An optimality principle for locomotor central pattern generators

Hansol X. Ryu[1], Arthur D. Kuo[1,2]

[1]Biomedical Engineering Program, [2]Faculty of Kinesiology
University of Calgary, Calgary, Alberta, CANADA
Corresponding author: H. X. Ryu, hansol.ryu@ucalgary.ca

Abstract

Legged locomotion is controlled by two neural circuits: *central pattern generators* (CPGs) that produce rhythmic motor commands (even in the absence of feedback, termed "fictive locomotion"), and *reflex circuits* driven by sensory feedback. Each circuit alone serves a clear purpose, and both are understood to be active normally. The difficulty is in how or why they work together, as there lacks an objective and operational criterion for combining the two. Here we propose that optimization in the presence of uncertainty can explain how feedback should be incorporated for locomotion. The key is to re-interpret the CPG as a state estimator: an internal model of the limbs that predicts their state, using sensory feedback to optimally balance competing effects of environmental and sensory uncertainties. We demonstrate use of optimally predicted state to drive a simple model of bipedal, dynamic walking, which thus yields minimal energetic cost of transport and best stability. The internal model may be implemented with classic neural half-center circuitry, except with neural parameters determined by optimal estimation principles. Fictive locomotion also emerges, but as a side effect of estimator dynamics rather than an explicit internal rhythm. Uncertainty could be key to shaping CPG behavior and governing optimal use of feedback.

Keywords: Central pattern generator; CPG; sensory feedback; locomotion; bipedal walking; walking model; simulation; optimization

Introduction

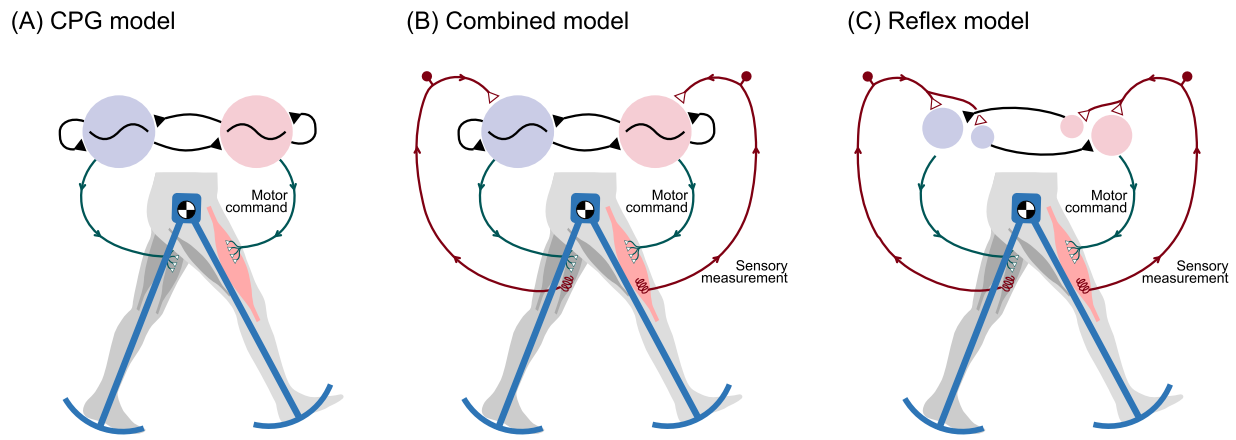

Figure 1. Three ways to control bipedal walking. (A) The central pattern generator (CPG) comprises neural oscillators that can produce rhythmic motor commands, even in the absence of sensory feedback. Rhythm can be produced by mutually inhibiting neural half-center oscillators (shaded circles). (B) The CPG can also incorporate sensory feedback, so that the periphery can influence the motor rhythm. (C) Sensory feedback can also control and stabilize locomotion through reflexes, without need for neural oscillators. The extreme of (A) CPG control without feedback is referred to here as pure feedforward control, and the opposite extreme (C) with no oscillators as pure feedback control. Animal locomotion is thought to incorporate (B) a combination of feedforward and feedback.

Animal locomotion appears to be controlled by two main types of neural circuitry. One type is the central pattern generator (CPG; Figure 1A), which generates pre-programmed, rhythmically timed, motor commands [1–3]. The other is the reflex circuit, which produces motor patterns triggered by sensory feedback (Figure 1C). A hierarchy of reflex loops act together, some integrating multiple sensory modalities for complex behaviors such as stepping and standing control [4,5]. Although reflexes alone seem sufficient to control locomotion, animal CPGs have also demonstrated fictive locomotion, in which rhythmic patterns are sustained even in the absence of sensory feedback [6,7]. In fact, within the intact animal, both types of circuitry work together for normal locomotion (Figure 1B; [8]). But this cooperation also presents a dilemma, of how and even why authority is shared between the two [9].

The combination of central pattern generators with sensory feedback has been explored in computational models. For example, the biologically-inspired Matsuoka oscillator [10] employs a network of mutually inhibiting neurons to intrinsically produce alternating bursts of activity. Sensory

input to the neurons can change network behavior based on system state, such as foot contact and limb or body orientation, to help respond to disturbances. The gain or weight of sensory input determines whether it slowly entrains the CPG [11], or whether it resets the phase entirely [12,13]. Controllers of this type have demonstrated legged locomotion in bipedal [14] and quadrupedal robots [15,16], and even swimming and other behaviors in others [17]. Feedback improves robustness, such as ability to traverse different terrains [18]. These models suggest how sensory feedback could improve the robustness of locomotion for animals.

This raises the question whether it might be optimal to use sensory feedback alone. Human-like models can learn reflexive control and produce quite complex and robust walking motions based on state feedback alone [19–21]. The bipedal (Atlas [22]) and quadrupedal (BigDog [23]) robots of highest performance and robustness are typically driven by feedback (of state of body and environment). In fact, reinforcement learning and other optimization approaches (e.g., dynamic programming [24,25]) are typically expressed in terms of state, rather than time. They have no need for, nor even benefit from, an internally generated rhythm. Thus, although CPGs may be modeled with feedback, operational performance seems to favor feedback alone.

There may nevertheless be a principled reason for a controller to have its own internal rhythm or dynamics. State feedback requires knowledge of state, which cannot be known perfectly but may be estimated from noisy and imperfect sensors. The state estimator [25] uses an internal model of the body to predict expected state and sensory information, despite two types of noise. One is due to uncertainty in environment and internal model, termed *process noise,* and the other to imperfect sensors, termed *sensor noise.* Error in predicted vs. actual sensory feedback is used to correct the state estimate. We previously proposed that these internal model dynamics can function like a CPG [26], albeit with its output interpreted not as the motor command *per se*, but as a state estimate that drives the motor command. A simple model of rhythmic leg motions demonstrates how such a scheme could produce the

equivalent of fictive locomotion [26]. But walking, as suggested by a preliminary model [27], is considerably more complex, with continuous-time dynamics, discrete and changing ground contact conditions, and risk of falling. Perhaps state estimation could apply to walking as well.

The purpose of the present study was to test an estimator-based CPG controller with a dynamic walking model. We devised a simple state feedback control scheme, producing stance and swing leg torques as a function of the leg states. Assuming noise acts on both the sensors and as disturbances to the system, we devised a state estimator for the linear, continuous-time dynamics of the legs, with a discrete switch between stance and swing dynamics. The combination of control and estimation thus define our version of a CPG controller that incorporates sensory feedback. In fact, this same controller may be realized in the form of a Matsuoka oscillator [10], with similar neuron-like dynamics. We expected that minimum state estimation error would allow this model to walk with optimal performance, in terms of measures such as mechanical cost of transport. Scaling the sensory feedback either higher or lower than optimal would be expected to yield poorer performance. Such a model may conceptually explain how CPGs could incorporate sensory feedback based on optimal estimation principles.

Results

Central pattern generator controls a dynamic walking model

The CPG controller produced a periodic gait with a model of human-like dynamic walking (Figure 2A). Much of the walking motion was due to the passive dynamics of pendulum-like legs, which can swing back and forth on their own. The legs were also influenced by active torque commands (T_1 and T_2, Figure 2B) from the CPG, which in turn could be influenced by sensory signals. The passive leg dynamics were sufficient to yield a periodic gait, if it were not for energy dissipation in each step's ground contact

collision [28,29]. Active control was therefore necessary to restore that energy through the torque commands. The result was an alternating motion of the two legs (Figure 2C), offset in phase by half a stride period. These leg angles and the ground contact condition ("GC", 1 for contact, 0 otherwise; Figure 2C) were treated as measurements to be fed back to the CPG. Each leg's states ($\boldsymbol{x}_i \triangleq [\theta_i, \dot{\theta}_i]^T$) described a periodic orbit or limit cycle (Figure 2D), which could be perturbed and made to fall (Figure 2E).

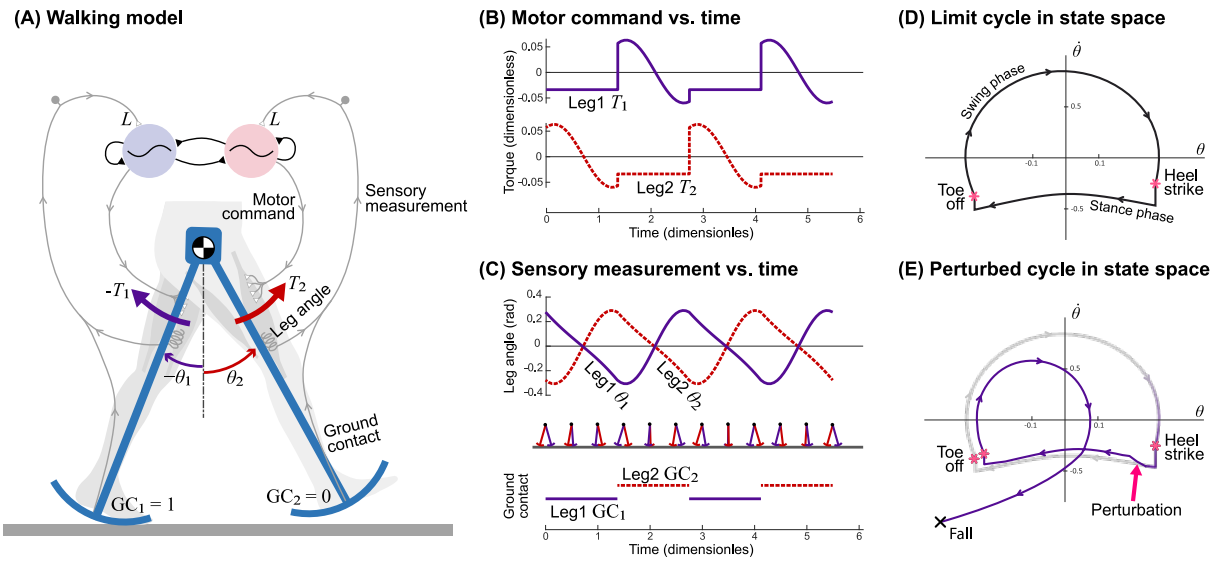

Figure 2. Dynamic walking model controlled by CPG controller with feedback. (A) Pendulum-like legs are controlled by motor commands for hip torques T_1 and T_2, with sensory feedback of leg angle and ground contact "GC" relayed back to controller. (B) Controller produces alternating motor commands vs. time, which drive (C) leg movement θ. Sensory measurements of leg angle and ground contact in turn drive the CPG. (D) Resulting motion is a nominal periodic gait (termed a "limit cycle") plotted in state space $\dot{\theta}$ vs θ. (E) Discrete perturbation to the limit cycle can cause model to fall.

