

# The International Journal of Robotics Research

<http://ijr.sagepub.com/>

---

## **Geometric motion planning: The local connection, Stokes' theorem, and the importance of coordinate choice**

Ross L Hatton and Howie Choset

*The International Journal of Robotics Research* 2011 30: 988 originally published online 3 June 2011

DOI: 10.1177/0278364910394392

The online version of this article can be found at:

<http://ijr.sagepub.com/content/30/8/988>

---

Published by:



<http://www.sagepublications.com>

On behalf of:



Multimedia Archives

**Additional services and information for *The International Journal of Robotics Research* can be found at:**

**Email Alerts:** <http://ijr.sagepub.com/cgi/alerts>

**Subscriptions:** <http://ijr.sagepub.com/subscriptions>

**Reprints:** <http://www.sagepub.com/journalsReprints.nav>

**Permissions:** <http://www.sagepub.com/journalsPermissions.nav>

**Citations:** <http://ijr.sagepub.com/content/30/8/988.refs.html>

>> [Version of Record](#) - Jul 13, 2011

[OnlineFirst Version of Record](#) - Jun 3, 2011

[What is This?](#)

# Geometric motion planning: The local connection, Stokes' theorem, and the importance of coordinate choice

Ross L Hatton and Howie Choset

## Abstract

*The locomotion of articulated mechanical systems is often complex and unintuitive, even when considered with the aid of reduction principles from geometric mechanics. In this paper, we present two tools for gaining insights into the underlying principles of locomotion: connection vector fields and connection height functions. Connection vector fields illustrate the geometric structure of the relationship between internal shape changes and the system body velocities they produce. Connection height functions measure the curvature of their respective vector fields and capture the net displacement over any cyclic shape change, or gait, allowing for the intuitive selection of gaits to produce desired displacements. Height function approaches have been previously attempted, but such techniques have been severely limited by their basis in a rotating body frame, and have only been useful for calculating planar rotations and infinitesimal translations. We circumvent this limitation by introducing a notion of optimal coordinates defining a body frame that rotates very little in response to shape changes, while still meeting the requirements of the geometric mechanics theory on which the vector fields and height functions are based. In these optimal coordinates, the height functions provide close approximations of the net displacement resulting from a broad selection of possible gaits.*

## Keywords

Geometric mechanics, kinematics, locomotion, nonholonomic motion planning, underactuated robots, undulatory systems

## 1. Introduction

Locomotion—that is, self-powered motion through the world—is one of the most basic actions performed by a robot or living organism. Driven wheels are the dominant paradigm for manmade locomotors, but their effectiveness is limited to relatively flat, smooth terrain. Nature provides many alternative modes of locomotion that can handle more challenging conditions: snakes crawl, fish swim, birds fly, and all manner of creatures walk. The facility with which animals use internal joint motions to move through their environments far exceeds that which has been achieved in artificial systems; consequently there is much interest in raising the locomotion capabilities of such systems to match or surpass those of their biological counterparts. The challenge in this task is in finding effective means of planning for these systems, where the controls are not neatly separable into ‘drive’ and ‘steering’ inputs.

A fundamental aspect of animal locomotion is that it is primarily composed of gaits: cyclic shape motions that efficiently transport the animal through its environment. Examples of such gaits include a horse's walking, trotting, and galloping, a fish's translation and turning strokes, and a snake's slithering and sidewinding. In all of these motions, the animals take advantage of an asymmetry in the reaction

forces they generate from the surrounding media to gain a net displacement over each cycle. The efficacy of such gaits, along with the abstraction that they allow from internal shape changes to position changes, suggests that they will form an equally important part of artificial locomotion, raising the question of how to identify useful asymmetries during motion planning.

Here, we are specifically interested in producing gait design tools for mechanical systems. Much prior work in gait design has taken the approach of choosing parameterized basis functions for gaits and simulating the motion of the system while executing the gaits, optimizing the input parameters to find gaits that meet the design requirements. Such optimization with forward simulation is computationally expensive and sensitive to the choice of starting parameters. Therefore, there is growing interest in finding gait

---

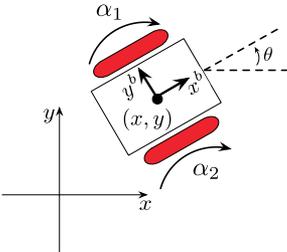
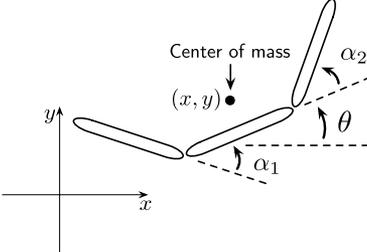
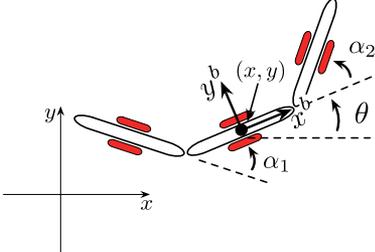
Robotics Institute and Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA

### Corresponding author:

Ross L Hatton, Robotics Institute and Department of Mechanical Engineering, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA.

Email: rlhatton@cmu.edu

**Table 1.** The three example systems.

Differential-drive car	Floating snake	Kinematic snake
		
<p>The differential-drive car has independently driven wheels on either side of the body. These wheels allow the system to drive in the direction in which it points, rotate around its center, or combine these motions to move along arcs of arbitrary curvature.</p>	<p>The links on the floating snake are only constrained by their joints. Because it does not interact with its environment, the translation and rotation of this system are dictated entirely by conservation of linear and angular momentum.</p>	<p>Passive wheelsets on the links of the kinematic snake prevent lateral translation while freely allowing rotation and longitudinal translation. The joint angles are actively controlled and the system can locomote by ‘pulling’ or ‘pushing’ against the wheel constraints.</p>

design tools that replace the simulation stage with a gait evaluation method more amenable to optimization.

Geometric mechanics offers a systematic means of approaching this optimization problem, providing rigorously defined quantities such as *curvature* for measuring the net displacement a system can achieve over cyclic internal shape motion. While methods based on these quantities are powerful, their dissemination has been unfortunately limited by the abstract nature of the differential geometry used to express them. Additionally, the curvature methods as previously developed have not been fully developed for several interesting systems; in their present state, either they are restricted to designing small, inefficient motions, or they provide incomplete information about the actual displacement of the system over the course of a gait.

In this paper, we address these limitations on two fronts. First, we review the key elements of the geometric mechanics formulations, but do so in the language of vector calculus rather than differential geometry and make broad use of concrete examples based on the systems in Table 1. By taking this approach, we seek to present this material to a broader audience. Second, using this simplified formulation, we develop the concept of a *body-velocity integral* (BVI), which provides an expanded and physically meaningful interpretation of previous curvature-based approaches to locomotion, and identify conditions under which the BVI is a good estimate of the true displacement resulting from a gait. We then introduce the notion that rather than being intrinsic to the system, the presence of these conditions depends on the choice of parameterization, and demonstrate that this choice of parameterization can be manipulated to ensure the conditions’ existence. We

finish by presenting a means for automatically finding the optimal choice of parameterization for any given system, such that the BVI approximates the displacement with minimal error over the widest domain of possible gaits. The theoretical developments are supported by experimental results, in which the BVI is shown to be a significantly more accurate approximation of the displacement under the new parameterization than those previously described in the literature.

## 2. Prior work

Our work builds on the body of locomotion literature that uses geometric mechanics to separate internal shape changes from the external motions they produce. The application of geometric mechanics to locomotion, pioneered by Shapere and Wilczek (1989) and further developed by Murray and Sastry (1993) and Kelly and Murray (1995), provides a powerful mathematical framework for analyzing locomotion. A key product of this work is the development of the *reconstruction equation* for non-holonomic systems, which relates body velocity to changes in internal shape for a broad class of locomoting systems. We will not rederive the reconstruction equation here; for a thorough treatment, see Bloch et al. (2003), Ostrowski and Burdick (1998), and Shammass et al. (2007a).

This reconstruction equation has been used in a variety of locomotion contexts. In Ostrowski and Burdick (1998) and Ostrowski et al. (2000) the reconstruction equation is combined with Lie-bracket theory to generate sinusoidal gaits that translate and rotate a variety of snake-like systems. Bullo and Lynch (2001) used the reconstruction equation to decouple the locomotion of kinodynamic

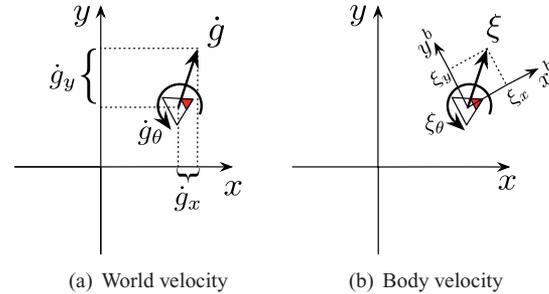
systems and design kinematic gaits. More recently, there has been interest in applying these techniques to swimming robots, such as McIsaac and Ostrowski's (2003) work on anguilliform (eel-like) robots and the work by Morgansen et al. (2007) on fish, both of which combine the geometric approach with biomimetic elements. In Hatton and Choset (2008), we introduced the *connection vector field* as a tool for visualizing the reconstruction equation differentially.

It is not generally possible to integrate the reconstruction equation in closed form, raising difficulties for the inverse problem of finding shape changes that result in desired translations. In special cases, however, Stokes' theorem can be used to find the net motion resulting from gaits (Kelly and Murray 1995). Mukherjee and Anderson (1993a,b) used this principle to analyze the motion of rolling disks, and Walsh and Sastry (1995) applied it to the case of an isolated three-link robot. Shamma et al. (2005, 2007a) combined this approach with the reconstruction equation to define height functions on the shape space of their three-link robots that allowed the design of gaits resulting in specified rotations. A similar technique was used by Melli et al. (2006) and later by Avron and Raz (2008) to generate gaits for swimming robots.

The approaches in Shamma et al. (2007a), Melli et al. (2006), and Avron and Raz (2008) all shared the limitation that for general macroscopic (non-infinitesimal) gaits, the height functions could only be used to determine the net rotations over the gaits, with the non-commutativity of translations and rotations precluding their use in determining the net translation. Melli et al. and Avron and Raz were able to extend the techniques so as to find the net translation over a limited set of macroscopic (but still small magnitude) gaits, by using a higher-order definition of the height functions. Shamma et al. took a different approach, using the height functions to generate coarse approximations of the net displacement that disregard non-commutativity. These approximations were, in general, not close enough to be used effectively in gait design, but in Hatton and Choset (2009), we observed that we can significantly reduce the error by making appropriate choices of system parameterization, and in Hatton and Choset (2010) we took the first steps towards optimizing the parameter choice.

### 3. Background

The present work makes use of several key techniques borrowed from geometric mechanics for representing the locomotion of mechanical systems, the salient points of which we review here. In reviewing this material, our intention is to present it in terms accessible to a reader familiar with vector calculus, rather than go into the details of the underlying differential geometry, which are thoroughly covered in the prior work, particularly Kelly and Murray (1995) and Bloch et al. (2003).



**Fig. 1.** Two representations of the velocity of a robot. The robot, represented by the triangle, is translating up and to the right, while spinning counterclockwise. In (a), the *world velocity*,  $\dot{g}$ , is measured with respect to the global frame. The *body velocity*,  $\xi$ , in (b) is the velocity represented in the robot's instantaneous local coordinate frame.

#### 3.1. The reconstruction equation

When analyzing a multi-body locomoting system, it is convenient to separate its configuration space  $Q$  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 arrangements of its bodies,<sup>1</sup> and then consider how manipulating the shape affects the position. The geometric mechanics community (Murray and Sastry 1993; Kelly and Murray 1995; Ostrowski and Burdick 1998; Bloch et al. 2003; Shamma et al. 2007a) has addressed this question with the development of the *reconstruction equation* and the *local connection*, tools for relating the body velocity of the system,  $\xi$ , to its shape velocity  $\dot{r}$ , and accumulated momentum  $p$ . The general reconstruction equation is of the form

$$\xi = -\mathbf{A}(r)\dot{r} + \mathbf{\Gamma}(r)p, \quad (1)$$

where  $\xi$  is the body velocity of the system,  $\mathbf{A}(r)$  is the *local connection*, a matrix that relates joint to body velocity,  $\mathbf{\Gamma}(r)$  is the *momentum distribution function*, and  $p$  is the *generalized non-holonomic momentum*, which captures how much the system is 'coasting' at any given time (Bloch et al. 2003). The body velocity, illustrated for a planar system in Figure 1(b), is the position velocity as expressed in the instantaneous local coordinate frame of the system, that is its forward, lateral, and turning velocities. For systems that translate and rotate in the plane, with an  $SE(2)$  position space, the body and world velocities are related to each other as

$$\xi = \begin{bmatrix} \xi_x \\ \xi_y \\ \xi_\theta \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \dot{g}, \quad (2)$$

where  $\theta$  is the system's orientation.

For systems that have as many velocity or momentum constraints as dimensions in the position space, the generalized momentum disappears (Shamma et al. 2007a),

and the system's behavior is dictated by the *kinematic reconstruction equation*,

$$\xi = -A(r)\dot{r}, \tag{3}$$

in which the local connection thus acts as a kind of Jacobian, mapping from velocities in the shape space to the corresponding body velocity.

The geometric mechanics literature upon which the present work is based includes a wide variety of models for multi-body locomoting systems that provide concrete examples of the reconstruction equation and local connection. As we develop the concepts in this paper, we will draw upon three of these models: the common *differential-drive car*, the *floating snake* or *planar space robot* (Kelly and Murray 1995; Shamma et al. 2005, 2007a), and the *kinematic snake* (Shamma et al. 2005, 2007a; Ishikawa 2009). Each of these systems moves in the plane, is kinematic, and has two internal configuration variables. The two snake-like models are simple enough to be tractable, but complex enough to pose interesting motion-planning problems. In contrast, the differential-drive car has been extensively studied through conventional means; we take advantage of its familiarity to provide an intuitive link between our new results and the prior work in this area. Table 1 illustrates these three systems and provides some notes on the principles of their locomotion, which are encoded in their reconstruction equations as follows:

**3.1.1. Differential-drive car** The coordinates for the differential-drive car are  $g = (x, y, \theta) \in SE(2)$ , and  $r = (\alpha_1, \alpha_2) \in \mathbb{S} \times \mathbb{S}$ . The reconstruction equation for this system is based on the no-slip and no-slide assumptions for its wheels, which provide three independent constraints relating the body velocity of the vehicle to the angular velocities of the wheels (Kelly and Murray 1995). Normalized for wheel radius and body width, the reconstruction equation for the differential-drive car is

$$\xi = - \begin{bmatrix} -1 & -1 \\ 0 & 0 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}. \tag{4}$$

This equation matches an intuitive understanding of how a differential-drive vehicle works: when the wheels are turned together, the vehicle moves forward; when they are turned oppositely, the vehicle rotates; and no combination of inputs can move the system laterally.

**3.1.2. Floating snake** For the floating snake, we take the position  $g = (x, y, \theta) \in SE(2)$  as the position of the center of mass and the orientation of the middle link and the shape  $r = (\alpha_1, \alpha_2) \in \mathbb{R}^2$  as the interlink joint angles. If the system starts at rest, conservation of linear momentum ensures that the velocity of the center of mass remains zero for all time, so the first two rows of the local connection (corresponding to translation) are zero. The third row of the local

connection identifies rotational velocities that preserve a net angular momentum of zero in response to specified joint velocities (Shamma et al. 2007a). Taking the rows together, the reconstruction equation for a floating snake with equal link lengths and inertias is

$$\xi = -\frac{1}{D} \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ a_{31} & a_{32} \end{bmatrix} \begin{bmatrix} \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}, \tag{5}$$

where

$$\begin{aligned} a_{31} &= 5 + 3 \cos(\alpha_2) + \cos(\alpha_1 - \alpha_2) \\ a_{32} &= 5 + 3 \cos(\alpha_1) + \cos(\alpha_1 - \alpha_2) \\ D &= 19 + 6(\cos(\alpha_1) + \cos(\alpha_2)) + 2 \cos(\alpha_1 - \alpha_2). \end{aligned}$$

**3.1.3. Kinematic snake** The position and orientation of the kinematic snake are similar to those of the floating snake, with  $g = (x, y, \theta) \in SE(2)$  and  $r = (\alpha_1, \alpha_2) \in \mathbb{R}^2$ , except that here we take the position as that of the middle link, rather than the center of mass. Outside of singularities, the no-slide conditions on the passive wheelsets constitute three independent constraints, equal in number to the position coordinates, so the system has no direction in which to coast and behaves kinematically. As detailed in Shamma et al. (2007a), the constraints define the local connection for the system, such that the reconstruction equation, normalized for link length, is

$$\xi = -\frac{1}{D} \begin{bmatrix} 1 + \cos(\alpha_2) & 1 + \cos(\alpha_1) \\ 0 & 0 \\ -\sin(\alpha_2) & -\sin(\alpha_1) \end{bmatrix} \begin{bmatrix} \dot{\alpha}_1 \\ \dot{\alpha}_2 \end{bmatrix}, \tag{6}$$

where  $D = \sin(\alpha_1) - \sin(\alpha_2) + \sin(\alpha_1 - \alpha_2)$ .

### 3.2. Shape changes, gaits, and image-families

To facilitate discussion of motion plans for locomoting systems, we define a vocabulary describing operations in their shape spaces.

**Definition 3.1 (Shape change).** A shape change  $\psi$  is a trajectory in the shape space  $M$  of the robot over an interval  $[0, T]$ , that is, the set of all shape changes is

$$\Psi = \{\psi \in C^1 \mid \psi : [0, T] \rightarrow M\} \tag{7}$$

where  $\psi(0), \psi(T) \in M$  are, respectively, the start and end shapes.

**Definition 3.2 (Gait).** A gait  $\phi$  is a cyclic shape change, that is, the set of all gaits is

$$\Phi = \{\phi \in \Psi \mid \phi(0) = \phi(T)\}. \tag{8}$$

Note that a gait has a defined start shape  $\phi(0)$ ; two gaits whose images in  $M$  are the same closed curve, but with different start points, are distinct.

**Definition 3.3 (Image-family).** The image-family  $\bar{\phi}$  of a gait  $\phi$  is the set of all gaits that share its image (i.e. trace out the same closed curve) in  $M$ ,

$$\bar{\phi} = \{\varphi \in \Phi \mid \text{Im}(\varphi) = \text{Im}(\phi)\}. \tag{9}$$

The image-family includes all continuous time-reparameterizations and changes in starting point of a gait.

### 3.3. Stokes' theorem

In the second half of this paper, we use Stokes' theorem (Boothby 1986) to approximate the net displacement of our robot over specified gaits. As a brief review, Stokes' theorem equates the line integral along a closed curve on a vector field  $V$  on a space  $U$  to the integral of the curl of that vector field<sup>2</sup> over a surface bounded by the curve,

$$\oint_{\partial\Omega} V(u) \cdot du = \iint_{\Omega} \text{curl } V(u) \, du, \tag{10}$$

where  $\Omega$  is a surface on  $U$  bounded by  $\partial\Omega$ .

