Biological Cybernetics (2019) 113:93–104
https://doi.org/10.1007/s00422-018-0772-z

Biological Cybernetics

ORIGINAL ARTICLE

Locomotion: exploiting noise for state estimation

John Guckenheimer[1] · Aurya Javeed[2]

Received: 15 January 2018 / Accepted: 19 July 2018 / Published online: 28 July 2018

Abstract
Running, walking, flying and swimming are all processes in which animals produce propulsion by executing rhythmic motions of their bodies. Dynamical stability of the locomotion is hardly automatic: millions of older people are injured by falling each year. Stability frequently requires sensory feedback. We investigate how organisms obtain the information they use in maintaining their stability. Assessing stability of a periodic orbit of a dynamical system requires information about the dynamics of the system *off* the orbit. For locomotion driven by a periodic orbit, perturbations that "kick" the trajectory off the orbit must occur in order to observe convergence rates toward the orbit. We propose that organisms generate excitations in order to set the gains for stabilizing feedback. We hypothesize further that these excitations are stochastic but have heavy-tailed, non-Gaussian probability distributions. Compared to Gaussian distributions, we argue that these are more effective for estimating stability characteristics of the orbit. Finally, we propose experiments to test the efficacy of these ideas.

Keywords Locomotion · Stochastic dynamical system · Floquet multiplier

1 Introduction

Basic principles of locomotion control remain enigmatic. This paper investigates one such issue, namely how organisms obtain information needed to maintain their stability. We study systems in which locomotion results from body movements organized around a periodic orbit of a dynamical system. Two different types of models can be used to represent the periodic orbits:

1. physical models based upon reconstructing the system from components, and
2. "data-driven" models that parameterize the dynamics of a general system near a periodic orbit in terms of the orbit's Floquet multipliers.

For both types of models, the organism must rely upon sensory inputs to estimate the system parameters.

This paper discusses the latter type of model (i.e., data-driven models), paying more attention to mathematics than to biomechanics. We explore how an organism might learn parameters that characterize its Floquet multipliers by monitoring its own behavior. We model the locomotion of the organism as a trajectory of a dynamical system that couples the organism with its environment. As part of this coupled system, the organism must be capable of real time responses to unexpected stimuli that threaten its stability.

The nervous system acts as a controller for locomotion. It obtains sensory information from both the external environment and the internal state of the organism, and it issues motor commands that stimulate muscles to produce forces. The agility and adaptability of many animals in comparison with robots is astonishing, but it is subject to disruption by seemingly minor disorders. Analyzing organisms as if they were machines yields insight into what is special about organisms and suggests potential improvements of machines. Doing so also sheds light on limitations of organisms, thereby focusing attention on aspects of animal locomotion that we still do not understand. This paper examines mechanisms used by organisms for state and parameter estimation in locomotion. We build machines and robots with devices that measure quantities needed for feedback control, but organisms rely

Communicated by Benjamin Lindner.

This article belongs to the Special Issue on *Control Theory in Biology and Medicine*. It derived from a workshop at the Mathematical Biosciences Institute, Ohio State University, Columbus, OH, USA.

✉ John Guckenheimer
jmg16@cornell.edu

[1] Mathematics Department, Cornell University, Ithaca, NY 14853, USA

[2] Center for Applied Mathematics, Cornell University, Ithaca, NY 14853, USA

on evolved sensory capabilities *and* their own behaviors to acquire the information they need for controlling locomotion.

Our analysis looks at the system as a whole without dissecting the relative effects of organism and environment on locomotion. A key observation is that excitations generated by the organism that perturb its movement affect its parameter estimation capabilities: the neuromuscular system is not just a passive observer. Our primary objective is to discover excitation strategies that organisms might use to estimate Floquet multipliers in a timely manner, with enough accuracy to set the gains of a stabilizing controller.

We assume that excitations used to parameterize a stabilizing controller are stochastic. This is hardly controversial: many studies have measured fluctuations in steady state activities like balancing a stick on the fingertip (Cabrera and Milton 2004a) or maintaining balance while standing (Collins and Luca 1994). Their probability distributions have been observed to be non-Gaussian and possess heavy tails. We propose an important functional role in this paper for these probability distributions in enabling the organism to estimate its controller parameters quickly. The time for estimation is critical in locomotor systems in which active control is required to maintain the organism's posture—especially when the controller must adapt to changing environments.

We use human running as an example to make our discussion concrete. Maus et al. (2015) analyzed data consisting of ground reaction forces and 31 body marker locations of runners on a treadmill. Their goals were

1. to obtain predictive models for the motion of the runner's center of mass (COM) one or two steps ahead, and
2. to create reduced models of locomotion with stable periodic orbits similar to those estimated from the motion capture data.

The first steps in this investigation were to compute a phase function (Revzen and Guckenheimer 2008) for the motion and an estimate for the center of mass trajectory (COM). Next, the data were reduced to a mechanical model called a spring-loaded inverted pendulum (SLIP) that has been widely applied to running animals (Blickhan and Full 1993; Holmes et al. 2006). Autonomous simulations of this SLIP model were unstable: trajectories did not remain close to the periodic orbit. This suggested that the runners employed active control of their locomotion. Maus et al. (2015) derived linear feedback controllers of enlarged SLIP models that successfully stabilized the periodic orbits and gave predictive information about observed step to step COM motion. While the conceptual framework for their analysis seems clear, it did little to shed light on how the runners learned control parameters that stabilized their motion. We explore this issue from a mathematical perspective.

2 Periodic orbits

A periodic orbit γ of period T for the dynamical system

$$\dot{x} = f(x)$$

is a trajectory, $x(t)$, that satisfies $x(t+T) = x(t)$. If $\phi_t(x)$ is the flow map of the system (i.e., as t varies, $\phi_t(x)$ is the trajectory starting at x), then the *Floquet multipliers* are the eigenvalues of the Jacobian $D\phi_T$ evaluated at a point $p \in \gamma$. The Jacobian depends upon the choice of p, but the eigenvalues do not because the Jacobians at different points of γ are matrices related by similarity transformations. Points of γ constitute a curve of fixed points of ϕ_T, so 1 is a Floquet multiplier. If the remaining multipliers have magnitude smaller than 1, then γ is asymptotically stable and it has a neighborhood filled with trajectories that approach γ at exponential rates.

Here is a prototypical example of a periodic orbit. Consider the state space $R^{n-1} \times S^1$ with coordinates (x, θ) and the system of equations

$$\dot{x} = Ax$$
$$\dot{\theta} = \frac{2\pi}{T},$$

where A is an $(n-1) \times (n-1)$ dimensional matrix. The θ axis, $0 \times S^1$, is a periodic orbit of period T, and the flow map is $\phi_t(x, \theta) = (\exp(tA)x, t\theta)$. The Floquet multipliers are the eigenvalues of $\exp(TA)$ together with 1. One might hope that systems with periodic orbits have coordinate systems near the orbit where the equations take this form, but that is not the case. However, *Floquet theory* (Hartman 2002) describes a useful, but weaker, result: if $Y(t)$ is a fundamental solution matrix for the variational equations $\dot{y} = Df(\gamma(t))y$, then there is a constant matrix, R, and a $2T$-periodic matrix, $Z(t)$, so that $Y(t) = Z(t)\exp(tR)$. In our example, the eigenvalues of $\exp(TR)$ are the Floquet multipliers of γ.

