Eur. Phys. J. Special Topics **224**, 3141–3174 (2015)

DOI: 10.1140/epjst/e2015-50085-y

THE EUROPEAN PHYSICAL JOURNAL SPECIAL TOPICS

Regular Article

Nonconservativity and noncommutativity in locomotion

Geometric mechanics in minimum-perturbation coordinates

R.L. Hatton[1,a] and H. Choset[2]

[1] School of Mechanical, Industrial, and Manufacturing Engineering, Oregon State University, USA
[2] Robotics Institute, Carnegie Mellon University, USA

Received 14 April 2015 / Received in final form 2 November 2015
Published online 15 December 2015

Abstract. Geometric mechanics techniques based on Lie brackets provide high-level characterizations of the motion capabilities of locomoting systems. In particular, they relate the net displacement they experience over cyclic gaits to area integrals of their constraints; plotting these constraints thus provides a visual "landscape" that intuitively captures all available solutions of the system's dynamic equations. Recently, we have found that choices of system coordinates heavily influence the effectiveness of these approaches. This property appears at first to run counter to the principle that differential geometric structures should be coordinate-invariant.

In this paper, we provide a tutorial overview of the Lie bracket techniques, then examine how the coordinate-independent *nonholonomy* of these systems has a coordinate-dependent separation into *nonconservative* and *noncommutative* components that respectively capture how the system constraints vary over the shape and position components of the configuration space. Nonconservative constraint variations can be integrated geometrically via Stokes' theorem, but noncommutative effects can only be approximated by similar means; therefore choices of coordinates in which the nonholonomy is primarily nonconservative improve the accuracy of the geometric techniques.

1 Introduction

The study of locomotion is an active field in robotics and biology. Understanding how changes in system shape interact with environmental constraints allows robot designers to produce systems that move through the world with agility and efficiency. In biology, understanding the links between shape and body motions provides baseline data on the potential capability of an organism, against which its observed behavior can be compared.

Geometric mechanics offers powerful tools for studying locomotion [1–8] that both simplify system models and provide deep links between physical motion and fundamental mathematical structures. These methods exploit mathematical symmetry in systems' equations of motion and highlight relationships between changes in the internal *shape* of a system and its

[a] e-mail: Ross.Hatton@oregonstate.edu

external *position* in the world. A key feature of these approaches is their basis in differential geometry, making use of techniques such as *Lie bracket averaging*, which approximate the net displacement over cyclic changes in shape with area integrals of the *curvature* of the system constraints. These approximations then aid the search for a system's most useful or effective cycles, which collectively describe the locomotion capabilities of the system and can be used as a library for planning longer-range motions.

The approximation error in Lie bracket averaging scales with the magnitude of the shape changes executed during a gait. Historically, this error scaling has restricted Lie bracket analysis to relatively small motions, which often do not include the most efficacious gaits available to the systems. In our recent work [9–11], however, we demonstrated that–contrary to the expectation that differential geometric expressions should be coordinate-invariant–the *rate* at which the error scales with gait amplitude is significantly affected by the choice of body frame used to parameterize the system. Based on this observation, we developed a method for optimizing the choice of coordinates, in which advantageous body frame selections are found by taking generalized *Hodge-Helmholtz decompositions* of the system constraints to yield the *minimum-perturbation* coordinates for the system, analogous to the *Coulomb gauge* in electromagnetic systems.

We initially developed this coordinate optimization in a robotics context, using simplified vector-calculus terminology and making explicit use of the planar nature of the example systems in our derivation. Given this genesis, questions naturally arise regarding rigorous mathematical justification for our results and the possibility of extending them to systems with different configuration spaces. In this paper, we address these questions by formulating our approach in a truly differential geometric framework, thus better relating it to key results in the field [1,8] and allowing a deeper examination of the underlying principles. In particular, we investigate how the *nonholonomy*,[1] or path-dependence, encoded by the constraint curvature has a coordinate-dependent separation into two parts: a *nonconservative* component that captures changes in the constraints across different system shapes, and can be integrated exactly in the averaging techniques; and a *noncommutative* component that captures how the order in which a system translates and rotates affects its net motion, and can only be approximated when averaging the Lie bracket.

Using this separation, we demonstrate that our optimal choice of coordinates is that for which the nonholonomy is minimally noncommutative, and thus maximally nonconservative. Additionally, considering the noncommutative effects (which we elided for simplicity in our previous work) allows us to characterize the efficacy of coordinate optimization for a given system and bound the error of our geometric tools. To facilitate comparison between these results and our prior formulation, we have selected the same set of example systems as in [11], shown in Fig. 1; as discussed in Sect. 3.2, these systems serve as both canonical examples in the geometric mechanics literature and exhibit special features that highlight aspects of our coordinate optimization.

2 Context and prior work

A defining feature of the geometric mechanics approach to locomotion is its rigorous mathematical distinction between the *shape* of a locomoting system and its *position* (including orientation) in the world. Formally, this separation is embodied in a differential geometric structure called a *fiber bundle*, in which the shape and position of the system are respectively

[1] Note that by *nonholonomy* we mean the degree to which the constraints on the system are nonholonomic (i.e., how different they are from a set of constraints that could be integrated to provide a lower-dimensional configuration space for the system), and therefore how much the integral along a path depends on the whole path and not just the endpoints. This is *not* an antonym for the term "holonomy", which refers to the result of integrating along a specific path.

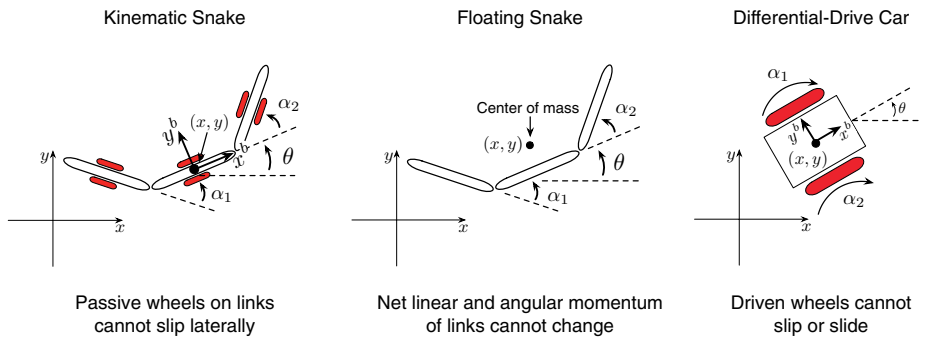

Fig. 1. The three example systems and their constraints.

referred to as elements of the *base* and *fiber* spaces. Combinations of shape and position velocities that accommodate constraints on the system (such as conservation of momentum, or interaction with a surrounding fluid) are encoded in *connections* on the fiber bundle.

In this paradigm, two assumptions about system behavior are often used to frame the analysis. First, that the system's shape is fully controlled (i.e. the shape can be regarded as an input), thus letting the connection induce a Jacobian from control inputs to the resulting position velocity, modulo any second-order momentum effects. Second, that when expressed in body coordinates the system constraints are invariant with respect to the position; intuitively, this leads to the concept that a shape change that pulls the system "forward" will always do so, regardless of how the system is located and oriented.

Much of the geometric locomotion framework can be traced to Shapere and Wilczek's seminal papers "The Geometry of Self Propulsion at Low Reynolds Number"and "Gauge Mechanics Of Deformable Bodies" [1,12]. The former paper considered the locomotion of circular and spherical objects immersed in highly viscous fluids, such as is the case for bacteria and other micro-organisms, while the latter considered the rotation of isolated bodies under the conservation of net angular momentum. The authors of those papers drew their terminology from gauge theory, and clearly articulated the ideas of using a fiber bundle to represent the swimmer configuration and of using a *local form of the connection* to map input shape velocities into position velocities. Interestingly, Shapere and Wilczek also posed a question that, to the best of our knowledge, remained open until our investigation: What is the best choice of position and orientation coordinates (in their terminology, choice of gauge) to use when modeling these systems? [1].

Following Shapere and Wilczek's investigations, geometric locomotion research focused more on the rigid body, momentum conservation approach in [12] than on the low Reynolds number formulation in [1]. Drawing on the mechanical symmetry work of Marsden et al. [13–15], researchers such as Krishnaprasad, Tsakiris, Kelly, Ostrowski, Burdick, Bloch, and Lewis [2–7] developed the *reconstruction equation*, which generalized the results of [1,12] to apply to systems with first-order nonholonomic constraints, such as ice skates or passive wheels. Among other benefits, this generalization included a momentum term capturing how much the system is coasting. This allows the study of systems which fall between the completely constrained low Reynolds number systems of [1], and those completely unconstrained except by momentum conservation, as in [12].

This reconstruction equation has been used in a variety of locomotion contexts. Ostrowski et al. [4,16] combined the reconstruction equation with Lie bracket theory to generate sinusoidal gaits which translate and rotate a variety of snake-like systems. Bullo and Lynch used the reconstruction equation to decouple the kinematic and dynamic elements of the locomotion of kinodynamic systems and thus to design kinematic gaits [17]. Shammas et al. modified the momentum term in the reconstruction equation into a *scaled momentum*, enabling motion planning techniques that guarantee at least the direction of that momentum [18].

To aid in visualizing the reconstruction equation, we developed the *connection vector field* [19]. Although most of the geometric locomotion literature uses a subset of the special Euclidean group $SE(3)$ to represent system position and orientation, Krishnaprasad and Tsakiris have also investigated using alternative group structures for the position space [2].

More recently, the geometric locomotion focus has returned to swimming systems, including Kelly's formulation of the reconstruction equation for swimming systems at various Reynolds numbers [20], and the respectively high and low Reynolds numbers studies of Melli et al. [8] and Avron and Raz [21]. Other geometric studies of swimming incorporate biomimetic elements along with mathematical formulations, such as McIsaac and Ostrowski's work on anguilliform (eel-like) robots [22] and Morgansen et al.'s work on fish [23].

It is not generally possible to integrate the reconstruction equation in closed form, raising difficulties for the inverse problem of finding shape changes that produce desired translations. In special cases, however, Stokes' theorem (as a special case of the Baker-Hausdorff-Campbell theorem) can be used to find the net motion resulting from gaits [3]. Mukherjee [24] used this principle to analyze the motion of rolling disks, and Walsh and Sastry [25,26] applied it to the case of an isolated three-link robot. Shammas et al. [27,28] 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. [8] and later by Avron and Raz [21] to generate gaits for swimming robots.

The approaches in [8,21,28] all shared the limitation that for general macroscopic (non-infinitesimal) gaits, the height functions could only be used to determine the net rotations over gaits; the noncommutativity of translations and rotations precludes the existence of an (inherently unordered) area integral to determine net translation. Melli et al. and Avron and Raz were able to extend the techniques to finding good approximation of the net translations induced by a limited set of macroscopic (but still small magnitude) gaits by using a higher-order definition of the height functions.

In [9], we observed that much of the limitation to small gaits was not inherent to the system, but instead depended on how the system was parameterized, with certain choices of coordinates greatly reducing the system noncommutativity. Building on this result, we presented in [10,11] a systematic optimization procedure for making this choice, which we extended in [29] to the swimming systems considered in [8,21]. These results were presented to the robotics community in a vector-calculus framework, and in this paper we generalize them into differential-geometric terms.

More recently, we have turned our attention to using the geometric paradigm to analyze real-world systems whose dynamics are approximated by a local connection (rather than being an exact match). In [30], we demonstrated that the geometric tools can provide useful insight about systems whose dynamics are analytically intractable, but the *output* of which can be fit to our geometric template. In that paper, which grew out of investigations of the burrowing behavior of the sandfish lizard, we considered the motion of a three-link robot immersed in a bed of small plastic spheres. Taking the gradient of this velocity relationship at each shape yielded a numerically-generated connection equivalent to that in the formal geometric model, from which we could calculate constraint curvature for the system and then proceed with the analysis.

3 Geometric mechanics

Geometric mechanics applies differential-geometric structures to problems in classical mechanics. In locomotion analysis, using Lie groups to represent the position and orientation of a system generates rigorously-defined body frames for the systems that simplify their representation. The relationship between changes in system shape and position is captured by differential forms; when combined with Lie brackets and Stokes' theorem these forms

provide information about the net displacements produced by cyclic changes in shape. This section relates physical aspects of locomotion to the geometric structures; we have provided a more general description of the structures themselves in the appendices.

3.1 Position, shape, and velocity

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 B, such that the position $g \in G$ locates the system in the world and the shape $r \in B$ gives the relative arrangements of its bodies. We can then consider how changes in shape induce changes in position. Typically, G is chosen to have a Lie group structure with the properties described in Sect. A.2. A natural choice of G for describing position and orientation is *SE*(3), the *special Euclidean group* of three-dimensional translations and rotations, or one of its subgroups, such as *SE*(2) for planar motion or the *special orthogonal group SO*(3) for spatial rotations. In differential geometric terms, this separation assigns a trivial, principal fiber bundle structure[2] $Q = G \times B$ to the system's configuration space, with G the fiber space and B the base space. Note that as well as denoting the system position, g also identifies a *body frame* rigidly attached to the system, and the transform associated with that coordinate frame: if a second object has position $h \in G$ with respect to the body frame, then its position with respect to the world is gh.

3.1.1 Position and shape

In this paper, we focus on systems that translate and rotate in the plane, and thus have a $G = SE(2)$ position space. Positions in this space are canonically parameterized as $g = (x, y, \theta)$, representing a position and orientation. The group action on *SE*(2) treats these positions as relative displacements, encoded as the multiplication of matrices of the form

$$g = \begin{bmatrix} \cos\theta & -\sin\theta & x \\ \sin\theta & \cos\theta & y \\ 0 & 0 & 1 \end{bmatrix}. \tag{1}$$

More details on such *homogeneous representations* can be found in [31].

The shape of a system specifies the relative positions of the points or particles that compose it. For a multi-body system made up of rigid elements, the space of shapes may be taken as the set of joint angles or extensions, interpreted as a subset of $\mathbb{R}^n$ bounded according to the joint limits. In some cases, such as when the joint is the axle of a wheel with no limits, the corresponding component of the shape space may be taken as a circle $\mathbb{S}$ rather than a subset of $\mathbb{R}$.

3.1.2 Velocity

The time derivative of the system position, i.e., its *world velocity*, $\dot{g}$, is the translational and rotational velocity of its body frame with respect to a fixed inertial frame, and can be expressed in coordinates as $(\dot{x}, \dot{y}, \dot{\theta})$ as illustrated in Fig. 2(a). World velocities at different positions that represent the same system *body velocity* or *spatial velocity* are related by *SE*(2)'s *lifted actions* (see Sect. A.2). The lifted actions can be combined into *adjoint actions* that convert between body and spatial representations of the velocity.