The special case of Stokes' theorem for vector fields on  $\mathbb{R}^2$  is also known as Green's theorem. In this case, the curl acts as a scalar derivative of the vector field that measures change in both magnitude and orientation; the curl of a vector field  $V$  on a space  $U \subset \mathbb{R}^2$  is

$$\text{curl } V = \frac{\partial v_2}{\partial u_1} - \frac{\partial v_1}{\partial u_2}. \tag{11}$$

The surface  $\Omega$  is just the region of  $U$  enclosed by  $\partial\Omega$ , and (10) specializes to the simple area integral

$$\int_{\partial\Omega} V(u) \cdot du = \iint_{\Omega} \left( \frac{\partial v_2}{\partial u_1} - \frac{\partial v_1}{\partial u_2} \right) du_1 \, du_2. \tag{12}$$

## 4. Connection vector fields

On their own, the reconstruction equation and local connection provide a simple, generalized process model for optimization-based approaches to motion planning (Sfakiotakis and Tsakiris 2007). They also provide a useful framework for applying Lie-bracket theory to the design of small-amplitude gaits (Ostrowski et al. 2000) (i.e. small-amplitude oscillations in the shape variables) that propel the systems in desired net directions. Unfortunately, the process-model approaches do not take advantage of the geometric structure of the reconstruction equation and the Lie-bracket techniques produce inefficient motions for which much of the input energy is expended in high-frequency oscillations.

We are interested here in developing motion-planning techniques that make efficient use of the structure of the local connection, both in the sense that the structure guides the optimization, and that the resulting gaits efficiently convert internal motions into net displacements. As the first step in developing these techniques, we have introduced

connection vector fields (Hatton and Choset 2008). These vector fields comprehensively illustrate the structure of the local connection in a form that intuitively captures the relationship between shape and position changes. The connection vector fields can be used directly to identify components of a given gait, such as the 'power stroke' and 'windup' phases, or, as in the following sections, as a basis for designing further motion-planning tools.

### 4.1. Definition

Each row  $\mathbf{A}(r)^i$  of the local connection can be considered as defining a vector field on the shape space whose dot product with the shape velocity produces the corresponding component of the body velocity,

$$\xi_i = \bar{\mathbf{A}}^{\xi_i}(r) \cdot \dot{r} \tag{13}$$

where the elements of  $\bar{\mathbf{A}}^{\xi_i}$  are taken from the  $i$ th row of the local connection and, for convenience, we wrap in the negative sign when extracting this row as a vector field.<sup>3</sup> Considering the local connection as a set of vector fields with the dot product operator provides a strong geometric intuition for understanding the relationship between shape and position motions. The geometric interpretation of the dot product in Equation (13) is

$$\xi_i = \bar{\mathbf{A}}^{\xi_i}(r) \cdot \dot{r} = \|\bar{\mathbf{A}}^{\xi_i}(r)\| \|\dot{r}\| \cos \Theta, \tag{14}$$

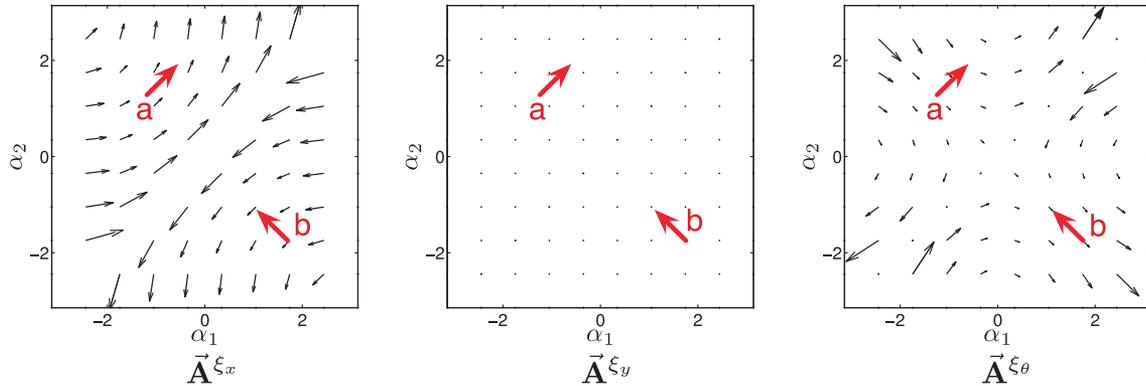
where  $\Theta$  is the angle between the vectors. Taking the  $\cos \Theta$  term as a measure of the alignment of  $\bar{\mathbf{A}}^{\xi_i}(r)$  and  $\dot{r}$ , the body velocity  $\xi_i$  is positive, negative, or zero when the two vectors have correspondingly positive, negative, or zero alignment, and is scaled by the magnitudes of  $\bar{\mathbf{A}}^{\xi_i}$  and  $\dot{r}$ .

For example, the two representative  $\dot{r}$  vectors  $\mathbf{a}$  and  $\mathbf{b}$  in Figure 2 are, respectively, aligned with or perpendicular to  $\bar{\mathbf{A}}^{\xi_x}$  (the 'longitudinal connection vector field') and perpendicular to or anti-aligned with  $\bar{\mathbf{A}}^{\xi_\theta}$  (the 'rotational connection vector field'). The 'lateral connection vector field'  $\bar{\mathbf{A}}^{\xi_y}$  is null, and thus has zero dot product with all vectors. Consequently, we can tell by inspection that the input vector  $\mathbf{a}$  produces a forward velocity with no lateral or rotational velocity, and that the input vector  $\mathbf{b}$  produces a negative rotational velocity with no longitudinal or lateral movement.

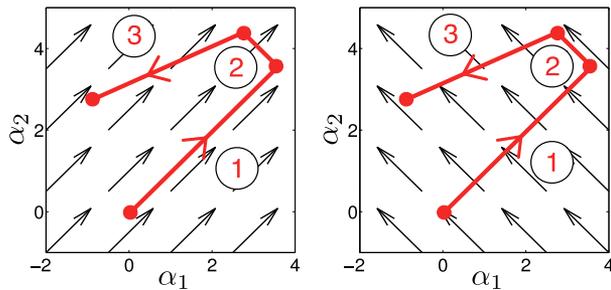
### 4.2. Applied examples

Beyond allowing for the characterization of the effect of input shape velocities, the connection vector fields also facilitate decomposing the displacements caused by shape changes and gaits into their constituent motions, as demonstrated in the examples below.

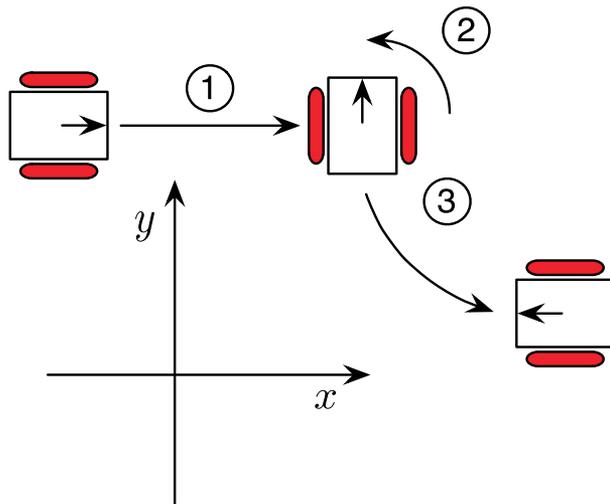
*4.2.1. Differential-drive car* As an introduction to working with connection vector fields, consider the differential-drive car, which has simple kinematics and two non-zero connection vector fields. These connection vector fields are plotted



**Fig. 2.** Connection vector fields for the three-link kinematic snake shown in Table 1. The  $\vec{A}^{\xi_y}$  field is null at all points in the shape space, as the middle wheelset constrains  $\xi_y$  to zero under all conditions. Due to singularities in the vector fields at the lines  $\alpha_1 = \pm\pi$ ,  $\alpha_2 = \pm\pi$ , and  $\alpha_1 = \alpha_2$ , the magnitudes of the vector fields have been scaled by their arctangents for readability. The input shape vectors at **a** and **b** produce pure forward translation and pure negative rotation, respectively.



(a)  $\vec{A}^{\xi_x}$  (left) and  $\vec{A}^{\xi_\theta}$  (right) for the differential-drive vehicle, along with an example shape change  $\psi$ .



(b) Position change of the differential-drive vehicle while executing  $\psi$ .

**Fig. 3.** Connection vector fields for the differential-drive vehicle, along with an example shape change  $\psi$  and the resulting position change.

with a shape change  $\psi$  in Figure 3. The connection vector fields are constant over the shape space and match our intuitive understanding of how the system works: turning the wheels together in the  $+\alpha_1, +\alpha_2$  direction moves the car

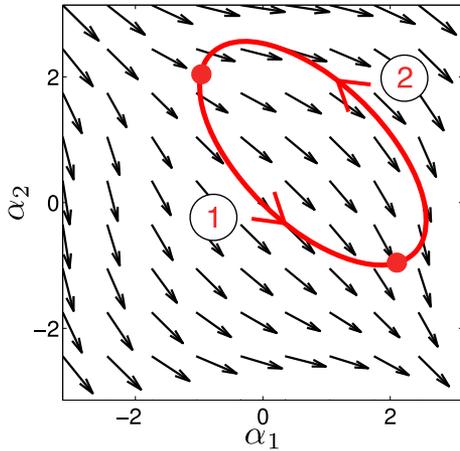
forward, while turning them oppositely in the  $-\alpha_1, +\alpha_2$  direction rotates the car.

The shape change  $\psi$  in Figure 3(a) is divided into three stages,  $\psi_1, \psi_2$ , and  $\psi_3$ . Over the first stage,  $\dot{\psi}_1$  is aligned with  $\vec{A}^{\xi_x}$  and orthogonal to  $\vec{A}^{\xi_\theta}$ ; the robot thus drives forward while executing  $\psi_1$ . In the second stage,  $\dot{\psi}_2$  is aligned with  $\vec{A}^{\xi_\theta}$  and perpendicular to  $\vec{A}^{\xi_x}$ ; the robot correspondingly rotates in place while executing  $\psi_2$ . In the third and final stage,  $\dot{\psi}_3$  has positive alignment with  $\vec{A}^{\xi_\theta}$  and negative alignment with  $\vec{A}^{\xi_x}$ ; the resulting position motion is that the robot rotates positively while reversing. The resulting position changes are illustrated in Figure 3(b).

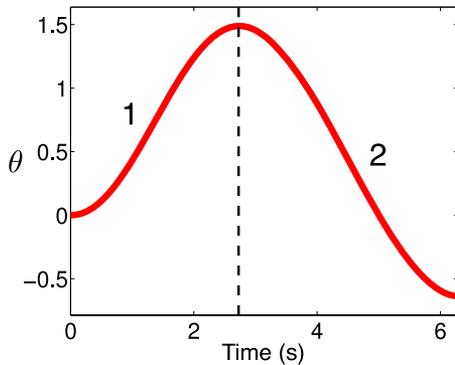
**4.2.2. Floating snake** The differential-drive car is a special case, with constant connection vector fields. Most systems have connection vector fields that vary with the system's shape. The floating snake provides a simple example of how this variation can provide information about the system's behavior. As the first two rows of its local connection in Equation (5) are null, all the information regarding this system's position change in response to a shape change is encoded in the third, rotational row of the local connection. The plot of the corresponding connection vector field,  $\vec{A}^{\xi_\theta}$ , in Figure 4(a) matches an intuitive understanding of how a system governed by conservation of angular momentum should behave: the general trend to point in the  $+\alpha_1, -\alpha_2$  direction corresponds to the central body counter-rotating against the outer links. The exact heading of the field varies with how much leverage the links exert on the system; the more extended an outer link is, the faster the central body needs to counter-rotate to nullify the outer link's angular momentum.

Figure 4(a) also overlays  $\vec{A}^{\xi_\theta}$  with a plot of a sample gait  $\phi(t) = (\alpha_1(t), \alpha_2(t))$ , with

$$\begin{bmatrix} \alpha_1(t) \\ \alpha_2(t) \end{bmatrix} = \begin{bmatrix} \frac{\pi}{4} \\ \frac{\pi}{4} \end{bmatrix} + \mathbf{R}(-\pi/4) \begin{bmatrix} \frac{9}{4} \cos(t + \frac{17}{16}\pi) \\ \frac{9}{8} \sin(t + \frac{17}{16}\pi) \end{bmatrix}, \quad (15)$$



(a)  $\vec{A}^{\xi\theta}$  for the floating snake, along with a gait  $\phi$ .



(b) Orientation  $\theta$  of the snake while executing  $\phi$ .

**Fig. 4.** Sample gait  $\phi$  for the floating snake. In the first segment, the shape velocity  $\dot{\phi}$  is positively aligned with  $\vec{A}^{\xi\theta}$ , so the snake rotates positively. In the second segment,  $\dot{\phi}$  is negatively aligned with the connection vector field, resulting in negative rotation.

where  $\mathbf{R}(\gamma)$  is a rotation by  $\gamma$  and the representation used is chosen to highlight the position, orientation, scale, and phase of the elliptical path. The shape velocity  $\dot{\phi}(t) = [\dot{\alpha}_1(t) \dot{\alpha}_2(t)]^T$  at time  $t$  is a tangent vector to  $\phi$  at  $\phi(t)$ . It is clear that  $\dot{\phi}(t)$  has positive alignment with  $\vec{A}^{\xi\theta}(\phi(t))$  over the first segment of the gait and negative alignment over the second segment, with zero alignment at the start, end, and transition between segments; we can thus immediately conclude that the robot rotates first positively, then negatively while executing the gait, as illustrated in Figure 4(b).

Closer observation of Figure 4(a) reveals that  $\vec{A}^{\xi\theta}$  for the floating snake has approximately constant magnitude and that its curvature results in the negatively aligned Segment 2 being longer than the positively aligned Segment 1, with the former acting as the ‘power stroke’ of the motion and the latter serving in a ‘windup’ role. This property results in the net negative rotation over the gait, and we will examine this property more closely in the following sections.

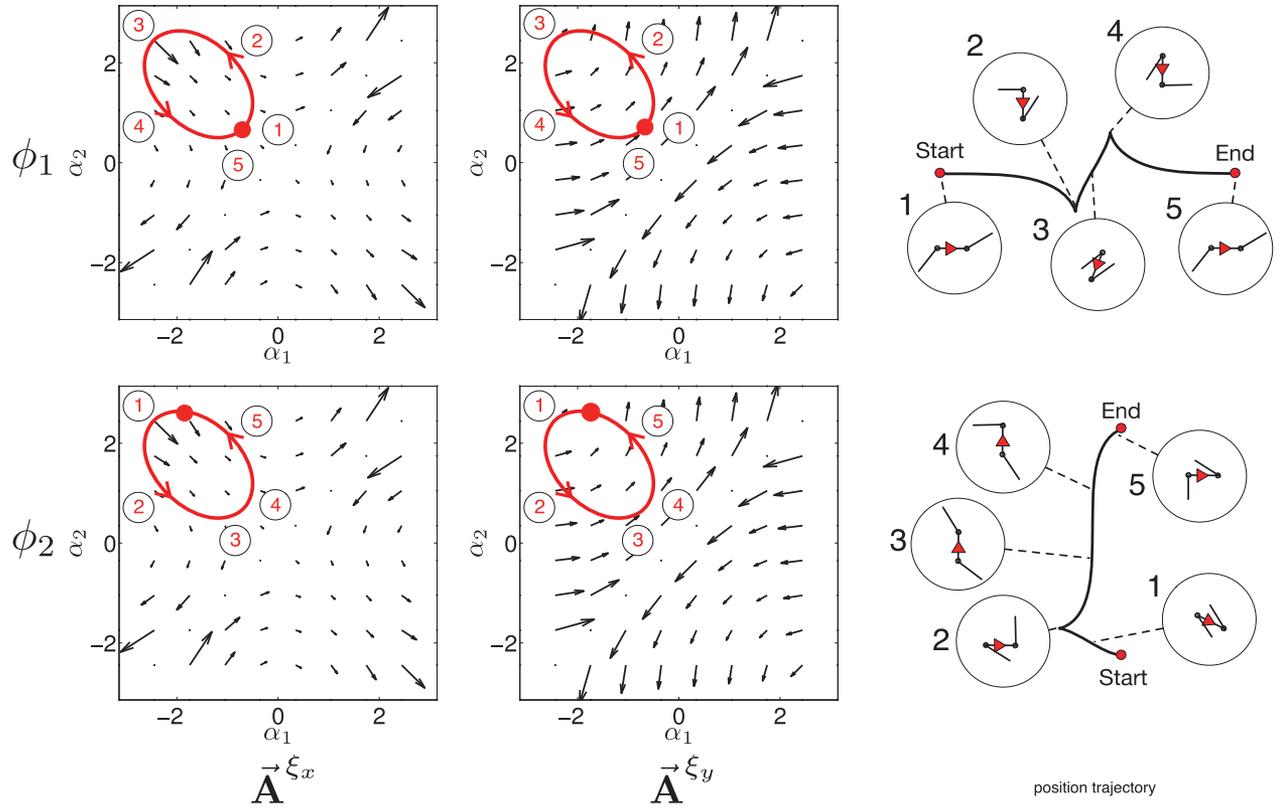
**4.2.3. Kinematic snake** The kinematic snake’s connection vector fields combine the interesting features from both the differential-drive car and the floating snake: there are multiple non-zero fields, and each varies richly with the shape of the robot. The vector fields also exhibit another interesting feature not seen in the previous examples: singularities. At shapes on the lines  $\alpha_1 = \pm\pi$ ,  $\alpha_2 = \pm\pi$ , and  $\alpha_1 = \alpha_2$ , the wheelset constraints are not independent and the local connection loses rank. For example, in Table 1, the kinematic snake is shown in a singular configuration with  $\alpha_1 = \alpha_2$ . The normal lines to the constraints all meet at a common center, allowing the system to rotate freely around that center. In such configurations, the magnitudes of the connection vector fields approach infinity, the constraints lock the joints, and the system is unable to change shape through internal torque. The system does gain a ‘coasting’ degree of freedom, but this state is uncontrollable, and the impracticality of reaching the singular state means that we need not consider gaits that pass through the singularities.