The locomotion models we investigate have two levels similar to the "plant" and "controller" of control theory. The lower-level plant is a piecewise-smooth vector field representing the motion of the musculoskeletal system of an organism relative to its COM.[1] We are interested in rhythmic gaits which are periodic orbits of this vector field. They may be unstable, in which case the model is enlarged to include a nervous system controller as well as the effects of external forcing. Noise is sure to be present, regardless of whether the orbit is stable or unstable. In many systems, delays are

[1] The vector fields are discontinuous at impacts, and their dimension may change due to different numbers of ground contacts. A slightly larger model would also describe the motion of the COM, but the periodic orbit would be replaced by a trajectory that undergoes a fixed horizontal translation each step.

also important, but they are not considered here. In what follows, we assume that the larger model has a stable periodic orbit in the absence of noise. We seek to characterize its Floquet multipliers from observations of trajectories in the noisy system.

Figure 1 shows the 3D motion of a foot marker and the 2D projection of the COM onto the horizontal plane for one of the runners in the data analyzed by Maus et al. (2015). They assumed that stochastic differential equations (SDEs) with Brownian noise increments were a good model for the locomotion, but did not test that hypothesis. Instead, they sought parameters of controlled, closed-loop SDE models that had stable limit cycles in the deterministic limit of no noise. When noise is restored to the model, the deviation of trajectories away from the limit cycle is roughly proportional to the magnitude of the noise term in the SDE. In these circumstances, innovations of the noise are confounded with the return of trajectories to the limit cycle. We discuss below the problem of estimating the deterministic periodic orbit and its Floquet multipliers from noisy trajectories.

Note that an alternative model for the running data does not require a stable limit cycle. Instead, it confines trajectories to a tubular region by creating thresholds at the boundary of the region where impulsive signals redirect a trajectory about to leave the region back into the region's interior. The impulses are analogous to the impact of a solid particle hitting a wall. As an example, Guckenheimer (1995) describes a controller for an inverted double pendulum on a cart. The target equilibrium of this four-dimensional vector field has a two-dimensional unstable manifold, but only a single direction actuated by the controller. Linear feedback control is still possible, but it stabilizes the system only in a narrow strip of the phase space. Impulses at selected boundaries of the phase space keep the pendulum upright in a much larger region. As shown in that paper, this type of impulsive controller can be used in conjunction with linear feedback in order to greatly enlarge the stability region of the closed-loop system.

3 Parameter and state estimation

We seek to characterize dynamical systems from trajectory data. In particular, we want to construct models that extrapolate from previous observations to predict still unobserved behavior. In the applications to locomotion, one goal is to create feedback controllers that modify models with unstable periodic orbits so that they become stable. For this, the locations of periodic orbits and their Floquet multipliers are key data. How do we find them? We first recall methods that are used to find periodic orbits and multipliers when a vector field is defined by formulas and then discuss analysis of time series data.

Stable periodic orbits of deterministic systems can be located as the limit sets of trajectories in their basin of attraction: the trajectories eventually lie on the orbit within the numerical accuracy of the simulation. Boundary value solvers yield additional methods that have the advantage that they may converge more rapidly, including to some unstable periodic orbits, but it is also more difficult to find successful starting data. The Floquet multipliers are eigenvalues of the Jacobian $D_x\phi_T$ evaluated at points of the periodic orbit. These Jacobians can be computed with numerical integration of the variational equations $\dot{y} = Df(\gamma(t))y$, or with finite-difference calculations of trajectories from carefully chosen initial conditions. These methods for estimating multipliers require additional numerical integrations that may not be feasible in the setting of data-driven analysis. Also, noise in empirical time series is a factor that complicates analysis of both the location of periodic orbits and their Floquet multipliers.

Extending the concepts of periodic orbits and Floquet multipliers to stochastic systems is problematic. As is common in many applications, we assume that our models are stochastic perturbations of a deterministic vector field and seek methods that estimate the periodic orbits and Floquet multipliers of that deterministic system. There are limitations that have only recently been recognized in applying this strategy to SDEs of the form $\mathrm{d}x = f(x)\mathrm{d}t + \varepsilon g(x)\mathrm{d}W$, with $\mathrm{d}W$ a Wiener process. In particular, the variance of the Jacobian $D_x\phi_T$ computed from fixed length sample paths of the SDE does not tend to zero as $\varepsilon \to 0$. Concretely, when $\varepsilon g \ll 1$ is constant and the periodic orbit is stable, Javeed (2017) argues that—for simple multipliers—the variance of unbiased estimators is at least

$$\frac{1-\lambda^2}{n}\int_0^T \left[\phi(t)^2 \int_0^T \phi(\tilde{t}+t)^{-2}\mathrm{d}\tilde{t}\right]^{-1} \mathrm{d}t. \tag{1}$$

Here, λ is the Floquet multiplier, n is the number of times the SDE cycles around the periodic orbit, and $\phi(t)$ is the flow of the $\varepsilon = 0$ system linearized about the periodic orbit and restricted to the invariant subspace associated with λ. Indeed, there is no dependence on the noise amplitude in this expression, consistent with numerical simulations (Guckenheimer 2014). This result (1) is a special case of an *information inequality* (Lehmann and Casella 1998). Heuristically, we have that fluctuations away from the deterministic limit cycle are of order ε so that the Jacobian is no longer defined on this scale: finite differences constructed from trajectories of stochastic systems have errors of $\mathcal{O}(\varepsilon)$ in the numerator and denominator, producing a net error that is $\mathcal{O}(1)$.

Estimating the Jacobian from a single simulation of a deterministic system is also problematic. Numerical trajectories are subject to both truncation and round-off errors. The optimal accuracy for computing a fixed time-length trajec-

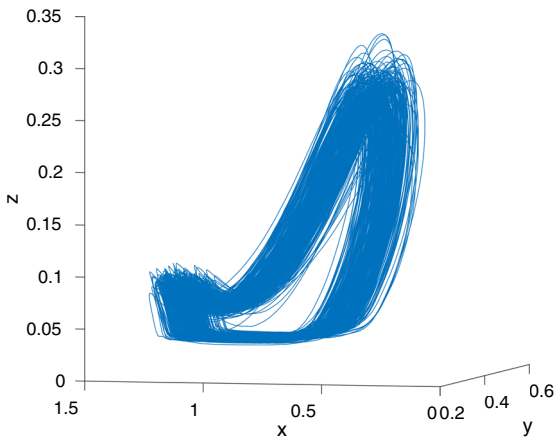

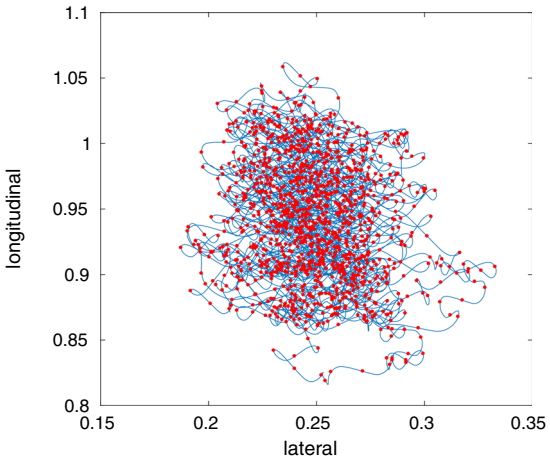

Fig. 1 (Left) The trajectory of a motion capture marker, in laboratory coordinates, placed on the right foot of an individual running on a treadmill for 30 min. The flat portions at the bottom of the trajectory occur during the stance phase of the right leg, as it slides with the treadmill belt. (Right) The center of mass trajectory projected onto the horizontal plane. The dots show the location at vertical local maxima and minima, i.e., intersections with the cross section $\dot{z} = 0$

tory segment is attained when there is a balance between these two sources of error. As the distance from a trajectory to the periodic orbit becomes smaller than this optimal accuracy, finite-difference calculations of the Jacobian and its Floquet multipliers become unreliable.