The resulting gait had approximately human-like parameters when walking without noisy disturbances. The nominal walking speed was equivalent to 1.25 m/s and step length 0.55 m (or normalized 0.4 $(gl)^{0.5}$ and 0.55 l, respectively; g is gravitational constant, l is leg length). The corresponding mechanical cost of transport was 0.053, comparable to other passive and active dynamic walking models (e.g., [30–32]).

Pure feedforward and pure feedback both susceptible to noise

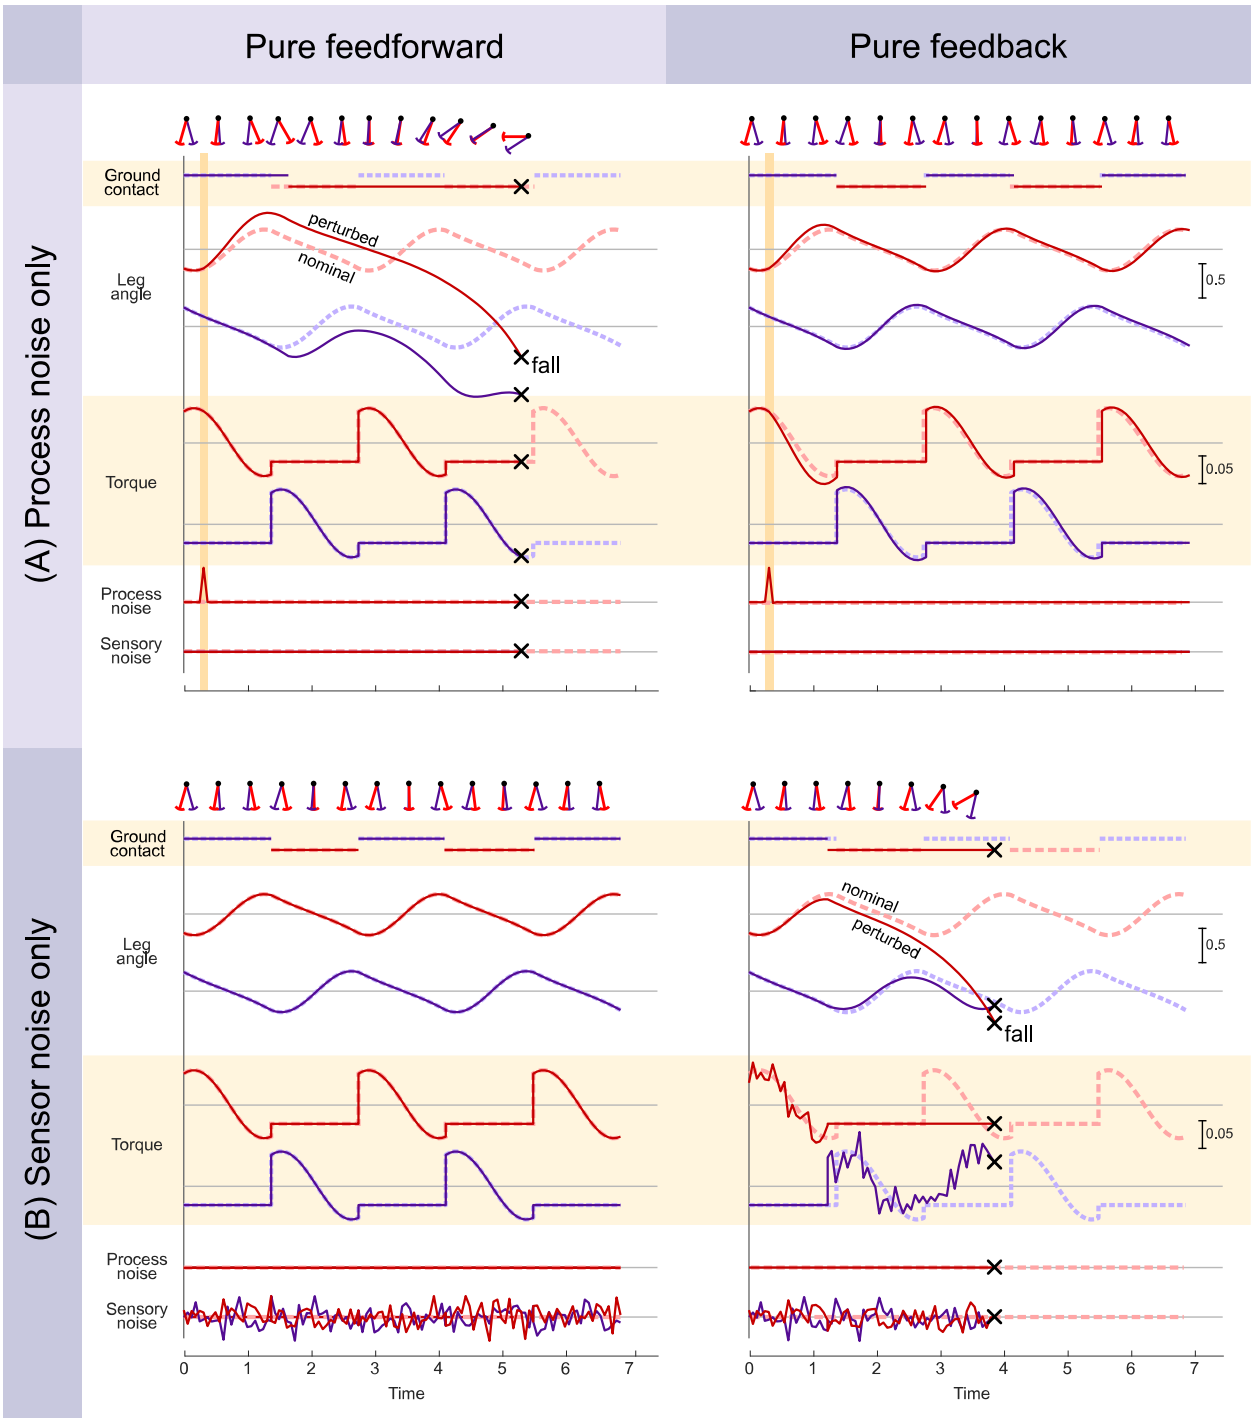

Figure 3. Effect of (A) process and (B) sensor noise on pure feedforward and pure feedback conditions (left and right columns, respectively). Plots show ground contact condition, leg angles, commanded leg torques, and noise levels vs. time, including both the nominal condition without noise (dashed lines), and the perturbed condition with noise (solid lines). With process noise alone, pure feedforward control tended to fall, whereas pure feedback was quite stable. With sensor noise alone, pure feedforward was unaffected, but pure feedback tended to fall.

The critical importance of sensory feedback in the presence of noise was demonstrated with the extremes of pure feedforward and pure feedback (Figure 3A). For both of these cases, we applied a process noise disturbance consisting of a single impulsive force acting on the body. The pure feedforward controller failed to recover (Figure 3A left), and would fall within about two steps. Its perturbed leg and ground contact states became mismatched to the nominal rhythm, which in pure feedforward does not respond to state deviations. In contrast, the feedback controller could recover from the perturbation (Figure 3A right) and return to the nominal gait. Feedback control is driven by system state, and therefore automatically alters the motor command in response to perturbations.

We also applied an analogous demonstration with sensor noise. Adding continuous noise to sensory measurements had no effect on pure feedforward control (Figure 3B left), which ignores sensory signals entirely. But pure feedback was found to be sensitive to noise-corrupted measurements, and would fall within a few steps (Figure 3B right). This is because erroneous feedback would trigger erroneous motor commands not in accordance with actual limb state. The combined result was that both pure feedforward and pure feedback control had complementary weaknesses. They performed identically without noise, but each was unable to compensate for its particular weakness, either process noise or sensor noise.

Equivalence between Matsuoka neural oscillator and state estimator

The same CPG model could be represented in two ways. The Matsuoka neural oscillator (Figure 4A) representation had two mutually inhibiting half-center oscillators, one driving each leg ($i = 1$ for left leg, $i = 2$ for right leg). Each half-center had a total of three neurons, one a primary Matsuoka neuron with standard second-order dynamics (states u and v). Its output drove the second neuron (α) producing the motor command to the ipsilateral leg. The third neuron was responsible for relaying ground contact ("c") sensory information, to both excite the ipsilateral Matsuoska neuron and inhibit the contralateral one.

The same CPG architecture was then re-interpreted in a control systems framework (Figure 4B), while changing none of the neural circuitry. Here, the structure was not treated as half-center oscillators, but rather as three neural stages from afferent to efferent. The first stage receiving sensory feedback signal was interpreted as a feedback gain L (upper rectangular block, Figure 4B), modulating the behavior of the second stage, interpreted as a state estimator (middle rectangular block, Figure 4B) acting as an internal model of leg dynamics. Its output was interpreted as the state estimate, which was fed into the third, state-based motor command stage (lower rectangular block, Figure 4B). In this interpretation, the three stages correspond with a standard control systems architecture for a state estimator driving state feedback control. In fact, the neural connection weights of the Matsuoka oscillator were determined by, and are therefore specifically equivalent to, a state estimator driving motor commands to the legs.