The vector fields for the kinematic snake are plotted in Figure 5 along with two example gaits  $\phi_1$  and  $\phi_2$ . The singularities in the connection vector fields also pose a display problem, as the large magnitudes near the singularities saturate the resolution of the vector magnitudes in the plot and obscure the field in the regions we are more interested in. To address this issue, we use a scaled arctangent normalization on the vector magnitudes while maintaining their headings. The non-linearity of the arctangent function preserves the differences between small values, while pulling large values down to a ceiling level. Along with the vector field plots, we present two gaits based on those in Shamma et al. (2007a), defined as

$$\begin{bmatrix} \alpha_1(t) \\ \alpha_2(t) \end{bmatrix} = \begin{bmatrix} -\frac{\pi}{2} \\ \frac{\pi}{2} \end{bmatrix} + \mathbf{R}(-\pi/4) \begin{bmatrix} \frac{5}{4} \cos(t + \tau_i) \\ \frac{17}{20} \sin(t + \tau_i) \end{bmatrix}, \quad (16)$$

with  $\tau_1 = 0$  for  $\phi_1$  and  $\tau_2 = \frac{49}{64}\pi$  for  $\phi_2$ . We note for future reference that they trace out the same closed curve in the shape space but from different starting points, and so are both in the same image-family (i.e.  $\phi_2 \in \phi_1$  or, equivalently,  $\phi_1 \in \phi_2$ ).

Unlike the differential-drive vehicle, the kinematic snake cannot be simply ‘driven forward’ and it is not immediately clear (outside of direct calculation) how executing the example gaits changes the position of the robot. With the aid of the connection vector fields, however, we can gain additional insight into the system’s behavior and identify characteristics of the position change during the gait. In  $\phi_1$ , the robot starts with  $\dot{\phi}_1$  positively aligned with  $\vec{A}^{\xi_x}$  and orthogonal to  $\vec{A}^{\xi_\theta}$ , and thus moves straight forward. As the gait moves from state 1 to state 2,  $\dot{\phi}_1$  becomes progressively anti-aligned with  $\vec{A}^{\xi_\theta}$  and orthogonal to  $\vec{A}^{\xi_x}$ , until at 2 the position change is purely negative rotation, generating a cusp in the position trajectory. From 2 to 3,  $\dot{\phi}_1$



**Fig. 5.** Connection vector fields for the kinematic snake robot, along with the second example gait  $\phi_2$  and the resulting position change.

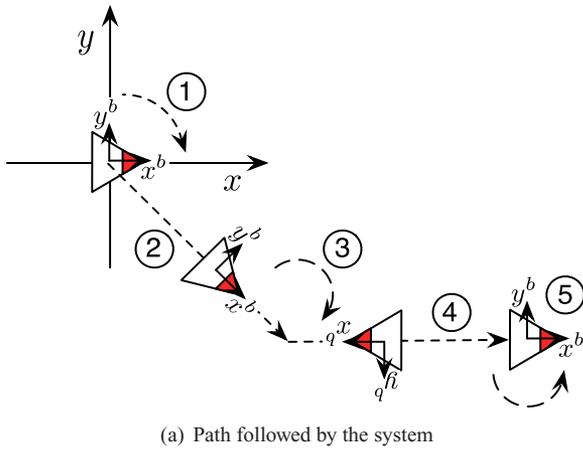
becomes negatively aligned with  $\vec{A}^{\xi_x}$  and perpendicular to  $\vec{A}^{\xi_\theta}$  generating a short stretch of backwards translation. As the robot reaches stage 4,  $\phi_1$  again becomes orthogonal to  $\vec{A}^{\xi_x}$ , but this time is aligned with  $\vec{A}^{\xi_\theta}$ , and the resulting cusp returns the robot to its initial orientation. Finally, at stage 5,  $\phi_1$  is once more aligned with  $\vec{A}^{\xi_x}$  and orthogonal to  $\vec{A}^{\xi_\theta}$ , producing a second stretch of forward motion.

The second gait,  $\phi_2$ , contains all the same segments of motion as  $\phi_1$ , but in a different order. The initial configuration of  $\phi_2$  is at a point of pure rotation, so the resulting position motion starts at what was the first cusp of  $\phi_1$  and follows the short backward translation into the second cusp. Coming out of the cusp, it has rotated by  $\pi/2$  radians from its original starting orientation as it starts the long translation segment from 3 to 4, sending this motion into the  $-y$  direction. The cusp at 5 serves to return the snake to its original orientation with no accrued  $x$  motion.

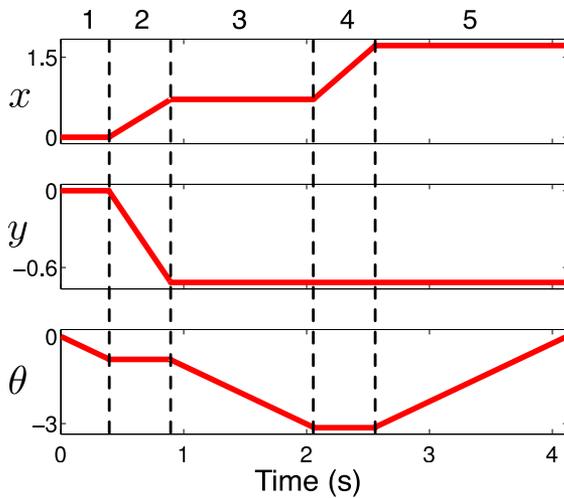
If we consider the cusps as turning maneuvers, the difference between the two gaits can roughly be characterized as  $\phi_1$  having a ‘move–turn–turnback–move’ pattern, while  $\phi_2$  has a ‘turn–move–turnback’ pattern. Inputs that fall between these two patterns (i.e. start at different points on the ellipse) produce net displacements that fall between the pure  $x$  and pure  $y$  displacements shown here.

### 5. Connection height functions

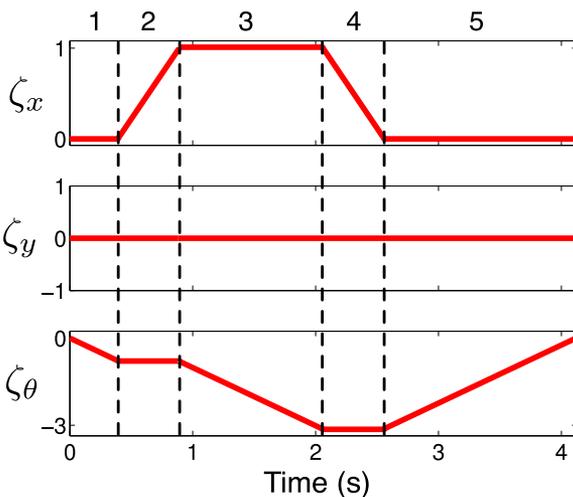
In Section 4.2.2, we noted that the ‘curvature’ of the vector field in Figure 4(a) resulted in the example gait opposing the ‘flow’ of the floating snake’s  $\vec{A}^{\xi_\theta}$  field more than it followed the flow. We then inferred a relationship between this property and the net negative rotation experienced by the system. To further examine this phenomenon, we turn to an approach based on Stokes’ theorem (Boothby 1986). By applying this theorem to the connection vector fields, we can generate scalar *connection height functions* on the shape space that directly capture the net motion resulting from arbitrary choices of gaits. These height functions have been investigated in several prior works (Walsh and Sastry 1995; Melli et al. 2006; Shamma et al. 2007a), but in each work were generated by direct mathematical processing of the local connection, with no physical intuition attached. Additionally, while the height functions were useful for designing gaits to produce desired net rotations, limitations on their applicability prevented them from accurately capturing system translations. Here, we incorporate intuition from the connection vector fields into the derivation and interpretation of the height functions; this both provides additional insight into their operation and lays the groundwork for the following section, in which we use the connection vector fields to circumvent previous limitations on the height functions.



(a) Path followed by the system



(b) Accrued displacement



(c) Accrued BVI

**Fig. 6.** Example motion showing the difference between displacement and BVI. Because the BVI is calculated in the body frame, the ‘backward’ motion in stage 4 cancels out the ‘forward’ motion in stage 2.

### 5.1. Displacement and the body-velocity integral

Before we can discuss tools for calculating the net motion over a gait at any technical level, we need to define what we mean by ‘net motion.’ While the net motion should clearly be an integral of the body velocity, there are two senses in which the phrase ‘integrating the body velocity over time’ can be interpreted. In the first, most natural sense, we integrate to find the resulting displacement.<sup>4</sup> This quantity is found by the non-linear iterative integral

$$g(t) = \int_0^t \begin{bmatrix} \cos \theta(\tau) & -\sin \theta(\tau) & 0 \\ \sin \theta(\tau) & \cos \theta(\tau) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \xi_x(\tau) \\ \xi_y(\tau) \\ \xi_\theta(\tau) \end{bmatrix} d\tau, \tag{17}$$

in which  $\xi_x$  and  $\xi_y$ , the translational components of the body velocity, are rotated into the world frame at each time.

The second sense of integrating  $\xi$  is to take its simple vector integral  $\zeta$ ,

$$\zeta(t) = \begin{bmatrix} \zeta_x(t) \\ \zeta_y(t) \\ \zeta_\theta(t) \end{bmatrix} = \int_0^t \begin{bmatrix} \xi_x(\tau) \\ \xi_y(\tau) \\ \xi_\theta(\tau) \end{bmatrix} d\tau, \tag{18}$$

where we term  $\zeta$  the *body-velocity integral*, or BVI. Physically, the BVI corresponds to the raw odometry for the robot (i.e. the net forwards minus backwards motion in each body direction) and it does not account for changes in the alignment between body and world directions, that is, the relative orientation of the body and world frames. While the integral in Equation (18) has appeared previously (Shammas et al. 2007a), its result  $\zeta$  has not previously been named, and its physical meaning has not been explored.

On first appearance, the BVI does not seem a particularly useful measure of the net motion of a system. For example, consider the path taken by the system in Figure 6. Over the course of the motion, the system clearly accrues a non-zero net displacement. However, because the ‘backward’ motion in stage 4 cancels out the ‘forward’ motion in stage 2, the net BVI is zero in all components: the same as if the system had not moved at all.

The BVI does have three saving graces that make it worthy of study. First, while the displacement integral in Equation (17) must be solved individually for any given gait, the BVI integral in Equation (18) can be replaced by an equivalent area integral via Stokes’ theorem, providing height functions that encode its solution over all possible gaits. Second, unlike the translational components, the rotational components of the displacement and BVI are identically equal, so the net rotation over a gait can be directly determined from the BVI. Third and most importantly, under the right conditions the translational components of the BVI can be taken as close approximations of the displacement, and these conditions can be expressed as simple functions on the shape space.

### 5.2. The BVI, Stokes' theorem, and height functions

Starting from the definition of the BVI in Equation (18), we can express the BVI over a shape change  $\psi$  as a set of line integrals on the connection vector fields,

$$\zeta(T) = \int_0^T \xi_i(\tau) \, d\tau \tag{19}$$

$$= \int_0^T \vec{\mathbf{A}}^{\xi_i}(\psi(\tau)) \cdot \dot{\psi}(\tau) \, d\tau \tag{20}$$

$$= \int_{\psi} \vec{\mathbf{A}}^{\xi_i}(r) \, dr. \tag{21}$$

For shape changes that are gaits, we can further apply Green's form of Stokes' theorem, Equation (12), to convert this line integral into an area integral,

$$\zeta(T) = \iint_{\phi_a} \text{curl } \vec{\mathbf{A}}^{\xi_i}(r) \, dr \tag{22}$$

$$= \iint_{\phi_a} \frac{\partial \vec{\mathbf{A}}_2^{\xi_i}}{\partial \alpha_1} - \frac{\partial \vec{\mathbf{A}}_1^{\xi_i}}{\partial \alpha_2} \, dr, \tag{23}$$

where  $\phi_a$  is the area on  $M$  enclosed by the gait. Note that as all gaits in an image-family enclose the same area, the integral in Equation (23), and thus the BVI, is the same for all gaits in an image-family.

By plotting the integrand of Equation (23) as a height function  $\mathbf{H}^{\xi_j}$  on the shape space, we can easily identify gaits that produce various displacements. Such height functions have featured in prior works, and rules of thumb for designing curves that produce desired values of the integral in Equation (23) are given in Shamma et al. (2007a):

1. *Non-zero net integral (I).* A loop in a region where the sign of  $\mathbf{H}^{\xi_j}$  is constant produces a non-zero net integral. The sign of the integral depends on the sign of  $\mathbf{H}^{\xi_j}$  and the orientation of the loop (the direction in which it is followed).
2. *Non-zero net integral (II).* A figure-eight across a line  $\mathbf{H}^{\xi_j}(r) = 0$ , where each half is in a region where the sign of  $\mathbf{H}^{\xi_j}$  is constant, produces a non-zero net integral.
3. *Zero net integral.* A loop that encloses equally positive and negative regions of  $\mathbf{H}^{\xi_j}(r)$ , such that the integrals over the positive and negative regions cancel each other out, produces a zero net integral.

In an automated gait design algorithm, these rules would serve to guide optimization of gait parameters, for example by indicating regions that the curve should grow to enclose or shrink to avoid.

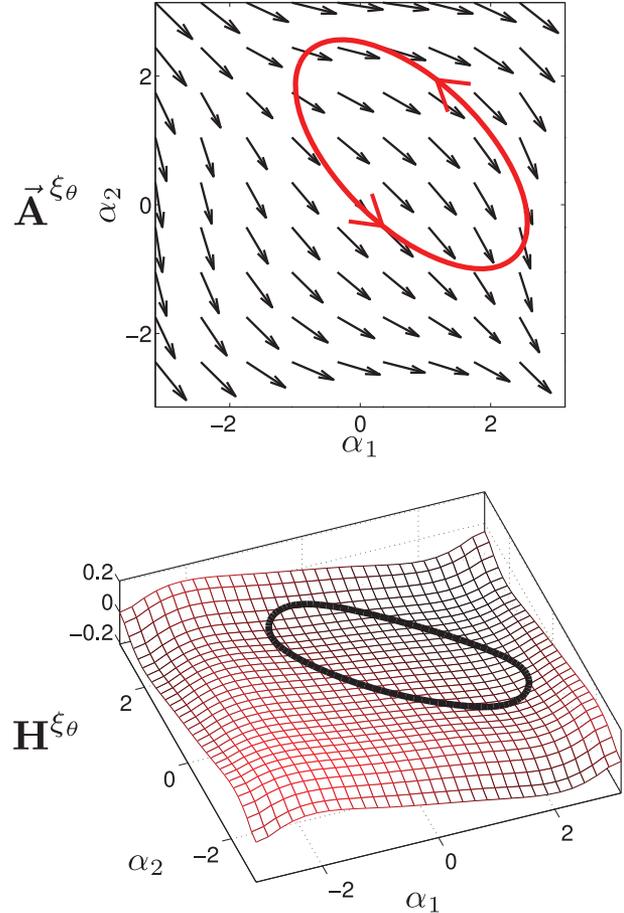


Fig. 7. Rotational connection vector field  $\vec{\mathbf{A}}^{\xi_\theta}$  and height function  $\mathbf{H}^{\xi_\theta}$  for the floating snake. The gait image-family encircles a negative region of the height function; correspondingly, any gait in the image-family produces a net negative rotation.

### 5.3. Finding net rotation from the height function

The floating snake is the classic example of using a connection height function (Kelly and Murray 1995; Shamma et al. 2005, 2007a). Taking the curl of  $\vec{\mathbf{A}}^{\xi_\theta}$  produces the height function  $\mathbf{H}^{\xi_\theta}$  shown in Figure 7. For any given gait, integrating  $\mathbf{H}^{\xi_\theta}$  over the area enclosed by that gait gives the net rotation. For instance, the gait image-family overlaid on  $\mathbf{H}^{\xi_\theta}$  in Figure 7 contains the gait shown in Figure 4(a). It positively encircles a negative region of  $\mathbf{H}^{\xi_\theta}$ , and thus produces the net negative rotation seen in Figure 4(b).

Note that in Section 4.2.2, we attributed the net negative rotation to the difference in length between the 'positive' and 'negative' segments of the gait, but we are now attributing it to the value of the height function in the region encircled by the gait. Some intuition on the workings of the height functions can be gained by realizing that these two reasons are in fact closely linked: the height function measures the curl of the vector field that bends it

to align more closely to the second half of the gait than the first.

### 5.4. The BVI as an approximation of displacement

Although the example in Figure 6 shows that we cannot *always* take the BVI as an indication of the net displacement, it does not show that we can *never* do so. To investigate the existence of gaits for which the BVI is a reliable measure of displacement, we consider the error  $\varepsilon_\zeta = \zeta - g$  between the BVI and displacement over an arbitrary gait. By subtracting the displacement integral in Equation (17) from the BVI integral in Equation (18) and combining the integrands, we can express this error as

$$\varepsilon_\zeta = \int_0^t \begin{bmatrix} 1 - \cos \theta & \sin \theta & 0 \\ -\sin \theta & 1 - \cos \theta & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \xi_x \\ \xi_y \\ \xi_\theta \end{bmatrix} d\tau, \quad (24)$$

with an implicit dependence on  $\tau$  for all variables in the integrand. From Equation (24), we can easily identify a condition that induces small error terms: if  $\theta(t)$  remains small for all  $t$ , then the matrix in the integrand remains close to a zero matrix for all  $t$ , and the error  $\varepsilon_\zeta$  is thus close to a zero vector. We will thus turn our attention to finding situations for which this small-angle condition holds true. Note that both the displacement and the BVI are calculated relative to the system's starting configuration, that is  $g(0) = (0, 0, 0)$  for the purposes of this calculation; any difference between the starting position and a preferred origin of the space results in a transform that applies equally to  $g$  and  $\zeta$  and drops out of the error calculation in Equation (24).

Trivially, small gaits (i.e. gaits described by small-amplitude loops in the shape space) offer  $\theta$  no opportunity to grow large; these gaits have been studied extensively in the controls literature with the aid of Lie-algebra analysis (Murray and Sastry 1993; Ostrowski and Burdick 1998; Melli et al. 2006; Morgansen et al. 2007). Small gaits, however, spend more energy in high-frequency oscillations than they do in producing net motion and are thus inefficient; we are interested in designing larger, more efficient gaits. To this end, we observe that as the orientation of the system over a gait is the integral of its rotational body velocity,

$$\theta(t) = \int_0^t \xi_\theta(\tau) d\tau = \int_{\phi(0)}^{\phi(t)} \vec{A}^{\xi_\theta}(r) dr, \quad (25)$$

the magnitude of its value at any time is bounded to the line integral of the magnitude of the rotational connection vector field,

$$|\theta(t)| \leq \int_{\phi(0)}^{\phi(t)} |\vec{A}^{\xi_\theta}(r)| dr. \quad (26)$$

For gaits in regions where  $\|\vec{A}^{\xi_\theta}(r)\|$  is small,  $\theta(t)$  and  $\varepsilon_\zeta(t)$  are thus guaranteed to be small, and the BVIs of those gaits are thus good approximations of their displacements.