Can an organism engaged in a periodic motion estimate the Floquet multipliers of that motion? If the trajectory lies strictly on the periodic orbit, the answer is clearly no. Perturbations are required to observe the return to the limit cycle. As described above, inferring the multipliers from an SDE with Brownian increments requires trajectories that may be too long to be feasible for stabilizing locomotion. The hypothesis that we explore is that different stochastic processes might be more effective in quickly generating sufficient data for this purpose. Informally, we reason as follows. The most common finite-difference calculation of the Jacobian begins with an initial condition p_0 on the periodic orbit and additional initial conditions $p_j = p_0 + \delta e_j$ with e_j the canonical basis vectors and δ chosen to balance truncation and round-off errors in the computation of the flow map. Then, $\frac{1}{T}(\phi_T(p_j) - p_0)$ is an approximation for the Jacobian that determines the multipliers. This algorithm can be modified to use random initial conditions q_j near p_0 so that the differences $q_j - p_0$ form a basis of the tangent space. In this case, $\frac{1}{T}(\phi_T(q_j) - p_0)$ would be expressed in terms of this basis. With more initial conditions and trajectory segments, we obtain over-determined systems of equations whose least-squares solution may improve the Jacobian estimate. We can implement this calculation in a single trajectory of a stochastic jump process with finite-size innovations that are small but larger than the variance of noise from other sources. When the jumps occur at times separated by intervals longer than $2T$, the trajectory segment in the intervening time approximates the deterministic flow map for long enough to estimate returns. With enough innovations to span the tangent space, we have data like those from the finite-difference algorithm with random initial conditions and can form Jacobian estimates of comparable accuracy.

We do not think that the algorithm in the last paragraph is a likely contender for how the locomotion controllers of organisms actually work. However, it illustrates the feasibility of obtaining accurate Jacobian estimates from a single trajectory on a shorter timescale than is possible from an SDE perturbation with Brownian increments. Since an organism needs to find stabilizing gains before it falls over or crashes, we do think that it is important to explore the trade-offs in time and accuracy for estimating multipliers with different types of stochastic processes.

4 Examples and numerical experiments

As a numerical example, we consider the van der Pol system

$$\begin{aligned} \dot{x} &= \mu\left(y - \frac{x^3}{3} + x\right) \\ \dot{y} &= \frac{1}{\mu}(a - x), \end{aligned} \tag{2}$$

with parameters $(\mu, a) = (1, 0.92)$. This planar system has a periodic orbit, γ, that is shown in Fig. 2. The Floquet multiplier of the orbit, $\lambda \approx 0.2173$, has modulus less than one, but the convergence of local trajectories is not monotonic: during each revolution, nearby solutions separate from the top

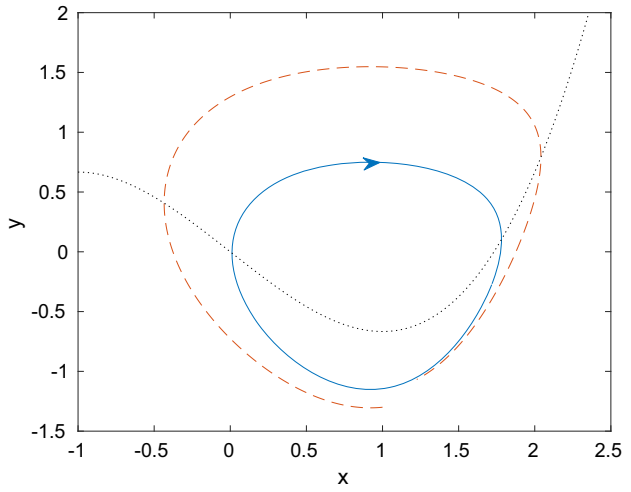

Fig. 2 The stable periodic orbit, γ, of the van der Pol system (2) is shown with a solid line. A nearby trajectory (dashed) and the $\dot{x} = 0$ nullcline (dotted) are also included

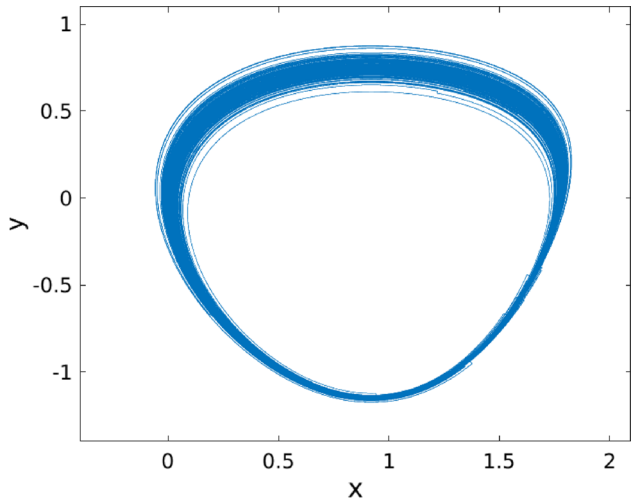

Fig. 3 A realization of the perturbed van der Pol system (approximately 10^3 cycles)

of the orbit before drawing closer to it as they cross beneath the nullcline $\dot{x} = 0$.

We perturb γ by a stochastic jump process like the one described in Sect. 3. Specifically, a small amount of Gaussian white noise is added to the equations of motion, yielding the SDE

$$\mathrm{d}x = f(x)\mathrm{d}t + g\mathrm{d}W. \tag{3}$$

Here, the coordinates of $W = W(t, \omega)$ are independent standard Brownian motions. This SDE is then perturbed by compound Poisson noise (meaning solutions are displaced by i.i.d. random variables at Poisson arrival times and subsequently evolved from the points at which they land). We consider displacements that are normally distributed, having zero mean and standard deviation $\sigma = 10^{-2}$. The amplitude of the Gaussian white noise is substantially smaller ($g = 10^{-4}$) so that the contributions of the Poisson noise are jumps whose size is on the order of the optimal finite-difference displacement: $\delta = \sqrt{g}$.

To simulate (3) during the time between jumps, we use a Wong–Zakai approximation whereby $W(\cdot, \omega)$ is replaced by a piecewise-linear interpolation of its values on the partition $\Pi = \{t_i = i\,\Delta t\}_{i \geq 0}$. Doing so reduces the SDE to an ODE, which we integrate with Runge–Kutta steps of size Δt. As the norm of Π tends to zero, Sussmann (1978) guarantees that solutions of the ODE converge uniformly to those of the SDE, almost surely.

Figure 3 shows a realization of the perturbed system that has been evolved until time $10^3 T$, where $T \approx 6.9755$ is the period of γ. The time step for the simulation is $\Delta t = 3 \times 10^{-3}$, and the intensity of the Poisson process that generates the jumps is $\nu = (3T)^{-1}$, implying an average of one jump every three cycles. Insofar as finite-difference calculations are concerned, we speculate that our value for ν is close to optimal: loosely speaking, $\lambda^{\frac{1}{\nu T}}\sigma \approx g$, so the jumps are spaced by the amount of time it takes the noise-free process to decay from a typical jump to the amplitude of the Gaussian noise.

Because of the noise, the flow map ϕ_T does not reliably map initial conditions through a cycle, so we estimate Floquet multipliers by considering returns to a cross section, Σ. Any local section that is transverse to the flow will do; we choose Σ to be the horizontal line segment intersecting the periodic orbit at $p_0 \approx (1.781628, 0.103452)$. Along this section, we introduce local coordinates that reverse the direction of increasing x and have p_0 as their origin (i.e., local coordinates that coincide with the Frenet–Serret frame of the orbit).

