

Geometric Swimming on a Granular Surface

Jin Dai*, Hossein Faraji[†], Chaohui Gong*, Ross L. Hatton[†], Daniel I. Goldman[‡], and Howie Choset*

*The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

jind, chaohuig, choset@cmu.edu

[†]Robotics Program
College of Engineering
Oregon State University
Corvallis, OR 97333

farajih, Ross.Hatton@oregonstate.edu

[‡]School of Physics
College of Sciences
Georgia Institute of Technology
Atlanta, GA 30332

daniel.goldman@physics.gatech.edu

Abstract—Snake robots can contact their environments along their whole bodies. This distributed contact makes them versatile and robust locomotors, but also makes controlling them a challenging problem involving high-dimensional configuration spaces, with no direct way to break their motion down into “driving” and “steering” actions. In this paper, we use concepts from geometric mechanics—e.g., expanded Lie bracket analysis—to simplify the problem of controlling a snake robot moving across a granular surface. Without needing force laws that model the interaction of the snake robot with the granular surface, the relationship between shape and body velocities can be experimentally derived by perturbing the robot’s shape from a sampling of initial configurations, which allows us to: 1. Generate an intuitive and visualizable relationship between gait cycles and the motion they induce; 2. Make accurate predictions about the most efficient gaits available to the robot; and 3. Identify an effective turning gait for the robot that to the best of our knowledge has not previously appeared in the snake robot literature. This geometric analysis of snake robot locomotion serves as a demonstration of how differential-geometric tools can provide insight into the motion of systems that do not have the analytic models often associated with such approaches.

I. INTRODUCTION

Locomotion over sand and other loose, flowing material is an ongoing challenge in mobile robotics. Snake robots, which can contact the ground along their whole body, have the potential to perform better in these environments than systems relying on wheels, tracks, or even legs for propulsion [1]. A major challenge in controlling snake robots, however, is coordinating their many degrees of freedom into useful motions; on granular media, this challenge is compounded by a lack of fundamental physical models for the robot-environment interaction.

In this paper, we investigate snake robot locomotion over granular surfaces via an empirical-geometric approach [2]: First, we project the high-dimensional shape space of the snake robot onto a two-dimensional set of *serpenoid waves*, which capture the body shapes seen in biological snakes [3]. Second, we use a series of small test motions to build up an empirical model of the local relationship between changes in the robot’s shape and changes in its position. Finally, we fit this model to a differential-geometric framework that is amenable to analysis and optimization.

In particular, adopting the geometric model allows us to use its *curvature*—which measures how the model changes



Fig. 1. A snake robot on a bed of small plastic spheres. The cameras mounted to the ceiling track the robot’s position changes as it changes shape. The joints on the snake are series-elastic actuators that allow for force measurements.

across the robot’s configuration space—to identify the gait patterns required to translate and rotate the robot in different directions. Combining this curvature with a *Riemannian metric* that captures the energy use of the system, then allows us to identify the the optimally-efficient amplitudes at which to execute these gaits.

By applying this procedure to the snake robot illustrated in Fig. 1, we:

- Derive an geometric model empirically, without having to have a force model of the environment;
- Represent this model via a set of *height functions*, whose topography dictates the effectiveness of different gait patterns;
- Use these height functions to identify the optimal wave amplitude for slithering across the medium; and
- Identify a quick-turn gait that we do not believe has been previously reported in the snake robot literature.

II. BACKGROUND

A. Geometric Locomotion Model

When analyzing a locomoting system, it is convenient to separate its configuration space Q (i.e. the space of its generalized coordinates q) into a position space G and a shape space M , such that the position $g \in G$ locates the system in the world, and the shape $r \in M$ gives the relative arrangement of the particles that compose it.¹

Taking the canonical three-link system in Fig. 2A as an example to illustrate the language that is used throughout the paper, the position is the location of center of mass and mean orientation of three links, $g = (x, y, \theta) \in SE(2)$, and its shape is parameterized by the two joint angles, $r = (\alpha_1, \alpha_2)$.

A useful model for locomotion in *kinematic* regimes where no gliding can occur, and which we employ in this paper, is that at each shape, there exists a linear relationship between changes in the system's shape and changes in its position,

$$\dot{g} = -\mathbf{A}(r)\dot{r}, \quad (1)$$

in which \dot{g} is the body velocity of the system (i.e., the world velocity \dot{g} expressed in the system's local coordinates), and the *local connection* \mathbf{A} acts like the Jacobian of a robotic manipulator, mapping from joint velocities to the corresponding body velocity.² Each row of \mathbf{A} can be regarded as a local-coordinates gradient of a position component with respect to the system shape, as illustrated for a swimmer in a viscous fluid in Fig. 2B.