[2] See Sect. A.5.

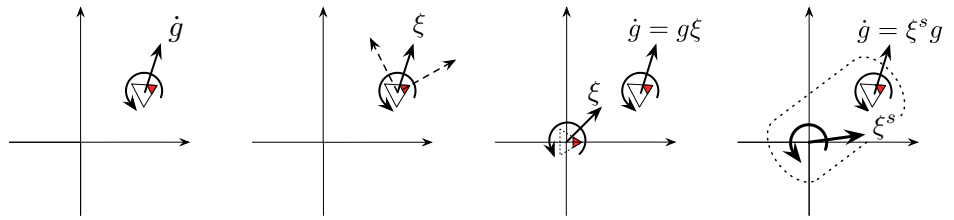

(a) World velocity (b) Body velocity I (c) Body velocity II (d) Spatial velocity

Fig. 2. Three 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*, ξ, in (b) is the velocity represented in the robot's instantaneous local coordinate frame. The body velocity is actually calculated by transporting the body back to the origin frame, as in (c), but by symmetry this is equivalent to bringing the world frame to the system. The *spatial velocity* in (d) is the velocity of the abstract rigid body to which the system is attached, measured at the origin.

Body velocity

The body velocity ξ in Fig. 2(b) is the position velocity expressed in the instantaneous body frame, i.e., its forward, lateral and turning velocities, and is calculated as

$$\xi = \begin{bmatrix} \xi^x \\ \xi^y \\ \xi^\theta \end{bmatrix} = \underbrace{\begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{T_g L_{g^{-1}}} \dot{g}, \tag{2}$$

where θ is the system's orientation and $T_g L_{g^{-1}}$ is calculated as described in Sect. A.3. Strictly speaking, (2) calculates the equivalent velocity to $\dot{g}$ at the group identity (origin selected for the position space) under left group actions, as shown in Fig. 2(c). General principles of coordinate invariance and group symmetry, however, mean that we can select any frame as the origin and that moving the body frame to the origin is equivalent to moving the origin to the body frame, making these two representations equivalent. Body velocities can be converted back into the world frame by inverting the lifted action from (2) as $T_e L_g = (T_g L_{g^{-1}})^{-1}$,

$$\dot{g} = \underbrace{\begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{T_e L_g} \xi. \tag{3}$$

Spatial velocity

Frames with the same spatial velocity ξ^s are moving as if rigidly attached to each other, and the spatial velocity is itself the velocity of the rigidly attached point that is currently over the origin. The spatial velocity of a system moving with given $\dot{g}$ at position g is calculated as the right-equivalent velocity at the origin of the space, and as depicted in Fig. 2(d) is equal to $\dot{g}$ plus a "cross-product" term that multiplies the rotational velocity by the translational displacement from the origin. This velocity, calculated as

$$\xi^s = \underbrace{\begin{bmatrix} 1 & 0 & y \\ 0 & 1 & -x \\ 0 & 0 & 1 \end{bmatrix}}_{T_g R_{g^{-1}}} \dot{g}, \tag{4}$$

Table 1. Interpretations of elements of T_eG, as used in this paper.

Symbol	Meaning	First Introduced
ξ	Body velocity	Sect. 3.1
z	Exponential coordinates	Sect. 3.1
ζ	Body velocity integral	Sect. 4.2.1
$\bar{\zeta}$	Corrected body velocity integral	Sect. 4.2.2

is seldom directly of interest as a physical quantity, but serves as a useful abstract representation of the system's motion–if an object's spatial velocity is known, then the world velocity of any frame attached to the body can be calculated as

$$\dot{g} = \underbrace{\begin{bmatrix} 1 & 0 & -y \\ 0 & 1 & x \\ 0 & 0 & 1 \end{bmatrix}}_{T_eR_g} \xi^s. \tag{5}$$

Velocities and Lie algebras

Because the body velocity ξ is an element of T_eG, it is also an element of the Lie algebra[3] $\mathfrak{g}$, producing physical interpretations for both the exponential map and the Lie bracket: In $SE(2)$ (for which the Lie algebra is denoted $\mathfrak{se}(2)$), the exponential map $g = \exp(z)$ of a vector $z \in \mathfrak{se}(2)$ finds the net displacement of a system starting at the origin and moving with body velocity $\xi = z$ for one unit of time, and takes the form

$$(x, y) = \begin{cases} (z^x, z^y), & \text{for } z^\theta = 0 \\ \dfrac{1}{z^\theta} \begin{bmatrix} \sin z^\theta & \cos z^\theta - 1 \\ 1 - \cos z^\theta & \sin z^\theta \end{bmatrix} \begin{bmatrix} z^x \\ z^y \end{bmatrix}, & \text{for } z^\theta \neq 0 \end{cases} \tag{6}$$
$$\theta = z^\theta.$$

The elements of z are the *exponential coordinates* of the position g. Note that both the body velocity and exponential coordinates are elements of the tangent space of G at the body frame, but have different physical interpretations–the body velocity encodes the rate at which the configuration is changing, and the exponential coordinates identify a configuration away from the origin. In the following discussion, we will maintain the orthographic distinction between these quantities, summarized in Table 1, adding two more interpretations in Sect. 4.2.

The Lie bracket[4] of two vectors $a, b \in \mathfrak{se}(2)$ finds the net effect of making sequential, infinitesimal moves in the $a, b, -a, -b$ directions and is calculated as

$$\left[\begin{pmatrix} a^x \\ a^y \\ a^\theta \end{pmatrix}, \begin{pmatrix} b^x \\ b^y \\ b^\theta \end{pmatrix} \right] = \begin{pmatrix} b^\theta a^y - a^\theta b^y \\ a^\theta b^x - b^\theta a^x \\ 0 \end{pmatrix}, \tag{7}$$

where the multiplications by a^θ and b^θ originate in a linearization (i.e., small angle approximation) of $T_{\mathbf{e}}L_g$ around $\theta = 0$. The x and y components of the $\mathfrak{se}(2)$ Lie bracket capture the "parallel parking" effect in which oscillating translation and rotation motions produce net translation in an orthogonal direction. The 0 value for the θ row of the Lie bracket product reflects that $SO(2)$ is an abelian (commutative) subgroup of the semi-direct group $SE(2)$ [8], so that oscillations in θ are always self-canceling, even in the presence of intermediate translations.

[3] See Sect. A.2.
[4] See Sect. C.4.

3.1.3 Reconstruction equation

Many locomoting systems' equations of motion are dictated by linear constraints defined in their body frame (i.e., which are symmetric with respect to left group actions). If there are as many of these constraints as there are dimensions in the position space, they define a *principal connection* (Sect. C.3) $\mathcal{A}$ on the configuration space. The geometric mechanics community [3,4,7,28,32] has exploited this connection structure with the development of the *kinematic reconstruction equation*, which makes use of the *local form* $\mathbf{A}$ of the connection to relate the body velocity of the system, ξ, to its shape velocity, $\dot{r}$, as

$$\xi = -\mathbf{A}(r)\dot{r}, \tag{8}$$

in which the local connection acts similarly to the Jacobian relating the end-effector velocity of a robotic arm to the configuration and velocity of its joints.

For these systems, the local connection is most easily calculated by identifying a set of constraints on the system's motion under which allowable combinations of body and shape velocity satisfy an equation of the form

$$\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} = \omega(\alpha) \begin{bmatrix} \xi \\ \dot{\alpha} \end{bmatrix}, \tag{9}$$

where ω is a 3×5 matrix. Once these constraints are identified and encoded in ω, they can be separated into two sub-blocks as

$$\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \omega_1^{3\times 3} & \omega_2^{3\times n} \end{bmatrix} \begin{bmatrix} \xi \\ \dot{\alpha} \end{bmatrix}, \tag{10}$$

(where n is the number of shape variables) and then recombined as

$$\omega_1 \xi = -\omega_2 \dot{\alpha} \tag{11}$$

and

$$\xi = -\omega_1^{-1}\omega_2 \dot{\alpha}. \tag{12}$$

Finally, setting $\mathbf{A} = \omega_1^{-1}\omega_2$ puts (12) into the form of (8).

3.2 Example systems

The geometric mechanics literature includes a wide variety of models for multi-body locomoting systems that provide concrete examples of the reconstruction equation and local connection. Of these systems, those that best illustrate the distinction between nonconservative and noncommutative effects have kinematic reconstruction equations of the form of (8). This category includes several broad classes of systems, including the *purely mechanical* and *purely kinematic* systems described in [4,28] and swimming systems at low and high Reynolds number [8,20,21,29].

Here, as in [11], we take the three systems in Fig. 1 as example models: the *kinematic snake* [27,28], the *floating snake* or *planar space robot* [3,27,28], and the common *differential-drive car*. The snake-like systems illustrate two fundamental physical regimes (nonholonomically and inertially constrained, respectively) and are simple enough to be tractable, but complex enough to pose interesting motion planning problems. Additionally, as we discuss in [29,30,33], they bear strong resemblance to the low and high Reynolds number swimming systems, making our system selection representative of a broad spectrum of kinematic systems. The differential-drive car is a simpler example of a nonholonomically-constrained system that highlights the principles under investigation.

Differential-drive car

The coordinates for the differential-drive car are $g = (x, y, \theta) \in SE(2)$ for the chassis, and $r = (\alpha_1, \alpha_2) \in \mathbb{S} \times \mathbb{S}$ for the wheel rotations. The reconstruction equation for this system is based on the no-slip and no-slide assumptions for the contact points of its wheels with the ground, which provide three independent constraints relating the body velocity of the vehicle to the angular velocities of the wheels [3]. These constraints can be encoded in the form of (9) as

$$\begin{bmatrix}0\\0\\0\end{bmatrix} = \begin{bmatrix}\xi_1^x\\ \xi_2^x\\ \xi_2^y\end{bmatrix} = \underbrace{\begin{bmatrix}1 & 0 & -w & -R & 0\\ 1 & 0 & w & 0 & -R\\ 0 & 1 & 0 & 0 & 0\end{bmatrix}}_{\omega(\alpha)} \begin{bmatrix}\xi\\ \dot{\alpha}\end{bmatrix}, \tag{13}$$

In which R and w are respectively the wheel radius and body half-width of the car. Normalized for wheel radius and body width, the reconstruction equation for the differential-drive car can then be found as

$$\xi = -\begin{bmatrix}-1 & -1\\ 0 & 0\\ 1 & -1\end{bmatrix}\begin{bmatrix}\dot{\alpha}_1\\ \dot{\alpha}_2\end{bmatrix}. \tag{14}$$

by applying (10)–(12). 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.

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 dictates that the linear and angular momentum remain zero, so the system's velocity vectors are subject to the constraint

$$\begin{bmatrix}0\\0\\0\end{bmatrix} = \begin{bmatrix}J_x\\ J_y\\ J_\theta\end{bmatrix} = \omega(\alpha)\begin{bmatrix}\xi\\ \dot{\alpha}\end{bmatrix}, \tag{15}$$

where the system kinematics provide the terms in ω.

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 [28]. Taking the rows together, the reconstruction equation for a floating snake with equal link lengths and inertias is

$$\xi = -\frac{1}{F}\begin{bmatrix}0 & 0\\ 0 & 0\\ a_1^3 & a_2^3\end{bmatrix}\begin{bmatrix}\dot{\alpha}_1\\ \dot{\alpha}_2\end{bmatrix}, \tag{16a}$$

where

$$a_1^3 = 5 + 3\cos(\alpha_2) + \cos(\alpha_1 - \alpha_2) \tag{16b}$$

$$a_2^3 = 5 + 3\cos(\alpha_1) + \cos(\alpha_1 - \alpha_2) \tag{16c}$$

$$F = 19 + 6(\cos(\alpha_1) + \cos(\alpha_2)) + 2\cos(\alpha_1 - \alpha_2). \tag{16d}$$

The constants in these expressions derive from the kinematics of the links and their positions in the center of mass frame [25,34].

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. Each link has a no-slide constraint, such as that imposed by a passive wheelset or ice skate, that acts as to prevent lateral motion while freely allowing longitudinal and rotational motion. These constraints can be encoded as

$$\begin{bmatrix}0\\0\\0\end{bmatrix} = \begin{bmatrix}\xi_1^y\\ \xi_2^y\\ \xi_3^y\end{bmatrix} = \omega(\alpha)\begin{bmatrix}\xi\\ \dot{\alpha}\end{bmatrix}, \tag{17}$$

where as with the floating snake, ω is constructed from the system kinematics. Normalized for unit body-length, these constraints can then be reworked into

$$\xi = -\frac{1}{F}\begin{bmatrix}(1+\cos(\alpha_2))/6 & (1+\cos(\alpha_1))/6\\ 0 & 0\\ \sin(\alpha_2) & \sin(\alpha_1)\end{bmatrix}\begin{bmatrix}\dot{\alpha}_1\\ \dot{\alpha}_2\end{bmatrix}, \tag{18}$$

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

The kinematic snake's singular configurations are those for which the no-slide constraints allow the system to move as a rigid body; these singularities correspond to configurations in which the constraint distribution contains a vector field aligned with the fiber direction, and so no longer acts as a connection (Sect. C.3). On the kinematic snake, singularities occur when $\alpha_1 = \alpha_2$, such that the normals to the passive wheels meet at a single point, or when α_1 or α_2 is equal to π, such that the two constraints are collocated.

3.3 Shape changes, gaits, and image-families

When discussing motion plans for locomoting systems, it is useful to have a vocabulary describing operations in their shape spaces. In this paper, we use the following definitions:

Definition 1 (Shape change). *A shape change ψ is a trajectory in the shape space B of the system over an interval $[0, T]$, i.e., the set of all shape changes is*

$$\Psi = \{\psi \in C^1 \mid \psi : [0, T] \to B\} \tag{19}$$

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

Definition 2 (Gait). *A gait ϕ is a cyclic shape change, i.e., the set of all gaits is*

$$\Phi = \{\phi \in \Psi \mid \phi(0) = \phi(T)\}. \tag{20}$$

Note that a gait has a defined starting shape $\phi(0)$ and pacing; two gaits whose images in B are the same closed curve, but that have different start points, are distinct.[5]

Definition 3 (Image-family). *The* image-family $\bar{\phi}$ *of a gait ϕ is the set of all gaits which share its image (i.e., trace out the same closed curve) in B,*

$$\bar{\phi} = \{\varphi \in \Phi \mid \mathrm{Im}(\varphi) = \mathrm{Im}(\phi)\}. \tag{21}$$

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

[5] Gaits that follow the same curve at different rates are also distinct, but the kinematic nature of our systems allows us to drop this distinction from the analysis.