We explore the distribution of returns to Σ by simulating our stochastic system from zero to $10^4 T$—a time interval ten times longer than the one used in Fig. 3. A histogram of the approximately 10^4 returns is plotted in the left pane of Fig. 4. The right pane shows the distribution of returns when the jumps are removed, leaving only Gaussian white noise to perturb γ. For each set of returns, the Gaussian distribution that is most likely to have generated the data is superimposed as a dashed line. In the absence of jumps, we see that the Gaussian fits the return data well, but this is not the case when jumps are present. Since the jump amplitude $\sigma \gg g$, the jumps produce a heavy-tailed distribution of returns that cannot be reasonably explained as Gaussian. As shown by the left pane of the figure, a more appropriate class of models appears to be the stable distributions, which are a superset of the Gaussians and have elements with heavy tails. (We have checked numerically that 10^4 is enough returns for these findings to be robust to sampling error.)

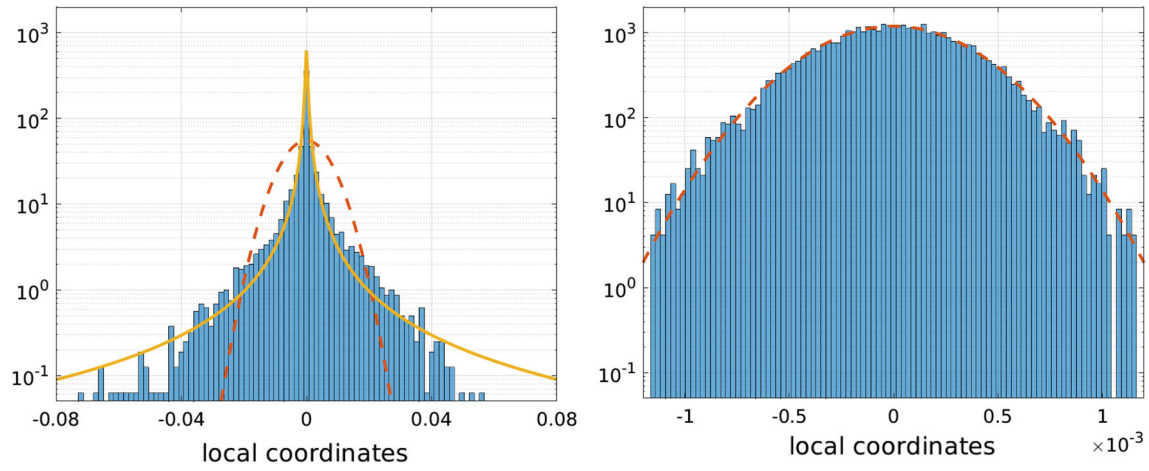

Fig. 4 The distribution of approximately 10^4 returns to Σ (left). The right frame shows the distribution of returns when jumps are removed from the stochastic process used to perturb γ. The dashed curves depict the most likely Gaussian density given the data, while the solid curve in the left frame is the most likely stable distribution (the shape parameter, α, describing the thickness of its tails is approximately 0.7)

At this point, it is natural to address whether the heavy-tailed return distribution generated by large jumps is preferable to a Gaussian density when estimating Floquet multipliers. Simulations (and intuition) suggest that the returns of a perturbed stable periodic orbit to a cross section are well-approximated by a stationary autoregressive process:

$$x_{i+1} = \lambda x_i + \varepsilon_{i+1}, \qquad |\lambda| < 1 \tag{4}$$

(if, of course, the stochastic perturbations to the dynamical system are small). Such processes are of fundamental importance in time series analysis, and there exist several theoretical results about the estimation of their parameters. For example,

- When the noise terms have finite second moments (which is technically the case for returns to Σ since periodic orbits have a finite size), the ordinary least-squares (OLS) estimator of λ has the same asymptotic properties as it does when the noise terms are Gaussian: it converges almost surely and is asymptotically normal with a limiting variance of $(1 - \lambda^2)$. The latter of these two facts implies that the estimator is $\sqrt{n}$-consistent, where n is the number of returns.
- If, on the other hand, the heavy tails of the noise terms are not truncated, the OLS estimator will still converge almost surely, and the rate of convergence will be faster than $\sqrt{n}$. The thicker the tails, the faster the convergence. This result is proved in Hannan and Kanter (1977), where they offer this heuristic explanation:

  "An isolated, very large, value of ε_i, which is more likely if the tails are thick, will result in relatively large x_i, and $x_{i+1} \approx \lambda x_i$. Such points, in the scatter diagram for x_{i+1} against x_i, will be almost on the line of slope λ through the origin and will be influential because of their large distance from the origin."

  (Here, their original notation has been altered to make it consistent with our own.)

For additional details and results, we recommend two standard references on time series analysis: Brockwell and Davis (1991) and Hamilton (1994).

In the spirit of the Hannan and Kanter (1977) excerpt—and using our aforementioned stochastic van der Pol system—we look at the scatter plot of return $i + 1$ versus return i for approximately 10^3 intersections with Σ. To simulate stationarity, we draw initial conditions from the best-fitting distributions in Fig. 4 (but it is the subsequent intersection, and not this initial condition, that is counted as the first return). As before, the left pane of Fig. 5 shows the returns of our stochastic system, while the right pane is the scatter plot when there are no jumps.

It is striking to see the collinearity of return pairs which do not have intervening jumps (see the circle-shaped markers in the left pane). As conjectured in Sect. 3, the jump process has generated a subset of data that conveys a great deal of information about the Floquet multiplier in question.

Consider the problem of estimating λ from these circle-shaped points, using OLS to fit the slope of the line that runs through them. If the number of returns is very large, there is an intuitive but simplistic lower bound for the variance of the estimated slope: Suppose, instead of $\nu = (3T)^{-1}$, a hypothetical jump intensity of $(2T)^{-1}$. This intensity prescribes an average of one jump every two cycles; hence, a coarse model for the number of jumps occurring between returns is the alternating sequence

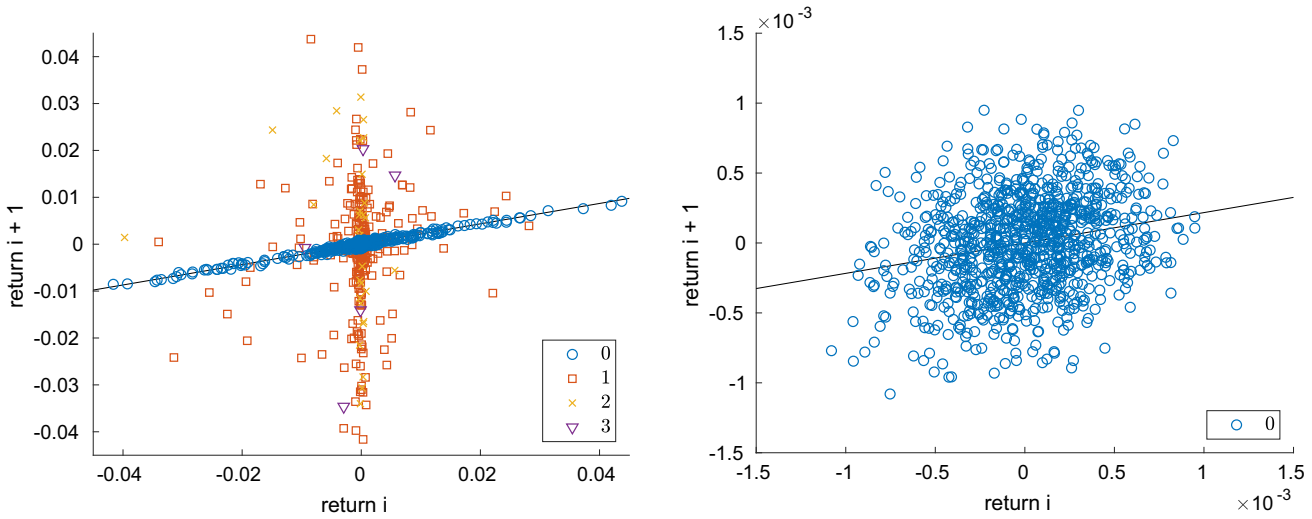

Fig. 5 A scatter plot of approximately 10^3 returns of the perturbed van der Pol system (left). The right frame recreates the scatter plot when the jumps are removed, leaving only Gaussian white noise to perturb γ. The type of marker indicates the number of intervening jumps, and the solid lines have slope equal to the Floquet multiplier, λ. Note the differences in scales of the two plots