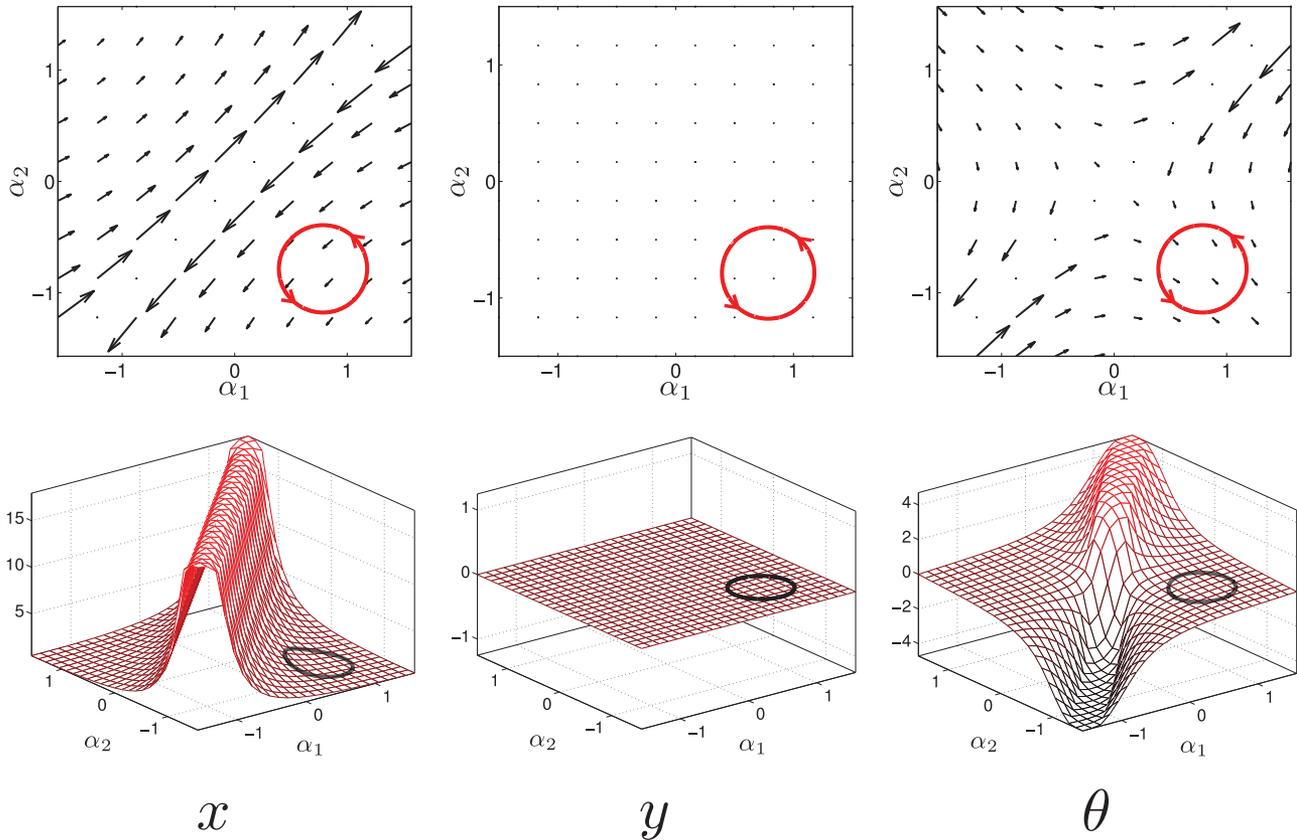
Unfortunately, this observation does not directly help us to apply the height function approach to translation gaits for systems like the kinematic snake. The  $\vec{A}^{\xi_\theta}(r)$  vector field for the kinematic snake in Figure 5 clearly does not vanish in any region, so we cannot expect the small-rotation condition to hold over gaits placed anywhere in the shape space. The most successful past attempt to address this limitation used an alternative definition of the height functions that includes second-order terms that account for rotation (Melli et al. 2006; Avron and Raz 2008). This higher-order approach allows for the design of gaits that are ‘small,’ rather than ‘differentially small,’ but does not extend to ‘large’ gaits. Another approach has been to withdraw from the goal of precise gait-evaluation via the height functions and to design gaits that have large BVIs, in the anticipation that they will produce correspondingly large displacements (Shammas et al. 2007a). This assumption that the magnitude of the displacement is correlated to that of the BVI has produced some usefully efficient gaits, but is unsatisfying as a provable motion-planning technique, as it relies on an assumption with direct counter-examples (such as in Figure 6). To bring the height functions into a form where they indicate displacement, a new approach is needed.

## 6. The importance of coordinate choice

As the final stage in developing a height function approach for systems like the kinematic snake, we now discuss the importance of coordinate choice. This constitutes a significant departure from previous work, in which the choice of body frame was considered largely unimportant, except that it facilitated expression of the system constraints. We have observed, however, that the choice of coordinate frame significantly affects the structure of the connection vector fields, including the presence or absence of the small-rotation condition described in Section 5.4. By carefully selecting our coordinates, then, we can manipulate the connection vector fields to meet these conditions, and thus produce useful height functions. Further, this selection can be accomplished by a simple change-of-coordinates operation after the local connection has been found in a convenient frame, thus preserving the advantages of the customary frame while gaining the benefits of the new frame.

### 6.1. BVI and displacement for a sample gait image-family

To demonstrate the benefits of changing the system's coordinates, we compare the performance of the connection height functions in approximating the displacement over a representative image-family of gaits  $\vec{\phi}$  under different choices of coordinates. As a baseline for comparison, we first consider the interaction of  $\vec{\phi}$  with the height functions corresponding to the original choice of coordinates.



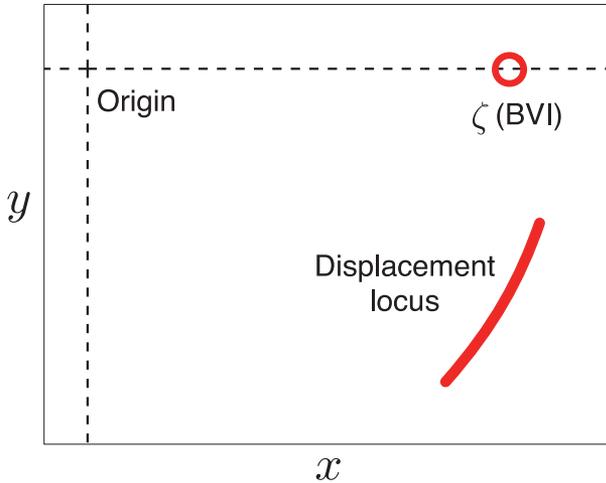
**Fig. 8.** Representative gait image-family overlaid on the connection vector fields and height functions of the kinematic snake in the original coordinate system depicted in Table 1. (Owing to singularities in the local connection, the magnitudes of the height functions and vector fields are scaled by their arctangents.)

Figure 8 shows the positive (counterclockwise) elements of  $\vec{\phi}$  overlaid on the connection vector fields and height functions. On the height functions, the gait family encircles a positive area on  $\mathbf{H}^{\zeta_x}$ , a zero area on  $\mathbf{H}^{\zeta_y}$ , and a net-zero area on  $\mathbf{H}^{\zeta_\theta}$ . As the rotational BVI and displacement are equal to each other, we can use the rules from Section 5.2, to conclude from the net-zero integral on  $\mathbf{H}^{\zeta_\theta}$  that the three-link robot undergoes no net rotation over the course of any gait from this family. However, because the  $\vec{\mathbf{A}}^{\zeta_\theta}$  field is non-vanishing the system will undergo intermediate rotation, forcing the BVI to diverge from the displacement.

The BVI for this image-family is shown in Figure 9 along with the locus of all of the net displacements generated by gaits in the image-family. The BVI is a point on the positive  $x$ -axis, with coordinates corresponding to the areas that the image-family encircles on the height functions, while the locus of net displacements is an arc with a significant  $-y$  component that is not captured by the BVI. Note that the arc of displacements is at constant radius from the origin: for a gait with zero net rotation, changing the starting shape only changes the *direction* of net translation as compared to the starting frame, and does not affect the magnitude of translation. A proof of this property is given in Appendix A.

### 6.2. BVI and displacement in a mean-orientation frame

While the  $\vec{\mathbf{A}}^{\zeta_\theta}$  field for the kinematic snake in Figure 8 had no region of small magnitude, our experience has shown that changing the definition of the system’s orientation substantially changes the structure of this field, to the degree that a good choice of coordinates can induce a null region in  $\vec{\mathbf{A}}^{\zeta_\theta}$ . For instance, instead of using the orientation of the middle link as the  $\theta$  component of  $g$ , we can derive the local connection with  $\theta$  representing the mean orientation of the three links, as shown in Figures 10 and 11. The connection vector fields and height functions corresponding to this coordinate choice are plotted with  $\vec{\phi}$  in Figure 12. With this new choice of coordinates,  $\vec{\mathbf{A}}^{\zeta_\theta}$  now has a region of small magnitude in the same region as  $\vec{\phi}$ , making the BVI a good approximation of the displacement for gaits in this image-family. The improvement of this estimate as compared to the estimate under the previous coordinate choice is shown in Figure 13, with both a significant reduction in the variance of net displacements and a much closer placement of the BVI to this locus. Figure 14 illustrates the underlying principle of this reduction in variance: Under the original coordinate choice, the two gaits selected from the example

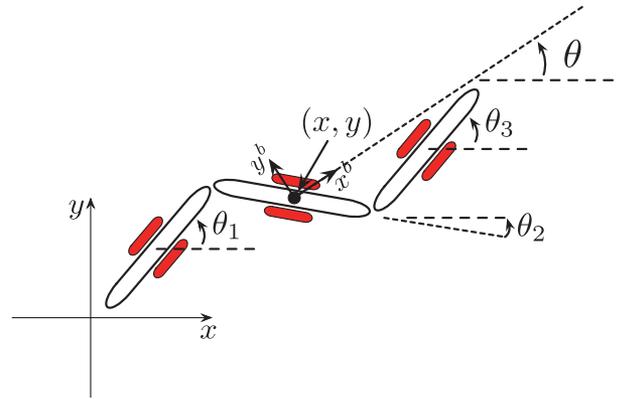


**Fig. 9.** The BVI and locus of displacements corresponding to the image-family of gaits for the kinematic snake robot, depicted in Figure 8, are represented, respectively, by the circle on the  $x$ -axis and the arc section. The BVI only depends on an area integral over the area enclosed by the gait and is thus the same for all gaits in the family. The displacement, however, also depends on the start/end shape of the gait, and there is, thus, a range of displacements, each corresponding to one starting point in the image-family. The displacements are of the order of one link length of the robot.

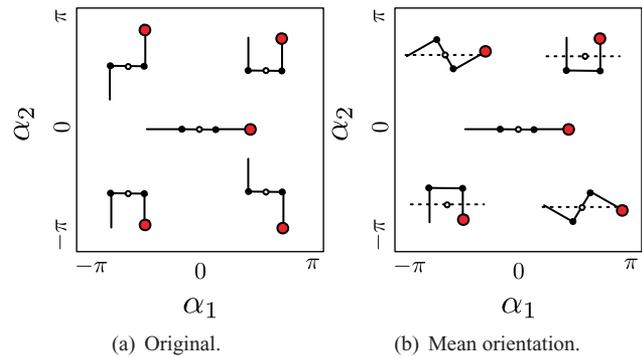
image-family produce displacements of the same magnitude, but in different directions, whereas the new choice of orientation realigns these displacements to point in the same direction.

### 6.3. Identifying valid coordinate choices

The increased utility of the connection height functions in the mean-orientation coordinate system raises two important questions: First, what properties of the mean-orientation coordinate system led to the improvement? Second, what other coordinate sets are available, and might they provide even more nullification of the  $\tilde{\mathbf{A}}^{\xi\theta}$  field? Starting with the second question, we offer the notion of a *valid coordinate set* as one that induces a body frame for which the body velocity is produced by a reconstruction equation of the form of Equation (1); for kinematic systems, this reconstruction equation should reduce to the form of Equation (3). This notion is compatible with the derivations of the reconstruction equation in Kelly and Murray (1995), Ostrowski (1998), and Shamma et al. (2007a), which also include a means of testing the validity of individual coordinate sets. The test used in those works is, unfortunately, only useful for validating given coordinate sets, and does not provide a means for finding new coordinate sets except by trial and error. However, if we have a valid coordinate set from which to start, such as one of the frames in Table 1, we can use the following lemma to directly find other valid coordinate sets.



**Fig. 10.** New orientation definition for kinematic snake. The orientation  $\theta$  is the mean of the individual link orientations  $\theta_1$ ,  $\theta_2$ , and  $\theta_3$ . Note that the body frame directions  $x^b$  and  $y^b$  are, respectively, aligned with and perpendicular to the orientation line, and not to the central axis of the middle link.

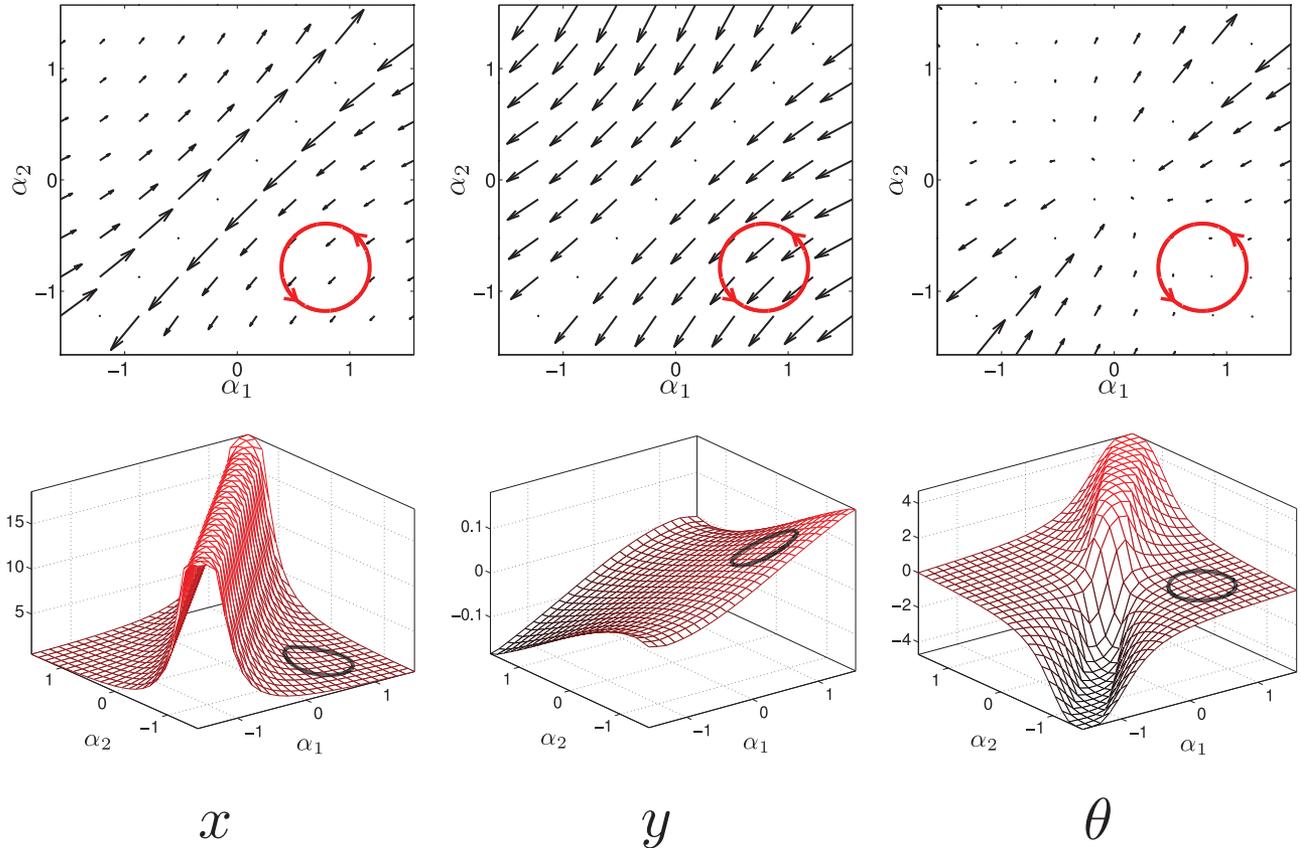


**Fig. 11.** A three-link system in various configurations. In both (a) and (b), the system is shown at five points on the shape space with reference orientation  $\theta = 0$ . In (a),  $\theta$  is the orientation of the middle link, while in (b),  $\theta$  is the mean orientation of the links.

**Lemma 6.1 (Relationship of valid coordinate sets).** *In any valid coordinate set (i.e. one for which a local connection can be derived), the body frame of the system described by the coordinates is related by a shape-dependent transform to the body frame in a coordinate set known to be valid.*

*Proof.* Consider the two coordinate frames shown in Figure 15. The frame designated by  $(x_{old}^b, y_{old}^b)$  represents an original, known-valid body frame for the system and the frame designated by  $(x_{new}^b, y_{new}^b)$  is a second frame that is a candidate to be a body frame under a new choice of coordinates. For any such frame, we can define its position and orientation in the original body frame to be  $\beta = (\beta_x, \beta_y, \beta_\theta)$ , as shown in Figure 15. The body velocity  $\xi_{new}$  of this frame, that is, its velocity as measured along its own instantaneous coordinates, is

$$\xi_{new} = B_{rot} \cdot (\xi + \dot{\beta} + B_\times \xi_\theta), \tag{27}$$



**Fig. 12.** Representative gait image-family overlaid on the connection vector fields and height functions of the kinematic snake in the mean-orientation coordinate system depicted in Figure 10. (Owing to singularities in the local connection, the magnitudes of the height functions and vector fields are scaled by their arctangents.)

where  $\dot{\beta}$  represents the relative velocity of the two frames,  $B_{\times}$  is the cross product term that gives the translational velocity of the new frame produced by rotation of the old frame,

$$B_{\times} = \begin{bmatrix} -\beta_y \\ \beta_x \\ 0 \end{bmatrix}, \tag{28}$$

and  $B_{\text{rot}}$  is the rotation matrix

$$B_{\text{rot}} = \begin{bmatrix} \cos \beta_{\theta} & \sin \beta_{\theta} & 0 \\ -\sin \beta_{\theta} & \cos \beta_{\theta} & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{29}$$

that handles the conversion of the velocity vector between the two frames.