4 Integrating the local connection

For a given shape change ψ, the resulting displacement can be calculated by passing the shape $\psi(t)$ and shape velocity $\dot{\psi}(t)$ at each time into the reconstruction Eq. (8) to find $\xi(t)$, the body velocity as a function of time, and then integrating $\dot{g} = T_e L_g \xi$ as a differential equation to find the net displacement. In coordinates, this displacement can be calculated via standard methods for ordinary differential equations,

$$g(T) = \int_0^T T_e L_{g(t)} \xi(t)\, dt = \int_0^T -T_e L_{g(t)} \mathbf{A}(r(t))\dot{r}(t)\, dt. \tag{22}$$

Finding a shape change that produces a desired net displacement or identifying the most effective gaits available to a system is considerably more difficult, as the recursive time integral in (22) is inherently resistant to inversion. In limited cases, it is possible to specify a desired position trajectory $g(t)$ and (pseudo)invert the local connection to find a corresponding control trajectory $\dot{r}(t) = -\mathbf{A}^{-1}(r(t))\xi(t)$, but joint limits and constraint singularities often prevent such techniques from succeeding over extended trajectories.

Gait-based approaches provide an attractive solution to the path planning and identification problems posed by these concerns. Rather than specifying a long-range path through position space and dictating that the system rigidly follow this path, a motion planner can design cyclic gaits with a range of short-term net displacements and then choose among them to build a motion plan that "on average" follows a specified path. This raises the question, then, of how to identify useful gaits. The geometric mechanics community has investigated principled answers to this question through *averaging* techniques [1–8, 23], which are related to the Lie bracket. Here, we briefly review averaging methods, while drawing connections between them that have not previously been made explicit.

4.1 Lie brackets and the local connection

A gait with very small amplitude in the shape space can be considered as a infinitesimal oscillation of the shape, and the net motion resulting from such an input can be found by the use of *Lie brackets* (see Sect. C.4). For a specified input shape velocity $\dot{r}$, a system at configuration $q = (r, g)$ has position velocity $\dot{g} = -T_e L_g \mathbf{A}(r)\dot{r}$. Moving with this velocity can be interpreted as flowing along the vector field $\mathbf{X}(q)$ defined over the configuration space as the combination of this shape velocity and the position velocity it induces,

$$\mathbf{X}(q) = \begin{pmatrix} \dot{r} \\ -T_e L_g \mathbf{A}(r)\dot{r} \end{pmatrix}. \tag{23}$$

If we define two unit-magnitude input shape velocities as $\dot{r}_1 = [1\ 0]^T$ and $\dot{r}_2 = [0\ 1]^T$, then the Lie bracket

$$\left[\begin{pmatrix} \dot{r}_1 \\ -T_e L_g \mathbf{A}(r)\dot{r}_1 \end{pmatrix}, \begin{pmatrix} \dot{r}_2 \\ -T_e L_g \mathbf{A}(r)\dot{r}_2 \end{pmatrix} \right] \Bigg|_{q_0}, \tag{24}$$

gives the average velocity vector achieved by infinitesimally flowing along the vector fields defined by $\dot{r}_1$, $\dot{r}_2$, $-\dot{r}_1$, and $-\dot{r}_2$ in order, i.e., infinitesimally oscillating the shape around an initial configuration q_0.[6]

The structure of the configuration space gives the Lie bracket in (24) a natural second representation (as derived in Sect. D),

$$\begin{pmatrix} \mathbf{0} \\ T_e L_{g_0} \left(-\mathbf{dA}(r_0) + \left[T_e L_{g_0^{-1} g} \mathbf{A}_1(r_0),\ T_e L_{g_0^{-1} g} \mathbf{A}_2(r_0) \right] \right) \end{pmatrix} \tag{25}$$

[6] Although we use coordinates here, the results hold for any selection of independent $\dot{r}_1$ and $\dot{r}_2$ vectors and so we do not incur any loss of generality.

where $\mathbf{A}_i$ is the ith column of $\mathbf{A}$ and $\mathbf{dA}$ is its exterior derivative with respect to $\dot{r}_1$ and $\dot{r}_2$. Note that in the Lie bracket in (24), $\mathbf{A}$ varies with r, but that in the Lie bracket term in (25), it is held to its value at r_0. In subsequent discussion, we will refer to this latter form as the *local* Lie bracket. The top half of (25) is null, as a cyclic motion in B by definition causes no net change in those coordinates. In the bottom half, the exterior derivative and local Lie bracket terms respectively represent how the system dynamics change with the shape and with motion through the position space, in a blockwise application of (C.4) to (24)

We can simplify the rather ungainly expressions in (24) and (25) by drawing on the Lie group structure of the position space: First, we focus our attention only on the position-space component in the lower half of the equation. Second, we set the origin of the space at the current location and orientation of the system (i.e., taking $g_0 = e$), such that (24) gives the averaged velocity in the body frame of the system, thereby eliminating the $T_e L_{g_0}$ term. Finally, we note that as $\mathbf{A}_1$ and $\mathbf{A}_2$ are both elements of $T_e G$, they are elements of $\mathfrak{g}$, and we can use the special definition for a Lie bracket of Lie algebra elements in (C.5) to make the lifted actions inside the Lie bracket implicit, as described in §C.4. These three changes reduce the lower halves of (24) and (25) to

$$\Big[\mathbf{A}(r)\dot{r}_1,\ \mathbf{A}(r)\dot{r}_2\Big]\Big|_{r_0} = \big(-\mathbf{dA} + \big[\mathbf{A}_1, \mathbf{A}_2\big]\big)\,(r_0) \tag{26}$$

$$= D\mathbf{A}(r_0), \tag{27}$$

in which the right-hand expression is the *local curvature* of the connection at r_0 [8]. An expanded treatment of this derivation is given in the Appendix.

The local curvature $D\mathbf{A}$ is the projection (Sect. C.3) of the curvature, or covariant exterior derivative (Sect. B.6), of the full connection one-form $\mathcal{A}$ into the space of allowable velocities. The $-\mathbf{dA}$ component measures the intrinsic change in the connection across the shape space; the $[\mathbf{A}_1, \mathbf{A}_2]$ component corresponds to the Christoffel symbols that capture the extrinsic changes in the connection as the space of allowable velocities rotates with the system body frame.

4.2 The body velocity integral (BVI) and the corrected body velocity integral (cBVI)

For larger-amplitude (i.e., non-infinitesimal) gaits, Radford and Burdick [35] have shown that the exponential coordinates $z(\phi)$ for the net displacement over a gait, $\phi(T)$, can be approximated as

$$z(\phi) = \iint_{\phi_a} -\mathbf{dA} + [\mathbf{A}_1, \mathbf{A}_2]\, dr_1\, dr_2 + \text{ higher-order terms}, \tag{28}$$

where ϕ_a is the oriented region of the shape space B enclosed by ϕ, and the two integrands sum to the right-hand side of (26). The expression in (28) is of an especially useful form, as the first two terms are both area integrals over the region of the shape space enclosed by the gait and the integrands can be obtained directly from the local connection. Plotting either the first integrand [10,28] or the sum of the two integrands [8,21], provides a visual representation of the displacements that different gaits produce, based on the values of the integrands in areas the gaits enclose.

This notion plays a key role in the results of this paper, but before we examine it in more depth, we make an aside to supply some insight as to the form of (28). In previous works, this equation has only been presented as resulting from the direct application of calculus to the Magnus expansion for a Lie group [36], but examining the components individually provides strong intuition as to their provenance and links them directly to the nonconservativity and noncommutativity of the system dynamics.

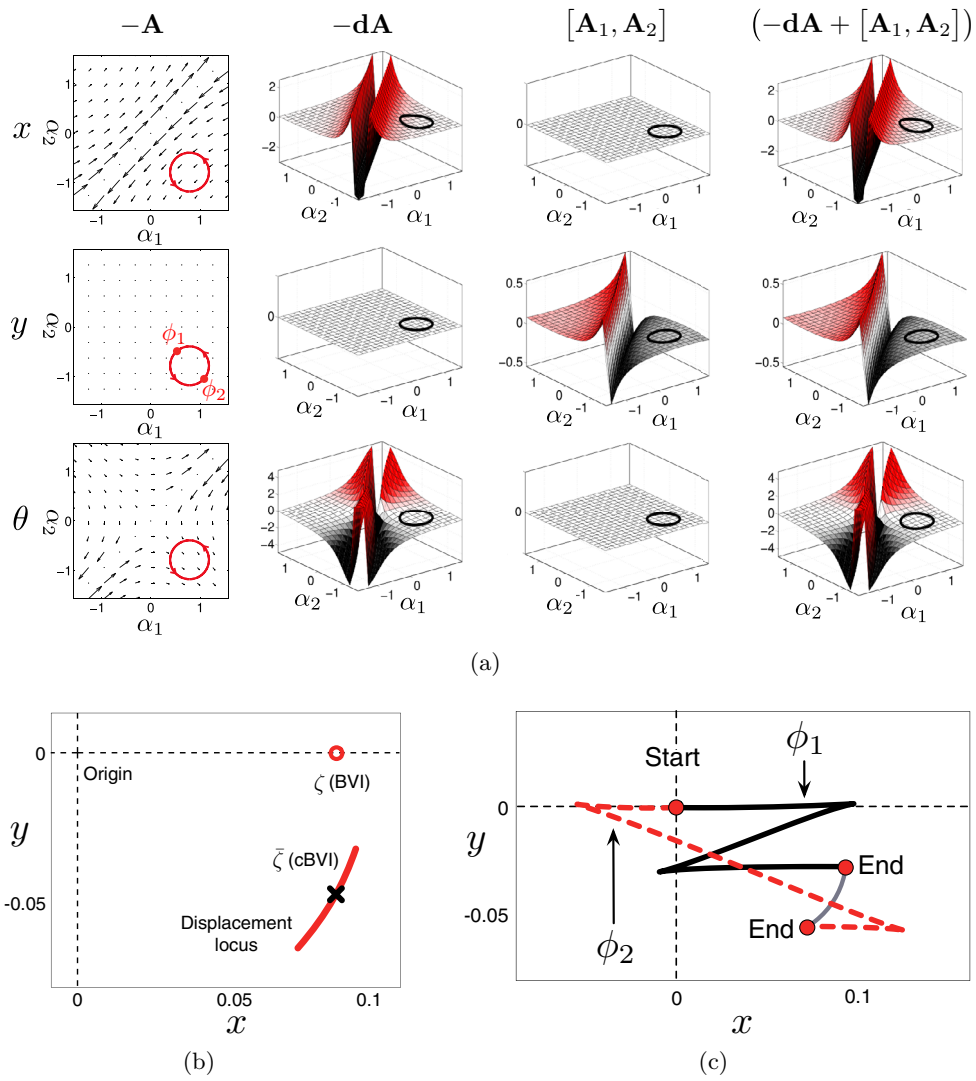

Fig. 3. (a) An example gait image-family for the kinematic snake, overlaid on the system's connection vector fields and curvature functions. The exterior derivative measures the nonconservativity of the local connection over the shape space, while the local Lie bracket measures its noncommutativity. Note x component of $\mathbf{dA}$ is positive definite, the y component of $[\mathbf{A}_1, \mathbf{A}_2]$ is sign-definite in each half of the shape space with a significant non-zero value, and that the θ component of $\mathbf{dA}$ has symmetric positive and negative regions in the area enclosed by the gaits. (b) BVI, cBVI, and displacements for the kinematic snake gaits in the image-family depicted in Fig. 3(a). (c) The intermediate motions of the system over the two gaits from the image-family whose displacements form the endpoints of the arc in (b).

To ground this discussion, we refer to the example kinematic snake gaits overlayed on the connection vector fields and their derivatives in Fig. 3(a). As discussed in [10,11], the counterclockwise circular path through the shape space in Fig. 3(a) corresponds to an image-family of gaits (one gait per starting point on the curve), for which the resulting locus of displacements is the arc in Fig. 3(b), with zero net rotation over any of the gaits. Note that for zero net rotation (i.e. $z_\theta = 0$), the exponential map in (6) is an identity map, and that for this

example, the z_x and z_y components of the exponential coordinates thus correspond directly to the x and y components of the displacement.

4.2.1 Body velocity integral

Previously, we have shown [9–11] that the first term of the integral in (28) forms the *body velocity integral* (BVI), ζ, of the system over ϕ. The BVI describes the net "forward minus backward" motion of the system over the gait in each body direction,

$$\zeta(T) = \int_0^T \xi(t)\, dt = -\int_\phi \mathbf{A}(r)\, dr = -\iint_{\phi_a} \mathbf{dA}\, dr_1\, dr_2, \tag{29}$$

where the last equality is given by Stokes' theorem. The BVI can equivalently be interpreted as a quantity for which $\zeta(T)/T$ is the time-averaged body velocity experienced by the system over the gait ϕ, making $\exp\zeta(T)$ a first-order approximation of $g(T)$, and $\zeta(T)$ thus a first-order approximation of the exponential coordinates $z(T)$ of the net displacement over the gait.

All the gaits in an image-family enclose identical regions of the shape space, and thus have the same BVI. By plotting the exterior derivative of the local connection, as in the second column of Fig. 3(a), we gain a visual representation of this integral. For the example image-family, the gaits enclose:

1. A positive region of the x exterior derivative function (which is positive-definite outside of the singularity);
2. A zero region of the lateral function (which is zero everywhere); and
3. Equal positive and negative regions of the rotational function (which is antisymmetric around the $\alpha_1 = -\alpha_2$ line).

These $(+, 0, 0)$ values exponentiate to the point on the x axis in Fig. 3(b), with an orientation of $\exp\zeta_\theta = \zeta_\theta = 0$.

The value of the BVI is entirely determined by the change in the local connection over the enclosed region of the shape space, as measured by the exterior derivative. As the exterior derivative measures how non-closed (Sect. B) the rows of the local connection are, or, equivalently via the dual formulation, how nonconservative the connection vector fields are, we take the BVI as the *nonconservative contribution* to the displacement over the gait in our subsequent discussion.