$1, 0, 1, 0, \ldots$

In this pattern of jumps, the returns preceding the zero-jump cycles (i.e., the horizontal coordinate of the circle-shaped points) will be approximately i.i.d., implying the OLS estimator is unbiased and efficient, with a variance on the order of $g^2(n\sigma^2)^{-1}$. This variance is the smallest one would hope for because—in the other extreme where all of the returns are correlated as in (4)—the OLS estimate of the slope has the limiting variance

$$\frac{1-\lambda^2}{n} \gg \frac{(g/\sigma)^2}{n}.$$

For our jump intensity of $\nu = (3T)^{-1}$, numerical simulations approximate the standard deviation of the estimated slope to be 2×10^{-3} when $n = 10^3$ and OLS is applied to the subset of returns having zero intervening jumps. In comparison, $\sqrt{(g/\sigma)^2/n} \approx 3 \times 10^{-4}$ and $\sqrt{(1-\lambda^2)/n} \approx 3 \times 10^{-2}$.

Figure 6 shows the accuracy of OLS in the small-n regime relevant to locomotion. As indicated in the legend, three different strategies are presented:

1. perturbing γ by Gaussian white noise ($g = 10^{-4}$) and using OLS to estimate λ from n returns to Σ (an intercept is included in the fit since p_0, the location of γ along Σ, may not be known *a priori*);
2. estimating λ in the same way, but perturbing γ by our stochastic jump process (Gaussian white noise and jumps);
3. the estimation procedure discussed in the preceding two paragraphs; i.e., filtering the n returns in Strategy 2 so that a subset of pairs without intervening jumps remain, and then estimating λ with OLS.

If, after filtering a realization, too few returns remain to uniquely specify the fit, the third strategy uses the Strategy 2 estimate.

The metric we choose to quantify accuracy is the root mean square error (RMSE) of the Floquet exponent, i.e., the logarithm of the multiplier, divided by T. [2] Letting $\widetilde{\lambda}$ denote the OLS estimator of λ, constructed using either Strategy 1, 2, or 3, we have that $\mathrm{RMSE}(\log\widetilde{\lambda}) = \mathrm{RMSE}(\log(\widetilde{\lambda}/\lambda)) = ||\log\widetilde{\lambda} - \log\lambda||_{L^2}$. Thus

$$-\mathrm{RMSE}\left(\log\frac{\widetilde{\lambda}}{\lambda}\right) < \log\widetilde{\lambda} - \log\lambda < \mathrm{RMSE}\left(\log\frac{\widetilde{\lambda}}{\lambda}\right)$$

for "typical" values of $\widetilde{\lambda}$. Roughly speaking then, applying the exponential to this inequality implies that $\lambda/\widetilde{\lambda}$ is between $x = \exp\circ\mathrm{RMSE}(\log(\widetilde{\lambda}/\lambda))$ and x^{-1} "most" of the time. This quantity, x, is plotted in Fig. 6, which is the main result of this section.

Figure 6 demonstrates that, relative to Brownian motion, other classes of stochastic perturbations have the potential to yield far better stability estimates from substantially less

[2] The RMSE is the L^2 error over the probability space of realizations. Its square has the appealing property of being the variance plus the square bias of the estimator. We use the RMSE of Floquet exponent estimates, as opposed to estimates of the Floquet multiplier, because the former has better invariance properties. In particular, the value of the exponent is the same when computed from trajectories which make different numbers of circuits around a periodic orbit. To simulate the expectation, 10^3 realizations are used.

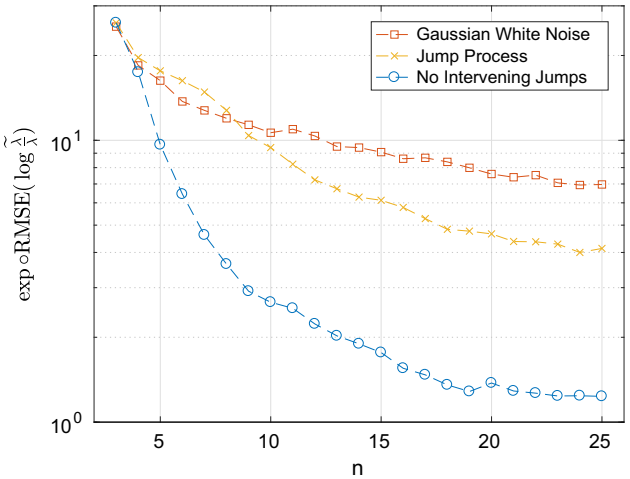

Fig. 6 The approximate fractional error $\tilde{\lambda}/\lambda$, where $\tilde{\lambda}$ is the OLS estimate constructed from n returns (or a subset of the n returns, in the case of the circle-shaped series)

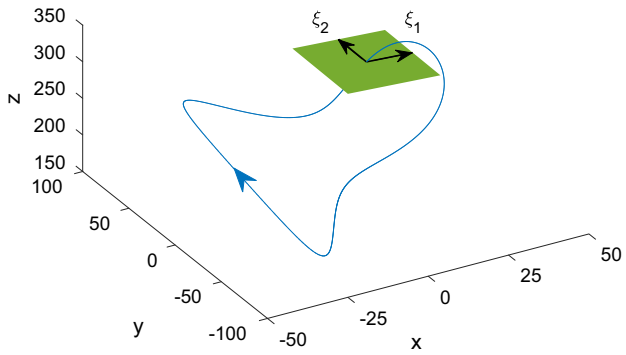

Fig. 7 A stable periodic orbit of the Lorenz system (5). The local cross section that we use is also included, along with its local coordinates, (ξ_1, ξ_2)

data. With just 13 returns or so, the figure shows that our simple jump process has the capacity to produce Floquet multiplier estimates that tend to be within a factor of 2 of the true multiplier estimate. When γ is perturbed only by Gaussian white noise, estimates of the multiplier are within a factor of 10. Though not shown in the figure, Strategy 1 estimates require approximately 130 returns to estimate the multiplier within a factor of 2.

4.1 Higher dimensions

To explore the generalizations of our numerical findings to higher-dimensional periodic orbits, we study the Lorenz system

$$\begin{aligned} \dot{x} &= \sigma(y - x) \\ \dot{y} &= rx - y - xz \\ \dot{z} &= -bz + xy, \end{aligned} \tag{5}$$

with $(\sigma, r, b) = (10, 240, 8/3)$. At these parameter values, the differential equations have the periodic orbit shown in Fig. 7. Henceforth, we refer to this orbit as γ. The Floquet multipliers of this new γ are $(\lambda_1, \lambda_2) \approx (-0.6161, -0.0026)$, and because $|\lambda_1/\lambda_2| \gg 1$, the local dynamics are stiff: nearby trajectories rapidly collapse to the invariant subspace of the slow multiplier, λ_1.

Guckenheimer (2014) analyzes OLS estimates of λ_1 and λ_2 when γ is perturbed by Gaussian white noise. We use a cross section that is essentially the same as the one in the reference, taking Σ to be a piece of the plane that intersects the periodic orbit at $p_0 \approx (29.418748, 69.200043, 239)$ and whose normal is $(\dot{x}, \dot{y}, \dot{z})|_{p_0}$ (see Fig. 7). Instead of perturbing γ only by Gaussian white noise, we use the same jump process we applied to the van der Pol system, except the process is now defined on $\mathbb{R}^3$ and has parameters

$$(\nu,\ \sigma,\ g) = \left((2T)^{-1},\ 3 \times 10^{-2},\ 3 \times 10^{-4}\right),$$

where $T \approx 0.4698$ is the period of the new γ.