For the kinematic systems we are considering, if this new frame is a valid body frame for the system, then  $\xi^{\text{new}}$  must be produced by a kinematic reconstruction equation of the form of Equation (3). We first address the sufficiency of our condition on  $\beta$ : when  $\beta$  is a function of  $r$ , the relative velocity between the frames is

$$\dot{\beta} = \frac{\partial \beta}{\partial r} \dot{r} \tag{30}$$

and the cross-product term is also a function of the shape, so Equation (27) thus becomes

$$\xi^{\text{new}} = B_{\text{rot}}(r) \cdot \left( \xi + \frac{\partial \beta}{\partial r} \dot{r} + B_{\times}(r) \xi_{\theta} \right). \tag{31}$$

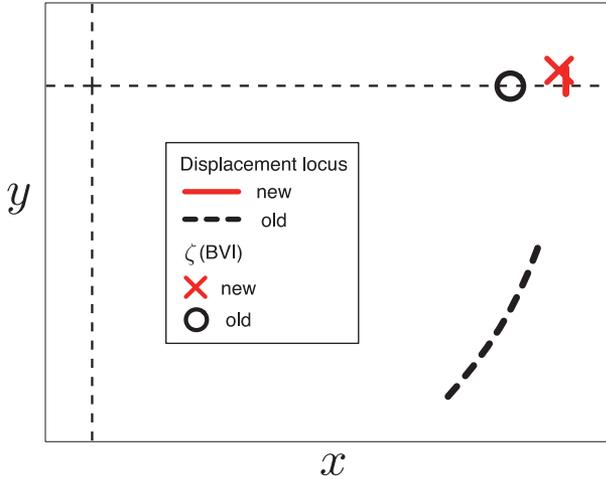
By substituting in the original kinematic reconstruction Equation (3) into Equation (31) and regrouping,

$$\xi^{\text{new}} = B_{\text{rot}}(r) \cdot \left( -\mathbf{A}(r) \dot{r} + \frac{\partial \beta}{\partial r} \dot{r} - B_{\times}(r) \mathbf{A}^{\xi_{\theta}}(r) \dot{r} \right) \tag{32}$$

$$= B_{\text{rot}}(r) \cdot \left( -\mathbf{A}(r) + \frac{\partial \beta}{\partial r} - B_{\times}(r) \mathbf{A}^{\xi_{\theta}}(r) \right) \dot{r} \tag{33}$$

$$= -\mathbf{A}_{\text{new}}(r) \dot{r}, \tag{34}$$

we see that  $\xi^{\text{new}}$  is produced by a matrix of the right form to be a local connection, and that having  $\beta$  as a function of  $r$  is sufficient to define a valid body frame. To see that it is also a necessary condition, consider the effect of defining  $\beta$  as a function of anything other than  $r$ . In this situation, the expansion of Equation (27) into Equation (31) would contain terms that cannot be collected into  $\mathbf{A}_{\text{new}}(r)$  in Equation (34), precluding such a  $\beta$  from meeting our definition of a valid coordinate system.  $\square$



**Fig. 13.** The BVI and displacements corresponding to the image-family of gaits in Figure 12 for the kinematic snake robot with the new measure of orientation depicted in Figure 10 are represented, respectively, by the cross above the  $x$ -axis and the short arc section, with the BVI and displacements as measured in the old coordinate choice presented for reference. With the new measure, the orientation of the system remains almost constant over the course of any of the gaits in the family, and the BVI  $\zeta$  is thus a good approximation of the displacements.

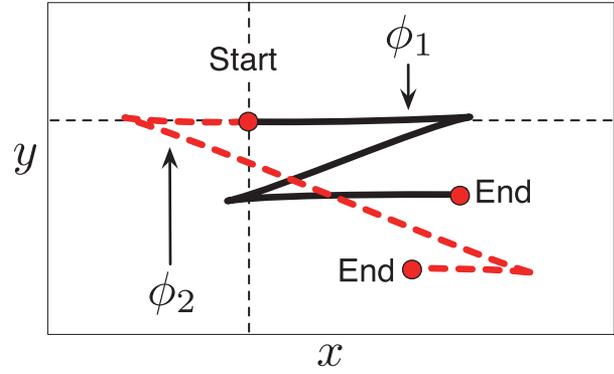
6.4. Finding an optimal coordinate choice

Having defined the valid body frames, and thus the available choices of coordinates, we can address the questions of why the mean-orientation body frame was a good choice and how to systematically find an optimal choice.

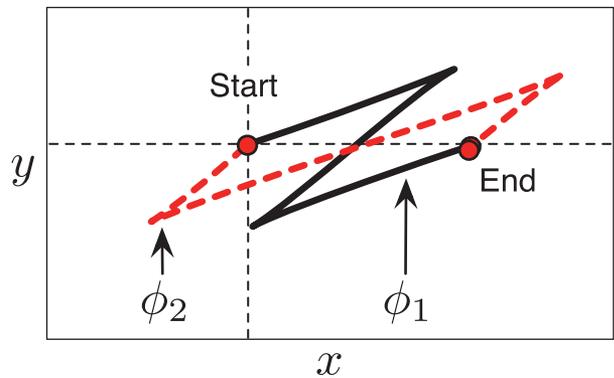
6.4.1. Optimal choice of orientation The key to finding the optimal choice of orientation lies in the third row of the  $\frac{\partial \beta}{\partial r}$  term in Equation (31),  $\frac{\partial \beta_\theta}{\partial r}$ . As  $\beta_\theta$  is a function of  $r$ ,  $\frac{\partial \beta_\theta}{\partial r}$  is its gradient with respect to the shape,  $\nabla_r \beta_\theta$ . Because  $B_{rot}$  and  $B_x$  only contribute to the translational components of the new local connection, any difference between  $\tilde{A}^{\xi_\theta}$  and  $\tilde{A}_{new}^{\xi_\theta}$  is due to the addition of  $\nabla_r \beta_\theta$ .

When moving from the center-link frame for the kinematic snake to the mean-orientation frame,  $\beta_\theta = (-\alpha_1 + \alpha_2)/3$  and its gradient is  $\nabla_r \beta_\theta = [-\frac{1}{3} \quad \frac{1}{3}]$ , as shown in Figure 16. Figure 17 visually represents the addition of the gradient field to the original  $\tilde{A}^{\xi_\theta}$  field. In the two circled regions, the fields cancel each other out, creating the null regions that we observed in Figure 12. In these regions, the new body frame counter-rotates with respect to the center-link frame, and thus rotates very little with respect to the world in response to shape changes.

Based on this observation of what makes the mean-orientation frame an effective choice of coordinates, we now have a criterion for an optimal choice of coordinates: It should induce a  $\nabla_r \beta_\theta$  field that most nullifies  $\tilde{A}^{\xi_\theta}$ . While the exact meaning of ‘most nullifies’ is open to interpretation, for now we will take it as a minimization of the average

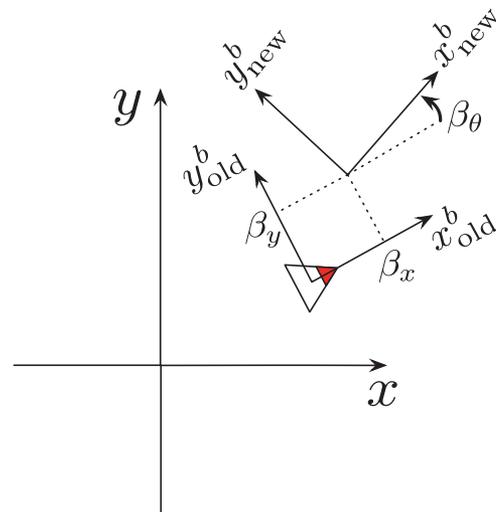


(a) Original coordinate choice.

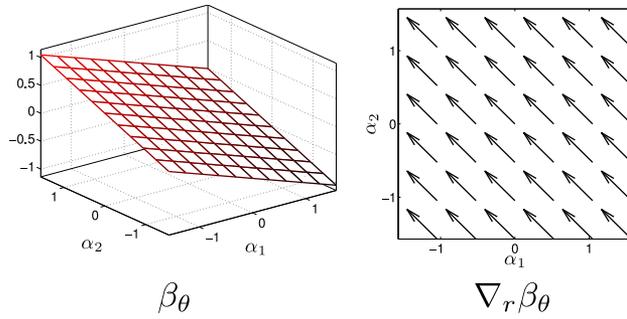


(b) Mean-orientation coordinate choice.

**Fig. 14.** Position trajectories for two gaits taken from the example image-family in Figure 12, as measured in the original and mean-orientation coordinate frames. In the mean-orientation coordinate frame, the trajectories are reoriented such that their net displacements are in the same direction.



**Fig. 15.** An original, known-valid body frame, and a new frame displaced from that frame by  $\beta$



**Fig. 16.** The relative orientation  $\beta_\theta$  and its gradient  $\nabla_r\beta_\theta$  between the mean-orientation frame and the original center-link frame for the kinematic snake, as a function of the shape  $\alpha$ .

squared magnitude  $\vec{\mathbf{A}}_{\text{new}}^{\xi_\theta}$  over the region of interest, that is, minimizing  $D$  for

$$D = \iint_{\Omega} \|\vec{\mathbf{A}}^{\xi_\theta} + \nabla_r\beta_\theta\|^2 \, d\Omega \quad (35)$$

for  $\Omega \subset M$ . Because  $\beta_\theta$  is a function of  $r$ , we can generate it directly from  $\nabla_r\beta_\theta$  (modulo an integration constant) and we can shift attention to directly finding  $\nabla_r\beta_\theta$  that minimizes Equation (35).

Identifying optimal coordinate choices with minimizations of  $D$  in Equation (35) provides the key to finding the best coordinate choices for arbitrary systems. The cost function  $D$  is minimized when  $\nabla_r\beta_\theta$  is the gradient field that most closely approximates  $-\vec{\mathbf{A}}^{\xi_\theta}$ , that is, it is the negative projection of  $\vec{\mathbf{A}}^{\xi_\theta}$  onto the space of conservative vector fields. This projection is identical to that in the *Hodge–Helmholtz decomposition* (Arfken 2005), which separates a vector field into a conservative component and a divergence-free remainder. Here, the conservative component is equal to  $-\nabla_r\beta_\theta$ , while the remainder becomes  $\vec{\mathbf{A}}_{\text{new}}^{\xi_\theta}$ . While an analytical solution for this projection requires that  $\vec{\mathbf{A}}^{\xi_\theta}$  meet certain regularity conditions and is not generally feasible, the *discrete Hodge–Helmholtz decomposition* (Guo et al. 2005) provides a numerical minimization of Equation (35) and is relatively easy to compute for arbitrary vector fields via finite-element methods, as explained in Appendix C.

In practice, we have found that an extra refinement to the decomposition is necessary when  $\vec{\mathbf{A}}^{\xi_\theta}$  has singularities. The large magnitudes of such fields in the vicinity of the singularities dominate the integral in Equation (35), producing a  $\nabla_r\beta_\theta$  field that closely approximates  $-\vec{\mathbf{A}}^{\xi_\theta}$  at the singularities, but at the expense of any accuracy over the rest of the field. We address this limitation by using the same arctangent scaling we use to display the vector fields, and instead solve the minimization problem with

$$D = \iint_{\Omega} \left\| \frac{\arctan(k\|\vec{\mathbf{A}}^{\xi_\theta}\|)}{k\|\vec{\mathbf{A}}^{\xi_\theta}\|} \vec{\mathbf{A}}^{\xi_\theta} + \nabla_r\beta_\theta \right\|^2 \, d\Omega, \quad (36)$$

where the scaling factor  $k = 0.2/\text{median}(\|\vec{\mathbf{A}}^{\xi_\theta}\|)$  modifies the arctangent function such that the center of its linear region corresponds to the median value of the range of input vector magnitudes. The arctangent scaling acts as a selective weighting function that emphasizes the quality of the fit in the low-control-force regions of the shape space, away from the singularities. Details of this scaling function are given in Appendix B.

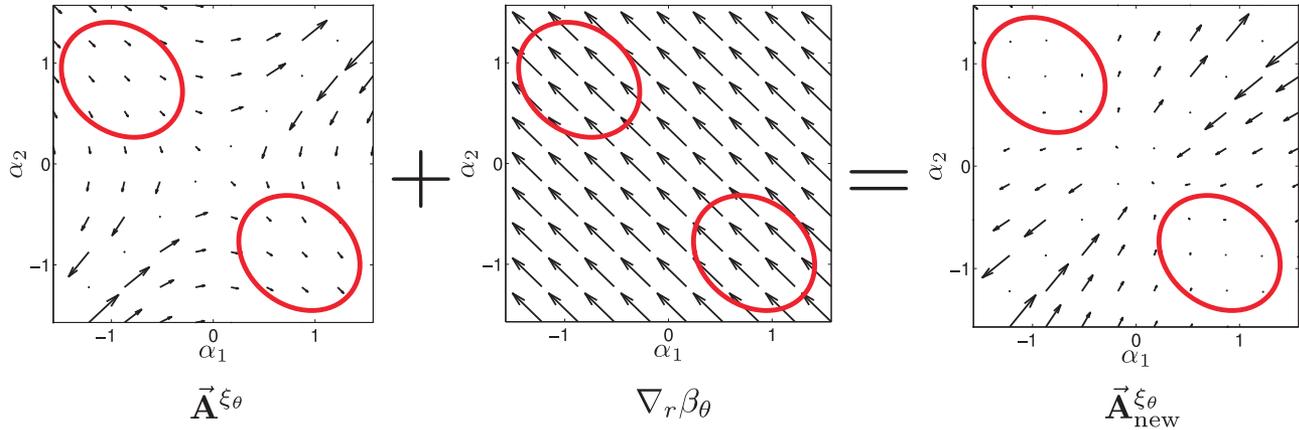
**6.4.2. Optimal choice of reference position** While the choices of  $\beta_x$  and  $\beta_y$  do not affect the magnitude of  $\vec{\mathbf{A}}^{\xi_\theta}$ , they do contribute to the goal of minimizing the error between the BVI and the displacement. Returning to Equation (24), we observe that the rotation matrix term is multiplied by the translational body velocity,  $[\xi_x \ \xi_y]^T$ . It, therefore, follows that minimizing this body velocity by choosing a coordinate system that minimizes  $\vec{\mathbf{A}}^{\xi_x}$  and  $\vec{\mathbf{A}}^{\xi_y}$  will further reduce the error beyond the reduction already accomplished by an optimal choice of  $\beta_\theta$ . We have already seen an ideal example of this principle in the floating snake: using the center of mass as the reference position induces zero-magnitude fields for  $\vec{\mathbf{A}}^{\xi_x}$  and  $\vec{\mathbf{A}}^{\xi_y}$ , and the translational BVI and displacement over a gait are both identically zero with no error. If another reference point, such as the center of the middle link, had been used,  $\vec{\mathbf{A}}^{\xi_x}$  and  $\vec{\mathbf{A}}^{\xi_y}$  would not be zero, and the BVI would not match the displacement.

Finding the optimal choice of position coordinates has a similar formulation to finding the optimal choice of orientation, in that we are looking for a change of coordinates that most nullifies the translational connection vector fields. The key difference from the orientation optimization lies in the  $B_\times$  term, which is now non-zero. Beyond adding a third term into the summations of vector fields that produce  $\vec{\mathbf{A}}_{\text{new}}^{\xi_x}$  and  $\vec{\mathbf{A}}_{\text{new}}^{\xi_y}$ ,  $B_\times$  also couples these two calculations, so we must simultaneously solve for  $\beta_x$  and  $\beta_y$  to minimize  $D$  for

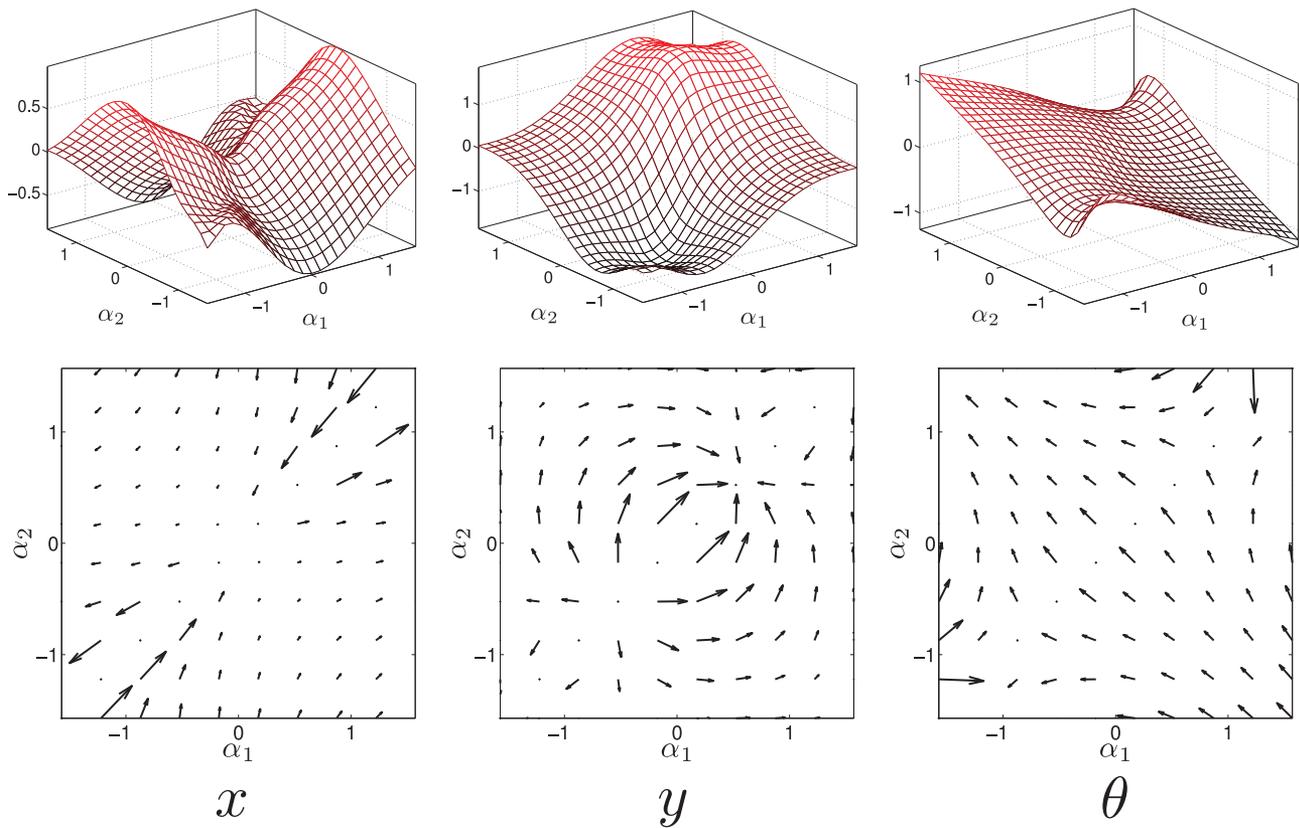
$$D = \iint_{\Omega} \|\vec{\mathbf{A}}^{\xi_x} + \nabla_r\beta_x - \beta_y\vec{\mathbf{A}}^{\xi_\theta}\|^2 + \|\vec{\mathbf{A}}^{\xi_y} + \nabla_r\beta_y + \beta_x\vec{\mathbf{A}}^{\xi_\theta}\|^2 \, d\Omega. \quad (37)$$

As with the discrete Hodge–Helmholtz decomposition, we take a finite-element approach to this minimization, the details of which appear in Appendix D. Additionally, when singularities are present, we make use of the arctangent scaling approach in Equation (36) to emphasize the regions of the shape space away from the singularities. Note that  $\vec{\mathbf{A}}^{\xi_\theta}$  in this equation is for the original choice of coordinates, and that  $B_{\text{rot}}$  does not appear in the optimization as it does not affect the magnitude of the connection vector fields.