Equation (1) is a *reconstruction equation*, in that it can be integrated to reconstruct a position trajectory from a known shape trajectory. By an extension of Stokes' theorem [15], the net displacement over a closed loop ϕ (i.e., a gait cycle) can be approximated³ by integrating the *curvature* of \mathbf{A} over a region enclosed by the loop,

$$\Delta g = \oint_{\phi} -\mathbf{A}(r) dr \approx \iint_{\phi} \overbrace{-\text{curl}\mathbf{A} + [\mathbf{A}_1, \mathbf{A}_2]}^{\text{curvature}} dr \quad (2)$$

where the curl of \mathbf{A} is taken row-wise and the local Lie bracket of the columns of \mathbf{A} evaluates as

$$[\mathbf{A}_1, \mathbf{A}_2] = \begin{bmatrix} \mathbf{A}_1^y \mathbf{A}_2^\theta - \mathbf{A}_2^y \mathbf{A}_1^\theta \\ \mathbf{A}_2^x \mathbf{A}_1^\theta - \mathbf{A}_1^x \mathbf{A}_2^\theta \\ 0 \end{bmatrix}. \quad (3)$$

Together, these terms measure how the constraints encoded by \mathbf{A} change during the cycle in ways that prevent the net integral from being self-canceling, and thus capture the net displacement induced by the cycle. As is also discussed

¹In the parlance of geometric mechanics, this assigns Q the structure of a (trivial, principal) *fiber bundle*, with G the *fiber space* and M the *base space*.

²This kinematic condition has been demonstrated for a wide variety of physical systems, including those whose behavior is dictated by conservation of momentum [5, 6], non-holonomic constraints such as passive wheels [6–9], and fluid interactions at the extremes of low [10–12] and high [11–14] Reynolds numbers.

³In this explanation, we have glossed over some formalisms regarding Lie group theory and exponential coordinates, see [15] for more details in this regard.

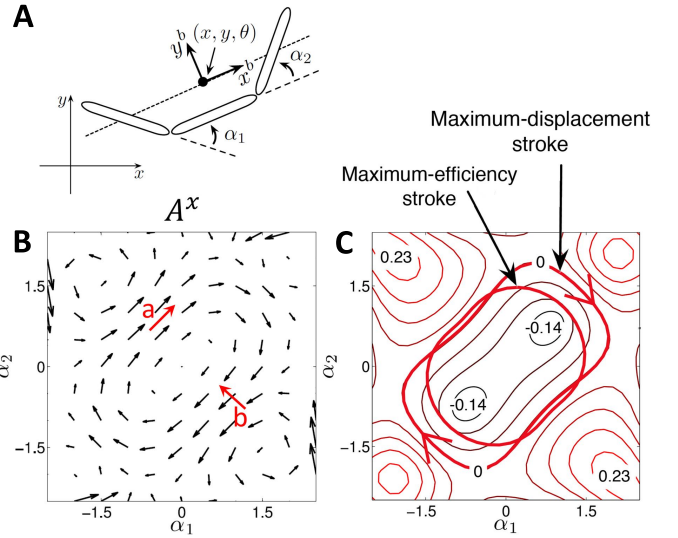


Fig. 2. (A) Three link kinematic vehicle with two actuated joints α_1, α_2 . (B) Local connection vector field corresponding to body velocity x component with a viscous-swimming locomotion model. The shape motion indicated by a is aligned with the field, and so induces positive x displacement, while shape motion b is orthogonal to the vector field, and so induces no x displacement. (C) The displacement- and efficiency-optimal strokes for the three-link viscous swimmer found in [4], overlaid on a contour plot of the x row of the curvature of the local connection. Units on the curvature plot are body-lengths per square radian.

in [15], this approximation is most accurate in *minimum perturbation* coordinates, which place the body frame roughly at the center of mass and aligned with the system's principle moment of inertia, and can be completely identified through a *Hodge-Helmholtz decomposition* of the system dynamics.

Over two-dimensional shape spaces, each component of the constraint curvature can be plotted as a *height function*, as illustrated in Fig. 2C for the x component of a viscous swimmer. Plotting the system dynamics makes the nature of effective gaits immediately apparent: gaits which enclose large sign-definite areas of the height functions produce large displacements, and the largest displacement-per-cycle gaits lie along the zero contours of the height functions [16], as shown in Fig. 2C.