4.2.2 Corrected body velocity integral

In abelian position spaces, exponentiating the average body velocity is sufficient to find the net displacement [8], and $z = \zeta$. Because of the semi-direct product rules discussed in Sect. A.4, this principle holds for the θ component of $SE(2)$ [8], and gaits can easily be selected for the net orientation change they produce by locating them in sign-definite or sign-balanced regions [28]. The group $SE(2)$ as a whole, however, is non-abelian, and the net displacement depends on the *order* in which the system moves forward, turns, and moves backward during the gait. Figure 3(b) illustrates the effect of this property on the example gaits, in that the BVI does not capture any of the lateral motion experienced by the system as a result of the intermediate turning motion apparent in Fig. 3(c). Including the second integral term in (28) (the local Lie bracket) incorporates some of the ordering information to correct for this deficiency. Specifically, it corresponds to the *cyclic ordering* of the translation and rotation motions in the gait, and thus encodes the average effects of turning on gaits in a given image-family.

In the example image-family, this cyclic ordering ensures that at least one of two conditions holds: either clockwise rotation precedes positive longitudinal (forward) translation, or counterclockwise rotation precedes negative longitudinal (backward) translation. Each of these conditions results in the system accruing a $-y$ displacement in a manner similar to parallel parking, the magnitude of which is approximated by the second integral in (28). The integrand of this expression is plotted in the third column of Fig. 3(a). Note that the x and θ components of the local Lie bracket term (calculated according to (7)) are zero; x because $\mathbf{A}^y = \mathbf{0}$, and θ as an inherent property of $SE(2)$ systems.

In contrast to the BVI, the local Lie bracket is entirely determined by the noncommutativity of the columns of $\mathbf{A}$ (which are vectors in the position space) at each point in the shape space, and is unaffected by any variation in the local connection. We therefore will refer to the area integral of the local Lie bracket as the *primary noncommutative contribution* to the displacement over a gait.

The sum of the exterior derivative and local lie bracket integrals features heavily in the averaging locomotion literature, most notably in [8,21], which developed motion planning techniques based on plots of the combined integrands similar to those in the fourth column of Fig. 3(a). This quantity, however, has yet to be explicitly named except as an approximation to the exponential coordinates of the net displacement [8]. In accordance with the physical intuition outlined above, we now identify it as the *corrected body velocity integral* (cBVI) $\bar{\zeta}$. As with the original BVI, as an area integral over the enclosed area, the cBVI is the same for all gaits in an image family.

4.2.3 Higher-order terms

The cBVI accounts for the two most significant terms in (28). To close out our discussion of the physical interpretation of this equation, we now turn to the higher-order terms that cause the displacements over the gaits in an image-family to form a distributed locus, rather than being at a single point. These higher-order terms capture the linearization error inherent in expanding the local Lie bracket to non-infinitesimal values and come from two main sources. First, each gait in the image-family starts at a different point in the cycle, so the specific order in which the system translates and turns is unique to each gait. Even though the cyclic order is the same for all the gaits, this change of start point can have a significant effect on the resulting displacement, such as that illustrated in Fig. 3(c) for two gaits chosen from the example image-family. Here, ϕ_1 moves the system forward at zero angle and backward with positive θ, while ϕ_2 moves the system backward at zero angle and forward with negative θ. As noted previously, both these patterns result in a net $-y$ displacement, but because the system moves further during the forward section of the gait than the backward section, ϕ_2 moves the system further in the $-y$ direction. The remaining gaits in the image family fall between the two selected motions, and fill in the arced displacement locus in Fig. 3(b).

Second, the magnitude of displacement for the gaits (which, for zero net rotation is the same across all gaits in the family [11]) can vary significantly from that of the cBVI. Both this error and the spread of the displacement locus can be viewed as linearization error in the Lie bracket, originating in small angle approximations for body-frame rotation during a gait. The spread error (orthogonal to the direction of motion) grows more quickly than the magnitude error (parallel to the direction of motion), much in the same way that a pendulum exhibits more side-to-side movement than vertical motion while swinging at small to moderate amplitudes.

As with the local Lie bracket, the higher-order terms arise as a consequence of the noncommutativity of the position space. The changes in the local connection across the shape space do influence these terms, but they only appear for non-abelian position spaces, and so we refer to them as the *secondary noncommutative contribution* to the displacement. While this contribution is not directly calculable from the local Lie bracket, the bracket (as a measure of noncommutativity) does provide an indication of its magnitude.

5 Minimum perturbation coordinates

Recently, we have shown that the span of the displacement locus and its distance from the cBVI (i.e. the magnitudes of the local Lie bracket and the higher-order terms) both depend strongly on the choice of coordinates used to describe the system [9], and that it is possible to optimize this choice of coordinates to minimize these two quantities [10,11], making the cBVI an accurate predictor of the displacement over non-infinitesimal gaits and removing the dependence on initial phase within the gait. In those works, we attributed the benefit of working in the new coordinates to having found a body frame for the system that rotated very little during locomotion, with the result that integrating velocities in the body frame or world frame was effectively equivalent. Using our notions of nonconservative and noncommutative contributions to locomotion, we can now extend this physical intuition into a differential geometric formulation, placing a firm theoretical foundation beneath these prior results.

5.1 Coordinate definition

The optimal coordinates identified in [10,11] are *minimum-perturbation coordinates*, in that they are the choice of body frame whose position and orientation moves the least in response to changes in the system shape, as measured by the norm of the local connection (in its role as the Jacobian from inputs to outputs). This choice of coordinates has two effects: First, it allows us to isolate the "true" motion of the system as a whole. Second, by reducing motion in the θ direction, we minimize the influence of the noncommutative terms in the equations of motion.

The motion of the body frame in response to shape changes is encoded in the local connection. At a given configuration, the natural norms for measuring this motion are the magnitudes of the velocity vectors produced in the translational and rotational subspaces of $SE(2)$ by the local connection,[7]

$$\|\mathbf{A}(r)\|_{xy}^2 = \sum_i \left((\mathbf{A}_i^x(r))^2 + (\mathbf{A}_i^y(r))^2\right) \tag{30a}$$

$$\|\mathbf{A}(r)\|_{\theta}^2 = \sum_i \left((\mathbf{A}_i^\theta(r))^2\right). \tag{30b}$$

Integrating[8] these norms over a region of interest in the shape space, $\mathcal{B} \subset B$, as

$$D_j = \int_{\mathcal{B}} \|\mathbf{A}\|_j^2 \, d\mathcal{B}, \tag{31}$$

finds the average scale of motion produced by changing shape within this region, allowing direct comparison between the local connections corresponding to different coordinate systems. Using the squared norms, rather than the norms themselves, gives the metric an energy-like aspect and means that a coordinate choice scores better (lower) if it spreads the magnitude of the local connection evenly over the region of interest, rather than concentrating it in a given region.

[7] For simplicity of notation, we assume here that unit infinitesimal changes in the shape variables are comparable inputs at different shapes. If a different metric (such as the energy dissipation metric in [37]) is preferable, it can be incorporated into these norms without loss of generality.

[8] If a metric is chosen for the norms in (30), it should also be incorporated into this integration; some applications may additionally suggest the incorporation of a second metric to weight the importance of different configurations, such as the arctangent scaling we discuss in [11] that de-emphasizes regions near kinematic singularities.

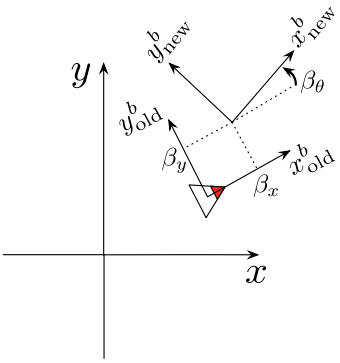

Fig. 4. An original, known-valid body frame, and a new frame displaced by that frame by β.

As described in [10,11], the minimum perturbation coordinates correspond to the solution of a modified Hodge-Helmholtz decomposition (Sect. B.5) of the local connection calculated in an initial, arbitrary choice of coordinates. This decomposition is predicated on the property that any new set of position coordinates defines a new body frame with a relative position $\beta \in SE(2)$ with respect to the original body frame, as illustrated in Fig. 4. Combined with the constitutive requirement that for any valid set of generalized coordinates, β is solely a function of the shape r [10,11], this leads to a general formulation for the local connection in new coordinates,

$$-\mathbf{A}_{\text{new}}(r) = T_{\beta(r)} L_{\beta(r)^{-1}} \big(-T_e R_{\beta(r)} \mathbf{A}(r) + \mathbf{d}\beta(r) \big), \tag{32}$$

where $T_{\beta(r)} R_{\beta(r)^{-1}}$ finds the motion of the new body frame induced by the motion of the original body frame, $\mathbf{d}\beta(r)$ adds its motion relative to the original frame, and $T_{\beta(r)} L_{\beta(r)^{-1}}$ changes the coordinates of these velocities from the original to new frames. As neither $T_\beta L_{\beta^{-1}}$ (whose structure follows that in (2)) nor $T_{\beta(r)} R_{\beta(r)^{-1}}$ (from (5)) affects the θ row of the local connection, the integrated norm for that row in the new coordinates is

$$D_\theta = \int_{\mathcal{B}} \| -\mathbf{A}^\theta + \mathbf{d}\beta_\theta \|^2 \, d\mathcal{B}. \tag{33}$$

The structure of equation (33) parallels that of (B.5), leading to our fundamental result in [10,11]: The orientation component of the minimum perturbation coordinates is that for which $\mathbf{d}\beta_\theta$ is the curl-free component of $-\mathbf{A}^\theta$. Applying the Hodge-Helmholtz decomposition to this one-form thus produces the coordinate transform that optimally minimizes rotation of the new body frame, leaving $-\mathbf{A}^\theta_{\text{new}}$ as its divergence-free component.

Finding the optimal xy coordinates that minimize translation is somewhat complicated by the cross-product term contributed by the right lifted action, but is essentially similar in principle. Expanding the right lifted action in (32),[9] the integrated norm of the local connection becomes

$$D_{xy} = \iint_{\mathcal{B}} \| -\mathbf{A}^x + \mathbf{d}\beta_x + \beta_y \mathbf{A}^\theta \|^2 + \| -\mathbf{A}^y + \mathbf{d}\beta_y - \beta_x \mathbf{A}^\theta \|^2 \, d\mathcal{B}, \tag{34}$$

for which β_x and β_y must be solved simultaneously, and in which $\mathbf{A}^x$ and $\mathbf{A}^y$ are in the original coordinates. We present a finite-element approach to solving both this decomposition and the standard Hodge-Helmholtz decomposition in [11].

For serial-chain systems, such as the kinematic and floating snake examples here, the optimal coordinates are roughly the position of the center of mass and the mean orientation of the links, with shape-dependent weightings of different links to account for nonlinearities

[9] Left lifted actions on $SE(2)$ rotate vectors without changing their magnitude, so the left action in (32) can be ignored when constructing the magnitude expression in (34).

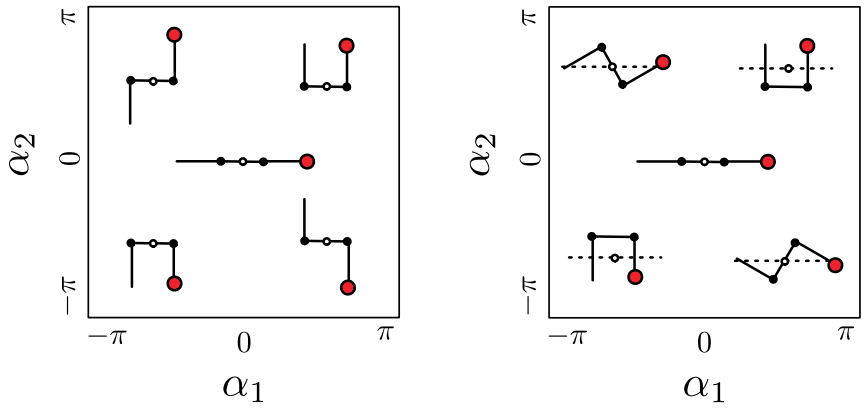

Fig. 5. Configuration of a three-link system with $\theta = 0$ in the original and minimum perturbation coordinates.

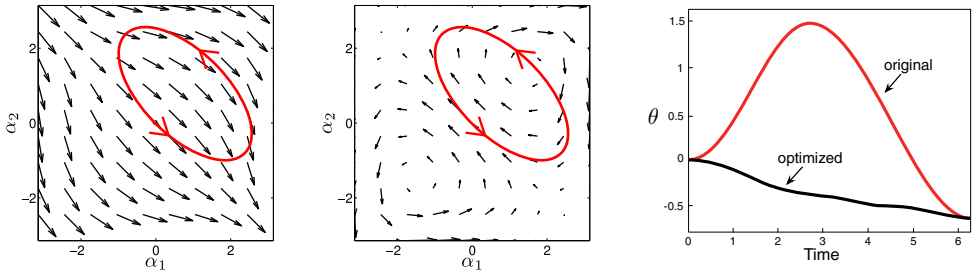

(a) Example gait overlaid on the original $-\mathbf{A}^{\theta}$ for the floating snake.

(b) Example gait overlaid on the optimized $-\mathbf{A}^{\theta}$ for the floating snake.

(c) Time history of the original and optimized orientations during the example gait.

Fig. 6. Optimizing the definition of θ for the floating snake strips out the gradient component from $\mathbf{A}^{\theta}$, 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 "wind up" and "power stroke" seen in the original orientation.

in the system constraints. Figure 5 compares the zero-orientation configuration for both the original (middle-link) coordinates and the optimized coordinates in a representative sampling of shapes.

5.2 Integrating in minimum-perturbation coordinates

Changing body frames naturally does not alter the physical motion of a system, but it does change how this motion is *represented*. We have just seen (in Sect. 5.1) how this affects infinitesimal motion via the transformation of the local connection. More directly relevant to locomotion analysis is how the representation of the *integrated* motion is modified.