Figure 8 shows 10^3 contiguous returns of the perturbed system to Σ. [3] The stiff dynamics confine the returns to the eigendirection of the slow multiplier—unless, of course, the solution is perturbed by a jump immediately prior to its return. Away from p_0 (the origin of the local coordinate frame), the returns lift away from the slow eigendirection by a slight amount. This is all but certain to be an artifact of the nonlinearity of the return map at large scales, as we later discuss in greater detail.

To characterize the distribution of returns, we again simulate a superset of about 10^4 contiguous returns. In this higher-dimensional setting, coordinates are then transformed to the principal components (PCs) of the return set. PC coordinates diagonalize the sample covariance, so the components of returns are approximately uncorrelated in this frame. Along each PC coordinate, the most likely Gaussian and stable distribution to have generated the data are computed, as before (Fig. 9). Our findings are consistent with the planar example: when γ is perturbed by the jump process (i.e., Gaussian white noise plus larger jumps at Poisson arrival times), the stable density is the better description of the heavy-tailed return distribution along each PC direction. Moreover, when the jumps are removed, the returns appear to be Gaussian (a fact also reported by Guckenheimer 2014).

[3] In this example and the last, the intersections of the numerical solution with the cross section of interest are approximated with cubic Hermite splines. Doing so is unnecessary in the van der Pol setting—a linear interpolant is sufficiently accurate. However, that is not the case for the perturbed Lorenz system studied here (time step $\Delta t = 2 \times 10^{-4}$).

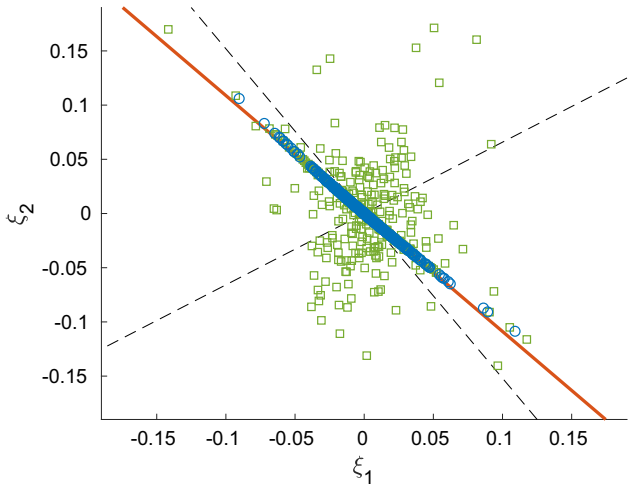

Fig. 8 A realization of approximately 10^3 returns to the Lorenz Σ. If the perturbed system has not experienced a jump since its previous return, the marker is circle-shaped. If there is at least one intervening jump, it is square-shaped. The eigendirection of λ_1 is shown with a solid line, and the dashed lines are the principal directions for a superset of 10^4 returns

Analogues of Fig. 6 for the slow and fast multipliers of the Lorenz system are presented as Fig. 10. To compute the multiplier estimates, we fit an affine return map (as before) and take the eigenvalues of the map's linear part.

Much like the van der Pol example, the learning rate of λ_1 is significantly faster for the circle-shaped strategy. This slow multiplier is significantly closer to one (i.e., the border between stability and instability) than the λ of the van der Pol system; however, under Strategy 3, typical estimates become smaller than one after the same 13 or so returns. Strategy 3 estimates for the fast multiplier are orders of magnitude worse than they are for λ_1, but this is unsurprising given the small magnitude of λ_2: informally, $g/\sigma = 10^2$, so the smallest multiplier that can be resolved with finite-difference calculations is on the order of 10^{-2}, which is roughly an order of magnitude larger than λ_2. This is essentially what we see in the right pane of Fig. 10: the circle-shaped strategy begins to level off at values that are approximately 20 times too large.

Finally, we return to Fig. 8 to assess the effectiveness of OLS in specifying the Jacobian, as opposed to its eigenvalues. We use OLS to fit a Jacobian to the 585 circle-shaped returns and those returns that immediately precede them. To each of these predecessors, we apply the fitted Jacobian and then compute the residuals between their images and the true values of the circle-shaped returns, measuring the residuals along the slow eigendirection (i.e., along the solid line in Fig. 8). The diamond-shaped points in Fig. 11 show these residuals as a function of the circle-shaped returns' location along the slow eigendirection.

Most residuals are as we would expect: comparable to the Gaussian noise amplitude $g = 3 \times 10^{-4}$. But away from p_0, the primary source of error appears to be the nonlinearity of the dynamics.

5 Discussion

We have studied periodic orbits of dynamical systems in the context of animal locomotion, asking how stochastic fluctuations enable nervous system controllers to stabilize the orbits and steer the organism. Consider the following hypotheses about the control of locomotion:

1. The targets of locomotion controllers are periodic orbits.
2. The controllers maintain the Floquet multipliers of these periodic orbits in suitable ranges.
3. The organism excites perturbations that probe its phase space in a neighborhood of the periodic orbit.
4. The probability distribution of the stochastic trajectories is heavy tailed: finite size jumps that move the trajectory farther from the orbit than other sources of noise yield improved estimates of the multipliers.

An alternate to the first hypothesis is that the control merely confines the motion to an annular region of the phase space. In particular, the control target might be a chaotic attractor or a bounded volume in phase space rather than a periodic orbit.

To test the hypotheses listed above, we call for new investigations of the stochastic dynamics of motor systems that characterize the probability distributions of their trajectories and distinguish among alternative controllers. This will be a multifaceted endeavor because motor tasks have differing mathematical structures. Consider five examples that have been studied extensively:

– balancing an upright stick on a finger
– bouncing a ball on a paddle
– flapping flight of insects and birds
– bipedal running and walking
– maintaining balance while standing still.

In stick balancing and ball bouncing, a person drives a physical object with little mechanical feedback from object to person. Experiments can be designed so that the actuation of the objects is done by a machine and physical contact between object and individual is totally eliminated. Stability of these tasks is an issue only for the object. Stick balancing uses feedback control to stabilize an unstable upright pendulum. The target state is an equilibrium, not a periodic orbit. Cabrera and Milton (2004a) have studied this task and report that the distribution of stick angles is heavy tailed, consistent with our hypotheses. Further experiments to characterize how finger movements excite fluctuations of the stick could help deepen