The maximum-efficiency gait, which minimizes the mechanical cost of transport or maximizes speed at a given power, is a contraction of this zero-contour, as shown in Fig. 2C. In a dissipative system the power consumed by a shape change is proportional to its pathlength through the shape space, and contracting the zero-contour minimally affects the area integral while significantly reducing the pathlength. The system is then able to execute more cycles in a given time or at a given power, moving further or faster at the same energetic cost [17].

B. Serpentine Mode Functions

Height functions provide an intuitive tool for designing gaits to achieve large movement in the workspace. However, because of the difficulty of visualizing geometric structures in high-dimensional spaces, they are currently most useful

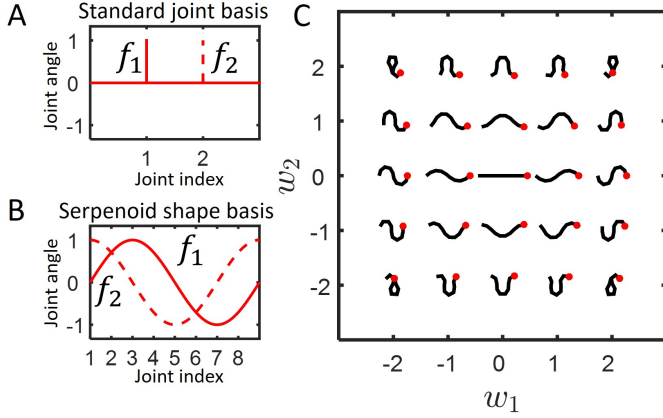


Fig. 3. (A) If all the joints of a locomoting system are independent, its configuration can be spanned by a standard basis of the same dimension as number of joints. (B) With a pair of continuous sine and cosine basis functions, the joints are coupled and the dimension of shape space for a continuous serpenoid swimmer can be reduced to 2. (C) The waveforms of a continuous swimmer in the serpenoid space are visualized in the right panel.

for two-dimensional shape spaces. By using mode functions, the tools of geometric mechanics can be readily applied to study behaviors of complex living systems, such as a snake’s slithering. Hirose proposed the serpenoid curve to model the wave geometry for a snake’s planar slithering motion, where the body curvature (or joint angle in a discrete system) varies sinusoidally along the backbone [3]. Such curves can be described as weighted sums of sine and cosine modes, shown as f_1 and f_2 in Fig. 3B, and angle α of the n th joint is

$$\alpha(n) = [\sin(\Omega n) \quad \cos(\Omega n)] \begin{bmatrix} w_1 \\ w_2 \end{bmatrix}, \quad (4)$$

where Ω is the spatial frequency of the curve, and the weights w serve as shape parameters that describe the amplitude and phase of the wave. The body shapes corresponding to different values of w are shown in Fig. 3C. Given this shape parameterization, we can then build a reconstruction equation of the form in (1) that maps changes in the shape parameters to changes in the body position as

$$\dot{g} = -\tilde{A}(w)\dot{w}, \quad (5)$$

where $\tilde{A}(w)$ is the local connection with respect to the serpenoid shape space.

III. EMPIRICAL DERIVATION OF LOCAL CONNECTION AND POWER METRIC

In this study, our goal was to characterize and optimize the motion of a snake robot over a bed of granular particles. Although we had no constitutive model for this medium, and certainly not one that met the theoretical standards normally encountered in geometric mechanics literature, we were able to empirically derive the local connection for an N-link snake robot swimming on a granular surface by (1) sampling a configuration w , (2) commanding a small shape velocity \dot{w} ,

TABLE I
PROPERTIES OF SNAKE ROBOT AND PLASTIC BALLS SUBSTRATE

number of joints	8
snake robot total length	0.89 m
snake robot intersection diameter	0.051 m
snake robot mass	4.61 kg
plastic ball diameter	0.006 m
plastic ball mass	0.12 g
robot-ball friction coefficient	0.50 ± 0.03

(3) measuring the resultant body velocity \dot{g} and power consumption, and (4) fitting the local connection $\tilde{A}(w)$ and power metric to the data. We repeated this procedure over a uniform sampling of the shape space to generate a comprehensive model of the system.

A. Experimental Setup

The granular bed was composed of 6mm plastic BB balls in a $2.6m \times 1.2m \times 0.23m$ pool as a test bed. Four OptiTrack Flex 13 cameras were installed to track the IR reflective markers uniformly attached on the snake robot, shown in Fig. 1. After each run, we manually smoothed out the surface of the pool to avoid any memory effects of the granular material.