As we note above, the first effect of the new coordinates is to more closely track the bulk motion of the system, with the new coordinates discarding extraneous motion that is canceled out over full gait cycles. A good example of this principle is the rotation of the three-link floating snake over the reorientation gait in Fig. 6, where the change in coordinates subtracts out the overall $(+\alpha_1, -\alpha_2)$ trend of $-\mathbf{A}^{\theta}$, leaving only the two true circulations [11]. Removing the conservative component means that only the net contribution of the shape change to the

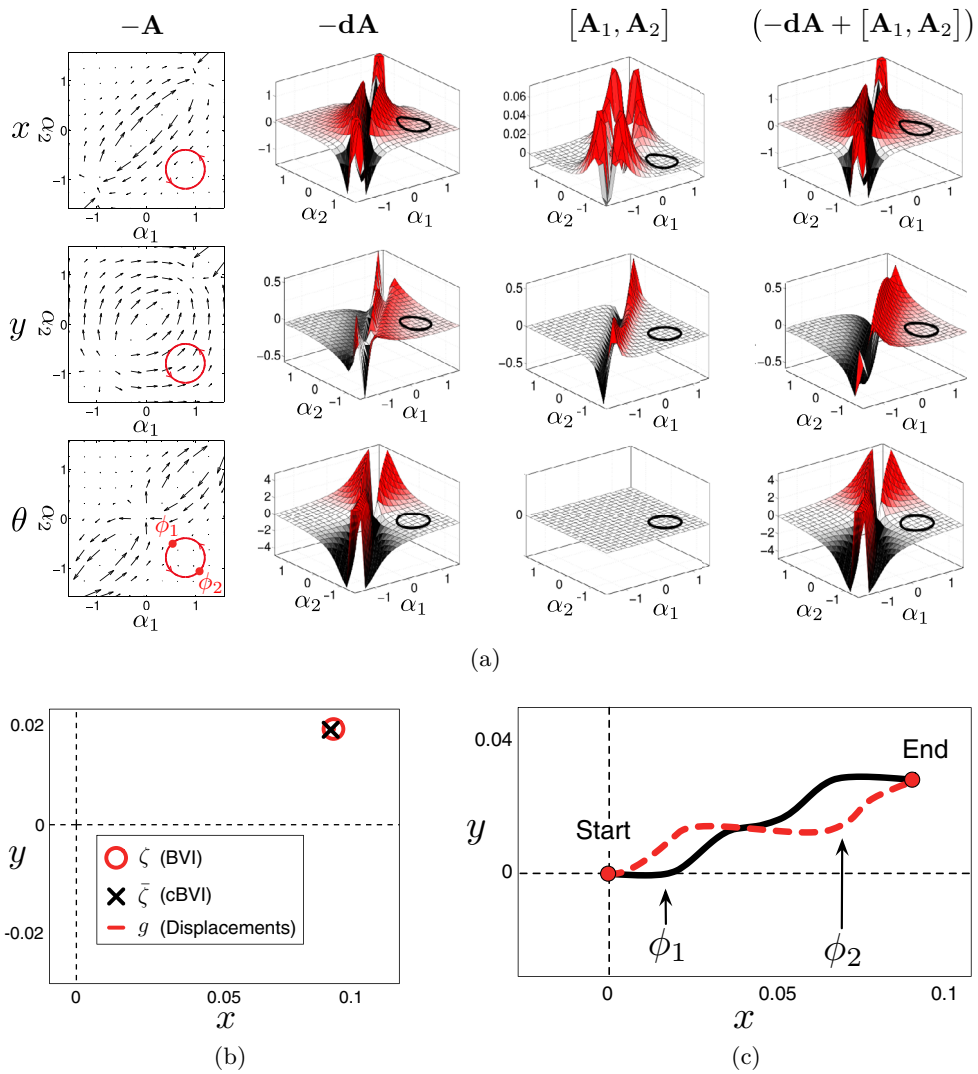

Fig. 7. Connection vector fields, curvature functions, BVI, cBVI, and displacements for the kinematic snake gaits in the image-family depicted in Fig. 3(a), as calculated in the minimum perturbation coordinates. (a) The exterior derivative measures the nonconservativity of the local connection over the shape space, while the local Lie bracket measures its noncommutativity. (b) The BVI and cBVI are now much closer and the displacement locus has effectively collapsed to a point collocated with ζ. (c) The intermediate motions the system over the two gaits from the image-family whose displacements form the endpoints of the arc in (b). The gaits now have almost-identical endpoints, and the change to a center-of-mass-like reference position has straightened the trajectories and removed the sharp cusps.

displacement is counted, clearly showing that the system monotonically approaches its final orientation. In the original coordinates, this trend is masked by the much larger oscillation of the middle link. The second, and more significant, effect of the new coordinates is that minimizing $\mathbf{A}^\theta$–and thus the θ terms in (7)–also minimizes the noncommutative terms in (28). This phenomenon is illustrated for the three-link kinematic snake in Fig. 7(a), in which the change to minimum-perturbation coordinates has effectively nullified $\mathbf{A}^\theta$ in the second and fourth quadrants. In these regions, $[\mathbf{A}_1, \mathbf{A}_2]^y$ has been reduced by two orders of magnitude

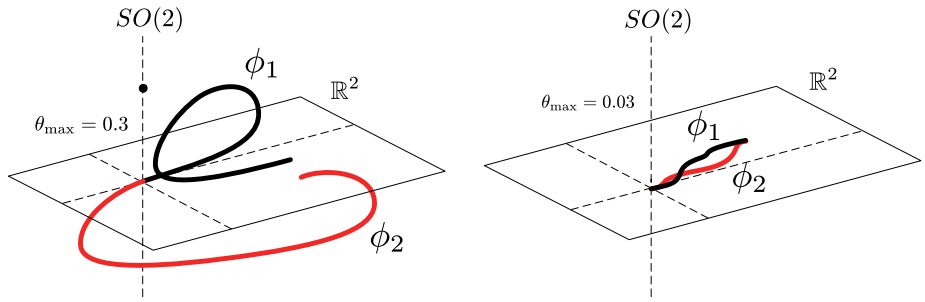

(a) Original coordinates (b) Optimized coordinates

Fig. 8. The example gaits from Figs. 3(c) and 7(c), plotted with the orientation θ in the third dimension.

from its values in the original coordinates, while $[\mathbf{A}_1, \mathbf{A}_2]^x$ remains an insignificant fraction of $\mathbf{dA}$. The net effect of these two properties is that away from the kinematic singularities at $\alpha_1 = \alpha_2$, the local curvature of $\mathbf{A}$ is approximately equal to its exterior derivative,

$$-\mathbf{dA} + [\mathbf{A}_1, \mathbf{A}_2] \approx -\mathbf{dA}, \tag{35}$$

and thus that the BVI and cBVI are approximately equal for gaits in these regions, as shown in Fig. 7(b). This equivalence was implicitly present in our earlier work, where we used the BVI rather than the cBVI, but the observations above give it a more solid foundation.

The higher-order terms in (28) have, like the local Lie bracket, noncommutative origins (and partly derive from the error accompanying the linearization inherent in the Lie bracket); they are thus also minimized by the change in coordinates, and go to zero more quickly than the local Lie bracket. As such, the exponential coordinates of the net displacement over a gait approach its (c)BVI,

$$z(\phi) \approx \left(\iint_{\phi_a} -\mathbf{dA}\, dr_1\, dr_2 + \iint_{\phi_a} [\mathbf{A}_1, \mathbf{A}_2]\, dr_1\, dr_2 \right) \approx \iint_{\phi_a} -\mathbf{dA}\, dr_1\, dr_2. \tag{36}$$

Critically, the disappearance in (36) of the higher-order terms means that the net displacement is (approximately) only a function of the area enclosed by the gait and no longer depends on the starting point. The net displacement is thus shared across all gaits in an image-family, as illustrated in Fig. 7(c) for the same pair of gaits as in Fig. 3(c), and the displacement locus for the image family effectively collapses into a point, as shown in Fig. 7(b).

5.3 The geometry of noncommutativity

Earlier, we provided a computational interpretation of the effect of working in minimum perturbation coordinates–that reducing the position velocity of the system, and in particular the rotational velocity, removes the noncommutative contributions to (28) by sending the local Lie bracket in (7) to zero. An alternative, geometric interpretation gives additional insight into this process and reconciles our results with the accepted notion that displacement should only coincide with the area integral of $\mathbf{dA}$ in abelian position spaces [8].

In Figs. 3(c) and 7(c), we plotted only the x and y trajectories of the system over the sample gaits, and noted the differences in smoothness and endpoint convergence between the two coordinate representations. If we include the θ components of these trajectories, as in Fig. 8, a third difference appears: under the original coordinates, the system makes significant excursions through the $SO(2)$ component of the position space, but in the optimized

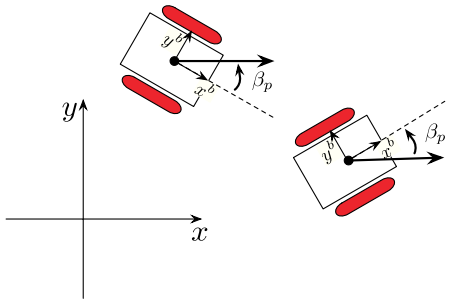

Fig. 9. The south-pointing chariot's pointer counter-rotates the main body by an angle $\beta_p = -\theta$, such that it maintains a constant angle with respect to the global coordinates.

coordinates it remains within a small neighborhood of the $\mathbb{R}^2$ plane in which it starts. As $\mathbb{R}^2$ is an abelian space, the system can be approximated as commutative.

Formally, the restriction to $\mathbb{R}^2$ corresponds to how the optimized coordinates affect the position elements of the system's constraint distribution (§C.2). At each configuration, the distributions for our locomoting systems span $T_r B$ (the shape space's tangent space), plus a two-dimensional subspace of $T_g G$ (the position space's tangent space). The position component $\mathcal{D}_G$ of the distribution is given by the columns of the local connection and takes the form

$$T_g L_{g^{-1}} \mathcal{D}_G(g, r) = \big(\mathbf{A}_1(r), \mathbf{A}_2(r)\big) \in T_e G \times T_e G, \tag{37}$$

i.e., is independent of the position when expressed in the body frame and is defined by the columns of the local connection. Selecting our coordinates to minimize the θ row of $\mathbf{A}$ thus aligns the constraint distribution closely with the $\mathbb{R}^2$ subspace, giving the system no opportunity to incur noncommutative effects by moving in the $SO(2)$ subspace; minimizing the x and y rows keeps the error generated by these excursions to a minimum.

This geometric argument can also be interpreted in terms of the relationship between the Lie bracket in (26) and the Christoffel symbols in the curvature of the connection, as discussed in Sect. 4.1. The Christoffel symbols correspond to the rotation of the horizontal space as the body frame rotates in response to shape changes; under a change of coordinates that minimizes body rotations, the information they contain is shifted into the exterior derivative.

5.4 The South-pointing chariot

The ability of the minimum-perturbation coordinates for the kinematic snake to *almost* eliminate $\mathbf{A}^\theta$ and thus *almost* align the constraint distribution with $\mathbb{R}^2$ raises the question as to whether it is ever possible to *completely* induce these effects. Returning to (33) suggests an answer to this question: if $-\mathbf{A}^\theta$ for a system is closed (i.e., conservative), then it can be completely nullified through a change of orientation β_θ for which $\mathbf{d}\beta_\theta = \mathbf{A}^\theta$.

In [11], we observed that $\mathbf{A}^\theta$ for the differential-drive car is constant (and therefore conservative) and calculated the change of coordinates that eliminates the θ row of the local connection. The new body frame we found coincides with the pointer on a *south-pointing chariot* [38]. This system is a two-wheeled vehicle topped by a horizontally-rotating pointer, synchronized to the wheels by a gear train. When the gear ratios are set correctly, (encoding the change of orientation function $\beta_\theta = \alpha_1 - \alpha_2$), the pointer exactly counter-rotates the body of the cart, thus maintaining its orientation with respect to the world as shown in Fig. 9. Chinese legend holds that such a device was used as a compass for coordinating military maneuvers prior to the discovery of magnetic needles, though dead-reckoning error would make this impractical and it is more likely that that the historical examples discovered were mechanical or mathematical curiosities [38].

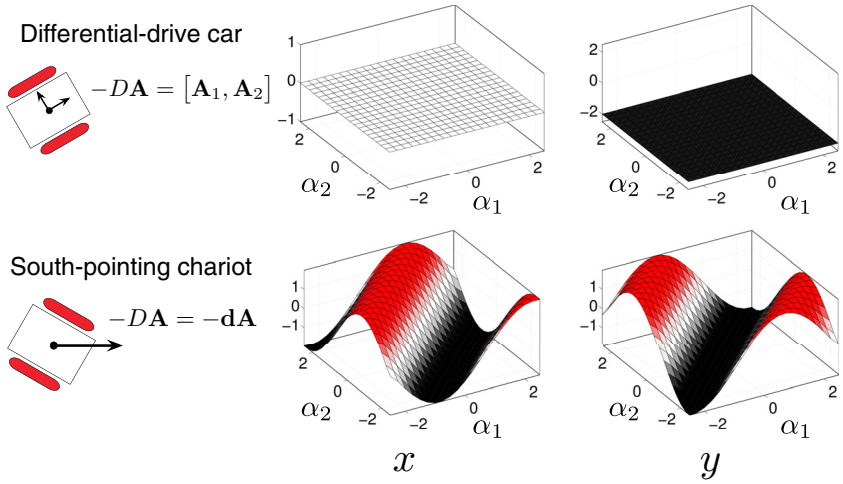

Fig. 10. The constraint curvature $D\mathbf{A}$ for the differential-drive car, shown for both the original coordinates (top row) and south-pointing chariot coordinates (bottom row), and broken into x and y components (left and right). In the original coordinates, the constraint curvature appears as a constant local Lie bracket whose only non-zero component is $[\mathbf{A}_1, \mathbf{A}_2]^y = -2$. In the south-pointing chariot coordinates, $D\mathbf{A}$ is equal to a shape-varying exterior derivative function that is non-zero in both x and y components, but has a constant magnitude equal to the magnitude of the Lie bracket from the original representation, $\|\mathbf{d}\mathbf{A}^{xy}\| = 2$.

Considering the differential-drive vehicle in the south-pointing chariot frame highlights some interesting properties of the system. First, it emphasizes the existence of a *holonomic* constraint between the system orientation and the wheel angles: for any given initial combination of θ, α_1, and α_2, the orientation is a function of the wheel angles. Mathematically, this divides the five-dimensional configuration space of the vehicle into a set of non-intersecting four-dimensional subspaces (known as *foliations*), each corresponding to a different orientation when $\alpha_1 = \alpha_2 = 0$. In the standard coordinates with the body frame locked to the vehicle chassis, these foliations are hard to visualize, but in the south-pointing chariot coordinates, each foliation corresponds to the full shape space combined with a single $\mathbb{R}^2$ slice of the $SE(2)$ position space.