**6.4.3. Summary of coordinate optimization** In summary, the coordinate optimization process is as follows. First, we use the Hodge–Helmholtz decomposition to find the



**Fig. 17.** The effect of adding  $\nabla_r \beta_\theta$  for the mean-orientation coordinate choice to  $\vec{A}^{\xi_\theta}$ . In the circled regions,  $\nabla_r \beta_\theta \approx -\vec{A}^{\xi_\theta}(r)$ , so  $\|\vec{A}_{\text{new}}^{\xi_\theta}\| \approx 0$  in these regions. Note that for visual clarity within each plot, the vectors in the different fields are not to scale.

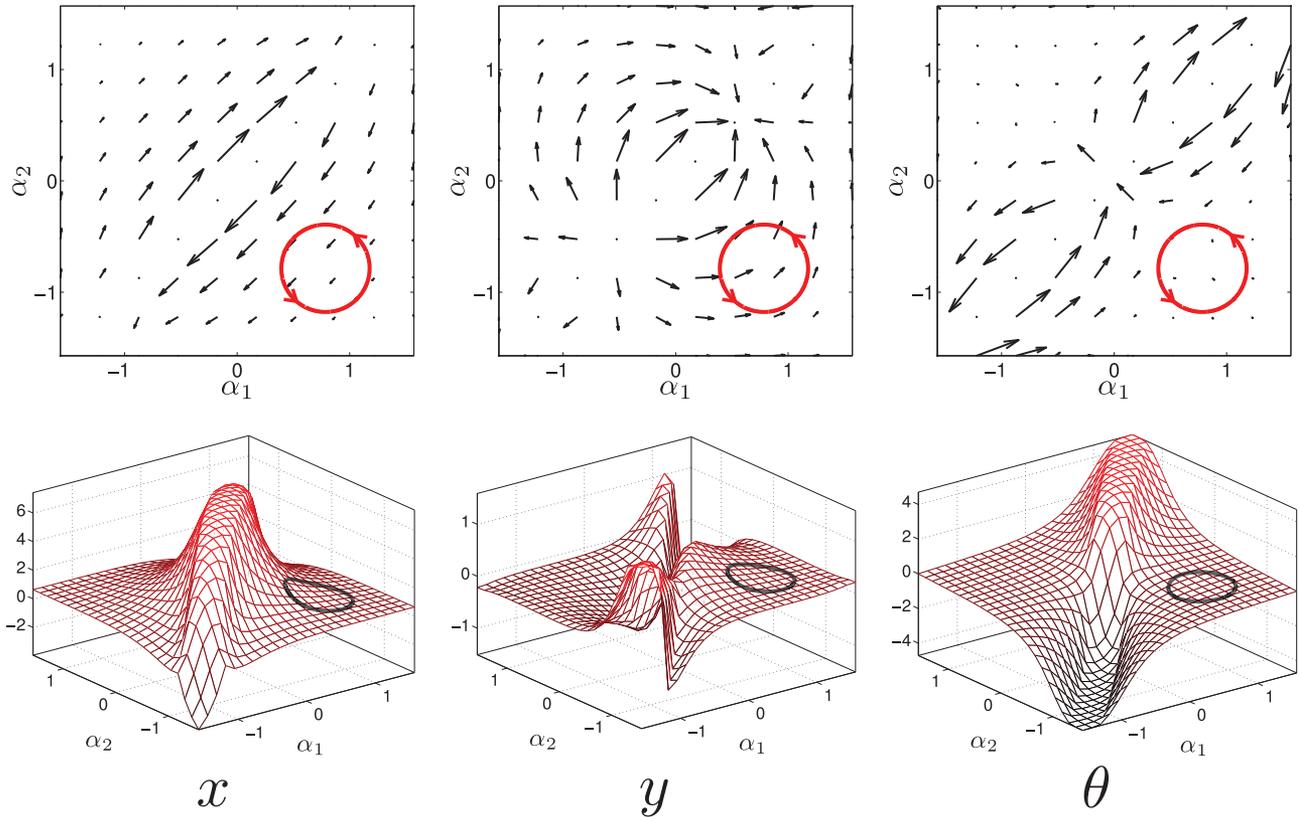


**Fig. 18.** Optimal  $\beta$  (top) and  $\nabla_r \beta$  (bottom) for the kinematic snake, using the arctangent metric in Equation (36).

$\nabla_r \beta_\theta$  field that minimizes  $\vec{A}^{\xi_\theta}$ . Second, we integrate this field to find  $\beta_\theta(r)$ , choosing the integration constant so that  $\beta_\theta(0,0) = 0$ ; from  $\beta_\theta$ , we generate the  $B_{\text{rot}}$  operator that converts the  $x$  and  $y$  components of the body velocity into the new frame. Third, solving the coupled optimization problem in Equation (37) provides us with  $\nabla_r \beta_x$ ,  $\nabla_r \beta_y$ , and  $B_\times$ . Finally, we insert these terms into Equation (33) to generate the new, optimized local connection.

### 6.5. BVI vs. displacement in optimal frame

Applying the discrete Hodge–Helmholtz decomposition and the position optimization to the connection vector fields for the kinematic snake produces the  $\beta$  and  $\nabla_r \beta$  functions shown in Figure 18. The first conclusion we can draw from these functions is that the mean-orientation coordinates are close to optimal for orientation: the plots of  $\beta_\theta$  and  $\nabla_r \beta$



**Fig. 19.** Representative gait image-family overlaid on the connection vector fields and height functions of the kinematic snake in the optimized coordinate system. Compared to the mean-orientation coordinate system,  $\vec{\mathbf{A}}^{\xi\theta}$  is nullified over a larger area, and  $\mathbf{H}^{\xi y}$  shows more variation. (Owing to singularities in the local connection, the magnitudes of the height functions and vector fields are scaled to their arctangents.)

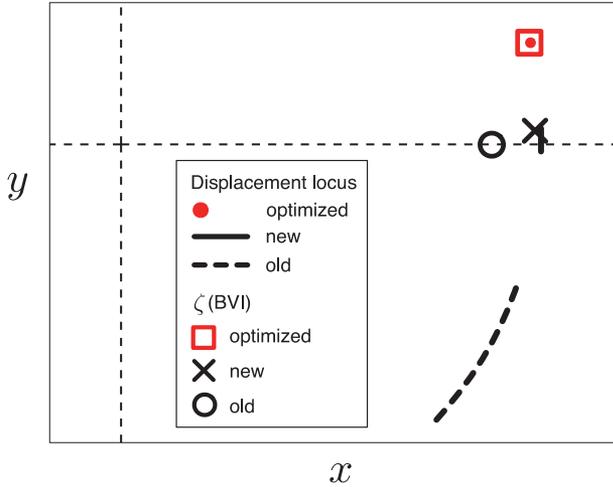
strongly resemble their mean-orientation counterparts in Figure 16. The chief difference between them is that the optimization process warps  $\beta_\theta$  so that  $\nabla_r \beta_\theta$  more closely approximates  $-\vec{\mathbf{A}}^{\xi\theta}$  near the edges of the region under consideration. As show in Figure 19, this change in  $\nabla_r \beta_\theta$  leads to the nullification of  $\vec{\mathbf{A}}^{\xi\theta}$  over a wider area than was the case for the mean-orientation coordinate system. Figure 20 shows the marked improvement of the BVI as an estimate of displacement in the optimized coordinate system. The span of the locus of displacements is reduced to the extent that it can be considered as a point, and the BVI calls out the same location as this point.

It is worth noting that the degree to which the new coordinate system nullifies  $\vec{\mathbf{A}}^{\xi\theta}$  is inversely proportional to the magnitude of  $\mathbf{H}^{\xi\theta}$ , and that  $\mathbf{H}^{\xi\theta}$  itself does not change. The former effect is explained by considering that as  $\mathbf{H}^{\xi\theta}$  is the curl of  $\vec{\mathbf{A}}^{\xi\theta}$ , it measures the non-conservativity of  $\vec{\mathbf{A}}^{\xi\theta}$ , so that the smaller the value of  $\mathbf{H}^{\xi\theta}$  in a region, the more we can nullify  $\vec{\mathbf{A}}^{\xi\theta}$  by the addition of a conservative field. The latter effect comes from the identity between the curl of the sum of two vector fields and the sum of their curls; the conservative field we added to  $\vec{\mathbf{A}}^{\xi\theta}$  during optimization has zero curl by definition, and so does not change the height function.

The  $\beta_x$  and  $\beta_y$  functions found by the coordinate optimization move the center of the new coordinate frame closer to a ‘true center’ of the system. For example, as illustrated in Figure 18,  $\beta_x$  has a saddle structure that increases with the distance of  $\alpha_1$  from 0 radians and decreases as  $\alpha_2$  moves away from 0 radians. This behavior corresponds to the system’s center of constraint<sup>5</sup> (roughly at the same location as the center of mass) being drawn forwards or backwards as it moves away from a fully extended configuration. Similarly,  $\beta_y$  increases and decreases proportionally to both  $\alpha_1$  and  $\alpha_2$  as the outer links move away from the centerline of the middle link.

### 6.6. Generality

Having demonstrated the accuracy of the BVI as an estimate of displacement for our example gait image-family, we now expand the scope of our analysis. To be an effective tool for motion planning, the BVI must closely approximate the displacement over a broad domain of gaits. One way to measure the size of the domain over which the approximation holds is to consider the ‘largest’ gaits that can be executed while keeping the error within acceptable

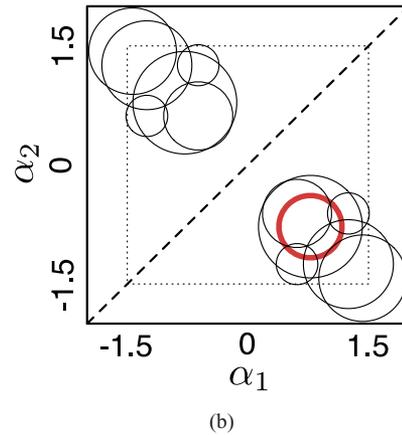
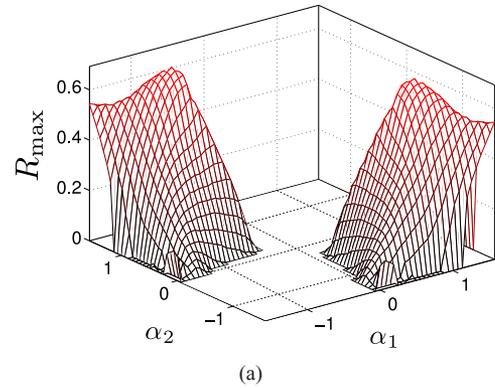


**Fig. 20.** The BVI and displacements, corresponding to the image-family of gaits in Figure 19 for the kinematic snake robot with the optimized measure of orientation from Figure 18, are represented, respectively, by the box above the  $x$ -axis and a very short arc section at the same location, with the BVI and displacements as measured in the original and mean-orientation coordinate choices shown for reference.

bounds. While it is naturally infeasible to find the BVI and displacement for all possible gaits, it is straightforward to consider the set of all circular gaits. From this set, we can then find the maximum radius for a circular gait image-family centered at each point in the shape space for which the error percentage is less than a given value, that is for which

$$\frac{\max_{\phi} \|\zeta - g_{\bar{\phi}}\|}{\min_{\phi} \|g_{\bar{\phi}}\|} < \epsilon_{\max}, \quad (38)$$

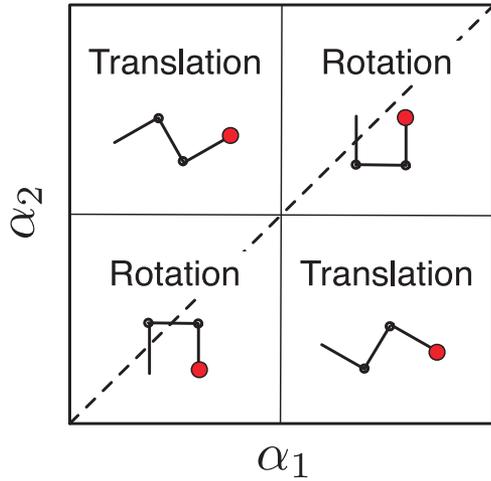
where  $g_{\bar{\phi}}$  is the locus of net translations corresponding to the image-family,  $\max \|\zeta - g_{\bar{\phi}}\|$  is the maximum distance from the BVI to a point on the net translation locus, and  $\min \|g_{\bar{\phi}}\|$  is the minimum magnitude of the net translation in the locus. Figure 21 shows this maximum radius for an acceptable error of  $\epsilon_{\max} = 0.1$ , both as a function on the shape space and as a representative collection of circular image-families. For reference, the gait used in the previous examples is drawn along with the 10% error gaits, showing it to be well within the range of allowable radii for its location. Note that by calculating this maximum radius, we are not suggesting that such calculations should form a part of motion planning, but rather are demonstrating the effects of the coordinate change on our example system. Also, the use of circular or elliptical gaits here and elsewhere should not be taken as a restriction on the applicability of the results; they appear here because they allow us to examine the underlying principles without adding in extra complicating factors in the form of gait parameterizations.



**Fig. 21.** Maximum radius for a circular gait image-family for  $\epsilon_{\max} = 0.1$ , (a) plotted as a function on the shape space, and (b) drawn as a representative collection of image-families. In (b), the dashed line denotes the  $\alpha_1 = \alpha_2$  singularity, while the dotted box marks the extent of the domain of the plot in (a). The thick-lined circle in the lower-right portion of the drawing is the gait image-family used in the main examples.

As expected, the regions of the shape space in which the BVI serves well as an approximation of the net displacement correspond to the regions in which  $\vec{A}^{\xi_{\theta}}$  has been nullified by the change of coordinates. It is worth noting that these regions are also inherently good regions in which to search for translation gaits, both because their distance from constraint singularities results in small required control forces, and because the low magnitude of  $\mathbf{H}^{\xi_{\theta}}$  in these regions produces a high density of gaits that produce pure net translation with zero net rotation, which is important when designing gaits that can be repeated to locomote the system in a straight line. The change of coordinates thus not only makes it possible to use the BVI to design translational gaits, but does so in the most useful regions of the shape space.

Taking this observation one step further, we can roughly partition the shape space into regions containing translational and rotational gaits, as illustrated in Figure 22. The translational regions have the properties noted above, while the rotational regions have large magnitudes of  $\mathbf{H}^{\xi_{\theta}}$ ,



**Fig. 22.** Partitioning of the shape space into regions that contain effective translational or rotational gaits. The translational regions correspond to S-shaped configurations of the robot, while the rotational regions correspond to C-shapes.

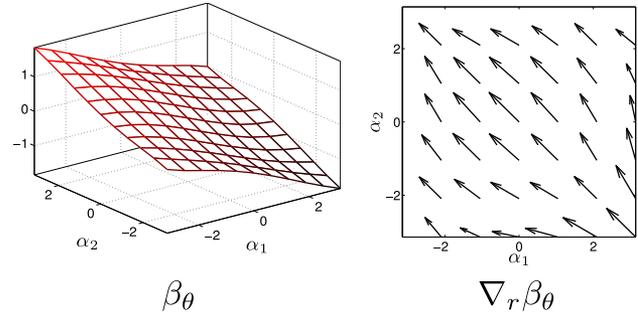
allowing for gaits that efficiently turn shape motion into net rotation. This partitioning bears strong similarity to the findings by Melli et al. (2006) that an S-shape for their three-link swimmer was most effective for producing translation, while a C-shape was better for rotation.

### 7. Optimal coordinates for other systems

While our chief interest in developing the coordinate optimization approach has been to analyze non-traditional locomotors such as the kinematic snake, applying coordinate optimization to the two other systems discussed earlier yields some interesting insights.

#### 7.1. Floating snake

As we noted previously, conservation of linear momentum dictates that the floating snake’s center of mass remains stationary when the system starts at rest. Accordingly, we can conclude that the center of mass is an inherently optimal choice of reference point according to Equation (37), and thus that optimizing the coordinate choice for the system reduces to finding the optimal choice of orientation. While our initial motivation for optimizing coordinate choice, making the BVI an approximation of displacement, is less relevant for the floating snake (as any gait for the system will still of course generate zero net displacement and zero translational BVI, along with a change in orientation that exactly matches the rotational BVI), working with an optimized choice of orientation still better captures the ‘true motion’ of the system. In this case, we regard the ‘true motion’ as the actual contribution of a shape change to the net rotation produced by a gait containing that shape change.



**Fig. 23.** The relative orientation  $\beta_\theta$  and  $\nabla_r \beta_\theta$  between the optimized orientation frame and the original center-link frame for the floating snake, as a function of the shape  $\alpha$ .

Similarly to the kinematic snake, the floating snake’s optimal orientation is a weighted average of the orientations of the three links, with  $\beta_\theta$  and  $\nabla_r \beta_\theta$  as plotted in Figure 23. Figure 24 shows the example gait for the floating snake overlaid onto the original and optimized  $\vec{A}^{\xi_\theta}$  fields, along with the time history of the orientation of the system using both measures. Using the original definition of orientation, the system exhibits significant positive ‘windup’ followed by a ‘power stroke’ that brings it to its final negative orientation. The new orientation, however, monotonically approaches the final negative value. The difference between the two plots highlights how the optimal orientation found by minimizing Equation (35) separates out the gradient component of  $\vec{A}^{\xi_\theta}$ : the ‘extra’ motion in the original coordinate choice corresponds to this (conservative) gradient component, and thus by definition cancels with itself over a gait. With the new choice of orientation, this self-canceling motion is stripped away, leaving only the actual contribution of a given change in shape to the net orientation change.