B. Empirical Height Functions

For every configuration, we first take a snapshot of the snake robot, logging the initial center of mass position and mean orientation in the workspace as a homogeneous representation $g(0) \in SE(2)$. Then we perturb the snake along both w_1 and w_2 directions independently, and log the new center of mass position and mean orientation $g(t) \in SE(2)$ in the workspace. With the knowledge of initial configuration $g(0)$ and final configuration $g(t)$ of snake robot, we can calculate the body velocity \dot{g} in a twist form [18]:

$$g(t) = g(0) \exp(\dot{g}t) \Leftrightarrow \dot{g} = \frac{\log(g^{-1}(0)g(t))}{t} \quad (6)$$

Then we can extract all three body velocity components for the planar locomotion from the twist:

$$\dot{g} = \begin{bmatrix} 0 & -\dot{\theta} & \dot{x} \\ \dot{\theta} & 0 & \dot{y} \\ 0 & 0 & 0 \end{bmatrix} \quad (7)$$

After evaluating body velocity at every sampled shape configuration, the local connection can be derived by linear regression as the best-fit linear map from shape to body velocity. In addition, as explained in §II we used *Hodge-Helmholtz Decomposition* to transform the local connection to an optimal body frame. By evaluating curvature of the vector field, we derived the empirical height functions, as shown in Fig. 4. Red areas indicate that the local connection vector field has positive curvature, while black areas indicate negative curvature. The intensity of either color represents the magnitude of the curvature of the local connection vector field and white regions are where the curvature is zero.

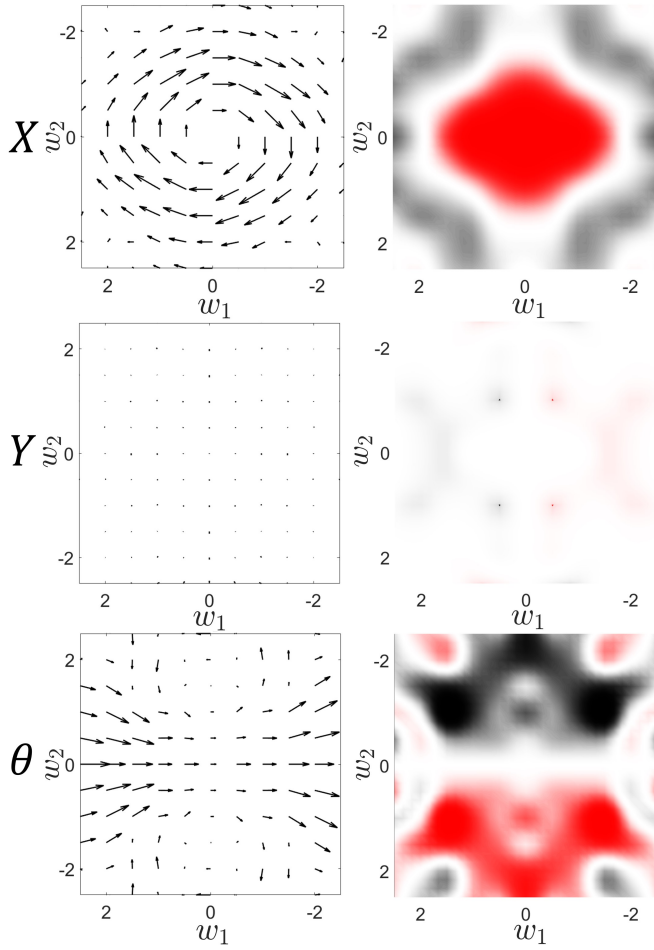


Fig. 4. Empirical local connection vector field and height function of an 8-link snake robot slithering on the surface of 6mm plastic particles, represented in minimal perturbation body frame. The range of the curvature plots is ± 0.0152 body-lengths or ± 0.0483 radians per square unit of shape. Red, white and black indicate positive, zero and negative values respectively.

C. Power Metric

Our model for the energy and power consumed by the snake robot as it changes shape is a dissipative Riemannian metric of the form

$$\left(\frac{dE}{dt}\right)^2 = P^2 = \dot{w}^T M_p \dot{w}, \quad (8)$$

where the metric tensor

$$M_P = \begin{bmatrix} a & \frac{1}{2}c \\ \frac{1}{2}c & b \end{bmatrix} \quad (9)$$

encodes the surrounding media's resistance to motion along different shape directions.⁴ The energetic cost of a gait cycle is then

$$E_{\text{cost}}^2 = \int dw^T M_p dw, \quad (10)$$

which is the pathlength of the cycle in the parameter space, weighted by M_P .