A second effect of working in the south-pointing chariot coordinates is that they emphasize the benefit of encoding the system nonholonomy as nonconservativity rather than as noncommutativity. In the original coordinates, the exterior derivative of the constraints is zero and the constraint curvature reduces to the local Lie bracket, $-D\mathbf{A} = [\mathbf{A}_1, \mathbf{A}_2]$, while in the new coordinates the local Lie bracket is zero and the constraint curvature is equal to the exterior derivative, $D\mathbf{A} = \mathbf{d}\mathbf{A}$. These two representations are plotted in Fig. 10. The magnitude of the constraint curvature is constant and equal across the two parameterizations with $\|D\mathbf{A}^{xy}\| = 2$, confirming that the curvature at each point is a fundamental property of the constraints, rather than of the parameterization. The area integrals of these two representations are, however, markedly different. The lack of variation of $D\mathbf{A}$ in the original coordinates means that its area integral always grows monotonically with the enclosed area, whereas separately integrating the x and y components of $\mathbf{d}\mathbf{A}$ incorporates the effects of placing a gait in different regions of the shape space. Fundamentally, this reflects that the local Lie bracket is a *linearization* of the system's nonholonomy with respect to the position space, while the exterior derivative is an *exact* representation of the nonholonomy with respect to the shape space. Choosing coordinates that force the nonholonomy into the exterior derivative thus provides a more exact representation of the system dynamics.

6 Conclusions

Separating nonholonomic effects into their nonconservative and noncommutative components highlights the role of coordinate choice in geometric analysis of locomoting systems. The nonholonomy is characterized by the curvature of the constraints, $D\mathbf{A}$. As befits a differential-geometric construct, this curvature is independent of the coordinates in which it is expressed, but its components $\mathbf{dA}$, representing the nonconservativity of the constraints over the shape space, and $[\mathbf{A}_1, \mathbf{A}_2]$, representing their noncommutativity with respect to the position space, are strongly coordinate-dependent. Choosing coordinates that keep the system in a relatively commutative region of the position space concentrates the curvature in the conservative component; as nonconservative effects can be added linearly, this allows for the generalization of differential-geometric results to macroscopic motions in a manner unavailable when the nonholonomy is heavily noncommutative.

More broadly, the coordinate optimization illustrates limits on the differential geometric notion of coordinate invariance. Conventional wisdom in this field holds that the parameterization selected for a system will at most affect the convenience and tractability of its equations of motion, and that the results of calculations in one parameterization can be transferred to another by a simple change of coordinates. In contrast, we have examined two ways in which the coordinates play a more significant role. First, the accuracy of approximations, such as the truncated series in (28) can be coordinate-dependent, even when the exact computations are coordinate-invariant. Second, good coordinate choices can reveal symmetries or near symmetries in the system dynamics that are not apparent in arbitrary parameterizations, such as the collapse of the systems' gait displacement loci.

Looking ahead, we aim to extend both the fundamental theory at hand and its applications. As a first step towards both of these goals, we have begun combining the constraint curvatures described here with distance metrics on the shape space to evaluate gait efficiencies as well as displacements [37]. We are also considering the the nature of optimal coordinates for three-dimensional translation and rotation, and for systems in new physical regimes, such as the sand-swimmer in [30].

With specific regards to three-dimensional translation and rotation, the general formulas we have provided here apply to fiber bundles built on Lie groups, and so transfer directly to systems on $SE(3)$. Most importantly for the extension of our results, these formulas include approximately integrating the reconstruction equation by taking an area integral of $D\mathbf{A}$ as in (28) and the fundamentally bilinear relationship between the norm of the connection and the Lie bracket which means that finding a minimum-perturbation body frame drives the system towards commutativity. To fully bring these results into three-dimensional motion, we chiefly need to introduce an appropriate objective function of the form in (34) for three-dimensional rotation, which we have derived for small angles in [39] and believe we can build on further.

Appendices: Concepts in differential geometry

The discussion in Sects. 4–5 rests on a body of differential geometry and geometric mechanics. In addition to the works cited in Sect. 2, several textbooks [7,40] provide an extensive overview of this subject. As an aid to the reader, we have collected a brief survey of the aspects of this material that bear most strongly on our results. This survey emphasizes general intuition over formal presentation; in particular we assume that the mathematical structures are defined over smooth, orientable spaces.

A Spaces and groups

The dynamics of our systems are heavily influenced by the structures of their configuration spaces, which take the form of *manifolds* and have *group structure*. Key notions for describing these spaces are *Lie groups*, *semi-direct product groups*, and *fiber bundles*:

A.1 Manifolds and tangent spaces

Manifolds are differentiable, locally Euclidean spaces that may have more complicated global structure. For example, a circle is locally like a line, but connects back to itself. Each point q in a manifold Q has an associated *tangent space*, T_qQ, which contains vectors originating at q, such as configuration velocities $\dot{q}$. Continuing with the circle analogy, the velocity of a point moving around a circle is a vector originating at the current configuration and lying tangent to the circle; the set of all such vectors forms a line tangent to the circle. A manifold's *tangent bundle* TQ is the collection of its tangent spaces, and a vector field is the assignment of a $\dot{q}$ vector to each T_qQ in TQ (or a subset of it corresponding to a region of Q).

A.2 Lie groups

A *Lie group* G is a manifold with group structure–i.e., in which each point has a dual interpretation as a transformation that applies to other points in the manifold. For example, the positive real numbers form a one-dimensional manifold, and multiplying any two numbers from this set produces a number that is also within the manifold. In standard notation, the group product of the two elements, $g, h \in G$ is $hg \in G$ and represents g transformed by h; if the group operation is noncommutative (e.g., matrix multiplication), this operation is specifically the *left* group action $L_h(g)$, and the *right* transformation of g by h is $R_h(g) = gh$.[10] The *identity element* of the group, $\mathbf{e} \in G$, leaves the element it acts on unchanged, $\mathbf{e}g = g = g\mathbf{e}$, and corresponds to the origin of the manifold.

As group actions, elements $h \in G$ also define *left* and *right lifted actions*, $T_gL_h = \partial(hg)/\partial g$ and $T_gR_h = \partial(gh)/\partial g$, which map vectors in T_gG to their equivalent vectors (encoding the same infinitesimal group action) in $T_{hg}G$ or $T_{gh}G$. For example, in the multiplicative group of real numbers the lifted actions scale velocities by the same factor as the group action, $T_gL_h = T_gR_h = h$; velocities on this group are therefore group-equivalent if they encode the same *proportional* rate of change in the system's configuration.

Two lifted actions that play a particularly important part of our analysis are $T_gL_{g^{-1}}$, which maps vectors into $T_{\mathbf{e}}G$, the tangent space at the identity/origin, and $T_eL_g = (T_gL_{g^{-1}})^{-1}$, which maps elements of $T_{\mathbf{e}}G$ into T_gG. The tangent space at the identity element, $T_{\mathbf{e}}G$, has a privileged role in Lie group analysis, and corresponds to $\mathfrak{g}$, the *Lie algebra* of the group. In this paper, we use elements of the Lie algebra in two ways. First, mapping vectors from different tangent spaces into a common space ($T_{\mathbf{e}}G$) allows them to be added and subtracted as a well-defined operation. Second, vectors $z \in \mathfrak{g}$ can be mapped to elements of G via the *exponential map* [7], which flows the system with velocity T_eL_gz for one unit of time.

A.3 Lifted actions on $SE(2)$

In Sect. 3.1.2, we make use of lifted actions on *SE*(2) to relate world, body, and spatial velocities. As an aid to the reader, we here include the specific calculations for their expressions:

[10] Note that hg and gh are also respectively the right and left actions of g on h.

Using the notation that group elements $g, h \in G$ have coordinate expressions $g = (x, y, \theta)$ and $h = (u, v, \beta)$, they compose as

$$hg = L_h(g) = \begin{pmatrix} x\cos\beta - y\sin\beta + u \\ x\sin\beta + y\cos\beta + v \\ \theta + \beta \end{pmatrix} \tag{A.1}$$

and

$$gh = R_h(g) = \begin{pmatrix} x + u\cos\theta - v\sin\theta \\ y + u\sin\theta + v\cos\theta \\ \theta + \beta \end{pmatrix}. \tag{A.2}$$

The lifted actions are the Jacobians of these functions, found by taking the derivatives of the expressions in (A.1) and (A.2) with respect to the coordinates of g,

$$T_g L_h = \frac{\partial(hg)}{\partial g} \tag{A.3}$$

$$= \begin{bmatrix} \frac{\partial(x\cos\beta - y\sin\beta + u)}{\partial x} & \frac{\partial(x\cos\beta - y\sin\beta + u)}{\partial y} & \frac{\partial(x\cos\beta - y\sin\beta + u)}{\partial\theta} \\ \frac{\partial(x\sin\beta + y\cos\beta + v)}{\partial x} & \frac{\partial(x\sin\beta + y\cos\beta + v)}{\partial y} & \frac{\partial(x\sin\beta + y\cos\beta + v)}{\partial\theta} \\ \frac{\partial(\theta+\beta)}{\partial x} & \frac{\partial(\theta+\beta)}{\partial y} & \frac{\partial(\theta+\beta)}{\partial\theta} \end{bmatrix} \tag{A.4}$$

$$= \begin{bmatrix} \cos\beta & -\sin\beta & 0 \\ \sin\beta & \cos\beta & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{A.5}$$

and

$$T_g R_h = \frac{\partial(gh)}{\partial g} \tag{A.6}$$

$$= \begin{bmatrix} \frac{\partial(x + u\cos\theta - v\sin\theta)}{\partial x} & \frac{\partial(x + u\cos\theta - v\sin\theta)}{\partial y} & \frac{\partial(x + u\cos\theta - v\sin\theta)}{\partial\theta} \\ \frac{\partial(y + u\sin\theta + v\cos\theta)}{\partial x} & \frac{\partial(y + u\sin\theta + v\cos\theta)}{\partial y} & \frac{\partial(y + u\sin\theta + v\cos\theta)}{\partial\theta} \\ \frac{\partial(\theta+\beta)}{\partial x} & \frac{\partial(\theta+\beta)}{\partial y} & \frac{\partial(\theta+\beta)}{\partial\theta} \end{bmatrix} \tag{A.7}$$

$$= \begin{bmatrix} 1 & 0 & -(u\sin\theta + v\cos\theta) \\ 0 & 1 & u\cos\theta - v\sin\theta \\ 0 & 0 & 1 \end{bmatrix}. \tag{A.8}$$

The specific forms used in Sect. 3.1.2 can then be found by evaluating (A.5) and (A.8) for the appropriate values of g and h.

A.4 Direct and semi-direct product groups

Groups can be combined to form larger groups. Often, the combination of two groups A and B into a new group C is taken to mean the creation of a *direct product group* $C = A \times B$. This means that in the new group operation, components that started in A or B only affect other components that started as elements of the same group, i.e., $c_1 c_2 = (a_1 a_2, b_1 b_2)$. Direct products preserve properties such as being abelian (commutative)–if A or B has this property, then so does the corresponding section of C.

In a semi-direct product group [8], $D = A \circledS B$, elements of A act not only on each other, but also on elements of B. For example, the rotation/translation group that appears in our kinematics discussion in Sect. 3, $SE(2)$, takes the form $d_1 d_2 = (b_1(a_1 b_2), a_1 a_2)$, with rotations both composing with other rotations and acting on translations. A key aspect of such groups is that even though they do not possess the full orthogonality of a direct product group, the A components do preserve their original properties, and thus results that depend on these properties can be applied to the corresponding elements of D.

A.5 Fiber bundles

A fiber bundle Q is a space that decomposes locally into a *fiber space* F and *base space* B [7]. The fiber bundles we discuss in this paper are additionally both *trivial* and *principal*. "Trivial" means that the decomposition is the same everywhere in Q, such that we can take $Q = F \times B$; "principal" means that the fiber space F has Lie group structure, so we will use the notation $Q = G \times B$ to refer to such bundles.

B Covectors and differential forms

The example systems' equations of motion are described in terms of *covectors* and other *differential forms*. In our analysis, we apply *Stokes' theorem* and the *Hodge-Helmholtz decomposition* to these forms and fields to identify their key characteristics:

B.1 Covectors

In vector calculus, directional quantities such as the velocity of a particle or the gradient of a function are both consiered "vectors." Differential geometry distinguishes between *contravariant vectors* (generally referred to simply as "vectors"), whose magnitude scales inversely with coordinate changes, and *covariant vectors* (or *covectors*) whose magnitude scales proportionally with changes of coordinates. Within this distinction, velocities are vectors–increasing the length of a unit distance reduces the number of units traveled in a given time at given speed–and function gradients are covectors–increasing the unit distance increases the rate at which the function changes per unit traveled. Note that that not all covectors are explicitly function gradients; as we discuss below, they need only *locally* represent the rate of change of a quantity across a space, and that quantity may not be explicitly integrable into a function on the space.

Covectors are elements of *cotangent spaces* T_q^*Q that are *dual* to vectors' tangent spaces, and covector fields are defined over sections of the *cotangent bundle* T^*Q. Vectors and covectors at a point have a natural product, based on their magnitudes and relative alignment; in coordinates, the product $\langle \omega, v \rangle$ of a vector $v \in T_qQ$ and a covector $\omega \in T_q^*Q$ is the sum of the products of their paired indices, $\sum_i \omega_i v^i$ (i.e., their "dot product"). When the vector and covector represent velocity and gradient terms, their product is the *directional derivative*, measuring the rate of change along the vector of the function associated with the covector.[11] The vector-covector product is computationally equivalent to the vector dot product, and identified with it in treatments where the distinction between vectors and covectors is not made. In practice, vectors and covectors are often expressed as "column" and "row" vectors (in a linear algebraic sense), such that their product is the left matrix product of the covector acting on the vector.

B.2 Visualizing covector fields

Three basic visual representations are available for covector fields, as illustrated in Fig. B.1. *Plane fields* [41] depict the covectors as surfaces to which they are locally the gradient, with the vector-covector product the rate at which a vector climbs or descends its

[11] Some sources use a definition of "directional derivative" that normalizes the vector-covector product by the magnitude of the velocity. By removing speed from the calculation, this definition better matches the notion of "directional", but is less informative as an analytical tool and undefined for a vector with zero magnitude.