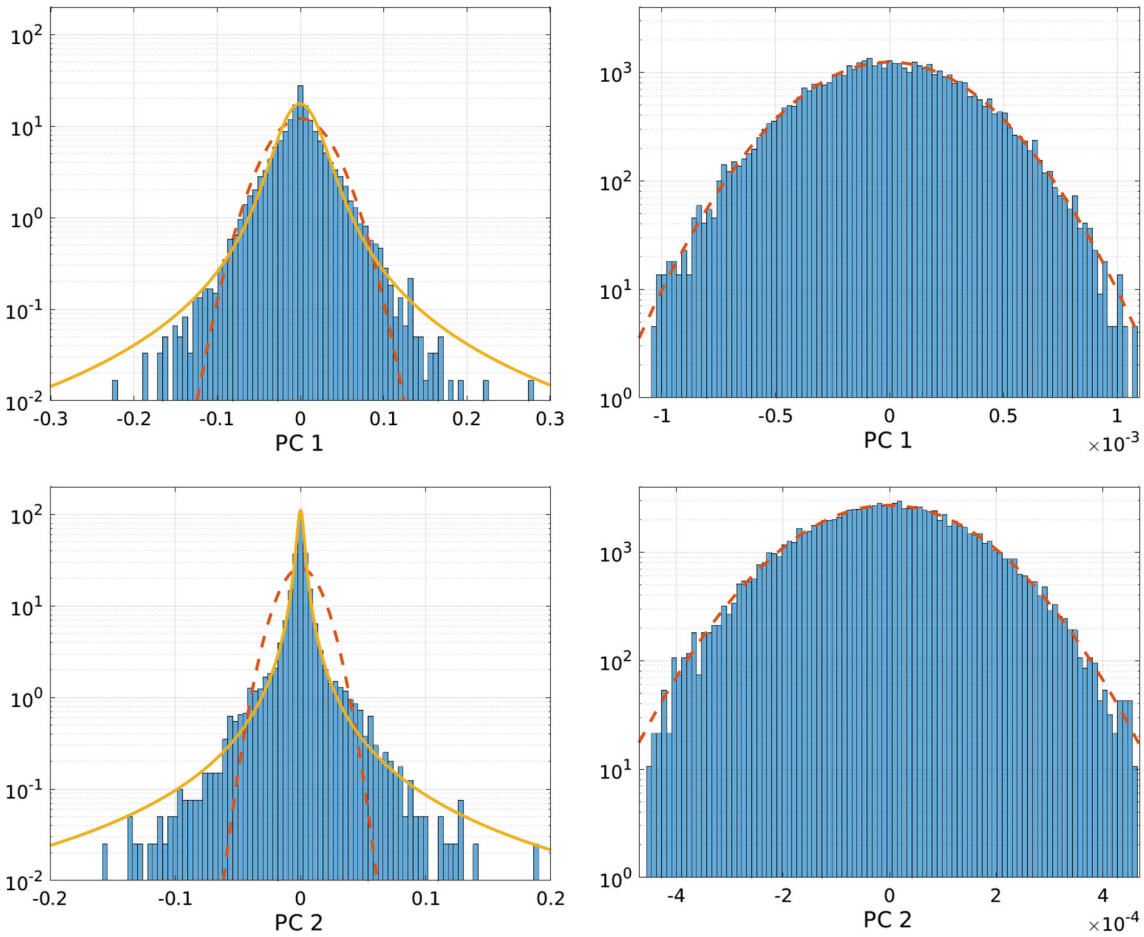

Fig. 9 The distribution of approximately 10^4 returns to the Lorenz Σ (left). The top and bottom frames show the returns projected onto the first and second principal components, respectively. The two frames on the right show the distribution of returns when jumps are removed from the stochastic process used to perturb γ. The dashed curves depict the most likely Gaussian density given the data, while the solid curves in the left frames are the most likely stable distributions

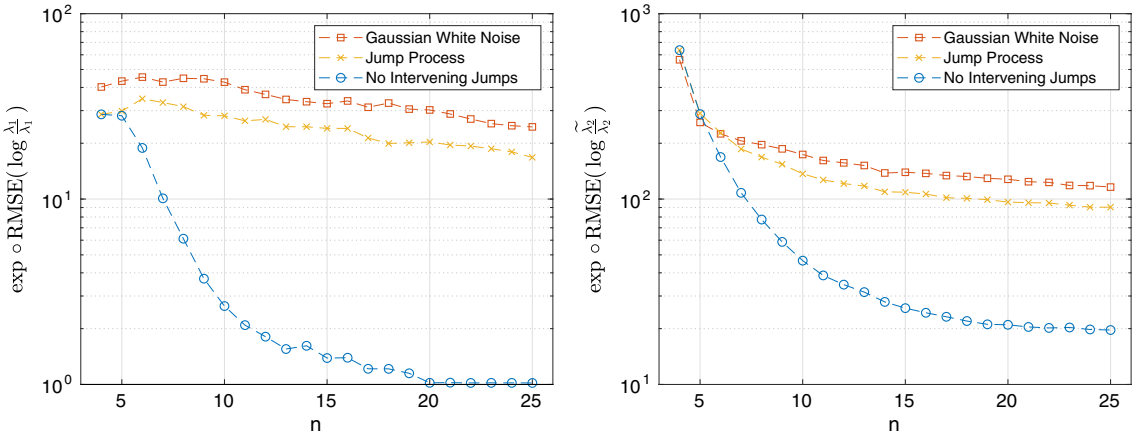

Fig. 10 The approximate fractional error of $\widetilde{\lambda}_i$, $i = 1, 2$; i.e., the approximate fractional error of the slow and fast multipliers of the Lorenz system (left and right, respectively)

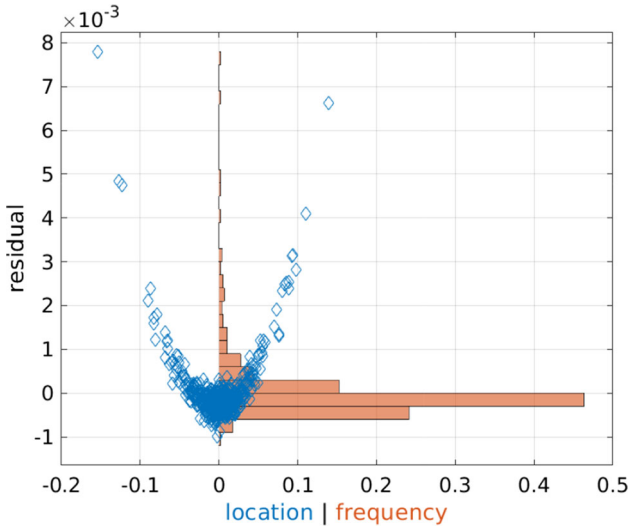

Fig. 11 The residuals when predicting the circle-shaped points in Fig. 8 with an OLS-fitted Jacobian. These residuals (diamond-shaped) are measured along $v_1 \approx (-0.6766, 0.7363)$, which is the eigendirection of the slow multiplier, λ_1. The abscissa of each residual is the location of its circle-shaped point along span(v_1). The histogram shows the relative frequencies of residuals

our understanding of effective means of control in this task. Note that sensory feedback delays are important and limit performance: people cannot balance short sticks (Cabrera and Milton 2004b).

Actuation of a bouncing ball occurs at times of impact between ball and paddle. These impacts are usually modeled as instantaneous events with velocities after impact determined by a coefficient of restitution. Schaal et al. (1996) analyzed the stability of ball bouncing and showed that there are strategies that make the open loop system stable. Nonetheless, control is exercised through the time, location, and velocity of impacts. Parameter and state estimation in ball bouncing requires that the individual plan the time of impact by estimating the motion of both ball and paddle. In addition to estimating these quantities, the person must also learn the coefficient of restitution to achieve a target height attained by the ball. When the motion is a stable periodic orbit, information about the state space off the orbit may not be needed. However, one can modify the experiment to make stabilization an essential part of this task. If the object hitting the ball is a cylinder (say a baseball bat), the open loop system will be unstable because transverse displacements of the impact point from the midline of the bat will result in larger sideways displacement of the ball. Learning the function that relates these displacements requires data from impacts that are not on the target periodic orbit.

Parallels can be made between stick balancing and insect flight. The wing-beat frequency of *Drosophila* flies is faster (about 250 Hz) than the timescale for sensory feedback and neural processing. Flight control is exercised by modifying the relative motions of the two wings. Stabilization and steering can be analyzed in terms of equilibria of vector fields that average over the flapping motion of their wings. Ristroph et al. (2010) measured the responses of flies to perturbations of their body orientation by gluing metal threads to the flies and generating torques with magnetic pulses. The flies were astonishing in their ability to restore their original flight direction. Reduced models based upon quasi-steady fluid dynamics with linear PI controllers were successfully fit to these data. We suggest that these perturbation methods could give a new level of insight into stick balancing. Magnetic fields could enable stochastic perturbations of the stick with arbitrary distributions of impulse magnitudes. These might have an effect on an individual's skill and the rate at which they learn to balance a stick well.