⁴This model is a Coulomb-friction equivalent to the viscous-friction metric discussed in detail in [17].

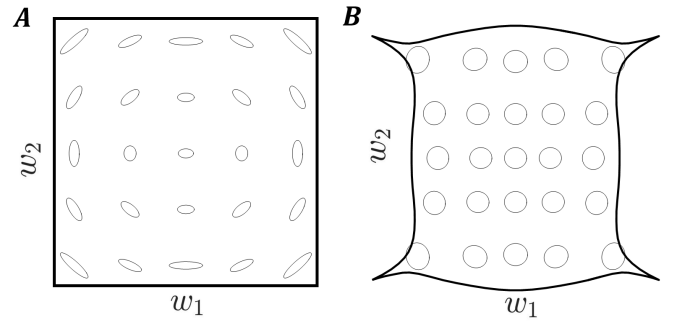


Fig. 5. (A) Power dissipation metric in shape coordinates is shown. The long axis in the ellipse corresponds to less energy consumption per change in shape. (B) The metric as it appears in under a power-normalizing cartographic reparameterization [17].

To evaluate M_P experimentally, we used the high-fidelity joint-torque sensors⁵ on the snake robot to measure the mechanical load on the joints during a sequence of test motions similar to those used to find the local connection in §III-B. Under our model, we would expect to see the energy consumed for small motions around a given starting shape to fit a conical function whose coefficients correspond to the elements in the metric tensor,

$$a(dw_1^2) + b(dw_2^2) + c(dw_1)(dw_2) = P^2 dt^2, \quad (11)$$

and we find the set of coefficients that best fit our observations using a linear regression of the data: For each set of n values for dw , P , and dt around a given starting shape, our model expectation is

$$\begin{bmatrix} (dw_1^2)_1 & (dw_1 dw_2)_1 & (dw_2^2)_1 \\ (dw_1^2)_2 & (dw_1 dw_2)_2 & (dw_2^2)_2 \\ \vdots & \vdots & \vdots \\ (dw_1^2)_n & (dw_1 dw_2)_n & (dw_2^2)_n \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} P^2 dt^2_1 \\ P^2 dt^2_2 \\ \vdots \\ P^2 dt^2_n \end{bmatrix}, \quad (12)$$

which is overconstrained but has a best fit solution via the pseudo-inverse,

$$X = A^+ Y, \quad (13)$$

giving a set of components for the metric tensor at that shape.

The metric tensors at different shapes can be visualized via the Tissot indicatrix [17, 20], which illustrates the extent to which lengths according to the metric are distorted by the working parameterization. In these plots, unit circles according to the metric are displayed as ellipses, with the long axes of the ellipses corresponding to directions in which motion is “easiest” (i.e., for which a unit amount of effort changes the parameters the most). In Fig. 5A, we can see that moving around the shape space (passing a wave down the body) is easier than crossing the center of the shape space (straightening out the body and then curving it again). We can also see that it is slightly easier to move in the w_1 sine mode direction than the w_2 cosine mode direction, which parallels the observations

⁵Series elastic mechanisms as described in [19].

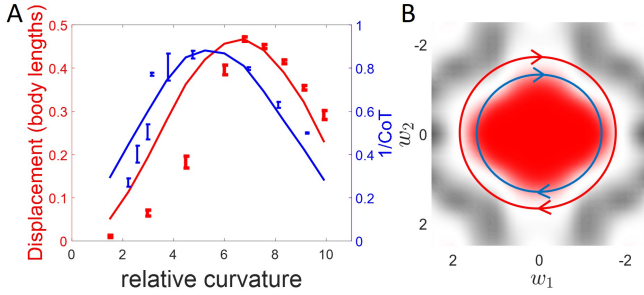


Fig. 6. (A) Body length normalized forward displacement and inverse of mechanical cost of transport per gait cycle vs. relative curvature. Both plots are predicted by empirical height function: displacement is estimated by area integral of height function and cost of transport is estimated by pathlength of the gait in the shape space normalized by area integral of the enclosed height function. Error bars represent robot swimming experiment data, where we sample joint angle amplitude for slithering on the surface of plastic balls. (B) Maximum forward displacement gait (red) and minimal mechanical cost of transport gait (blue) is visualized in serpenoid shape space. Colormap of the height function is the same as in Fig. 4.

in [17] that bending into a “C” shape costs more energy than bending into an “S” shape because the “C” motion sweeps the body through more of the surrounding medium.