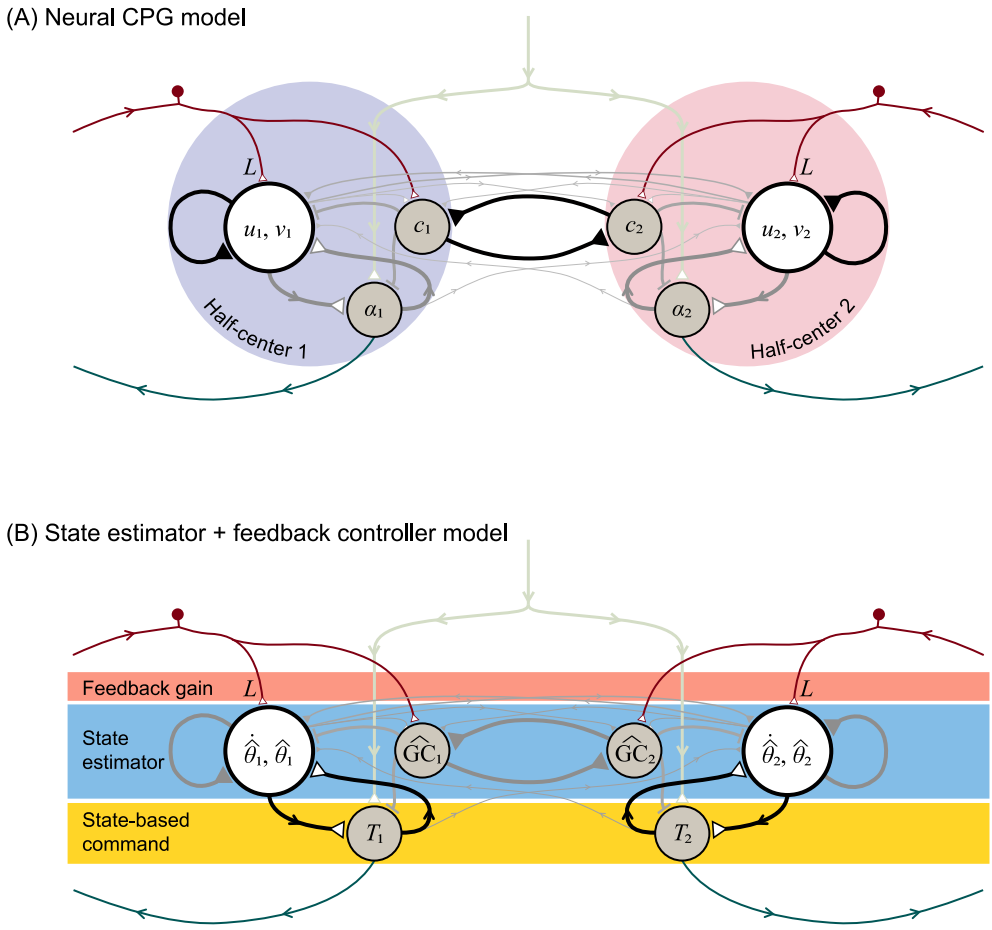

Figure 4. Locomotion control circuit interpreted in two ways: (A) Neural central pattern generator with mutually inhibiting half-center oscillators, and as (B) state estimator with feedback control. Each half-center has a primary neuron with two states representing membrane potential and fatigue (u and v, respectively), an auxiliary neuron c for registering ground contact, and an alpha motoneuron α driving leg torque commands. Inputs include a tonic descending drive, and afferent sensory data with gain L. State estimator acts as second-order internal model of leg dynamics to estimate leg states $\hat{\theta}$ (hat symbol denotes estimate) and ground contact $\widehat{GC}$, which drive state-based command T. All estimator parameters including sensory feedback L are designed for minimum mean-square estimation error. Leg dynamics have nonlinear terms (see Methods) of small magnitude (thin grayed lines).

Sensory feedback gain L optimized by state estimation

We next examined walking performance in the presence of noise, while varying sensory gain L above and below optimal (Figure 5). We applied a combination of both process and sensor noise (with fixed covariances), which made sensory gain L critical to walking performance, unlike the noiseless case. Both pure feedforward and pure feedback control yielded poor performance, as quantified by mechanical cost of transport, step variability, mean time between falls, and state estimator error (Figure 5). Better performance was achieved by varying sensory feedback L continuously between these extremes. The combination of feedforward and feedback, where the CPG rhythm was modulated by sensory information, performed better than either extreme alone.

Best performance was found for the gain L equal to L^*_{lqe} predicted theoretically by linear quadratic estimation (LQE) principles (Figure 5). We had designed L^*_{lqe} based on the covariances of process and sensor noise. Using that gain in nonlinear simulation, the mechanical cost of transport was 0.077, somewhat higher than the nominal 0.053 without noise. Step length variability was 0.046 l, and the model experienced occasional falls, with MTBF (mean time between falls) of about 9.61 $g^{-0.5}l^{0.5}$ (or about 7.1 steps). This optimal case served as a basis for comparisons with other values for gain L.

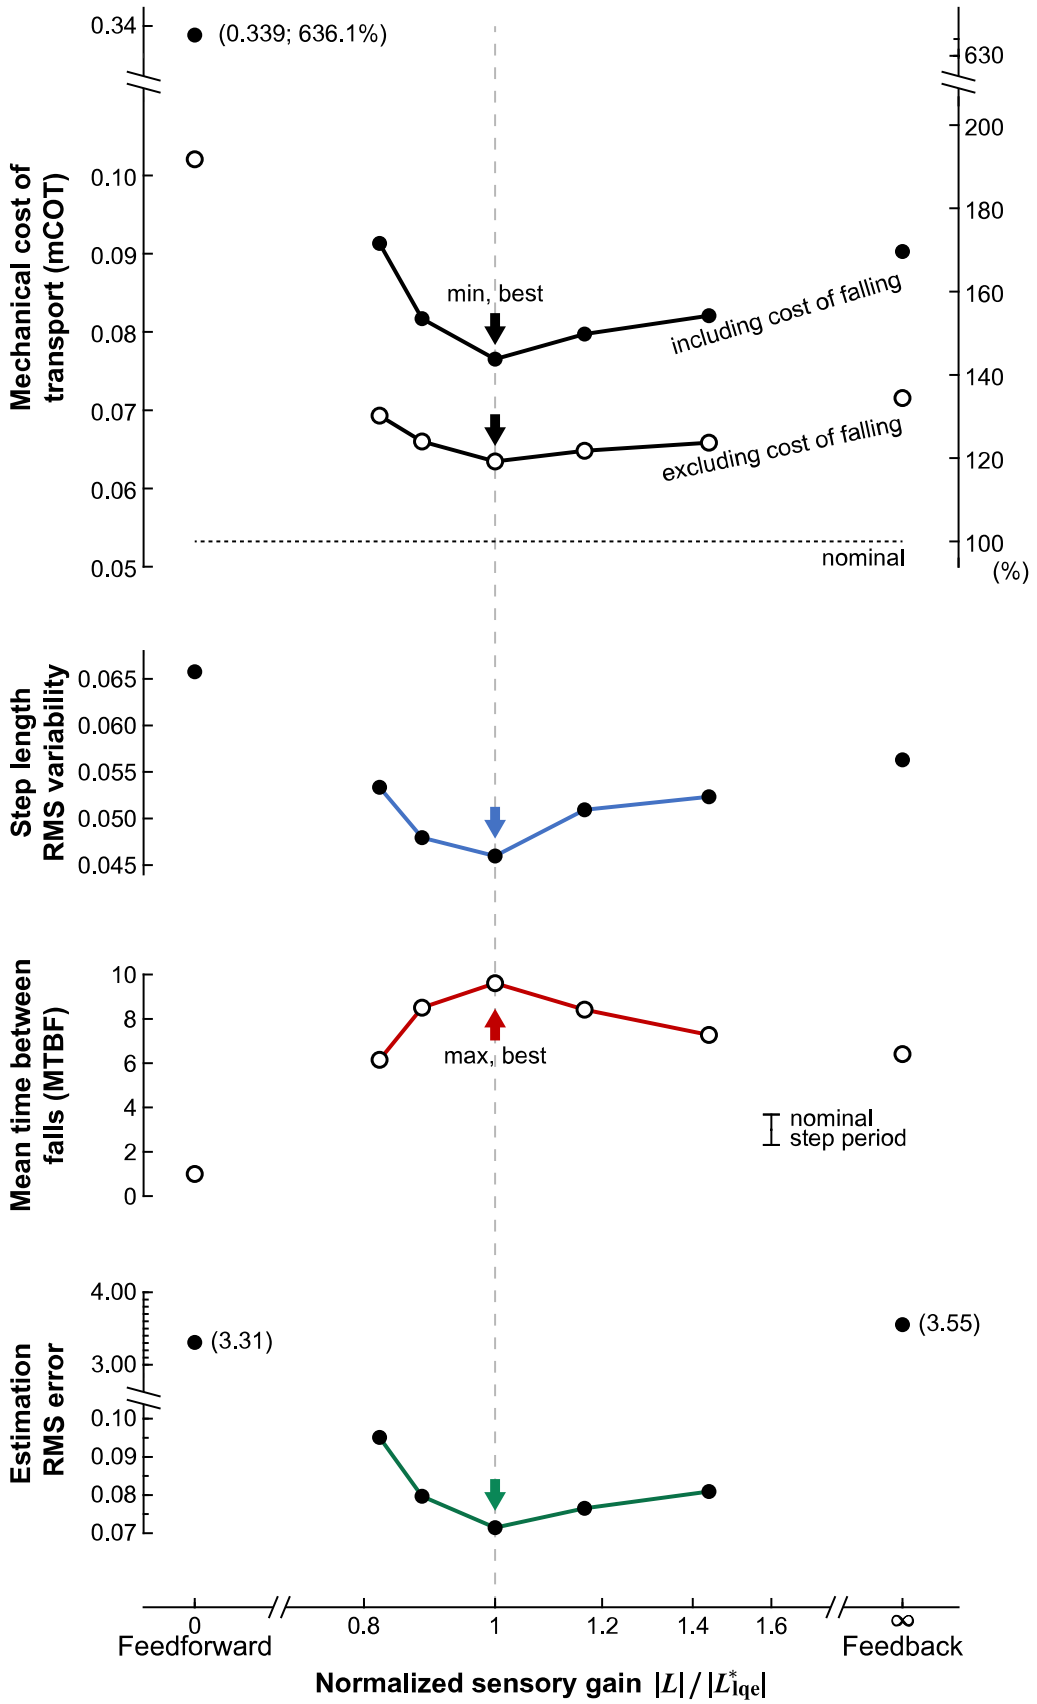
Mechanical cost of transport (mCOT)
0.34
0.10
0.09
0.08
0.07
0.06
0.05
(0.339; 636.1%)
min, best
including cost of falling
excluding cost of falling
nominal
630
200
180
160
140
120
100
(%)
Step length RMS variability
0.065
0.060
0.055
0.050
0.045
Mean time between falls (MTBF)
10
8
6
4
2
0
max, best
nominal step period
Estimation RMS error
4.00
3.00
0.10
0.09
0.08
0.07
(3.31)
(3.55)
0
0.8
1
1.2
1.4
1.6
∞
Feedforward
Feedback
Normalized sensory gain $|L| / |L^*_{\mathrm{lqe}}|$