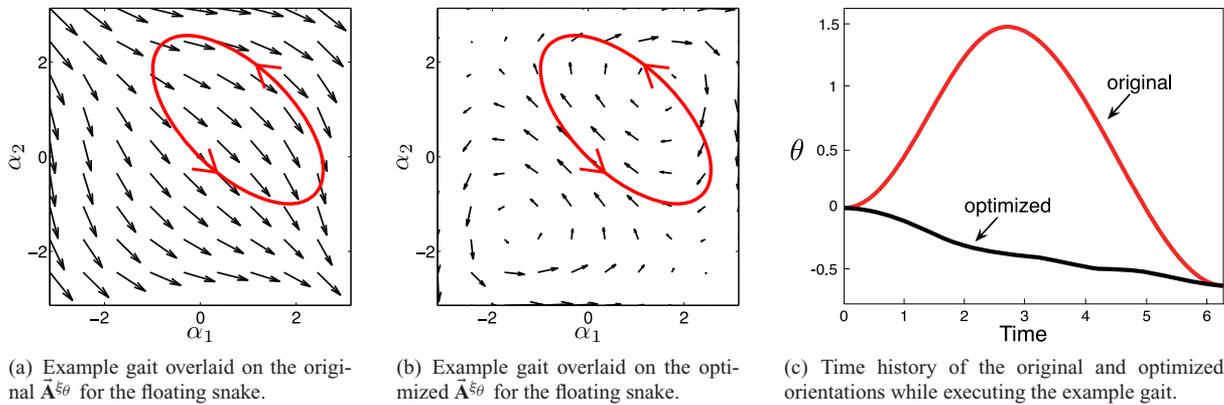
#### 7.2. Differential-drive car

The differential-drive car presents a second special case of coordinate optimization. For this system  $\vec{A}^{\xi_\theta}$  is a constant (and, thus, conservative) vector field, and can, therefore, be exactly nullified by the addition of  $\nabla_r \beta_\theta = -\vec{A}^{\xi_\theta}$ . In the coordinates defined by the implied choice of

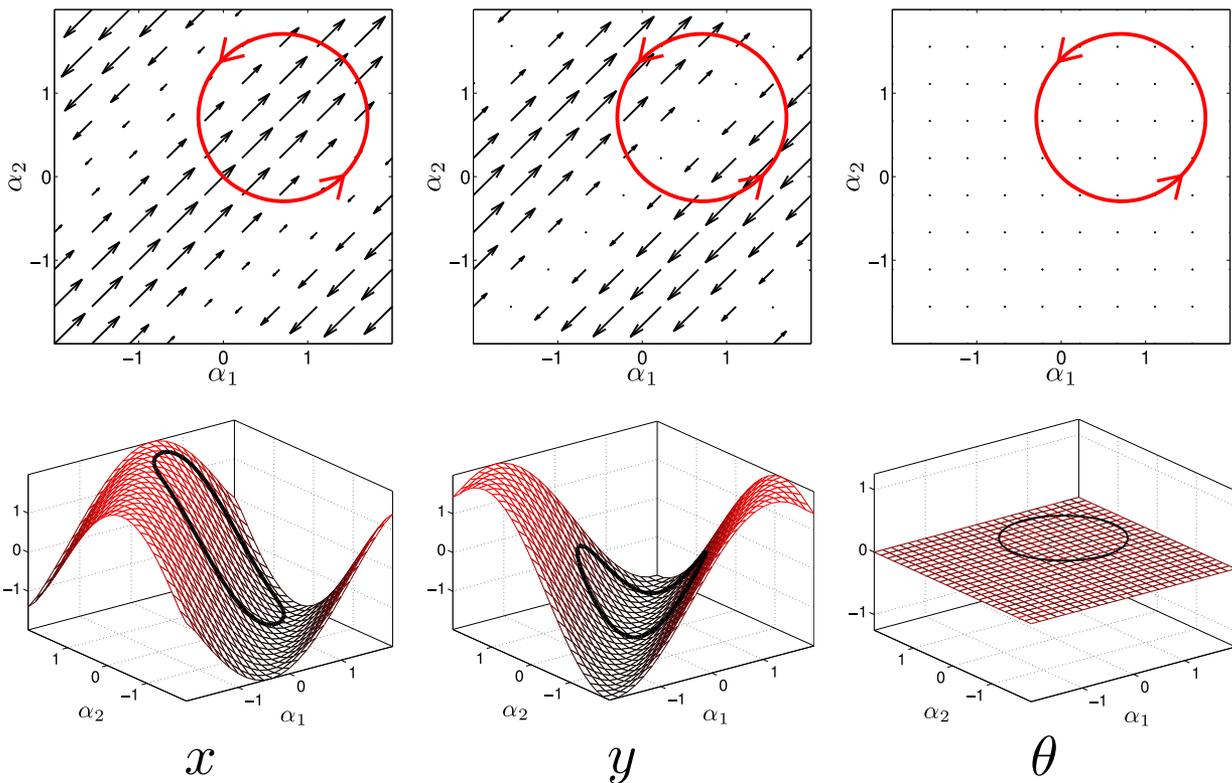
$$\beta_\theta = \alpha_1 - \alpha_2, \tag{39}$$

$\dot{\theta}$ ,  $\theta$ , and  $\varepsilon_\zeta$  are all identically zero, so there is no need to optimize the choice of reference point. Figure 25 shows the optimized connection vector fields and height functions for the differential-drive car, along with a parallel-parking gait that produces the trajectory in Figure 26. The gait encircles a net-zero region on the  $\mathbf{H}^{\zeta_x}$  height function and a negative region on  $\mathbf{H}^{\zeta_y}$ ; as  $\zeta_x$  and  $\zeta_y$  are, respectively, equal to the  $x$  and  $y$  displacements with no error, the resulting motion has net-lateral displacement.

The optimal coordinates for the differential-drive car reveal some interesting relationships between that system



**Fig. 24.** Optimizing the definition of  $\theta$  for the floating snake strips out the gradient component from  $\vec{A}^{\xi\theta}$ , leaving only the portion that actively contributes to net rotation over a gait. For the example gait, this means that the optimized orientation changes monotonically towards its final value, without the large ‘windup’ and ‘power stroke’ seen in the original orientation.



**Fig. 25.** Connection vector fields and height functions for the differential-drive car in the optimized coordinate choice. For this choice of coordinates, the  $\vec{A}^{\xi\theta}$  field is identically zero, so  $\zeta_x$  and  $\zeta_y$  are, respectively, equal to the  $x$  and  $y$  displacements with no error. The example gait shown here encircles a net-zero region on  $\mathbf{H}^{\zeta_x}$  and a negative region on  $\mathbf{H}^{\zeta_y}$ , producing the net-lateral ‘parallel-parking’ motion shown in Figure 26.

and two other systems that appear in the geometric mechanics literature: the south-pointing chariot (Santander 1992) and the rolling disk (Mukherjee and Anderson 1993a; Bloch et al. 2003). The south-pointing chariot is a two-wheeled cart with a pointer that counter-rotates with respect to the vehicle by an angle  $\beta_p = -\theta$  such that the pointer maintains a constant angle with respect to the global

coordinate frame, as illustrated in Figure 27. The pointer is implemented mechanically with a differential gear set that takes the difference between the two wheels’ axle angles; for unit wheel spacing and radius, the pointer angle is

$$\beta_p = \alpha_1 - \alpha_2. \tag{40}$$

A comparison of Equation (39) with Equation (40) clearly shows that  $\beta_\theta = \beta_p$ , and, thus, that optimizing the coordinates for the differential-drive car produces a mathematical description of the south-pointing chariot.

The rolling disk is an elementary non-holonomic system, often used in tutorial examples. Its kinematic model is depicted in Figure 28, with position  $g = (x, y) \in \mathbb{R}^2$  and ‘shape’  $r = (\alpha_a, \alpha_b) \in \mathbb{S} \times \mathbb{S}$ . The connection vector fields and height functions for the rolling disk are the same as for the optimized differential-drive, but rotated clockwise by 45 degrees; the reason for this becomes immediately clear when we consider that rolling the disk forward is equivalent to driving both wheels of the car together, while changing the disk’s heading corresponds to driving the wheels in opposite directions to each other. This relationship between the  $\alpha$  parameters for the disk and car is illustrated in Figure 29.

While a deeper examination of the relationships among the differential-drive car, the south-pointing chariot, and the rolling disk is beyond the scope of the current work, their equivalence under coordinate optimization suggests that they hold the seeds to additional insights into the coordinate optimization process. The differential-drive car in its classical coordinate system has zero-valued height functions and thus represents one end of the spectrum of height-function analysis, with no translation information available from the functions, while the south-pointing chariot and rolling disk represent the other extreme, with the height functions providing an exact measure of the displacement over gaits. A key element of our future work will be an examination of this spectrum, with an eye towards incorporating it into a more fully differential geometric treatment of the coordinate change.

### 8. Demonstration

We applied the gaits from the image-family in Figures 8 and 19 to a physical instantiation of the kinematic snake, shown in Figure 30, and plotted the resulting displacements in the original and optimized coordinates in Figure 31. During the experiments, we observed some backlash in the joints and slip in the constraints, to which we attribute the differences between the calculated and experimental loci in Figures 20 and 31. Even with this error, the BVI was a considerably more effective estimate of the displacement with the new coordinate choice than with the old choice. Under the old coordinate choice, the error in the the net direction of travel between the BVI estimate and the actual displacement ranged from 0.72 to 0.25 radians. With the new choice, the error in the estimate ranged from 0.08 to  $-0.05$  radians, reducing the maximum error magnitude by 89% and eliminating the minimum error.

### 9. Conclusions

In this paper, we have shown that the choice of coordinates used to represent a mobile system directly affects

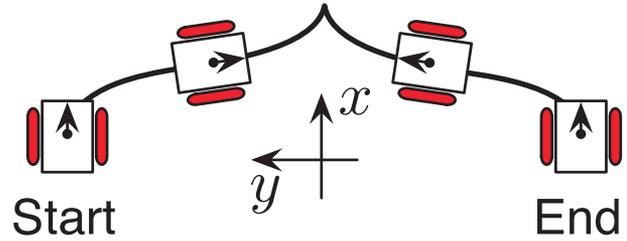


Fig. 26. ‘Parallel-parking’ motion for the differential-drive car produced by the gait in Figure 25.

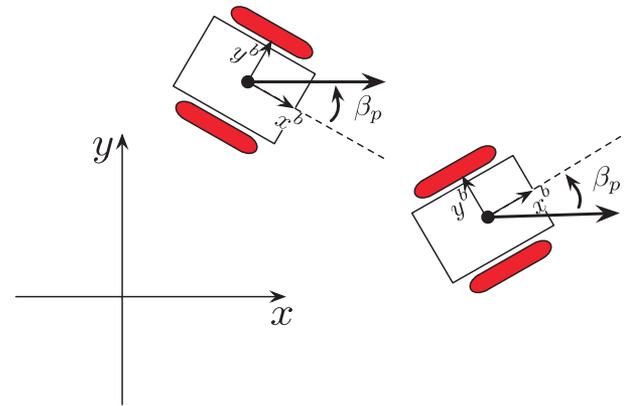


Fig. 27. The south-pointing chariot’s pointer counter-rotates with respect to the main body by an angle  $\beta_p = -\theta$ , such that it maintains a constant angle with respect to the global coordinates.

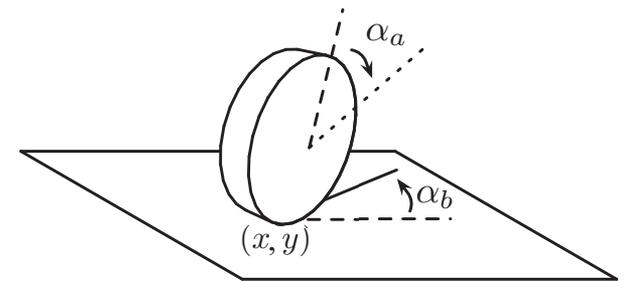


Fig. 28. Kinematic model for a rolling disk. The disk’s position is the  $(x, y)$  location of its contact point on the plane, and its ‘shape’ is the roll angle  $\alpha_a$  and the heading angle  $\alpha_b$ .

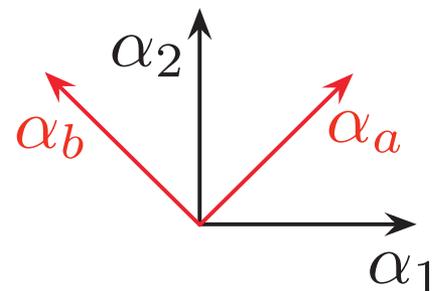
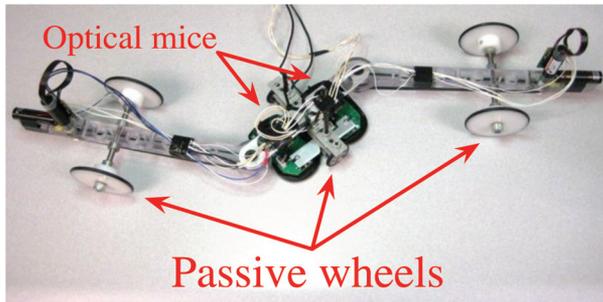
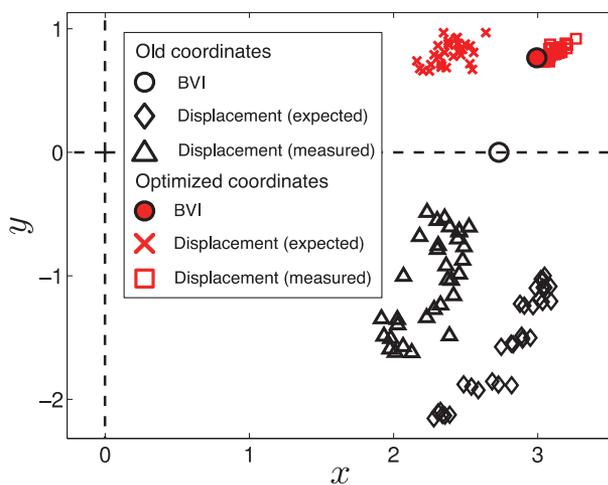


Fig. 29. Relationship between  $\alpha_1$  and  $\alpha_2$  for the differential-drive car and  $\alpha_a$  and  $\alpha_b$  for the rolling disk.



**Fig. 30.** Physical instantiation of the kinematic snake. Two optical mice on the middle link provide odometry data, including an estimation of constraint slip.



**Fig. 31.** BVI and displacements for the physical kinematic snake while executing the example image-family of gaits. The boxes and diamonds show the expected displacements for the robot, calculated from the actual trajectories followed by the joints, while the crosses and triangles show the actual measured displacement, which differs from the expected displacement because of backlash and constraint slip. As in Figure 20, the BVI is a much better estimate of the displacement with the new, optimized coordinate choice than with the original choice, and the variation in the displacement for different starting shapes is much smaller.

the analytical tools that can be applied to the system. By choosing a set of coordinates that minimizes the rotation of the body frame in response to changes in system shape, we have advanced a technique for approximating the net displacement resulting from a gait by means of Stokes' theorem, rather than by direct integration of the equations of motion. In turn, this approximation provides both an intuitive visualization of the system's kinematics and a means of designing gaits by encircling sign-definite regions of well-defined height functions. The technique constitutes a significant improvement over previous efforts in this area, as it is applicable to efficient macroscopic translational gaits, while previous work was limited to macroscopic rotational gaits or inefficient small-amplitude translational gaits.

By considering the effect of coordinate choice, we have added a new dimension to analysis of the motion-planning problem. Beyond the specific advances presented here, this widened perspective opens the door for new analysis outside the scope of this paper. A key example of such future work is the unification of our observations with the geometric mechanics theory on which we have built. This foundational work has typically striven to be independent of the coordinates used, so it is likely that new insights can be attained by lifting our coordinate-based approach from the systems examined here to a more abstract level.

Developing a more formal geometric formulation of the coordinate change will also provide the necessary structure to expand the classes of systems considered in our analysis. Thus far, our examples have been limited to planar systems with two shape inputs and governed by conservation of momentum or a single set of non-holonomic constraints. Concepts such as curl are inherently dimensionally limited, but differential geometry encompasses their generalization to higher dimensions, allowing us a route to include more shape parameters. Similarly, the more formal approach will inform the construction of objective functions for optimizing coordinate choice for fully-spatial motion in  $SE(3)$ . While extending our analysis to higher dimensions, we will also consider systems governed by other constraints that give rise to a fiber bundle structure, such as swimming systems (Melli et al. 2006; Avron and Raz 2008), the mixed kinodynamic systems of Shammass et al. (2007b), and potentially to hybrid systems that cycle through multiple sets of constraints over the course of a gait.

A second area that the present work leaves open to further exploration is finding more refined objective functions for optimizing the coordinate choice. Weighting the Hodge–Helmholtz projection to emphasize the quality of the fit between  $\nabla_r \beta_\theta$  and  $-\vec{A}^{\xi_\theta}$  in targeted regions, as in Equation (36), improves the accuracy of the height functions in regions of the shape space where useful gaits will exist. While the scaled arctangent scaling we use here serves this purpose adequately by emphasizing areas of the shape space away from singularities, there is still much room to investigate alternative weighting schemes. Such approaches might, for example, de-emphasize regions of the shape space that correspond to self-collision of the robot and which are thus unusable for gait design.

## Notes

1. In the parlance of differential geometry, this separation assigns a (trivial, principal) fiber bundle structure  $Q = G \times M$  to the system's configuration space, with  $M$  the base space and  $G$  the fiber space.
2. A more formal differential geometric treatment of this subject matter would, respectively, replace 'vector field' and 'curl' with 'one-form' and 'exterior derivative,' but here (without loss of generality) we choose our language for broader accessibility.

3. In strict differential geometric language, each row  $\mathbf{A}^i$  of  $\mathbf{A}$  is a one-form over  $M$  acting on  $\dot{r}$ , and  $\vec{\mathbf{A}}^{\xi_i}$  is the negative dual of that one-form.
4. In the geometric mechanics literature, the displacement over a gait is often referred to as the ‘geometric phase.’
5. The intersection point of the lines perpendicular to the wheelsets. When there are multiple intersection points, we regard the center of constraint as their mean position.

### Acknowledgements

The authors would like to thank Scott Kelly and Matthew Tesch for many insightful discussions.

### Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

### Conflict of interest statement

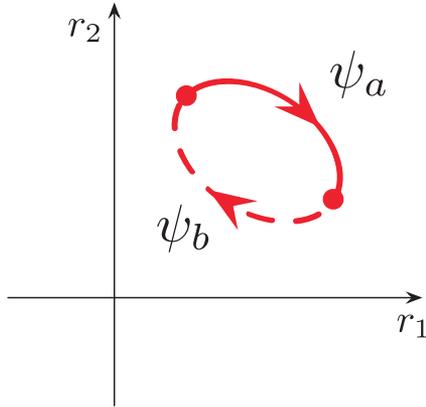
The authors declare that they have no conflicts of interest.