To remove some of the metric distortion captured by the Tissot indicatrices from the shape space, we can use a *cartographic reparameterization* [17] to find a set of coordinates that match the metric as much as is possible, in the same way that a conical map projection is a best-fit flat parameterization of the globe. This operation treats the metric distortions as pre-strains in an elastic sheet and relaxes these strains as much as possible while maintaining continuity in the space. As illustrated in Fig. 5B, applying this reparameterization to the shape space significantly reduces the distortion.

IV. DESIGNING GAITS

A. Forward locomotion

A slithering gait of constant curvature amplitude can be represented by a circle centered at the origin in the serpenoid shape space, with the radius representing the wave curvature amplitude. From the empirical height function corresponding to forward displacement in Fig. 6B, we observe that as the wave curvature amplitude increases from zero, forward displacement per gait cycle increases as the circle encloses more red (positive) area. After a certain point, forward displacement is maximized and then begins decreasing as the radius keeps increasing, because the gait starts enclosing black (negative) area. This observation of how forward displacement changes as a function of body curvature matches what we saw when conducting a manual sweep of the parameter space, as shown in Fig. 6A. Moreover, the maximum forward displacement gait found in the latter exactly encloses the most signed area in the empirical height function as the red circle marked in Fig. 6B. The relatively poor agreement in the magnitude of displacement for relative curvature less than 6 is discussed in section V.

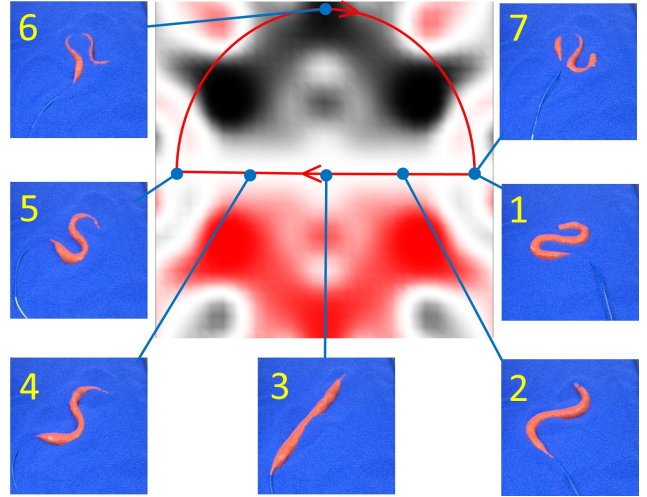


Fig. 7. Periodic turning gait in serpenoid space. Seven key frames are extracted (index is marked in yellow) and linked with shape configurations. The first frame and the seventh frame have the same shape configuration but rotated 90 degrees clockwise, showing the success of the turning maneuver. Colormap of the height function is the same as in Fig. 4.

In the experiments, the robot’s mechanical efficiency⁶ peaked at a lower gait amplitude than did the net displacement. This data matched our prediction from the geometric model that the efficiency should be equal to dividing the area integral of the height function by its metric-weighted perimeter length. As illustrated by the inner blue curve in Fig. 6B, the experiments also matched our prediction that the most-efficient gait would be a contraction of the maximum-displacement cycle, giving up the low-yield regions at the edge of the red central area.

B. Turning gait

As shown in Fig. 7, we can use the θ height function to intuitively identify good turning gaits. Semi-circles, on either the upper half or lower half of the shape space, can enclose a close-to-maximal signed area. Moreover, we observed that as the stroke amplitude (radius of the semi-circle) increases from zero, turning rate increases little at first as the circle encloses light colored area. After a certain point, the turning rate begins to increase significantly, because the semi-circle gait begins to enclose dark colored area. This observation of how turning rate changes as a function of body curvature matches what we saw when conducting a manual sweep of the parameter space, as shown in Fig. 8. The discrepancy in the magnitude of turning is discussed in section V. In Fig. 7, the semi-circle turning gait achieves about 90° per gait cycle, which is an effective motion and to the best of our knowledge has never been found by optimal motion planning or exploring the parameter space, which usually requires a clever heuristic parameterized gait model.

⁶The robot’s weight multiplied by average velocity divided by average mechanical power; reciprocal to the mechanical cost of transport.