179

Figure 5. Walking performance as a function of sensory gain, in terms of mechanical cost of transport (mCOT), step length variability, mean time between falls (MBTF), and state estimator error. Normalized sensory gain is defined as $|L|/|L^*_{\mathrm{lqe}}|$, with 1 corresponding to theoretical optimum L^*_{lqe} ($|\cdot|$ denotes matrix norm). Gain was varied between extremes of pure feedforward ($L = 0$) to pure feedback ($L \to \infty$), as well as intermediate values. Vertical arrow indicates best performance (minimum for all measures except maximum for MTBF). For all gains, model was simulated with a fixed combination of process and sensor noise as input to multiple trials, yielding ensemble average measures. Mechanical cost of transport (mCOT) was defined as positive work divided by body weight and distance travelled, and step variability as root-mean-square (RMS) variability of step length. Falling takes time and dissipates mechanical energy, and so mCOT was computed both including and excluding losses from falls (work, time, distance.

Other values for sensory gain generally resulted in poorer performance (Figure 5). Over the range of gains examined ($|L|/|L^*_{\mathrm{lqe}}|$ ranging 0.82 – 1.44), the performance measures worsened on the order of about 10%. This suggests that, in a noisy environment, a combination of feedforward and feedback is important for achieving precise and economical walking, and for avoiding falls. Moreover, the optimal combination for performance can be designed using control and estimation principles.

Fictive locomotion emerges

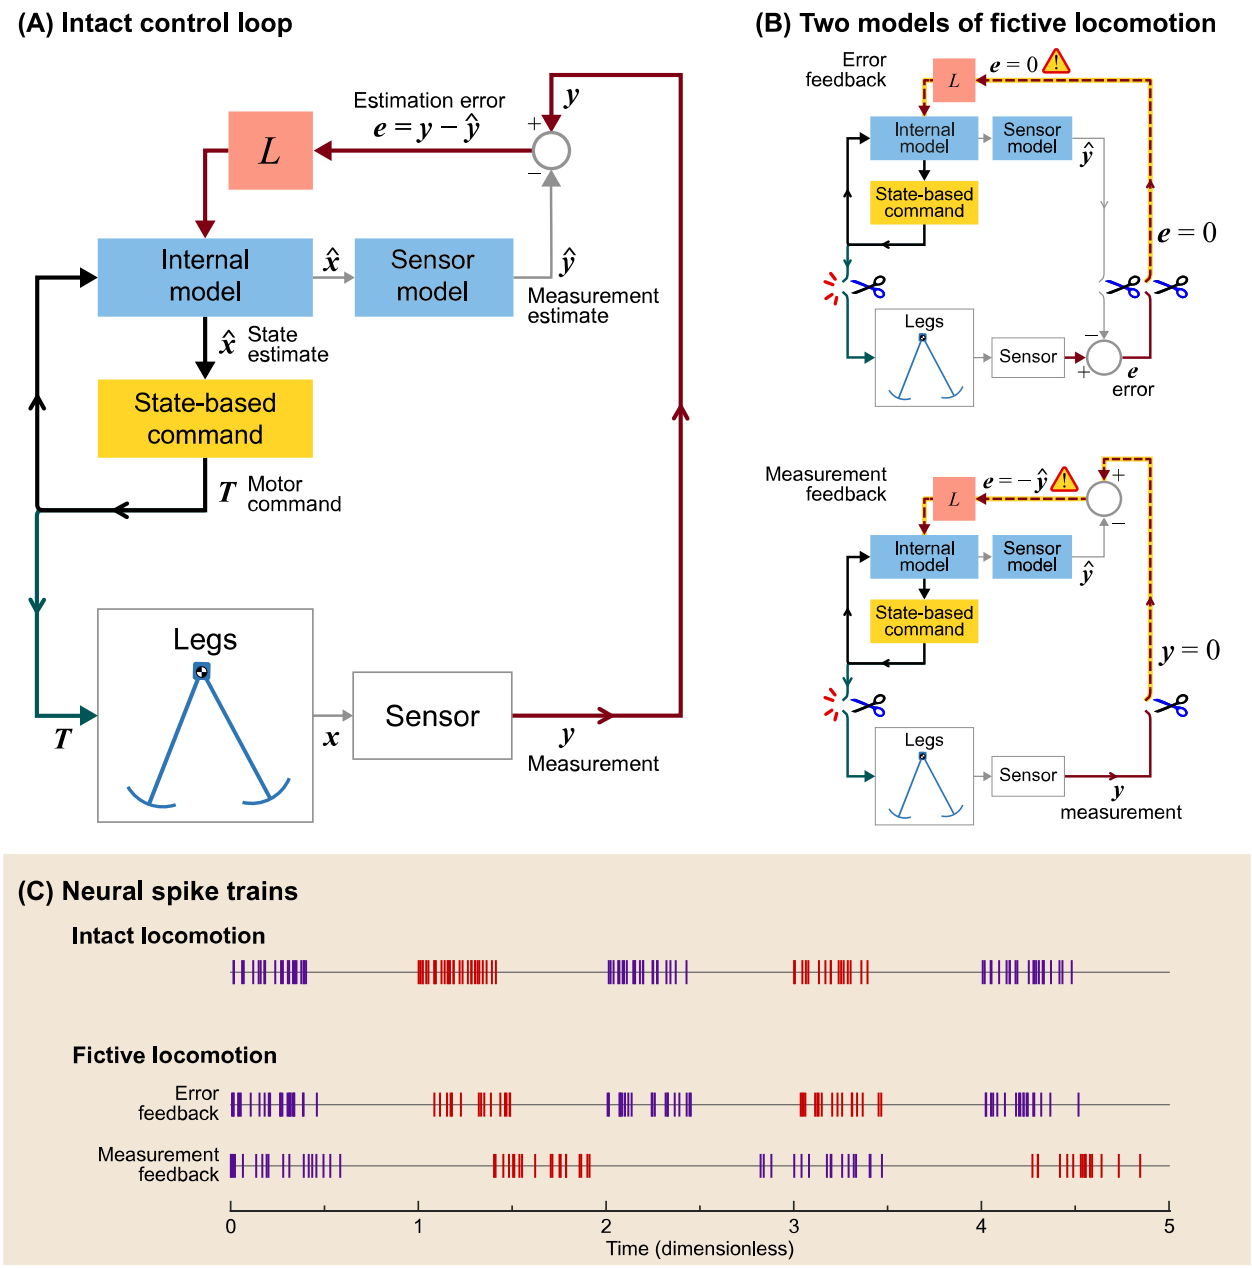

Figure 6. Emergence of fictive locomotion from CPG model. (A) Block diagram of intact control loop, where sensory measurements y and estimation error e are fed into internal model. Motor command T drives the legs and (through efference copy) the internal model of legs. (B) Two models of fictive locomotion, where sensory feedback is removed in two ways. Error feedback refers to sensors that receive efferent copy as inhibitory drive (e.g., some muscle spindles). Removal of error (dashed line) results in sustained fictive rhythm, due to feedback between internal model and state-based command. Measurement feedback refers to other, more direct sensors of limb state $\boldsymbol{x}$. Removal of such feedback can also produce sustained rhythm from internal model of legs and sensors, interacting with state-based command. (C) Simulated motor spike trains show how fictive locomotion can resemble intact. Measurement feedback case produces slower and weaker rhythm than error feedback.

Although the CPG model normally interacts with the body, it was also found to produce fictive locomotion even with peripheral feedback removed (Figure 6). Here we considered two types of biological sensors, referred to as "error feedback" and "measurement feedback" sensors. Error feedback refers to sensors that can distinguish unexpected perturbations from intended movements [33]. For example, some muscle spindles and fish lateral lines [34] receive corollary efferent signals (e.g. gamma motor neurons in mammals, alpha in invertebrates [35]) that signify intended movements, and could be interpreted as effectively computing an error signal within the sensor itself [34]. Movement feedback refers to sensors without efferent inputs, such as nociceptors, golgi tendon organs, cutaneous skin receptors, and other muscle spindles [36], that feed back information more directly related to body movement. Both types of sensors are considered important for locomotion, and so we examined the consequences of removing either type.

These cases were modeled by disconnecting different components of the closed-loop system. This is best illustrated by redrawing the CPG (Figure 4) more explicitly as a traditional state estimator block diagram (Figure 6A). The case of removing error feedback (Figure 6B) was modeled by disconnecting error signal e, so that the estimator would run in an open-loop fashion, as if the state estimate were always correct. Despite this disconnection, there remained an internal loop between the estimator internal model and the state-based command generator, that could potentially sustain rhythmic oscillations. The case of removing measurement feedback (Figure 6C) was modeled by disconnecting afferent signal *y,* and reducing estimator gain by about half, as a crude representation of highly disturbed conditions. There remained an internal loop, also potentially capable of sustained oscillations. We tested whether either case would yield a sustained fictive rhythm, illustrated by transforming the motor command T into neural firing rates using a Poisson process.