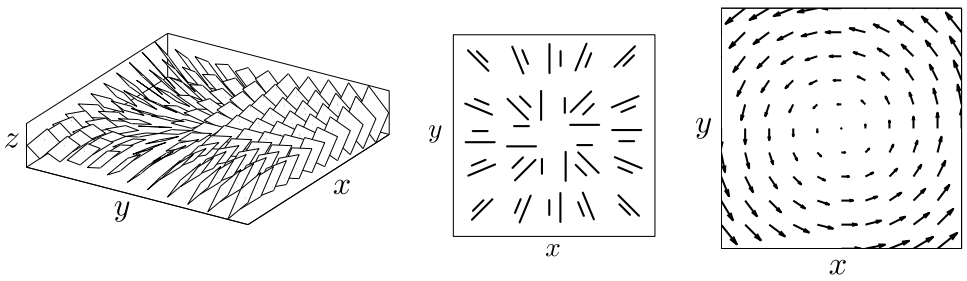

(a) Plane field (b) Local contours (c) Vector field

Fig. B.1. Three visual representations of the covector field $\omega(x, y) = (-y, x)$. All three renderings illustrate that clockwise vectors have a positive product with this covector field and radial vectors have a zero product with it, but this property is is most clearly discernible from the vector field representation.

corresponding surface. *Local contour* [42] plots project these surfaces onto the xy plane, rendering each plane via a pair of contour lines separated by a fixed z distance. "Vector field" plots [19] similarly project the surfaces onto the xy plane, but represent them via their gradient vectors. Note that plotting a covector field in this matter does not conflate vectors and covectors; rather, it makes use of their natural duality when represented in a shared basis [7], such as that of the page. In particular, it allows us to visually identify the vector-covector product with the vector dot product $\langle \omega, v \rangle = \|\omega\| \|v\| \cos\theta$, i.e., the product of the vector and covector magnitudes and alignment.

Many sources advocate the use of plane field or local contour representations, on the grounds that they emphasize the gradient-like nature of covectors and minimize confusion between the vector-covector product and the inner product of two vectors in a shared tangent space. We agree that these plots serve well as pedagogical tools for teaching differential geometry, but for examining system behavior we find that the vector field representations are more convenient. In particular, the vector representation conveniently illustrates the addition operation on one-forms, which plays a key part in our coordinate optimization approach. Further, the local contour representation has the disadvantage of requiring a pair of contour lines infinitely far apart to represent a covector of zero magnitude, whereas in our representation we would plot this as a zero magnitude vector.

B.3 Differential forms

Covector fields are an example of a broader class of differential geometric constructs, *differential forms*. In general, a differential k-form can be thought of as a position-dependent operator that linearly maps sets of k vectors in the tangent space at that point to an output. Unless otherwise specified, the output of a differential form is assumed to be a scalar value, but other classes of outputs that can be meaningfully added, such as vectors, are also possible. For the purposes of this paper, we can restrict our attention to the simplest of such structures: zero-forms, one-forms, and two-forms.

Zero-forms are functions whose outputs depend only on the input position. One-forms are covector fields; their product depends on both position in the manifold and the orientation and magnitude of an input vector. Two-forms take in a pair of vectors representing a infinitesimal parallelogram and return a value based on its orientation and scale (i.e., the projections of the parallelogram onto each pair of basis vectors). It is often useful to think of differential forms in terms of their integrals. For instance, one-forms have *line integrals* along one-dimensional paths that sum their products with the paths' tangent vectors. Similarly, two-forms have two-dimensional *surface integrals* that intuitively correspond to "fluxes" through the surfaces.

The *exterior derivative* $\mathbf{d}\omega$ of a k-form ω measures how components of ω change in directions normal to themselves, and is itself a $(k+1)$-form. For the examples we consider in this paper, which are restricted to k-forms over two-dimensional spaces, the exterior derivative is intuitively linked to two concepts in vector calculus–gradient and curl. As discussed above, the gradient of a function (zero-form) is a covector field (one-form). Similarly, the curl of a vector field (dual to a one-form) corresponds to the exterior derivative of a one-form. In vector calculus the curl is treated as a vector field in its own right (defined in terms of components along basis vectors), rather than as a structure with components along pairs of basis vectors; this is a convenient artifact of working in three dimensions, where each pair of basis vectors is uniquely defined by their shared normal vector.

Two important classifications of differential forms are *closed* and *exact*. Closed forms have null exterior derivatives, $\mathbf{d}\omega = 0$, and exact forms are themselves the exterior derivatives of other forms, $\omega = \mathbf{d}\Omega$. An especially useful result from differential geometry [7] is that exact forms are automatically closed, $\mathbf{d}(\mathbf{d}\omega) = 0$, which, in vector calculus terms is equivalent to saying that the curl of a gradient field is always zero (and thus that gradient fields are conservative). We will return to this point in our discussion of optimal coordinates.

B.4 Stokes' theorem

Stokes' theorem [43] equates the integral of a differential k-form ω along a closed k-dimensional manifold $\partial\Omega$ embedded in a space U to the integral of the exterior derivative of that k-form over a $(k+1)$-dimensional manifold Ω bounded by the original manifold,

$$\int_{\partial\Omega} \omega(u)\, du = \int_{\Omega} \mathbf{d}\omega(u)\, du. \tag{B.1}$$

In two dimensions and with ω a one-form, the surface Ω is simply the region of U enclosed by the closed curve $\partial\Omega$, and (B.1) specializes to the simple area integral

$$\oint_{\partial\Omega} \omega(u)\, du = \iint_{\Omega} \left(\frac{\partial \omega_2}{\partial u^1} - \frac{\partial \omega_1}{\partial u^2} \right) du^1\, du^2. \tag{B.2}$$

This equation is commonly encountered as "Green's theorem" in vector calculus, where it is described as equating a line integral along a closed loop on a vector field with an area integral of the field's curl.

B.5 Hodge-Helmholtz decomposition

The fundamental theorem of calculus states that for a one-dimensional function, if we know dy/dx as a function of x, then we know $y(x)$ up to a constant of integration. Expanding this principle to higher dimensions leads to the principle that if we know the gradient of a function, we know the function up to the addition of a scalar, and then to the fundamental theorem of vector calculus (also known as Helmholtz's theorem), which states that if we know the curl of a vector field, we know the vector field up to the addition of a gradient (conservative) field. In terms of differential forms, these statements are all aspects of a more general theorem that states that if we know the exterior derivative $\mathbf{d}\omega$ of a form ω, then we know ω up to the addition of a closed form.

In single-variable calculus, many functions have canonical anti-derivatives, such as $\int \cos = \sin$, to which the constant of integration is added. Similarly, it is often useful to designate a canonical form Ω as the anti-exterior-derivative of a form ω. This separation is achieved by the *Hodge decomposition* [44], referred to as the *Helmholtz decomposition* in

vector calculus contexts. The Hodge-Helmholtz decomposition splits a k-form ω into three components,

$$\omega = \mathbf{d}\mathcal{A} + \delta\mathcal{B} + \mathcal{C}, \tag{B.3}$$

where $\mathbf{d}\mathcal{A}$ is an exact, and therefore closed, k-form (corresponding to a curl-free vector field), $\delta\mathcal{B}$ is a k-form orthogonal to the set of closed k-forms (a divergence-free field in vector calculus), and $\mathcal{C}$ is a *harmonic remainder* (satisfying boundary conditions), generally much smaller than either of the first two components.

Because the exterior derivative is a linear operator and the double-exterior derivative is always zero, a one-form's exterior derivative independent of $\mathcal{A}$,

$$\mathbf{d}\omega = \mathbf{d}(\mathbf{d}\mathcal{A} + \delta\mathcal{B} + \mathcal{C}) = \mathbf{d}(\delta\mathcal{B} + \mathcal{C}), \tag{B.4}$$

with $\mathbf{d}(\mathbf{d}\mathcal{A}) = 0$. In this context we can consider $\delta\mathcal{B}+\mathcal{C}$ as the anti-exterior-derivative of $\mathbf{d}\omega$, and the set of closed forms $\mathbf{d}\mathcal{A}$ as the set of constants of integration.

The $\mathbf{d}\mathcal{A}$ term can also be considered as the projection of ω onto the space of conservative k-forms. For one-forms, the sum $\delta\mathcal{B} + \mathcal{C}$ therefore has the interesting property of being the *smallest* one-form with an exterior derivative equal to that of ω,

$$\underset{\mathbf{d}\Omega=\mathbf{d}\omega}{\operatorname{argmin}}\left(\int \|\Omega\|^2\right) = \omega - \mathbf{d}\mathcal{A} = \delta\mathcal{B} + \mathcal{C}, \tag{B.5}$$

where the "size" of Ω is its squared norm (steepness) integrated over the domain under consideration. In [10,11], we demonstrated a useful correspondence between this sum (which, due to the relatively small contribution of $\mathcal{C}$, we refer to as the "divergence free component" for brevity) and an optimal choice of coordinates for locomotion analysis, and presented a numerical algorithm for finding it over a finite domain, based on that in [45]. In Sect. 5, we examine this application more closely, and from a differential geometric standpoint.

B.6 Covariant derivatives

In spaces defined by Cartesian coordinates, the directional derivatives of vector fields and differential forms are taken as the individual directional derivatives of their components. This definition inherently makes use of the "parallel" nature of the basis vectors at different points in these spaces: the difference between two nearby vectors is the difference in their components. The notion of a directional derivative does not, however, extend cleanly to non-Cartesian spaces, such as the plane in polar coordinates or the special Euclidean group we consider below. In these spaces the basis vectors are position-dependent, and the change in these bases affects the component-wise differentiation of a vector field or form.

Covariant derivatives generalize directional derivatives to such spaces in a way that captures the changes in the basis vectors. When represented in coordinates, they augment the directional derivative with a set of *Christoffel symbols* that act as differential operators on the vector field or form and encode the rotation and dilation of the basis vectors. These symbols obey transformation laws that ensure coordinate-independence of the covariant derivative–if the coordinates change such that a given vector field or form's directional derivative shrinks, its Christoffel symbol differentiation grows proportionally, and vice versa. The covariant derivative can be combined with other forms of differentiation, such as the exterior derivative, to ensure that they correctly handle changes in the underlying basis vectors.

C Constraints

For the system we are investigating, the equations of motion are dictated by *holonomic and nonholonomic constraints*, which together generate a *constraint distribution* containing the

allowable directions in which the system can move. The constraints on the locomoting systems considered here take the form of *principal connections* on their configuration spaces. *Lie brackets* provide a means of examining how these constraints change as the system moves through its configuration space.

C.1 Holonomic and nonholonomic constraints

Differential geometry offers many categorizations for the constraints on a system's configurations, with the different categories corresponding to tools that can be applied to their respective systems. In this work, we make use of *holonomic* and *nonholonomic* constraints. Broadly speaking, holonomic constraints restrict a system to a lower-dimensional subspace of its configuration space (such as the end of a pendulum being constrained to a constant-radius arc around the pivot), whereas nonholonomic constraints only constrain the instantaneous velocity (a car's wheels prevent it from moving sideways, but a parallel-parking maneuver will let it move in a net lateral direction).

C.2 Constraint distributions

A constraint distribution $\mathcal{D}$ is a set of vector fields encoding the directions in which a system can move through its configuration space without violating its constraints. The distribution for an unconstrained system with an n-dimensional configuration space contains n independent vector fields; each constraint added to the system removes one of these vector fields. In general, $\mathcal{D}(q)$, the set of vectors in the distribution at point q, forms an independent basis at each q. A point where these fields do not form an independent basis is a type of singularity, in which the system has "extra" freedom to move in some directions, while losing the ability to move in others.

C.3 Principal connections

A *principal connection* $\mathcal{A}$ is a constraint distribution on a principal fiber bundle (Sect. A.5) that is both complementary to the space of pure fiber velocities and symmetric with respect to group actions on the fiber space [3]. The first property means that the distribution contains as many independent vector fields as there are dimensions in the base space, and that no vectors in these fields are aligned purely in the fiber direction. The second property means that the connection needs only to be defined for a single fiber value (such as the group identity $\mathbf{e}$), after which it can be generated at any other fiber value via the group's lifted actions (A.2).

Principal connections are encoded as vector-valued one-forms on the fiber bundle that output elements of $T_{\mathbf{e}}G$. The kernels (null-spaces) of these one-forms contain the velocities allowed by the system constraints,

$$\mathcal{D}_{\mathcal{A}}(q) = \{\dot{q} \in T_qQ \mid \mathcal{A}(q)\,\dot{q} = \mathbf{0}\}. \tag{C.1}$$

Because this distribution is complementary to the space of pure fiber velocities, vectors in $\mathcal{D}_{\mathcal{A}}$ are fully specified by their base-space components: at a configuration $q = (g, r)$, for any base velocity $\dot{r} \in T_rB$ there is a unique fiber velocity $\dot{g} \in T_gG$ for which the total velocity $\dot{q} = (\dot{g}, \dot{r})$ satisfies the null-space condition in (C.1) [3]. This attribute allows the connection to be expressed in a *local form* as a one-form $\mathbf{A}$ that maps from base velocities to their associated fiber velocities.

The symmetry in a principal connection means that their local forms can be further reduced to depend on only the base variables. For a system whose constraints are symmetric with respect to left group actions, the relationship between base and fiber velocities is

$$T_gL_{g^{-1}}\dot{g} = -\mathbf{A}(r)\dot{r}, \tag{C.2}$$

where the left-hand side is the left group-equivalent velocity to $\dot{g}$ in $T_{\mathbf{e}}G$, and the negative sign at right is a longstanding convention in the study of connections [3].

C.4 Lie bracket of vector fields

The *Lie bracket* [31] of two vector fields X and Y measures their rates of change with respect to each other. For fields defined on an n-dimensional space U, the Lie bracket is a vector field defined as

$$[X, Y] = (\nabla Y \cdot X) - (\nabla X \cdot Y), \tag{C.3}$$

or componentwise as

$$[X, Y]^i = \sum_{j=1}^{n} \left(x^j \frac{\partial y^i}{\partial u^j} - y^j \frac{\partial x^i}{\partial u^j} \right) = -[Y, X]^i. \tag{C.4}$$

There are several applications of the Lie bracket in the context of control systems, but the interpretation most germane to the present work is that flowing infinitesimally along X, Y, $-X$, and $-Y$ is equivalent to flowing infinitesimally along $[X, Y]$.

A special definition of the Lie bracket applies to elements of the Lie algebra $\mathfrak{g}$ corresponding to a Lie group G. In this case, the Lie bracket of two vectors $a, b \in \mathfrak{g}$ is the Lie bracket of the two *left-invariant* vector fields on G generated by applying the left lifted action to a and b,

$$[a, b] \equiv [T_e L_g a, T_e L_g b]\big|_{g=e}. \tag{C.5}$$

Note that this class of Lie bracket is a vector-valued two-form: the derivatives from (C.4) are fully determined here by the group structure, so the bracket takes in a pair of vectors and bilinearly maps them to an output vector.