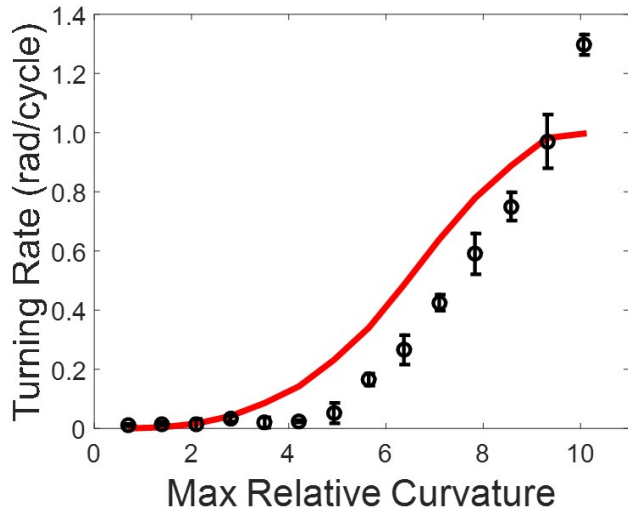


Fig. 8. Turning rate vs. max relative curvature for fast turning gait shown in Fig. 7. The red curve is predicted by area integral of empirical height function corresponding to rotation, whereas the points with error bar represent robot swimming experiment data, where we sample joint angle amplitude for turning on the surface of plastic balls.

C. Lateral translation

The average magnitude of the empirical height function corresponding to the lateral direction (y) is about 10^{-3} times smaller than that of the forward direction. Therefore, it would appear that we cannot design an effective periodic planar gait that achieves significant lateral translation in serpenoid shape space. However, with effective forward motion and in-place turning motion, we can represent the snake robot as a differential drive car, and then take advantage of noncommutativity to design a lateral translation gait that consists of a sequence of simple motions: in-place turn 90° , slither forward, and in-place turn -90° . It is also worth noting that by commanding two orthogonal independent serpenoid waves along the body, the snake robot can sidewind laterally [1, 21].

V. DISCUSSION AND CONCLUSION

Snake robots are potentially versatile and robust locomotion platforms, capable of traversing many kinds of terrain. The challenge in making effective use of this potential is in managing their many degrees of freedom and the complexities of whole-body interaction with the environment. In this paper, we have explored how geometric control theory can provide a path through these difficulties. In particular, fitting a geometric model to a set of small test motions allowed us to:

- Bypass the lack of governing force models of environments;
- Garner a deep understanding of the relationship between gait amplitude and the size and efficiency of the displacement it induces;
- Reduce the space of gaits we needed to consider for translation to simple traveling waves along the body;

- Identify a turn-in-place gait that to the best of our knowledge has not previously appeared in the snake robot literature; and
- Rule out any chance of finding an effective “simple” gait for lateral motion, but identify a set of forward-and-rotate gait cycles that will achieve this motion.

We achieved these insights despite the lack of an ideal differential geometric model for the system dynamics, emphasizing that the principles of geometric gait analysis apply beyond the physics models that they are normally paired with.

Further robot swimming experiments show strong agreement with the predictions of empirical height function for the locomotion performance versus stroke amplitude. We chiefly attribute the discrepancy in the magnitude to a memory effect in the granular substrate. When experimentally determining the local connection, the surface is flattened after each “twitch” to ensure uniform surface contact. However, for robot swimming experiment, the granular surface has already been disturbed by the previous motions and the contact pattern for the robot can hardly be guaranteed to be uniform during the whole swimming motion.

In our ongoing work, we are expanding on these ideas. Avenues we are exploring include incorporating these techniques into model-based predictive control as in [22], the use of higher-dimensional equivalents to height functions, so as to allow the use of more shape modes, and the application of these techniques to animal models as well as robots. To address the “memory” effect in the granular medium, we are developing a sampling protocol that extracts the connection along a standardized test trajectory (instead of resetting the experiment for each sample); this procedure will thus produce a model that better-reflects the experimental protocol.

ACKNOWLEDGMENTS

The authors thank all the members in the Biorobotics Lab, especially Xuesu Xiao for assistance with robot experiments, Tony Dear for comments that greatly improved the manuscript, and Dr. Matthew Travers for insightful discussions. We also thank Perrin E. Schiebel from CRAB Lab for beneficial discussions about biological snake behaviors and RFT models. This research is supported by NSF Grant 1462555 and Robotics Technology Collaborative Alliance.