We found that removal of both types of sensors still yielded sustained neural oscillations (Figure 6C), equivalent to fictive locomotion. In the case of error feedback, the motor commands from the isolated

CPG were equivalent to the intact case without noise in terms of frequency and amplitude. In the case of measurement feedback, simulations still produced periodic oscillations, albeit with slower frequency and reduced amplitude compared to intact. The state estimator tended to drive estimate $\hat{\theta}$ toward zero, and how this altered state estimation affected the final motor commands was quite dependent on the specifics of the state-based motor command.

Discussion

We have examined how central pattern generators could optimally integrate sensory information to control locomotion. Our CPG model offers an adjustable gain on sensory feedback, to allow for continuous adjustment between pure feedback control to pure feedforward control, all with the same nominal gait. The model is compatible with previous neural oscillator models, while also being designed through optimal state estimation. Simulations reveal how sensory feedback becomes critical under noisy conditions, although not to the exclusion of intrinsic, neural dynamics. In fact, a combination of feedforward and feedback is generally favorable, and the optimal combination can be designed through standard estimation principles. Estimation principles apply quite broadly, and could be readily applied to other models, including ones far more complex than examined here. The state estimation approach also suggests new interpretations for the role of CPGs in animal or robot locomotion.

One of our most basic findings was that the extremes of pure feedforward or pure feedback control each performed relatively poorly in the presence of noise (Figure 3). Pure feedforward control, driven solely by an open-loop rhythm, was highly susceptible to falling as a result of process noise. The general problem with feedforward or time-based control is that a noisy environment can disturb the legs from their nominal motion, so that the nominal command pattern is mismatched for the perturbed state. Under noisy conditions, it is better to trigger motor commands based on feedback of actual limb state,

253 rather than time. But feedback also has its weaknesses, in that noisy sensory information can lead to
254 noisy commands. The solution is to combine both feedforward and feedback together, modulated by
255 sensory feedback gain L. A more uncertain environment (higher process noise) would favor higher gain,
256 and noisier sensors would favor reduced gain. And for a given combination of noise, the theoretically
257 optimal gain L^*_{lqe} would be expected to minimize estimation error, and in turn yield best gait
258 performance (Figure 5). This is expected because theoretically, optimal control also typically calls for
259 optimal state estimation [25,37]. And empirically, imprecise visual information can induce variability in
260 foot placement [38] and poorer walking economy [39]. As expected, the present model walks best with
261 the optimal trade-off between noise effects.

262 Our model also explains how neural oscillators can be interpreted as state estimators (Figure 4).
263 Previous oscillator models (e.g., [10]) have demonstrated how neural half-centers could be modulated
264 by sensory feedback, but not how the feedback gain should operationally be determined, nor how
265 mechanistic principles can be used to combine feedforward and feedback. We have re-interpreted
266 neural oscillator circuits in terms of state estimation (Figure 4), and shown how the gain can be
267 determined in a principled manner, to minimize estimation error (Figure 5). The nervous system has
268 long been interpreted in terms of internal models, for example in central motor planning and control
269 [40–42] and in peripheral sensors [33]. Here we apply internal model concepts to CPGs, for better
270 locomotion performance.

271 This interpretation also explains fictive locomotion as an emergent behavior. We observed persistent
272 CPG activity despite removal of sensors (and either error or measurement feedback; Figure 6), but this
273 was not because the CPG was in any way intended to produce rhythmic timing. Rather, fictive
274 locomotion was a side effect of a state-based motor command, in an internal feedback loop with a state
275 estimator, resulting in an apparently time-based rhythm (Figure 6B). Others have cautioned that CPGs

should not be interpreted as generating decisive timing cues [43–45], especially given the critical role of peripheral feedback in timing [9,46,47]. In normal locomotion, central circuits and periphery act together in a feedback loop, and so neither can be assigned primacy. The present model operationalizes this interaction, demonstrates its optimality for performance, and shows how it can yield both normal and fictive locomotion.

This study argues that it is better to control with state rather than time. The kinematics and muscle forces of locomotion might appear to be time-based trajectories, and therefore require an internal clock to drive them. But another view is that the body and legs comprise a dynamical system dependent on state (described e.g. by phase-plane diagrams, Figure 2), such that the motor command should also be a function of state. Indeed, this is generally the case in optimal control [25], dynamic programming [24], and related methods (e.g., iterative linear quadratic regulators [48] and deep reinforcement learning [21]). As also demonstrated in robots [49,50], state-based control typically also calls for state estimation in realistic conditions [22,51]. Thus, robots with state-driven optimal control and estimation might also exhibit CPG-like fictive behavior, despite having no explicit time-dependent controls.

State estimation may also be applicable to movements other than locomotion. The same circuitry employed here (Figure 4) could easily contribute a state estimate $\hat{x}$ for any state-dependent movements, whether rhythmic [26], non-rhythmic, or discrete. In our view, persistent oscillations were the outcome of state estimation with an appropriate state-based command for the α motoneuron (see Methods). But the same half-center circuitry could be active and contribute to other movements that use non-locomotory, state-based commands. It is certainly possible that biological CPGs are indeed specialized purely for locomotion alone, but the state estimation interpretation suggests the possibility of a more general, and perhaps previously unrecognized, role in other movements.

The present optimization approach may offer insight on neural adaptation. Although we have explicitly designed a state estimator here, we would also expect a generic neural network, given an appropriate objective function, to be able to learn the equivalent of state estimation. That objective could be to minimize error of predicted sensory information, or simply locomotion performance such as cost of transport. Moreover, our results suggest that the eventual performance and control behavior should ultimately depend on body dynamics and noise. A neural system adapting to relatively low process noise (and high sensor noise) would be expected to learn and rely heavily on an internal model. Conversely, relatively high process noise (and low sensor noise) would rely more heavily on sensory feedback. A limitation of our model is that it places few constraints on neural representation, because there are many ways (or "state realizations" [37]) to achieve the same input-output function for estimation. But the importance and effects of noise on adaptation are hypotheses that might be testable with artificial neural networks or animal preparations.

There are, however, cases where state estimation may not apply. State estimation applies best to systems with inertial dynamics or momentum. Examples include inverted pendulum gaits with limited inherent (or passive dynamic; [31]) stability and pendulum-like leg motions [26]. The perturbation sensitivity of such dynamics make state estimation more critical. But other organisms may have well-damped limb dynamics and inherently stable body postures, and thus benefit less from state estimation. There may also be task requirements that call for fast reactions with short synaptic delays, or organismal, energetic, or developmental considerations that limit the complexity of neural circuitry. Such concerns might call for reduced-order internal models [37], or even their elimination altogether, in favor of faster and simpler pure feedforward or feedback. A more holistic view would balance the principled benefits of internal models and state estimation against the practicality, complexity, and organismal costs.

There are a number of limitations to this study. The "Anthropomorphic" walking model does not capture three-dimensional motion and multiple degrees of freedom in real animals. We used such a simple model because it is unlikely to have hidden features that could produce the same results for unexpected reasons. We also modeled extremely simple sensors, without representing the complexities of actual biological sensors. The estimator also used a constant, linear gain, and could be improved with nonlinear estimator variants. We also used a particularly simple, state-based command law, which was designed more for robustness than for economy. Better economy could be achieved by powering gait with precisely-triggered, trailing-leg push-off [26], rather than the simple hip torque applied here. However, the timing is so critical that feedforward conditions ($L < 1$) would fall too frequently to yield meaningful economy or step variability measures. We therefore elected for more robust control to allow a range of feedforward through feedback to be compared (Figure 5). But even with more economical control, we would still expect optimal performance to correspond with optimal gain.

Our principal contribution has been to reconcile the biological evidence for CPGs with the principles of feedback control and state estimation. The evidence of fictive locomotion has long been suggestive that neural oscillators produce the definitive timing and amplitude cues for locomotion. But pre-determined timing is also problematic for control in unpredictable situations [44], making it questionable whether CPG oscillators should dictate timing [43]. To our knowledge, previous CPG models have not included process or sensor noise in control design. Such noise is simply a reality of non-uniform environments and imperfect sensors. But it also yields an objective criterion for uniquely defining control and estimation parameters. The resulting neural circuits resemble previous oscillator models and can produce and explain nominal, noisy, or fictive locomotion. In our interpretation, there is no issue of primacy between CPG oscillators and sensory feedback, because they interact optimally to deal with a noisy world.

Method

Details of the model and testing are as follows. The CPG model is first described in terms of neural, half-center circuitry, which is then paired with a walking model with pendulum-like leg dynamics. The walking gait is produced by a state-based command generator, which governs how state information is used to drive motor neurons. The model is subjected to process and sensor noise, which tend to cause the gait to be imprecise and subject to falling. The CPG is then re-interpreted as an optimal state estimator, for which sensory feedback gain and internal model parameters may be designed, as a function of noise characteristics. The model is then simulated over multiple trials to computationally evaluate its walking performance as a function of sensory gain. It is also simulated without sensory feedback, to test whether it produces fictive locomotion.

CPG architecture based on Matsuoka oscillator

The CPG consists of two, mutually-inhibiting half-center oscillators, receiving a tonic descending input (Figure 4A). Each half-center has a primary Matsuoka neuron with second-order dynamics, described by states u_i for membrane potential and v_i for adaptation or fatigue. The membrane potential also produces an output q_i that can be fed to other neurons. In addition, we included two types of auxiliary neurons (for a total of three neurons per half-center): one for accepting the ground contact input (c_i, with value 1 when in ground contact and 0 otherwise for leg i), and the other to act as an alpha (α_i) motoneuron to drive the leg. We used a single motoneuron to generate both positive and negative (extensor and flexor) hip torques, as a simplifying alternative to including separate rectifying motoneurons.

Each half-center receives a descending command and two types of sensory feedback. The descending command is a tonic input s, which determines the walking speed. Sensory input from the corresponding leg includes continuous and discrete information. The continuous feedback contains information about leg angle from muscle spindles and other proprioceptors [52], which could be modeled as leg angle y_i for measurement feedback, or error e_i for error feedback sensors. The discrete information is about ground contact c_i sent from cutaneous afferents [53].

