpiral, and K. Tumer, "Flop and roll: Learning robust goal-directed locomotion for a tensegrity robot," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2014, pp. 2236–2243.

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

[4] K. Caluwaerts, J. Despraz, A. Işçen, A. P. Sabelhaus, J. Bruce, B. Schrauwen, and V. SunSpiral, "Design and control of compliant tensegrity robots through simulation and hardware validation," *Journal of The Royal Society Interface*, vol. 11, no. 98, p. 20140520, 2014.

[5] 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.

[6] E. Shammas, H. Choset, and A. Rizzi, "Natural Gait Generation Techniques for Principally Kinematic Mechanical Systems," in *Robotics: Science and Systems (RSS)*, 2005.

[7] M. Tesch, "Improving Robot Locomotion Through Learning Methods for Expensive Black-Box Systems," Ph.D. dissertation, Carnegie Mellonw University, 2013.

[8] S. Lessard, J. Bruce, A. Agogino, V. Sunspiral, and M. Teodorescu, "Robust Monte Carlo Control Policies to Maneuver Tensegrity Robots out of Obstacles," in *ARMS*, 2015, pp. 1–16.

[9] G. Walsh and S. Sastry, "On Reorienting Linked Rigid Bodies using Internal Motions," *IEEE Transactions on Robotics and Automation*, vol. 11, no. 1, pp. 139–146, 1995.

[10] J. Ostrowski and J. Burdick, "The Mechanics and Control of Undulatory Locomotion," *International Journal of Robotics Research (IJRR)*, vol. 17, no. 7, pp. 683–701, 1998.

[11] J. Ostrowski, J. Desai, and V. Kumar, "Optimal Gait Selection for Nonholonomic Locomotion Systems," *International Journal of Robotics Research (IJRR)*, 2000.

[12] S. Hirose, *Biologically Inspired Robots: Snake-like Locomotor and Manipulator*. Oxford University Press, 1993.

[13] C. Wright, A. Johnson, A. Peck, Z. McCord, A. Naaktgeboren, P. Gianfortoni, M. Gonzalez-Rivero, R. Hatton, and H. Choset, "Design of a modular snake robot," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2007, pp. 2609–2614.

[14] J. Hopkins and S. K. Gupta, "Design and Modeling of a New Drive System and Exaggerated Rectilinear-gait for a Snake-Inspired Robot," *ASME Journal of Mechanisms and Robotics*, vol. 6, no. 2, 2014.

[15] M. Sfakiotakis and D. Tsakiris, "Neuromuscular Control of Reactive Behaviors for Undulatory Robots," *Neurocomputing*, no. 70, pp. 1907–13, 2007.

[16] F. Herreo-Carron, F. B. Rodriguez, and P. Varona, "Bio-inspired Design Strategies for Central Pattern Generator Control in Modular Robotics," *Bioinspiration & Biomimetics*, 2011.

[17] Z. Bing, L. Cheng, G. Chen, F. Röhrbein, K. Huang, and A. Knoll, "Towards Autonomous Locomotion: CPG-based control of smooth 3D slithering gait transition of a snake-like robot," *Bioinspiration & Biomimetics*, 2017.

[18] A. J. Ijspeert and A. Crespi, "Online trajectory generation in an amphibious snake robot using a lamprey-like central pattern generator model," *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 262–268, 2007.

[19] A. Crespi and A. J. Ijspeert, "Online optimization of swimming and crawling in an amphibious snake robot," *IEEE Transactions on Robotics*, vol. 24, no. 1, pp. 75–87, 2008.

[20] M. Oliveira, V. Matos, C. P. Santos, and L. Costa, "Multi-objective Parameters CPG Optimization for Gait Generation of a Biped Robot," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2013.

[21] M. S. Branicky, R. A. Knepper, and J. J. Kuffner, "Path and trajectory diversity: Theory and algorithms," *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 1359–1364, 2008.

[22] R. A. Knepper and M. T. Mason, "Path diversity is only part of the problem," *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 3224–3229, 2009.

[23] C. K. Williams and C. E. Rasmussen, "Gaussian processes for regression," *Neural Information Processing Systems (NIPS)*, pp. 514–520, 1996.

[24] M. P. Deisenroth, D. Fox, and C. E. Rasmussen, "Gaussian processes for data-efficient learning in robotics and control," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 2, pp. 408–423, 2015.

[25] R. Calandra, A. Seyfarth, J. Peters, and M. P. Deisenroth, "Bayesian optimization for learning gaits under uncertainty," *Annals of Mathematics and Artificial Intelligence*, vol. 76, no. 1-2, pp. 5–23, 2016.

[26] M. Deisenroth and C. E. Rasmussen, "Pilco: A model-based and data-efficient approach to policy search," in *28th Int. Conf. on Machine Learning (ICML)*, 2011, pp. 465–472.

[27] D. Surovik and K. E. Bekris, ""Symmetric Reduction Of Tensegrity Rover Dynamics For Efficient Data-Driven Control," in *ASCE Earth and Space Conference*, 2017.

[28] F. W. Rauskolb, K. Berger, C. Lipski, M. Magnor, K. Cornelsen, J. Effertz, T. Form, F. Graefe, S. Ohl, W. Schumacher, *et al.*, "Caroline: An autonomously driving vehicle for urban environments," *Journal of Field Robotics*, vol. 25, no. 9, pp. 674–724, 2008.

[29] F. Von Hundelshausen, M. Himmelsbach, F. Hecker, A. Mueller, and H.-J. Wuensche, "Driving with tentacles: Integral structures for sensing and motion," *J. of Field Robotics*, vol. 25, no. 9, pp. 640–673, 2008.

[30] R. L. Hatton, R. a. Knepper, H. Choset, D. Rollinson, C. Gong, and E. Galceran, "Snakes on a plan: Toward combining planning and control," *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 5174–5181, 2013.

[31] J. Močkus, "On bayesian methods for seeking the extremum," in *Optimization Techniques IFIP Technical Conference*. Springer, 1975, pp. 400–404.

[32] D. J. Lizotte, T. Wang, M. H. Bowling, and D. Schuurmans, "Automatic gait optimization with gaussian process regression." in *IJCAI*, vol. 7, 2007, pp. 944–949.

[33] Z. Littlefield, K. Caluwaerts, J. Bruce, V. SunSpiral, and K. E. Bekris, "Integrating Simulated Tensegrity Models with Efficient Motion Planning for Planetary Navigation," in *Int. Symp. on AI, Robotics and Automation in Space (i-SAIRAS)*. June, 2016.

[34] Z. Littlefield, D. Surovik, W. Wang, and K. E. Bekris, "From Quasi-Static To Kinodynamic Planning For Spherical Tensegrity Locomotion," in *International Symosium on Robotics Research (ISRR)*, 2017.