REFERENCES

- [1] Henry C Astley, Chaohui Gong, Jin Dai, Matthew Travers, Miguel M Serrano, Patricio A Vela, Howie Choset, Joseph R Mendelson, David L Hu, and Daniel I Goldman. Modulation of orthogonal body waves enables high maneuverability in sidewinding locomotion. *Proceedings of the National Academy of Sciences*, 112(19):6200–6205, 2015.
- [2] Ross L Hatton, Yang Ding, Howie Choset, and Daniel I Goldman. Geometric visualization of self-propulsion in a complex medium. *Physical review letters*, 110(7):078101, 2013.

- [3] Yoji Umetani and Shigeo Hirose. Biomechanical study of serpentine locomotion. In *Proc. 1st RoManSy Symp*, volume 73, pages 171–184. Springer, 1974.
- [4] Daniel Tam and Anette E. Hosoi. Optimal stroke patterns for Purcell’s three-link swimmer. *Phys. Review Letters*, 98(6):068105, 2007.
- [5] Gregory C. Walsh and Shankar Sastry. On reorienting linked rigid bodies using internal motions. *Robotics and Automation, IEEE Transactions on*, 11(1):139–146, January 1995.
- [6] Elie A. Shamma, Howie Choset, and Alfred A. Rizzi. Geometric motion planning analysis for two classes of underactuated mechanical systems. *Int. J. of Robotics Research*, 26(10):1043–1073, 2007.
- [7] Richard M. Murray and Shankar S. Sastry. Nonholonomic motion planning: Steering using sinusoids. *IEEE Transactions on Automatic Control*, 38(5):700–716, Jan 1993.
- [8] James P. Ostrowski and Joel Burdick. The mechanics and control of undulatory locomotion. *International Journal of Robotics Research*, 17(7):683–701, July 1998.
- [9] Anthony M. Bloch et al. *Nonholonomic Mechanics and Control*. Springer, 2003.
- [10] Joseph E. Avron and Oren Raz. A geometric theory of swimming: Purcell’s swimmer and its symmetrized cousin. *New Journal of Physics*, 9(437), 2008.
- [11] Scott D. Kelly. The mechanics and control of driftless swimming. In press.
- [12] Ross L. Hatton and Howie Choset. Connection vector fields and optimized coordinates for swimming systems at low and high Reynolds numbers. In *Proceedings of the ASME Dynamic Systems and Controls Conference (DSCC)*, Cambridge, Massachusetts, USA, Sep 2010.
- [13] Juan B. Melli, Clarence W. Rowley, and Dzhelil S. Rufat. Motion planning for an articulated body in a perfect planar fluid. *SIAM Journal of Applied Dynamical Systems*, 5(4):650–669, November 2006.
- [14] Eva Kanso. Swimming due to transverse shape deformations. *Journal of Fluid Mechanics*, 631:127–148, 2009.
- [15] RL Hatton and H Choset. Nonconservativity and non-commutativity in locomotion. *The European Physical Journal Special Topics*, 224(17-18):3141–3174, 2015.
- [16] Ross L Hatton and Howie Choset. Geometric swimming at low and high reynolds numbers. *Robotics, IEEE Transactions on*, 29(3):615–624, 2013.
- [17] Ross L Hatton and Howie Choset. Kinematic cartography for locomotion at low reynolds numbers. *Proc. Robot. Sci. Syst. VII*, 2011.
- [18] Richard M Murray, Zexiang Li, S Shankar Sastry, and S Shankara Sastry. *A mathematical introduction to robotic manipulation*. CRC press, 1994.
- [19] David Rollinson, Yigit Bilgen, Ben Brown, Florian Enner, Steven Ford, Curtis Layton, Justine Rembisz, Mike Schwerin, Andreas Willig, Prasanna Velagapudi, et al. Design and architecture of a series elastic snake robot. In *Intelligent Robots and Systems (IROS 2014)*, 2014
- IEEE/RSJ International Conference on*, pages 4630–4636. IEEE, 2014.
- [20] Auguste Tissot. *Mémoire sur la représentation des surfaces et les projections des cartes géographiques*. Gauthier-Villars, 1881.
- [21] Ross L Hatton, Ross A Knepper, Howie Choset, David Rollinson, Chaohui Gong, and Enric Galceran. Snakes on a plan: Toward combining planning and control. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 5174–5181. IEEE, 2013.
- [22] Kristi Morgansen, Benjamin Triplett, Daniel J Klein, et al. Geometric methods for modeling and control of free-swimming fin-actuated underwater vehicles. *Robotics, IEEE Transactions on*, 23(6):1184–1199, 2007.