The primary neuron's dynamics are as follows. The membrane potential u_i has first-order dynamics, and is mainly affected by its own adaptation, a mutually inhibiting connection from other neurons, sensory input, and efference copy of the motor commands. Adaptation or fatigue v_i decays with first-order dynamics, driven by the same neuron's membrane potential as well as sensory input. This is described by the following equations, inspired by [10] and previous robot controllers designed for rhythmic arm movements (e.g., [54] and walking [18]):

$$\dot{u}_i + a_i u_i = -b_i v_i + \sum_{j=1}^{2} -w_{ij} q_j + \sum_{j=1}^{2} h_{ij} e_j + \sum_{j=1}^{2} r_{ij} \alpha_j\left(s, v_j, c_j\right) + f_i(\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{c}) \tag{1}$$

$$q_i = g(u_i) \tag{2}$$

$$\dot{v}_i + a_i' v_i = q_i + \sum_{j=1}^{2} h_{ij}' e_j \tag{3}$$

where there are several synaptic weightings: membrane potential decay a_i and a_i', adaptation gain b_i, mutual inhibition strength w_{ij} (weighting of neuron i's input from neuron j's output, where $w_{ii} = 0$), sensory input gains h_{ij} and h_{ij}', and efference copy strength r_{ij}. The neuron also receives efference copy of its associated motor command $\alpha_j\left(s, v_j, c_j\right)$, which depends on neuron state, descending drive and ground contact. There are also secondary, higher-order influences summarized by the function $f_i(\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{c})$, which have a relatively small effect on membrane potential but are part of the state estimator. The network parameters for Matsuoka oscillators are traditionally set through a combination

of design rules of thumb and hand-tuning, but here nearly all of the parameters will be determined from an optimal state estimator, as described below.

Walking model with pendulum dynamics

The system being controlled is a simple bipedal model walking in the sagittal plane (Fig. 3A). The passive dynamics of pendulum-like legs [31] are actively actuated by added torque inputs (the "Anthropomorphic Model," [30]), and energy is dissipated mainly with the collision of leg with ground in the step-to-step transition. The dissipation determines the amount of positive work required each step. In humans, muscles perform much of that work, which in turn accounts for much of the energetic cost of walking [28].

The walking model is described mathematically as follows. The equations of motion may be written in terms of vector $\boldsymbol{\theta} \triangleq [\theta_1, \theta_2]^T$ as

$$\mathrm{M}(\boldsymbol{\theta}, \mathbf{GC})\ddot{\boldsymbol{\theta}} + \mathrm{C}(\boldsymbol{\theta}, \dot{\boldsymbol{\theta}}, \mathbf{GC})\dot{\boldsymbol{\theta}} + \mathrm{G}(\boldsymbol{\theta}, \mathbf{GC}) = \mathrm{T}(\mathrm{s}, \boldsymbol{\theta}, \mathbf{GC}) \tag{4}$$

where M is the mass matrix, C describes centripetal and Coriolis effects, G contains position-dependent moments such as from gravity, $\mathbf{GC} \triangleq [\mathrm{GC}_1, \mathrm{GC}_2]^T$ contains ground contact, and $\mathrm{T} \triangleq [T_1, T_2]^T$ contains hip torques exerted on the legs (State-based control, below). The equations of motion depend on ground contact because each leg alternates between stance and swing leg behaviors, inverted pendulum and hanging pendulum, respectively. We define each matrix to switch the order of elements at heel-strikes, so that equation of motion can be expressed in the same form.

At heelstrike, the model experiences a collision with ground affecting the angular velocities. This is modeled as a perfectly inelastic collision. Using impulse-momentum, the effect may be summarized as the linear transformation

$$\begin{bmatrix} \dot{\theta}_1^+ \\ \dot{\theta}_2^+ \\ \mathrm{GC}_1^+ \\ \mathrm{GC}_2^+ \end{bmatrix} = \mathrm{S}(\boldsymbol{\theta}, \mathbf{GC}) \cdot \begin{bmatrix} \dot{\theta}_1^- \\ \dot{\theta}_2^- \\ \mathrm{GC}_1^- \\ \mathrm{GC}_2^- \end{bmatrix} \tag{5}$$

where the plus and minus signs ('+' and '–') denote just after and before impact, respectively. The ground contact states are switched such that the previous stance leg becomes and swing leg, and vice versa.

State-based motor command generator

The model produces state-dependent hip torque commands to the legs. Of the many ways to power a dynamic walking model (e.g., [30,50,55,56]), we apply a constant extensor hip torque against the stance leg, for its parametric simplicity and robustness to perturbations. The torque normally performs positive work (Fig. 3B) to make up for collision losses, and could be produced in reaction to a torso leaned forward (not modeled explicitly here; [31]). The swing leg experiences a hip torque proportional to swing leg angle (Fig. 3B, 3C), with the effect of tuning the swing frequency [26].

The overall torque command T_i for leg i is used as the motor command α_i, and may be summarized as

$$T_i(s, \theta_i, \mathrm{GC}_i) = -\,(k_{\mathrm{st}} + \mu_{\mathrm{st}}\, s\,) \cdot\ \mathrm{GC}_i - (k_{\mathrm{sw}}\, \theta_i) \cdot (1 - \mathrm{GC}_i) \tag{6}$$

where the stance phase torque is increased from the initial value k_{st} by the amount proportional to the descending command s with gain μ_{st}. The swing phase torque has gain k_{sw} for the proportionality to leg angle θ_i.

There are also two higher level types of control acting on the system. One is to regulate walking speed, by slowly modulating the tonic, descending command s (Eqn. 6). An integral control is applied on s, so keep attain the same average walking speed despite noise, which would otherwise reduce average speed. The second type of high-level control is to restart the simulation after falling. When falling is

detected (as a horizontal stance leg angle), the walking model is reset to its nominal initial condition, except advanced one nominal step length forward from the previous footfall location. No penalty is assessed for this re-set process, other than additional energy and time wasted in the fall itself. We quantify the susceptibility to falling with a mean time between failures (MTBF), and report overall energetic cost in two ways, including and excluding failed steps. The wasted energy of failed steps is ignored in the latter case, resulting in lower reported energy cost.

Noise model with process and sensor noise

The walking dynamics are subject to two types of noisy disturbances, process and sensor noise (Figure 3). Both are modeled as zero-mean, Gaussian white noise. Process noise n_x (with covariance N_x) acts as an unpredictable disturbance to the states, due to external perturbations or noisy motor commands. Sensor or measurement noise n_y (with covariance N_y) models imperfect sensors, and acts additively to the sensory measurements y. The errors induced by both types of noise are unknown to the CNS controller, and so both tend to reduce performance.

The noise covariances were set so that the model would be significantly affected by both types of noise. We sought levels sufficient to cause significant risk of falling, so that good control would be necessary to avoid falling while also achieving good economy. Process noise was described by covariance matrix N_x, with diagonals filled with variances of noisy accelerations, which had standard deviations of $0.015\ (g/l)$ for stance leg, $0.16\ (g/l)$ for swing leg. Sensor noise covariance N_y was also set as a diagonal matrix with both entries of standard deviation 0.1. Noise was implemented as a spline interpolation of discrete white noise sampled at frequency of 16 $(g/l)^{0.5}$ (well above pendulum bandwidth) and truncated to no more than ±3 standard deviations.

State estimator with internal model of dynamics

A state estimator is formed from an internal model of the leg dynamics being controlled (see block diagram in Fig. 4), to produce a prediction of the expected state $\hat{x}$ and sensory measurements $\hat{y}$ (with the hat symbol '^' denoting an internal model estimate). Although the actual state is unknown, the actual sensory feedback y is known, and the expectation error $e = y - \hat{y}$ may be fed back to the internal model with negative feedback (gain L) to correct the state estimate. Estimation theory shows that regulating error e toward zero also tends to drive the state estimate towards actual state (assuming system observability, as is the case here; e.g., [37]). This may be formulated as an optimization problem, where gain L is selected to minimize the mean-square estimation error. Here we interpret the Matsuoka oscillator network as such an optimal state estimator, the design of which will determine the network parameters.

The estimator equations may be described in state space. The estimator states are governed by the same equations of motion as the walking model (Eqns. 4, 5), with the addition of the feedback correction. Again using hat notation for state estimates, the nonlinear state estimate equations are

$$\begin{bmatrix} \dot{\hat{\boldsymbol{\theta}}} \\ \ddot{\hat{\boldsymbol{\theta}}} \end{bmatrix} = \begin{bmatrix} \dot{\hat{\boldsymbol{\theta}}} \\ \mathrm{M}^{-1}\left(-\mathrm{C}\dot{\hat{\boldsymbol{\theta}}} - \mathrm{G} + \mathrm{T}\right) \end{bmatrix} + \mathrm{L}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \tag{7}$$

We used standard state estimator equations to determine a constant sensory feedback gain L. This was done by linearizing the dynamics about a nominal state, and then designing an optimal estimator based on process and sensor noise covariances (N_x and N_y) using standard procedures ("lqe" command in Matlab, The MathWorks, Natick, MA). This yields a set of gains that minimize mean-square estimation error ($\boldsymbol{x} - \hat{\boldsymbol{x}}$), for an infinite horizon and linear dynamics. The constant gain was then applied to the nonlinear system in simulation, with the assumption that the resulting estimator would still be nearly optimal in behavior. Another sensory input to the system is ground contact $\mathrm{GC_i}$, a boolean variable. The state estimator ignores measured $\mathrm{GC_i}$ for pure feedforward control (zero feedback gain L), but for all other conditions (non-zero L), any sensed change in ground contact overrides the estimated ground

contact $\widehat{\mathrm{GC}}_{\mathrm{i}}$. When the estimated ground contact state changes, the estimated angular velocities are updated according to the same collision dynamics as the walking model (eqn. 5 except with estimated variables).

The state estimate is applied to the state-based motor command (Eqn. 6). Although the walking control was designed for actual state information $(\theta_i, \mathrm{GC}_i)$, for walking simulations it uses the state estimate instead:

$$T_i(\mathrm{s}, \hat{\theta}_i, \widehat{\mathrm{GC}}_{\mathrm{i}}) = -\left(k_{\mathrm{st}} + \mu_{\mathrm{st}}\, s\right) \cdot \widehat{\mathrm{GC}}_{\mathrm{i}} - \left(k_{\mathrm{sw}}\, \hat{\theta}_i\right) \cdot \left(1 - \widehat{\mathrm{GC}}_{\mathrm{i}}\right) \tag{8}$$

As with the estimator gain, this also requires an assumption. In the present nonlinear system, we assume that the state estimate may replace the state without ill effect, a proven fact only for linear systems (certainty-equivalence principle, [25,57]). Both assumptions, regarding gain L and use of state estimate, are tested in simulation below.