### References

- Arfken GB (2005) *Mathematical Methods for Physicists*, 6th edition. Elsevier.
- Avron J and Raz O (2008) A geometric theory of swimming: Purcell’s swimmer and its symmetrized cousin. *New Journal of Physics* 9: 063016.
- Becker EB, Carey GF and Oden JT (1981) *Finite Elements: An Introduction*, volume 1. Texas Institute for Computational Mathematics.
- Bloch AM with Baillieul J, Crouch PE and Marsden JE (2003) *Nonholonomic Mechanics and Control*. Berlin: Springer.
- Boothby WM (1986) *An Introduction to Differentiable Manifolds and Riemannian Geometry*. New York: Academic Press.
- Bullo F and Lynch KM (2001) Kinematic controllability for decoupled trajectory planning in underactuated mechanical systems. *IEEE Transactions on Robotics and Automation* 17: 402–412.
- Guo Q, Mandal MK and Li MY (2005) Efficient Hodge–Helmholtz decomposition of motion fields. *Pattern Recognition Letters* 26: 493–501.
- Hatton RL and Choset H (2008) Connection vector fields for underactuated systems. In: *Proceedings of the IEEE BioRobotics Conference*, Scottsdale, AZ, USA, October.
- Hatton RL and Choset H (2009) Approximating displacement with the body velocity integral. In: *Proceedings of Robotics: Science and Systems*, Seattle, WA, USA, June.
- Hatton RL and Choset H (2010) Optimizing coordinate choice for locomoting systems. In: *Proceedings of the IEEE International Conference on Robotics and Automation*, Anchorage, AK, USA, May.
- Ishikawa M (2009) Iterative feedback control of snake-like robot based on principal fibre bundle modelling. *International Journal of Advanced Mechatronic Systems* 1: 175–182.
- Kelly SD and Murray RM (1995) Geometric phases and robotic locomotion. *Journal of Robotic Systems* 12: 417–431.
- McIsaac KA and Ostrowski JP (2003) Motion planning for anguilliform locomotion. *IEEE Transactions on Robotics and Automation* 19: 637–652.
- Melli JB, Rowley CW and Rufat DS (2006) Motion planning for an articulated body in a perfect planar fluid. *SIAM Journal of Applied Dynamical Systems* 5: 650–669.
- Morgansen KA, Triplett BI and Klein DJ (2007) Geometric methods for modeling and control of free-swimming fin-actuated underwater vehicles. *IEEE Transactions on Robotics* 23: 1184–1199.
- Mukherjee R and Anderson DP (1993a) Nonholonomic Motion Planning Using Stokes’ Theorem. In: *IEEE International Conference on Robotics and Automation*, pp. 802–809.
- Mukherjee R and Anderson DP (1993b) A surface integral approach to the motion planning of nonholonomic systems. In: *American Control Conference, 1993*, pp. 1816–1823.
- Murray RM and Sastry SS (1993) Nonholonomic motion planning: steering using sinusoids. *IEEE Transactions on Automatic Control* 38: 700–716.
- Ostrowski J (1998) Reduced equations for nonholonomic mechanical systems with dissipative forces. *Reports on Mathematical Physics* 42: 185–209.
- Ostrowski J and Burdick J (1998) The mechanics and control of undulatory locomotion. *International Journal of Robotics Research* 17: 683–701.
- Ostrowski JP, Desai JP and Kumar V (2000) Optimal gait selection for nonholonomic locomotion systems. *The International Journal of Robotics Research* 19: 225–237.
- Santander M (1992) The Chinese south-seeking chariot: A simple mechanical device for visualizing curvature and parallel transport. *American Journal of Physics* 60: 782–787.
- Sfakiotakis M and Tsakiris DP (2007) Biomimetic centering for undulatory robots. *International Journal of Robotics Research* 26: 1267–1282.
- Shammas E, Schmidt K and Choset H (2005) Natural gait generation techniques for multi-bodied isolated mechanical systems. In: *IEEE International Conference on Robotics and Automation*, April, pp. 3664–3669.
- Shammas EA, Choset H and Rizzi AA (2007a) Geometric motion planning analysis for two classes of underactuated mechanical systems. *The International Journal of Robotics Research* 26: 1043–1073.
- Shammas EA, Choset H and Rizzi AA (2007b) Towards a unified approach to motion planning for dynamic underactuated mechanical systems with non-holonomic constraints. *The International Journal of Robotics Research* 26: 1075–1124.
- Shapere A and Wilczek F (1989) Geometry of self-propulsion at low Reynolds number. *Journal of Fluid Mechanics* 198: 557–585.
- Walsh G and Sastry S (1995) On reorienting linked rigid bodies using internal motions. *IEEE Transactions on Robotics and Automation* 11: 139–146.

### Appendix A: Magnitude of displacement for gaits in an image-family

In Section 6.1, we commented that the locus of net displacements resulting from an image-family of gaits that produces no net rotation will be an arc of constant radius, centered on the origin. To see that this must be the case, consider the two



**Fig. 32.** Two shape changes  $\psi_a$  and  $\psi_b$  that form gaits when concatenated in either order.

shape changes

$$\psi_a : [0, T_a] \rightarrow M \tag{41}$$

$$\psi_b : [0, T_b] \rightarrow M, \tag{42}$$

arranged such that  $\psi_a(T_a) = \psi_b(0)$  and  $\psi_b(T_b) = \psi_a(0)$ , as illustrated in Figure 32. By concatenating these two shape changes, we can generate two distinct gaits from a common image-family  $\bar{\phi}$ ,

$$\phi_1 = \psi_a \psi_b \tag{43}$$

$$\phi_2 = \psi_b \psi_a. \tag{44}$$

The net displacements over these two gaits are found by composing the net displacements over each section,

$$g_{\phi_1}(T) = g_{\psi_a}(T_a) g_{\psi_b}(T_b) \tag{45}$$

$$g_{\phi_2}(T) = g_{\psi_b}(T_b) g_{\psi_a}(T_a), \tag{46}$$

where  $T = T_a + T_b$ . Evaluating Equation (45) as  $g_{\phi_1}(T) = (x_1, y_1, \theta_1)$  gives the net displacement over  $\phi_1$  as

$$g_{\phi_1}(T) = (x_a + x_b \cos \theta_a - y_b \sin \theta_a, y_a + x_b \sin \theta_a + y_b \cos \theta_a, \theta_a + \theta_b). \tag{47}$$

The squared magnitude of the  $xy$  displacement over  $\phi_1$  is thus

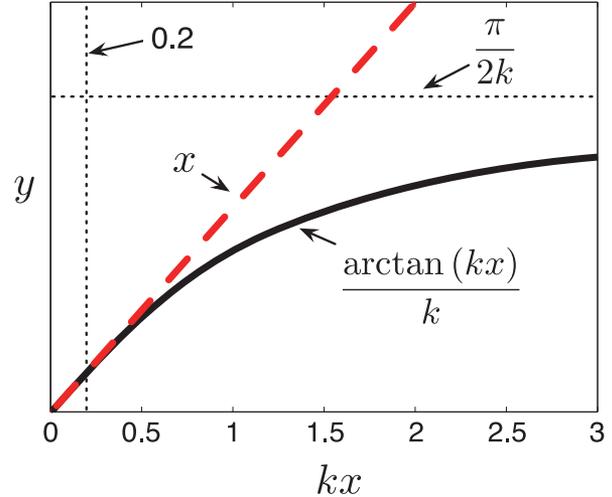
$$\|g_{\phi_1}(T)\|_{xy}^2 = x_a^2 + y_a^2 + x_b^2 + y_b^2 + \cos \theta_a(x_a x_b + y_a y_b) + \sin \theta_a(x_a y_b - x_b y_a); \tag{48}$$

by a similar calculation, the squared magnitude of the  $xy$  displacement over  $\phi_2$  is

$$\|g_{\phi_2}(T)\|_{xy}^2 = x_a^2 + y_a^2 + x_b^2 + y_b^2 + \cos \theta_b(x_a x_b + y_a y_b) - \sin \theta_b(x_a y_b - x_b y_a), \tag{49}$$

and the difference between these two values is

$$\|g_{\phi_1}(T)\|_{xy}^2 - \|g_{\phi_2}(T)\|_{xy}^2 = (\cos \theta_a - \cos \theta_b)(x_a x_b + y_a y_b) + (\sin \theta_a + \sin \theta_b)(x_a y_b - x_b y_a). \tag{50}$$



**Fig. 33.** Scaled arctangent function. For small  $kx$ ,  $\arctan(kx)/k \approx x$ , but for larger values it is limited to  $\pi/2k$ . In this paper, we choose  $k$  to place the median value of  $kx$  at 0.2, in the center of the distortion-free domain.

For image-families that produce no net rotation, that is, for which  $\theta_a + \theta_b = 0$ , the two trigonometric terms in Equation (50) go to zero, and the magnitude of displacement is the same for  $\phi_1$  and  $\phi_2$ . As the start and end points of  $\psi_a$  and  $\psi_b$  can be arbitrarily chosen along the curve defined by  $\bar{\phi}$ , the results of Equation (50) apply to all gaits in  $\bar{\phi}$ . This approach can be generalized to  $SO(n)$  to show that in an image-family that produces no net rotations, the magnitude of displacement for all gaits must be the same.

### Appendix B: Scaled arctangent

When plotting and optimizing the connection vector fields for the kinematic snake, we modify their magnitudes by a scaled arctangent function

$$y = \frac{\arctan(kx)}{k} \tag{51}$$

to prevent the large magnitudes of the fields near the singularities from saturating the plot or optimization process. Figure 33 illustrates the behavior of this function. For  $kx < 0.5$ ,  $y \approx x$ , but for larger values of  $kx$ ,  $y$  increases much more slowly, asymptotically approaching an upper limit of  $\pi/2k$ . The scaling parameter  $k$  can be chosen based on the values of  $x$  to minimize the distortion of ‘reasonable’ values, while filtering out overly large outliers. In this paper, we use  $k = 0.2/\text{median}(x)$  to place  $\text{median}(kx)$  at 0.2, in the center of the distortion-free domain.

### Appendix C: Discrete Hodge–Helmholtz decomposition

In Section 6.4.1, we use the discrete Hodge–Helmholtz decomposition to find the optimal choice of orientation, following the algorithm in Guo et al. (2005). The central

operation in this decomposition is finding a potential function  $E$  that minimizes the cost function

$$D(E) = \int_{\Omega} \|\nabla E - V\|^2 d\Omega \tag{52}$$

for an input vector field  $V$ . Here, we present an expanded derivation of the finite-element approach to solving this minimization. This derivation contains some details of the formulation that were not presented in Guo et al. (2005) and serves as an introduction for the minimization in Appendix D, which has a similar form, but several additional terms.

For generality, the derivation here uses generic expressions for the potential function and input vector fields; when we apply the results to Section 6.4.1, we take  $E = -\beta_{\theta}$  and  $V = \mathbf{A}^{\xi_{\theta}}$ .

### C1 Finite-element formulation

There is no general analytical solution for minimizing Equation (52), but a finite-element approach readily provides a numerical solution. The first stage of this optimization is to take the output potential function  $E$  as being the weighted sum of a set of (as yet unspecified) basis functions  $\phi$ ,

$$E = \sum_i \phi_i e_i. \tag{53}$$

An optimal set of weights  $e$  that minimizes Equation (52) must result in  $D$  being stationary with respect to each weight,

$$\forall i, \frac{\partial D}{\partial e_i} = 0. \tag{54}$$

The partial derivative of  $D$  with respect to the weights is found by first expanding the right-hand side of Equation (52),

$$D(E) = \int_{\Omega} (\nabla E - V) \cdot (\nabla E - V) d\Omega \tag{55}$$

$$= \int_{\Omega} (\nabla E \cdot \nabla E) + (V \cdot V) - 2(\nabla E \cdot V) d\Omega, \tag{56}$$

and then differentiating under the integration sign,

$$\frac{\partial D}{\partial e_i} = \int_{\Omega} \underbrace{\frac{\partial(\nabla E \cdot \nabla E)}{\partial e_i}}_A + \underbrace{\frac{\partial(V \cdot V)}{\partial e_i}}_0 - 2 \underbrace{\frac{\partial(\nabla E \cdot V)}{\partial e_i}}_B d\Omega. \tag{57}$$

The second term of the summation is clearly zero, as the input vector field does not depend on the weights, and the first and third terms evaluate to

$$A : \frac{\partial(\nabla E \cdot \nabla E)}{\partial e_i} = 2 \frac{\partial \nabla E}{\partial e_i} \cdot \nabla E = 2 \nabla \phi_i \cdot \nabla E \tag{58}$$

$$B : 2 \frac{\partial(\nabla E \cdot V)}{\partial e_i} = 2 \nabla \phi_i \cdot V, \tag{59}$$

allowing the consolidation of Equation (57) into

$$\frac{\partial D}{\partial e_i} = 2 \int_{\Omega} \nabla \phi_i \cdot (\nabla E - V) d\Omega. \tag{60}$$

Applying the necessary condition from Equations (54) to (60) produces an alternate condition on the weights  $e$  to minimize  $D$ : They should generate  $e$  such that

$$\int_{\Omega} \nabla \phi_i \cdot \nabla E d\Omega = \int_{\Omega} \nabla \phi_i \cdot V d\Omega. \tag{61}$$

It may interest a reader with knowledge of other finite-element approaches to note the resemblance between Equation (61) and an alternate formulation of the problem, in which we seek to approximate  $V$  with an element of the set of gradient fields, using  $\nabla \phi$  as a test function.

### C2 Piecewise-linear solution

To solve Equation (61), we need to specify the basis functions. While many classes of basis functions are available, here we use piecewise-linear bases on a triangular mesh. One such function is depicted in Figure 34; it has a value of 1 at its central node and zero at all other nodes, with linearly interpolated values in the interiors of the triangles. This is a canonical choice of basis, and as noted in Guo et al. (2005), several aspects make it especially convenient. First, the property of  $\phi_i$  being 1 at the  $i$ th node and zero everywhere else means that  $E$  at node  $i$  is equal to  $e_i$ . Second, this property also means that for each  $i$ , we need only evaluate Equation (61) over the six (or fewer) triangles containing node  $i$ , making the minimization condition

$$\sum_k \int_{T_k} \nabla \phi_i \cdot \nabla E dT_k = \sum_k \int_{T_k} \nabla \phi_i \cdot V dT_k, \tag{62}$$

where  $T_k$  is the  $k$ th triangle containing the  $i$ th node, as numbered in Figure 34. From the quadrature rules for integrating over linear triangular regions (Becker et al. 1981), the integrals in Equation (62) are equal to the areas of the triangles multiplied by the integrands evaluated at the triangles' geometric centers,

$$\sum_k \nabla \phi_{ik} \cdot \nabla E_k A_k = \sum_k \nabla \phi_{ik} \cdot V_k A_k. \tag{63}$$

For odd values of  $k$ , the terms in the integrands evaluate to

$$\nabla \phi_{ik} = [\phi_{id} - \phi_{ic}, \phi_{ia} - \phi_{ic}]^T \tag{64}$$

$$\nabla E_k = [E_d - E_c, E_a - E_c]^T \tag{65}$$

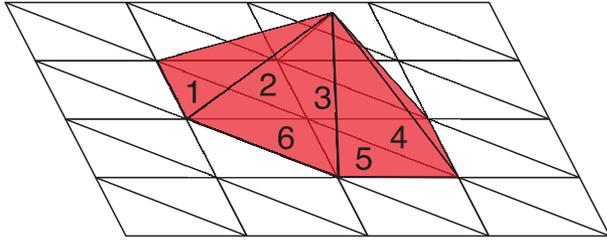
$$V_k = (V_a + V_c + V_d) / 3, \tag{66}$$

and for even values of  $k$ ,

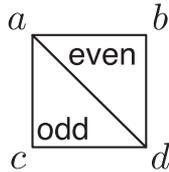
$$\nabla \phi_{ik} = [\phi_{ib} - \phi_{ia}, \phi_{ib} - \phi_{id}]^T \tag{67}$$

$$\nabla E_k = [E_b - E_a, E_b - E_d]^T \tag{68}$$

$$V_k = (V_a + V_b + V_d) / 3, \tag{69}$$



**Fig. 34.** Triangular mesh with piecewise-linear basis function. The basis function has a value of 1 at its central node and 0 at all other nodes, with linearly interpolated values in the interiors of the triangles.



**Fig. 35.** Geometry of odd- and even-numbered triangles in the grid. The lettering *a* through *d* denotes a local coordinate system for the nodes in each triangle.

where the nodes are lettered as in Figure 35.

As *E* and *V* each appear linearly in  $\nabla E_k$  and  $V_k$ , Equation (63) can be reformed into a linear equation

$$S_1 E = S_2 V, \tag{70}$$

where the *i*th rows of  $S_1$  and  $S_2$  encode the summations over the triangles containing the *i*th node. These two matrices depend only on the grid spacing (and the definition of the basis functions), and so can be precalculated as  $S = S_1^{-1} S_2$ , such that

$$E = S V \tag{71}$$

produces the potential function whose gradient best approximates the input vector field. In practice, one row and column of  $S_1$  and one row of  $S_2$  must be eliminated while precalculating  $S$ ; as  $\nabla(E + C) = \nabla E$  for any constant  $C$ , there is a whole family of potential functions that are solutions to Equation (71), and eliminating the row and column locks the problem down to a single solution.

### Appendix D: Reference-point optimization

We take a similar approach to optimizing the reference-point position as we did for the orientation. The main difference is that rather than a single input vector field and output potential function, there are now three input vector fields,

$U, V, W$ , and two output potential functions,  $E, F$ , which correspond, respectively, to  $\vec{A}^{\xi_x}, \vec{A}^{\xi_y}, \vec{A}^{\xi_\theta}$ , and  $\beta_x, \beta_y$ , and a cost function

$$D(E, F) = \int_{\Omega} \|U + \nabla E - FW\|^2 + \|V + \nabla F + EW\|^2. \tag{72}$$

As in the Hodge–Helmholtz decomposition, we minimize  $D$  by first expanding the squared vector norms,

$$\begin{aligned} D(E, F) = \int_{\Omega} & (U \cdot U) + (\nabla E \cdot \nabla E) + (F^2 W \cdot W) \\ & + 2(U \cdot \nabla E) - 2(FW \cdot \nabla E) \\ & + (V \cdot V) + (\nabla F \cdot \nabla F) + (E^2 W \cdot W) \\ & + 2(V \cdot \nabla F) + 2(EW \cdot \nabla F) \, d\Omega, \end{aligned} \tag{73}$$

then differentiating by the basis coefficients  $e$  and  $f$ ,

$$\begin{aligned} \frac{\partial D}{\partial e_i} = 2 \int_{\Omega} & \nabla \phi_i \cdot (\nabla E - FW + U) \\ & + \phi_i W \cdot (EW + V + \nabla F) \, d\Omega \end{aligned} \tag{74}$$

and

$$\begin{aligned} \frac{\partial D}{\partial f_i} = 2 \int_{\Omega} & \nabla \phi_i \cdot (\nabla F - EW + V) \\ & + \phi_i W \cdot (FW - U - \nabla F) \, d\Omega. \end{aligned} \tag{75}$$

Applying the minimization conditions of  $\partial D/\partial e_i = 0$  and  $\partial D/\partial f_i = 0$  produces the equated integrals

$$\begin{aligned} \int_{\Omega} & \nabla \phi_i \cdot (\nabla E - FW) + \phi_i W \cdot (EW + \nabla F) \, d\Omega \\ = \int_{\Omega} & \nabla \phi_i \cdot (-U) - \phi_i (V \cdot W) \, d\Omega \end{aligned} \tag{76}$$

and

$$\begin{aligned} \int_{\Omega} & \nabla \phi_i \cdot (\nabla F + EW) + \phi_i W \cdot (FW - \nabla E) \, d\Omega \\ = \int_{\Omega} & \nabla \phi_i \cdot (-V) + \phi_i (V \cdot W) \, d\Omega, \end{aligned} \tag{77}$$

which can be operated on by the same piecewise-linear approach as in Section C2, resulting in a matrix equation

$$\begin{bmatrix} E \\ F \end{bmatrix} = S(W) \begin{bmatrix} U \\ V \end{bmatrix}. \tag{78}$$

Unlike Equation (61), Equations (76) and (77) contain non-gradient terms, and so admit only a single solution. Consequently, the  $S$  matrix here is full rank and does not need to have a row and column removed in order to find a solution.