D Derivation of the local curvature

The expression for the local curvature of the connection in (26) was presented in [8], based on earlier work in [14]. To the best of our knowledge, however, the relationship between this curvature and the Lie bracket in (24) has not been explicitly shown in modern notation. Here, we offer a derivation linking the two expressions, i.e., showing that

$$\left[\begin{pmatrix} \dot{r}_1 \\ -T_e L_g \mathbf{A}(r)\dot{r}_1 \end{pmatrix}, \begin{pmatrix} \dot{r}_2 \\ -T_e L_g \mathbf{A}(r)\dot{r}_2 \end{pmatrix} \right]\Bigg|_{q_0} = \begin{pmatrix} \mathbf{0} \\ \left(-\mathbf{dA} + [\mathbf{A}_1, \mathbf{A}_2]\right)(r_0) \end{pmatrix}. \tag{D.1}$$

First, a general Lie bracket in which the vector components can be grouped into two subvectors, i.e., of the form

$$[\dot{q}_1, \dot{q}_2] = \left[\begin{pmatrix} \dot{r}_1 \\ \dot{g}_1 \end{pmatrix}, \begin{pmatrix} \dot{r}_2 \\ \dot{g}_2 \end{pmatrix} \right], \tag{D.2}$$

can be evaluated by applying the Lie bracket formula in (C.4) blockwise to (D.2), producing the expression

$$\left[\begin{pmatrix} \dot{r}_1 \\ \dot{g}_1 \end{pmatrix}, \begin{pmatrix} \dot{r}_2 \\ \dot{g}_2 \end{pmatrix} \right]\Bigg|_{(r_0, g_0)} = \begin{pmatrix} \left(\dfrac{\partial \dot{r}_2}{\partial r}\dot{r}_1 - \dfrac{\partial \dot{r}_1}{\partial r}\dot{r}_2 \right) + \left(\dfrac{\partial \dot{r}_2}{\partial g}\dot{g}_1 - \dfrac{\partial \dot{r}_1}{\partial g}\dot{g}_2 \right) \\ \left(\dfrac{\partial \dot{g}_2}{\partial r}\dot{r}_1 - \dfrac{\partial \dot{g}_1}{\partial r}\dot{r}_2 \right) + \left(\dfrac{\partial \dot{g}_2}{\partial g}\dot{g}_1 - \dfrac{\partial \dot{g}_1}{\partial g}\dot{g}_2 \right) \end{pmatrix}\Bigg|_{(r_0, g_0)}. \tag{D.3}$$

Second, for the problem at hand we are considering a pair of vector fields in which the $\dot{r}$ vectors are each a unit vector aligned with the chosen basis, and so have components (using Kronecker delta notation),

$$\dot{r}_i^j = \delta_{ij}. \tag{D.4}$$

The position velocities in the vector fields are functions of the shape velocity,

$$\dot{g}_i = -T_{\mathbf{e}}L_g\mathbf{A}(r)\dot{r}_i = -T_{\mathbf{e}}L_g\mathbf{A}_i(r), \tag{D.5}$$

where the final equality is due to the ith shape velocity serving to select the ith column of the local connection.

Given these two definitions, the partial derivatives of the shape velocity fields are clearly zero,

$$\frac{\partial \dot{r}_i}{\partial r} = \mathbf{0} \qquad \text{and} \qquad \frac{\partial \dot{r}_i}{\partial g} = \mathbf{0}, \tag{D.6}$$

as they are constant fields. The position velocity fields' derivatives with respect to the shape fields can be found by applying the product rule for differentiation,

$$-\frac{\partial \dot{g}_i}{\partial r}\bigg|_{(r_0,g_0)} = \frac{\partial\, T_{\mathbf{e}}L_g\mathbf{A}_i(r)}{\partial r}\bigg|_{(r_0,g_0)} = \frac{\partial\, T_{\mathbf{e}}L_g}{\partial r}\mathbf{A}_i(r_0) + T_{\mathbf{e}}L_{g_0}\frac{\partial \mathbf{A}_i(r)}{\partial r}, \tag{D.7}$$

and noting that $T_{\mathbf{e}}L_g$ is independent of r, so

$$-\frac{\partial \dot{g}_i}{\partial r}\bigg|_{(r_0,g_0)} = T_{\mathbf{e}}L_{g_0}\frac{\partial \mathbf{A}_i(r)}{\partial r}. \tag{D.8}$$

Further, each derivative of $\dot{g}_i$ in the lower-left term of (D.3) is multiplied by the shape velocity $\dot{r}_j$ with $i \neq j$, selecting out the jth derivative and giving these terms the form

$$-\frac{\partial \dot{g}_i}{\partial r}\dot{r}_j\bigg|_{(r_0,g_0)} = -\frac{\partial \dot{g}_i}{\partial r^j}\bigg|_{(r_0,g_0)} = T_{\mathbf{e}}L_{g_0}\frac{\partial \mathbf{A}_i(r)}{\partial r^j}. \tag{D.9}$$

Differentiating the position velocities with respect to the position space follows the same pattern,

$$-\frac{\partial \dot{g}_i}{\partial g}\bigg|_{(r_0,g_0)} = \frac{\partial\, T_{\mathbf{e}}L_g\mathbf{A}_i(r)}{\partial g}\bigg|_{(r_0,g_0)} = \frac{\partial\, T_{\mathbf{e}}L_g}{\partial g}\mathbf{A}_i(r_0) + T_{\mathbf{e}}L_{g_0}\frac{\partial \mathbf{A}_i(r)}{\partial g} \tag{D.10}$$

$$= \frac{\partial\, T_{\mathbf{e}}L_g}{\partial g}\mathbf{A}_i(r_0). \tag{D.11}$$

Inserting the results of (D.6), (D.9), and (D.11) into (D.3) and factoring out a $T_{\mathbf{e}}L_{g_0}$ term from each subexpression in the bottom row reduces the Lie bracket to

$$\left[\begin{pmatrix}\dot{r}_1\\ \dot{g}_1\end{pmatrix}, \begin{pmatrix}\dot{r}_2\\ \dot{g}_2\end{pmatrix}\right]\bigg|_{(r_0,g_0)}$$

$$= \begin{pmatrix} \mathbf{0} \\ T_{\mathbf{e}}L_{g_0}\left(-\left(\frac{\partial \mathbf{A}_2(r)}{\partial r^1} - \frac{\partial \mathbf{A}_1(r)}{\partial r^2}\right) + \begin{pmatrix}\frac{\partial\, T_{\mathbf{e}}L_g\,\mathbf{A}_2(r_0)}{\partial g}T_{g_0}L_g\mathbf{A}_1(r_0) \\ -\frac{\partial\, T_{\mathbf{e}}L_g\,\mathbf{A}_1(r_0)}{\partial g}T_{g_0}L_g\mathbf{A}_2(r_0)\end{pmatrix}\right)\end{pmatrix}. \tag{D.12}$$

The left-hand term in this new expression is the exterior derivative of the local connection,

$$\left.\frac{\partial \mathbf{A}_2(r)}{\partial r^1} - \frac{\partial \mathbf{A}_1(r)}{\partial r^2}\right|_{r_0} = \mathbf{dA}(r_0). \tag{D.13}$$

Taking $g_0 = \mathbf{e}$ (placing the origin at the initial position of the system) eliminates the $T_{\mathbf{e}}L_{g_0}$ factor applied to the whole row and makes the right-hand term in the new expression

$$\left(\frac{\partial\, T_{\mathbf{e}}L_g\, \mathbf{A}_2(r_0)}{\partial g} T_{\mathbf{e}}L_g \mathbf{A}_1(r_0) - \frac{\partial\, T_{\mathbf{e}}L_g\, \mathbf{A}_1(r_0)}{\partial g} T_{\mathbf{e}}L_g \mathbf{A}_2(r_0)\right). \tag{D.14}$$

Using the definition of the Lie bracket in (C.3), this expression is clearly the bracket

$$[T_{\mathbf{e}}L_g \mathbf{A}_1(r_0), T_{\mathbf{e}}L_g \mathbf{A}_2(r_0)] = [\mathbf{A}_1, \mathbf{A}_2](r_0), \tag{D.15}$$

where the expression to the right makes use of the simplified notation for Lie brackets on left-invariant fields given in (C.5).

Finally, combining the expressions in (D.13) and (D.15) provides an equation matching that in (D.1):

$$\left[\begin{pmatrix} \dot{r}_1 \\ -T_e L_g \mathbf{A}(r)\dot{r}_1 \end{pmatrix}, \begin{pmatrix} \dot{r}_2 \\ -T_e L_g \mathbf{A}(r)\dot{r}_2 \end{pmatrix}\right]\Bigg|_{q_0} = \begin{pmatrix} \mathbf{0} \\ \left(-\mathbf{dA} + [\mathbf{A}_1, \mathbf{A}_2]\right)(r_0) \end{pmatrix}. \tag{D.16}$$

References

1. A. Shapere, F. Wilczek. J. Fluid Mech. **198**, 557 (1989)
2. P.S. Krishnaprasad, D.P. Tsakiris, G-snakes: Nonholonomic kinematic chains on lie groups. In *33rd IEEE Conference on Decision and Control* (Lake Buena: Vista, Florida, 1994)
3. S.D. Kelly, R.M. Murray, J. Robotic Syst. **12**, 417 (1995)
4. J.P. Ostrowski, J. Burdick, Int. J. Robotics Res. **17**, 683 (1998)
5. A.D. Lewis, IEEE Trans. Automatic Control. **45**, 1420 (2000)
6. P.S. Krishnaprasad, D.P. Tsakiris, Dynamical Syst. **16**, 347 (2001)
7. M. Anthony, Bloch, et al., *Nonholonomic Mechanics and Control* (Springer, 2003)
8. J.B. Melli, C.W. Rowley, D.S. Rufat, SIAM J. Appl. Dynamical Syst. **5**, 650 (2006)
9. R.L. Hatton, H. Choset, Approximating displacement with the body velocity integral. In *Proceedings of Robotics: Science and Systems V* (Seattle, WA USA, 2009)
10. R.L. Hatton, H. Choset, Optimizing coordinate choice for locomoting systems. In *Proc. IEEE Int. Conf. Robotics and Automation* (Anchorage, AK USA, 2010), p. 4493
11. R.L. Hatton, H. Choset, Int. J. Robotics Res. **30**, 988 (2011)
12. A. Shapere, F. Wilczek, Am. J. Phys. **57**, 514 (1989)
13. R. Abraham, J.E. Marsden, *Foundations of Mechanics* (Addison Wesley, 1985)
14. J.E. Marsden, R. Montgomery, T.S. Ratiu, Mem. Am. Math. Soc. **436** (1990)
15. J.E. Marsden, *Introduction to Mechanics and Symmetry* (Springer-Verlag, 1994)
16. J.P. Ostrowski, J.P. Desai, Vijay Kumar. Int. J. Robotics Res. **19**, 225 (2000)
17. F. Bullo, K.M. Lynch, IEEE Trans. Robotics Automation **17**, 402 (2001)
18. E.A. Shammas, H. Choset, A.A. Rizzi, The Int. J. Robotics Res. **26**, 1075 (2007)
19. R.L. Hatton, H. Choset, Proc. IEEE BioRobotics Conf. 451 (2008)
20. S.D. Kelly, The mechanics and control of driftless swimming (in press)
21. J.E. Avron, O. Raz, New J. Phys. **9** (2008)
22. K. McIsaac, J.P. Ostrowski, Robotics Automation 637 (2003)
23. K.A. Morgansen, B.I. Triplett, D.J. Klein, IEEE Trans. Robotics **23**, 1184 (2007)
24. R. Mukherjee, D.P. Anderson, IEEE Int. Conf. Robotics Automation, 802 (1993)
25. G.C. Walsh, S. Sastry, Robotics and Automation, IEEE Transactions **11**, 139 (1995)
26. G.C. Walsh, S. Sastry, In Proc. 33th Conf. Decision Control, 1190 (1991)

27. E.A. Shammas, K. Schmidt, H. Choset, In IEEE Int. Conf. Robotics Automation (2005)
28. E.A. Shammas, H. Choset, A.A. Rizzi, Int. J. Robotics Res. **26**, 1043 (2007)
29. R.L. Hatton, H. Choset. In *Proceedings of the ASME Dynamic Systems and Controls Conference (DSCC)* (Cambridge, Massachusetts, USA, 2010)
30. R.L. Hatton, H. Choset, Y. Ding, D.I. Goldman, Phys. Rev. Lett. **110**, 078101 (2013)
31. R.M. Murray, Z. Li, S.S. Sastry, A Mathematical Introduction to Robotic Manipulation (CRC Press, 1994)
32. R.M. Murray, S.S. Sastry, IEEE Trans. Automatic Control. **38**, 700 (1993)
33. R.L. Hatton, H. Choset, IEEE Trans. Robotics **29**, 615 (2013)
34. E.A. Shammas, Ph.D. thesis, Carnegie Mellon University, 2006
35. J.E. Radford, J.W. Burdick, In Proc. Int. Symposium on Mathematical Theory of Networks and Syst. (Padova, Italy, 1998)
36. W. Magnus, Comm. Pure Appl. Math. **VII**, 649 (1954)
37. R.L. Hatton, H. Choset, Proc. Robotics: Sci. Syst. VII (Los Angeles, CA USA, 2011)
38. M. Santander, Am. J. Phys., 782 (1992)
39. M.A. Travers, R.L. Hatton, H. Choset, Proc. Am. Controls Conf. (ACC) (2013)
40. F. Bullo, A.D. Lewis, *Geometric Control of Mechanical Systems: Modeling, Analysis, and Design for Simple Mechanical Control Systems* (Springer, 2004)
41. D. Bachman, *A Geometric Approach to Differential Forms* (Birkhäuser, 2006)
42. W.L. Burke, Appl. Differential Geometry (Cambridge University Press, 1985)
43. W.M. Boothby, *An Introduction to Differentiable Manifolds and Riemannian Geometry* (Academic Press, 1986)
44. G.B. Arfken. Mathematical Methods for Physicists, 6th edition (Elsevier, 2005)
45. Q. Guo, M.K. Mandal, M.Y. Li, Pattern Recognition Lett. **26**, 493 (2005)