Theoretical equivalence between neural oscillator and state estimator

Having fully described the walking model in terms of control systems principles, the equivalent Matsuoka oscillator may be determined (Figure 4B). The identical behavior is obtained by re-interpreting the neural states in terms of the dynamic walking model states,

$$u_1 \triangleq \dot{\hat{\theta}}_1, v_2 \triangleq \hat{\theta}_1, u_2 \triangleq \dot{\hat{\theta}}_2, v_2 \triangleq \hat{\theta}_2 \tag{9}$$

along with the neural output function defined as identity,

$$q_i = g(u_i) \triangleq u_i\ . \tag{10}$$

In addition, motor command and ground contact state are defined to match state-based variables (Eqn. 6):

$$\alpha_1 \triangleq T_1, \alpha_2 \triangleq T_2, c_1 \triangleq \widehat{\mathrm{GC}}_1, c_2 \triangleq \widehat{\mathrm{GC}}_2 \ . \tag{11}$$

The synaptic weights and higher-order functions (Eqns. 1 – 3) are defined according to the internal model equations of motion (Eqn. 7),

$$\begin{bmatrix} a_1 & w_{12} \\ w_{21} & a_2 \end{bmatrix} = \mathrm{M}^{-1}\mathrm{C} \tag{12}$$

$$\begin{bmatrix} b_1\hat{\theta}_1 - f_1\left(\widehat{\boldsymbol{\theta}}, \dot{\widehat{\boldsymbol{\theta}}}, \widehat{\mathbf{GC}}\right) \\ b_2\hat{\theta}_2 - f_2\left(\widehat{\boldsymbol{\theta}}, \dot{\widehat{\boldsymbol{\theta}}}, \widehat{\mathbf{GC}}\right) \end{bmatrix} = \mathrm{M}^{-1}\mathrm{G} \tag{13}$$

$$\begin{bmatrix} h'_{11} & h'_{12} \\ h'_{21} & h'_{22} \\ h_{11} & h_{12} \\ h_{21} & h_{22} \end{bmatrix} = \mathrm{L} \tag{14}$$

$$\begin{bmatrix} r_{11} & r_{12} \\ r_{21} & r_{22} \end{bmatrix} = \mathrm{M}^{-1} \tag{15}$$

$$a'_1 = a'_2 = 0 \tag{16}$$

Because the mass matrix and other variables are state dependent, the weightings above are state dependent as well. The functions f_1 and f_2 are higher-order terms, which could be considered optional; omitting them would effectively yield a reduced-order estimator.

The result of these definitions is that the Matsuoka neuron equations (Eqns. 1 – 3) may be rewritten in terms of $\hat{\theta}_i$ and $\dot{\hat{\theta}}_i$, to illustrate how the network models the leg dynamics and receives inputs from sensory feedback and efference copy:

$$\ddot{\hat{\theta}}_i + a_i\dot{\hat{\theta}}_i = -b_i\hat{\theta}_i + \sum_{j=1}^{2} -w_{ij}\dot{\hat{\theta}}_j + \sum_{j=1}^{2} h_{ij}e_j + \sum_{\mathrm{j}=1}^{2} r_{ij}\alpha\left(s, \hat{\theta}_j, \widehat{\mathrm{GC}}_j\right) + f_i(\widehat{\boldsymbol{\theta}}, \dot{\widehat{\boldsymbol{\theta}}}, \widehat{\mathbf{GC}}) \tag{17}$$

$$\dot{\hat{\theta}}_i + a'_i\hat{\theta}_i = \dot{\hat{\theta}}_i + \sum_{j=1}^{2} h_{ij}e_j \tag{18}$$

The above may be interpreted as an internal model of the stance and swing leg as pendulums, with pendulum phasing modulated by error feedback e_j and efference copy of the motor command (plus small nonlinearities due to inertial coupling of the two pendulums).

Parametric effect of varying sensory feedback gain L

The sensory feedback gain is selected using state estimation theory, according to the amount of process noise and sensor noise. High process noise, or uncertainty about the dynamics and environment, favors a higher feedback gain, whereas high sensor noise favors a lower feedback gain. The ratio between the noise levels determines the optimal linear quadratic estimator gain L^*_{lqe} (Matlab function "lqe"). A constant gain was determined based on a linear approximation for the leg dynamics, an infinite horizon for estimation, and a stationarity assumption for noise [37]. In simulation, the state estimator was implemented with nonlinear dynamics, assuming this would yield near-optimal performance.

It is thus instructive to evaluate walking performance for a range of feedback gains. Setting L too low or too high would be expected to yield poor performing. Setting L equal to the optimal LQE gain L^*_{lqe} would be expected to yield approximately the least estimation error, and therefore the most precise control (e.g. [58]). In terms of gait, more precise control would be expected to reduce step variability and mechanical work, both of which are related to metabolic energy expenditure in humans (e.g., [39]). The walking model is also prone to falling when disturbed by noise, and optimal state estimation would be expected to reduce the frequency of falling.

We performed a series of walking simulations to test the effect of varying the feedback gain. The model was tested with 20 trials of 100 steps each, subjected to pseudorandom process and sensor noise of fixed covariance (W and V, respectively). In each trial, walking performance was assessed with

mechanical cost of transport (mCOT, defined as positive mechanical work per body weight and distance travelled; e.g., [32]), step length variability, and mean time between falls (MTBF) as a measure of walking robustness (also referred to as Mean First Passage Time [59]). The sensory feedback gain L^*_{lqe} was first designed in accordance with the experimental noise parameters, and then the corresponding walking performance was evaluated. Additional trials were performed, varying sensory feedback gain L with lower and higher than optimal values to test for a possible performance penalty. These sub-optimal gains were determined by re-designing the estimator with sensor noise ρV (ρ between 10^-4 and 10^0.8, with smaller values tending toward pure feedforward and larger toward pure feedback). This procedure guarantees stable closed-loop estimator dynamics, which would not be the case if the matrix L^*_{lqe} were simply scaled higher or lower. For all trials, the redesigned L was tested in simulations using the fixed process and sensor noise levels. The overall *sensory gain* was quantified with a scalar, defined as the L2 norm (largest singular value) of matrix L, normalized by the L2 norm of L^*_{lqe}.

We expected that optimal performance in simulation would be achieved with gain L close to the theoretically optimal LQE gain, L^*_{lqe}. With too low a gain ($L = 0$, feedforward Figure 1A), the model would perform poorly due to sensitivity to process noise, and with too a high gain ($L \to \infty$, feedback Figure 1C), it would perform poorly due to sensor noise. And for intermediate gains, we expected performance to have an approximately convex bowl shape, centered about a minimum at or near L^*_{lqe}. These differences were expected from noise alone, as the model was designed to yield the same nominal gait regardless of gain L. Simulations were necessary to test the model, because its nonlinearities do not admit analytical calculation of performance statistics.

Evaluation of fictive locomotion

We tested whether the model would produce fictive locomotion with removal of sensory feedback. Disconnection of feedback in a closed-loop control system would normally be expected to eliminate any

persistent oscillations. But estimator-based control actually contains two types of inner loops (Figure 6A), both of which could potentially allow for sustained oscillations in the absence of sensory feedback. However, the emergence of fictive locomotion and its characteristics depend on what kind of sensory signal is removed. We considered two broad classes of sensors, referred to producing error feedback and measurement feedback, with different expectations for the effects of their removal.

Some proprioceptors relevant to locomotion, including some muscle spindles and lateral lines [34], could be regarded as producing error feedback. They receive corollary discharge of motor commands, and appear to predict intended movements, so that the afferents are most sensitive to unexpected perturbations. The comparison between expected and actual sensory output largely occurs within the sensor itself, yielding error signal e (Figure 6B). Disconnecting the sensor would therefore disconnect error signal e, and would isolate an inner loop between state-based command and internal model. The motor command normally sustains rhythmic movement of the legs for locomotion, and would also be expected to sustain rhythmic oscillations within the internal model. Fictive locomotion in this case would be expected to resemble the nominal motor pattern.

Sensors that do not receive corollary discharge could be regarded as direct sensors, in that they relay measurement feedback related to state. In this case, disconnecting the sensor would be equivalent to removing measurement y. This isolates two inner loops, both the command-and-internal-model loop above, as well as a sensory prediction loop between sensor model and internal model. The interaction of these loops would be expected to yield a more complex response, highly dependent on parameter values. Nonetheless, we would expect that removal of y would substantially weaken the sensory input to the internal model, and generally result in a weaker or slower fictive rhythm.

We tested for the existence of sustained rhythms for both extremes of error feedback and measurement feedback. Of course, actual biological sensors within animals are vastly more diverse and

complex than this model. But the existence of sustained oscillations in extreme cases would also indicate whether fictive locomotion would be possible with some combination of different sensors within these extremes.

Acknowledgements

This work was supported in part by NSERC (Discovery Award and CRC Tier I) and Dr. Benno Nigg Research Chair.

Reference

1. Brown TG. On the nature of the fundamental activity of the nervous centres; together with an analysis of the conditioning of rhythmic activity in progression, and a theory of the evolution of function in the nervous system. The Journal of Physiology. 1914;48: 18–46. doi:10.1113/jphysiol.1914.sp001646

2. Wilson DM. The Central Nervous Control of Flight in a Locust. Journal of Experimental Biology. 1961;38: 471–490.

3. Wilson DM, Wyman RJ. Motor Output Patterns during Random and Rhythmic Stimulation of Locust Thoracic Ganglia. Biophys J. 1965;5: 121–143.