Bipedal running and walking are modes of locomotion prone to instability. Indeed, falls of the elderly are an enormous health risk. According to the US National Council on Aging, [4] falls result in more than 2.8 million injuries treated in emergency departments annually, including over 800,000 hospitalizations and more than 27,000 deaths. In 2014, the total cost of fall injuries was $31 billion. Clearly, we would benefit from a better understanding of control of walking and running. Structurally, the mechanical and neural subsystems are apparent, but the dynamic consequences of their coupling and feedbacks are not.

Collins and Luca (1993) measured postural sway of people standing upright with "stabilogram diffusion functions." They modeled observed center of pressure fluctuations as fractional Brownian motions. Their analysis suggested that there were at least two control systems, targeting different timescales. They asserted that the fast controller was open loop, while the slow controller was closed loop. Further analysis of the trajectories (Collins and Luca 1994) indicated that their correlation dimensions (Grassberger and Procaccia 1983) were large, suggesting that the feedback mechanisms for maintaining balance are "sloppy," allowing the center of pressure to drift in a neighborhood of its dynamical equilibrium. Peterka (2000) obtained similar stabilogram diffusion functions with a single PID controller that employed time delays.

It seems apparent from the five examples we have cited that additional data and/or analysis is required to infer the structure of their stabilizing controllers. Our goal has been complementary to the quest to identify the architecture of the controllers and build mechanistic models for them. Instead, we analyzed consequences of the assumption that their mathematical architecture consists of a stochastic perturbation of rhythmic motion of the coupled mechanical, neural and environmental system around a (closed-loop) stable limit cycle. In this general mathematical setting, we find that estimating

[4] https://www.ncoa.org/news/resources-for-reporters/get-the-facts/falls-prevention-facts/.

Floquet multipliers which characterize system stability from trajectory data is a subtle problem that presents previously unsuspected issues. Specifically, there are bounds depending upon the probability distribution of the stochastic process (but not its amplitude) as to how accurately the multipliers can be estimated with finite-length trajectories. Subsequent work of Javeed (2017) gives some mathematical explanation for these observations.

There are obstacles to the use of our theoretical and numerical analysis of stochastic perturbations of periodic orbits to provide further insight into the control of motor skills. The first involves the quality of data. Measurement errors are always present, and they are confounded with the effects of abrupt stimuli initiated by the subject. If these two sources of stochasticity in the data have comparable magnitude, then it may be difficult to disentangle them to test our hypotheses about heavy-tailed probability distributions. In particular, we have not yet found a way to do this in the running data analyzed by Maus et al. (2015). A second obstacle comes from the lengthy transients that occur while an individual learns a new motor skill. In all of the human examples we have discussed, there is a preliminary training period in which the subjects practice the task before data are collected. The subject might learn the location and Floquet multipliers of a periodic orbit during that period, thereby facilitating the system identification problem it faces subsequently. Bayesian models may aid in the mathematical representation of acquired knowledge, but further investigation is needed in order to understand the nature of the prior and the estimation problem.

Summarizing, we propose that experiments on motor control give renewed attention to (1) the learning process that occurs during training, (2) measuring responses to event-driven perturbations and (3) fitting probability distributions to data.

Our data-driven analysis relies upon the identification of generic properties of classes of models that are representative of the activities we study. This is a difficult task. For example, mechanical impacts of feet with the ground are a structural aspect that make walking and running more complex than flapping flight. The equations that govern walking by an organism or robot with rigid limbs change discontinuously when a leg transitions between stance and flight. However, the material properties of the human musculoskeletal system preclude accurate models of its impacts as instantaneous events. Our bodies may have sufficient compliance that data-driven models of running and walking as smooth dynamical systems may work better than we initially expect that they will. Careful analysis of (1) time series data informed by further research on stochastic dynamical systems and (2) experiments designed to compare organisms with the stochastic models should yield significant progress in our understanding of locomotion and other motor skills.

References

Blickhan R, Full RJ (1993) Similarity in multilegged locomotion: bouncing like a monopode. J Comp Physiol A 173(5):509–517. https://doi.org/10.1007/BF00197760

Brockwell PJ, Davis RA (1991) Time series: theory and methods, 2nd edn. Springer, New York. https://doi.org/10.1007/978-1-4419-0320-4

Cabrera JL, Milton JG (2004a) Human stick balancing: tuning Lévy flights to improve balance control. Chaos 14(3):691–698. https://doi.org/10.1063/1.1785453

Cabrera JL, Milton JG (2004b) Stick balancing: on-off intermittency and survival times. Nonlinear Stud 11(3):305–318

Collins JJ, De Luca CJ (1993) Open-loop and closed-loop control of posture: a random-walk analysis of center-of-pressure trajectories. Exp Brain Res 95(2):308–318. https://doi.org/10.1007/BF00229788

Collins JJ, De Luca CJ (1994) Random walking during quiet standing. Phys Rev Lett 73(5):764–767. https://doi.org/10.1103/PhysRevLett.73.764

Grassberger P, Procaccia I (1983) Characterization of strange attractors. Phys Rev Lett 50(5):346–349. https://doi.org/10.1103/PhysRevLett.50.346

Guckenheimer J (1995) A robust hybrid stabilization strategy for equilibria. IEEE Trans Automat Control 40(2):321–326. https://doi.org/10.1109/9.341802

Guckenheimer J (2014) From data to dynamical systems. Nonlinearity 27(7):R41. https://doi.org/10.1088/0951-7715/27/7/R41

Hamilton JD (1994) Time series analysis. Princeton University Press, Princeton

Hannan EJ, Kanter M (1977) Autoregressive processes with infinite variance. J Appl Probab 14(2):411–415. https://doi.org/10.2307/3213015

Hartman P (2002) Ordinary differential equations, 2nd edn. Society for Industrial and Applied Mathematics, Philadelphia, corrected reprint of the second (1982) edition. https://doi.org/10.1137/1.9780898719222

Holmes P, Full RJ, Koditschek D, Guckenheimer J (2006) The dynamics of legged locomotion: models, analyses, and challenges. SIAM Rev 48(2):207–304. https://doi.org/10.1137/S0036144504445133

Javeed A (2017) An uncertainty principle for estimates of Floquet multipliers. arXiv:1711.10992

Lehmann EL, Casella G (1998) Theory of point estimation, 2nd edn. Springer, New York. https://doi.org/10.1007/b98854

Maus HM, Revzen S, Guckenheimer J, Ludwig C, Reger J, Seyfarth A (2015) Constructing predictive models of human running. J R Soc Interface. https://doi.org/10.1098/rsif.2014.0899

Peterka RJ (2000) Postural control model interpretation of stabilogram diffusion analysis. Biol Cybern 82(4):335–343. https://doi.org/10.1007/s004220050587

Revzen S, Guckenheimer JM (2008) Estimating the phase of synchronized oscillators. Phys Rev E 78:051907. https://doi.org/10.1103/PhysRevE.78.051907

Ristroph L, Bergou AJ, Ristroph G, Coumes K, Berman GJ, Guckenheimer J, Wang ZJ, Cohen I (2010) Discovering the flight autostabilizer of fruit flies by inducing aerial stumbles. Proc Natl Acad Sci USA 107(11):4820–4824. https://doi.org/10.1073/pnas.1000615107

Schaal S, Sternad D, Atkeson CG (1996) One-handed juggling: a dynamical approach to a rhythmic movement task. J Mot Behav 28(2):165–183. https://doi.org/10.1080/00222895.1996.9941743

Sussmann HJ (1978) On the gap between deterministic and stochastic ordinary differential equations. Ann Probab 6(1):19–41. https://doi.org/10.1214/aop/1176995608