4. Sherrington CS. Flexion-reflex of the limb, crossed extension-reflex, and reflex stepping and standing. The Journal of Physiology. 1910;40: 28–121. doi:10.1113/jphysiol.1910.sp001362

5. Pringle JWS. The Reflex Mechanism of the Insect Leg. Journal of Experimental Biology. 1940;17: 8–17.

6. Grillner S. Locomotion in vertebrates: central mechanisms and reflex interaction. Physiological Reviews. 1975;55: 247–304. doi:10.1152/physrev.1975.55.2.247

7. Feldman AG, Orlovsky GN. Activity of interneurons mediating reciprocal 1a inhibition during locomotion. Brain Research. 1975;84: 181–194. doi:10.1016/0006-8993(75)90974-9

8. Büschges A. Sensory control and organization of neural networks mediating coordination of multisegmental organs for locomotion. J Neurophysiol. 2005;93: 1127–1135. doi:10.1152/jn.00615.2004

9. Bässler U, Büschges A. Pattern generation for stick insect walking movements--multisensory control of a locomotor program. Brain Res Brain Res Rev. 1998;27: 65–88. doi:10.1016/s0165-0173(98)00006-x

10. Matsuoka K. Mechanisms of frequency and pattern control in the neural rhythm generators. Biol Cybernetics. 1987;56: 345–353. doi:10.1007/BF00319514

11. Iwasaki T, Zheng M. Sensory Feedback Mechanism Underlying Entrainment of Central Pattern Generator to Mechanical Resonance. Biol Cybern. 2006;94: 245–261. doi:10.1007/s00422-005-0047-3

12. Nassour J, Hénaff P, Benouezdou F, Cheng G. Multi-layered multi-pattern CPG for adaptive locomotion of humanoid robots. Biol Cybern. 2014;108: 291–303. doi:10.1007/s00422-014-0592-8

13. Tsuchiya K, Aoi S, Tsujita K. Locomotion control of a biped locomotion robot using nonlinear oscillators. Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat No03CH37453). Las Vegas, NV, USA: IEEE; 2003. pp. 1745–1750. doi:10.1109/IROS.2003.1248896

14. Morimoto J, and, Hyon S, Cheng G, Bentivegna D, Atkeson CG. Modulation of simple sinusoidal patterns by a coupled oscillator model for biped walking. Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006 ICRA 2006. 2006. pp. 1579–1584. doi:10.1109/ROBOT.2006.1641932

15. Fukuoka Y, Kimura H, Cohen AH. Adaptive Dynamic Walking of a Quadruped Robot on Irregular Terrain Based on Biological Concepts. : 16.

16. Righetti L, Ijspeert AJ. Pattern generators with sensory feedback for the control of quadruped locomotion. 2008 IEEE International Conference on Robotics and Automation. Pasadena, CA, USA: IEEE; 2008. pp. 819–824. doi:10.1109/ROBOT.2008.4543306

17. Bliss T, Iwasaki T, Bart-Smith H. Central Pattern Generator Control of a Tensegrity Swimmer. IEEE/ASME Trans Mechatron. 2013;18: 586–597. doi:10.1109/TMECH.2012.2210905

18. Endo G, Morimoto J, Nakanishi J, Cheng G. An empirical exploration of a neural oscillator for biped locomotion control. IEEE International Conference on Robotics and Automation, 2004 Proceedings ICRA '04 2004. 2004. pp. 3036-3042 Vol.3. doi:10.1109/ROBOT.2004.1307523

19. Geyer H, Herr H. A Muscle-Reflex Model That Encodes Principles of Legged Mechanics Produces Human Walking Dynamics and Muscle Activities. IEEE Transactions on Neural Systems and Rehabilitation Engineering. 2010;18: 263–273. doi:10.1109/TNSRE.2010.2047592

20. Heess N, TB D, Sriram S, Lemmon J, Merel J, Wayne G, et al. Emergence of Locomotion Behaviours in Rich Environments. arXiv:170702286 [cs]. 2017 [cited 16 Feb 2018]. Available: http://arxiv.org/abs/1707.02286

21. Peng XB, Berseth G, Yin K, Van De Panne M. DeepLoco: Dynamic Locomotion Skills Using Hierarchical Deep Reinforcement Learning. ACM Trans Graph. 2017;36: 41:1–41:13. doi:10.1145/3072959.3073602

22. Kuindersma S, Deits R, Fallon M, Valenzuela A, Dai H, Permenter F, et al. Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot. Autonomous Robots. 2016;40: 429–455.

23. Wooden D, Malchano M, Blankespoor K, Howardy A, Rizzi AA, Raibert M. Autonomous navigation for BigDog. 2010 IEEE International Conference on Robotics and Automation. 2010. pp. 4736–4741. doi:10.1109/ROBOT.2010.5509226

24. Bellman R. The theory of dynamic programming. Bull Amer Math Soc. 1954;60: 503–515.

25. Bryson AE, Ho Y-C, Siouris GM. Applied Optimal Control: Optimization, Estimation, and Control. IEEE Transactions on Systems, Man, and Cybernetics. 1979;6: 366–367. doi:10.1109/TSMC.1979.4310229

26. Kuo AD. The relative roles of feedforward and feedback in the control of rhythmic movements. Motor Control. 2002;6: 129–145.

27. O'Connor SM. The relative roles of dynamics and control in bipedal locomotion. University of Michigan. 2009.

28. Donelan JM, Kram R, Kuo AD. Mechanical work for step-to-step transitions is a major determinant of the metabolic cost of human walking. Journal of Experimental Biology. 2002;205: 3717–27.

29. Kuo AD. Energetics of actively powered locomotion using the simplest walking model. Journal of Biomechanical Engineering. 2002;124: 113–20.

30. Kuo AD. A Simple Model of Bipedal Walking Predicts the Preferred Speed--Step Length Relationship. J Biomech Eng. 2001;123: 264–269. doi:10.1115/1.1372322

31. McGeer T. Passive dynamic walking. The International Journal of Robotics Research. 1990;9: 62.

32. Collins S, Ruina A, Tedrake R, Wisse M. Efficient bipedal robots based on passive-dynamic walkers. Science. 2005;307: 1082–1085. doi:10.1126/science.1107799

33. Dimitriou M, Edin BB. Human muscle spindles act as forward sensory models. Current Biology. 2010;20: 1763–1767.

34. Straka H, Simmers J, Chagnaud BP. A New Perspective on Predictive Motor Signaling. Current Biology. 2018;28: R232–R243. doi:10.1016/j.cub.2018.01.033

35. Delcomyn F. Reflexes and pattern generation, Ch. 16. Foundations of Neurobiology. New York: W. H. Freeman; 1998. pp. 383–400.

36. Iggo A. Handbook of sensory physiology. Volume II. Somatosensory system. 1973; 851pp.

37. Kailath T. Linear Systems. Prentice-Hall; 1980.

38. O'Connor SM, Kuo AD. Direction-dependent control of balance during walking and standing. J Neurophysiol. 2009;102: 1411–1419. doi:10.1152/jn.00131.2009

39. O'Connor SM, Xu HZ, Kuo AD. Energetic cost of walking with increased step variability. Gait & Posture. 2012;36: 102–107. doi:10.1016/j.gaitpost.2012.01.014

40. Hwang EJ, Shadmehr R. Internal Models of Limb Dynamics and the Encoding of Limb State. J Neural Eng. 2005;2: S266–S278. doi:10.1088/1741-2560/2/3/S09

41. Kawato M. Internal models for motor control and trajectory planning. Current Opinion in Neurobiology. 1999;9: 718–727. doi:10.1016/S0959-4388(99)00028-8

42. Uno Y, Kawato M, Suzuki R. Formation and control of optimal trajectory in human multijoint arm movement. Biol Cybern. 1989;61: 89–101. doi:10.1007/BF00204593

43. Bässler U. On the definition of central pattern generator and its sensory control. Biol Cybern. 1986;54: 65–69. doi:10.1007/BF00337116

44. Cruse H. The functional sense of central oscillations in walking. Biol Cybern. 2002;86: 271–280. doi:10.1007/s00422-001-0301-2

45. Pearson KG. Central Pattern Generation: A Concept Under Scrutiny. In: McLennan H, Ledsome JR, McIntosh CHS, Jones DR, editors. Advances in Physiological Research. Boston, MA: Springer US; 1987. pp. 167–185. doi:10.1007/978-1-4615-9492-5_10

46. Donelan JM, Pearson KG. Contribution of sensory feedback to ongoing ankle extensor activity during the stance phase of walking. Canadian journal of physiology and pharmacology. 2004;82: 589–598.

47. Pearson KG. Proprioceptive regulation of locomotion. Current opinion in neurobiology. 1995;5: 786–791.

48. Li W, Todorov E. Iterative linear quadratic regulator design for nonlinear biological movement systems. ICINCO (1). 2004. pp. 222–229.

49. Atkeson C, Wisely Babu BP, Banerjee N, Berenson D, Bove C, Cui X, et al. What Happened at the DARPA Robotics Challenge Finals. Springer Tracts in Advanced Robotics. 2018. pp. 667–684. doi:10.1007/978-3-319-74666-1_17

50. Westervelt ER, Grizzle JW, Koditschek DE. Hybrid zero dynamics of planar biped walkers. IEEE Transactions on Automatic Control. 2003;48: 42–56. doi:10.1109/TAC.2002.806653

51. Barfoot TD. State Estimation for Robotics. Cambridge University Press; 2017.

52. Proske U, Gandevia SC. The Proprioceptive Senses: Their Roles in Signaling Body Shape, Body Position and Movement, and Muscle Force. Physiological Reviews. 2012;92: 1651–1697. doi:10.1152/physrev.00048.2011

53. Trulsson M. Mechanoreceptive afferents in the human sural nerve. Experimental Brain Research. 2001;137: 111–116. doi:10.1007/s002210000649

54. Williamson MM. Neural control of rhythmic arm movements. Neural Networks. 1998;11: 1379–1394. doi:10.1016/S0893-6080(98)00048-3

55. Srinivasan M, Ruina A. Computer optimization of a minimal biped model discovers walking and running. Nature. 2005;439: 72–75.

56. Spong MW. Passivity based control of the compass gait biped. IFAC Proceedings Volumes. 1999;32: 506–510. doi:10.1016/S1474-6670(17)56086-3

57. Simon HA. Dynamic Programming Under Uncertainty with a Quadratic Criterion Function. Econometrica (pre-1986); Evanston. 1956;24: 74.

58. Rebula JR, Ojeda LV, Adamczyk PG, Kuo AD. The stabilizing properties of foot yaw in human walking. J Biomech. 2017;53: 1–8. doi:10.1016/j.jbiomech.2016.11.059

59. Byl K, Tedrake R. Metastable Walking Machines. The International Journal of Robotics Research. 2009;28: 1040–1064. doi:10.1177/0278364909340446