Analysis and Control Techniques for the Compass Gait with a Torso Walking on Stochastically Rough Terrain

Min-Yi Chen and Katie Byl

Abstract—Dynamic walking gaits which exploit inverted pendulum dynamics have demonstrated significant promise for biped robot locomotion. For example, these gaits can reduce the energy expended and the number and complexity of actuators required for level-ground walking. However, robot walkers employing dynamic gaits are, in general, also notoriously sensitive to terrain variations. In this paper, we focus on new methods for developing improved control strategies for and analyzing resulting stability of a simple yet effective model for biped walking on rough terrain. Our primary contributions are as follows. (1) We quantify the stabilizing value of adding a torso to the standard compass gait model; (2) we optimize a class of simple controllers on this walker to be robust to unsensed changes in upcoming terrain height; and (3) we develop improved numerical tools for estimating the statistics of fall events for rough terrain walking. Our results indicate that the torso walker can handle unanticipated step changes in terrain of approximately 14% of leg length, and that our statistical tools are effective for a 6-dimensional state space system, indicating promise in the challenge of addressing the curse of dimensionality when applying machine learning techniques to rough terrain walking.

I. INTRODUCTION

Both efficiency and stability are central goals in bipedal robot locomotion. Achieving both goals simultaneously is challenging, however. In meeting these challenges, we suggest that it is first practical to identify models that are simple enough to yield general principles – yet not so simple that they seriously limit the control capabilities of the system. Second, we suggest appropriate tools are required to investigate the "mixing effects" of step-to-step variability in terrain on the dynamics of the system and, correspondingly, on its stability. In particular, we hypothesize that understanding the effects of a single perturbation on a limit-cycle behavior is not sufficient in predicting failure rates nor in providing guarantees of stability for walking models on stochastically rough terrain. This paper investigates the twin issues of identifying simple yet capable models for biped walking and of quantifying and optimizing their stability on rough terrain. A third challenge is to use machine learning techniques for high dimensional systems, i.e., to address the "curse of dimensionality" [1].

Toward addressing the issue of efficiency, researchers have demonstrated stable biped gaits in purely passive mechanical devices [2, 3]. The principles of passive dynamic walking (PDW) have since been extended to design a variety

of powered robots based upon dynamic walking principles [4], and investigations of the energetics of various powered robot [5, 6] certainly seem to indicate that there are significant energetic benefits in exploiting passive dynamic elements in bipedal robots.

A fundamental model used in understanding the dynamics of biped gaits is the compass gait (CG) walker [7, 8, 9]. Although this model demonstrates a stable limit cycle, it is notoriously sensitive to variations in terrain [10]. Several researchers have begun to investigate the stability of the compass gait model on rough terrain, studying passive models [11] and developing control strategies for actuated CG models to improve stability [10, 12, 13]. Recent work on this problem has also been tested on a real-world robot designed to capture the essential dynamics of the CG model [14, 15, 16]. Despite these efforts, the compass gait with a passive ankle (no ankle torque) can negotiate only mildly rough terrain, e.g., height changes on the order of 1-2% of leg length for blind walking [15, 17].

Dynamic programming (DP) has shown some potential in developing effective tools for analysis and optimization of control policies for low-dimensional system such as the compass gait walker. For example, in [10], an approximate optimal control policy is developed for an actuated CG walker with a torque at the hip, and in [17], the dynamics of walking are analyzed using a step-to-step transition matrix. In both cases, however, a mesh is employed to discretize state space, and this presents a potential limitation: meshing is traditionally only practical for very low-dimensional system, typically three dimensions or fewer. [18] present one approach to extend meshed DP techniques to higher-dimensional models of walking, dividing the control problem into a set of nearly-decoupled subsystems.

In this paper, we address the issue of improving stability on rough terrain by studying a walking model with an additional degree of freedom – a torso – based on a hypothesis that effectively decoupling the twin problems of adding energy at each step and of achieving adequate ground clearance for uphill obstacles will require at least two control actuations. Our model of the torso walker can handle step changes in terrain of greater than 10% of leg length during blind walking on stochastically rough terrain.

Increasing the complexity of our walking model makes analysis of the step-to-step dynamics more challenging. We present an approach to extend the analyses developed in [17] toward a sparse set of discrete points in a five-dimensional state space. This is possible because our low-level controllers for swing leg and torso angle tend to restrict the step-to-step

M.-Y. Chen and K. Byl are with the Electrical and Computer Engineering Department, University of California, Santa Barbara, CA 93106 USA.
minyichen@umail.ucsb.edu, katiebyl@ece.ucsb.edu

dynamics of the system to a region near a thin, 2-dimensional manifold within this larger space.

The rest of the paper proceeds as follow. Section II introduces the compass gait both with and without a torso and compares the performance of heuristically-tuned controllers for each. Section III describes methods and results for a hill-climbing optimization to maximize the magnitude of single-step perturbation the torso walker can handle during blind walking. In Section IV, we outline the Poincaré analysis techniques we use, and analyze the metastable dynamics of the system when walking blindly on stochastically rough terrain, via an extension of the techniques in [17] to higher-order systems. We also briefly present experimental simulation results, toward verifying the validity of various simplifications made throughout. Finally, we conclude and discuss directions for future work in Sections V and VI.

II. MODIFIED MECHANICAL STRUCTURE

From [9], it seems reasonable to hypothesize that an upper body may be used to improve stability of rough terrain walking. In this work, we present techniques to quantify the extent to which this is true. As an initial step, we compare the performance of the CG with and without a torso to determine the maximum single-step change in height that each can successfully be negotiated during blind walking on otherwise flat terrain. Based on these single-perturbation tests, we will later focus on the torso walker as a more promising candidate for rough terrain walking. First, the CG models and results of our initial simulation tests are explained in more detail below.

A. Models: Compass Gait With and Without a Torso

The classic (non-torso) compass gait model is depicted on the left side of Figure 1. It consists of two rigid links, connected by a frictionless, revolute joint at the common point called the hip. Additionally, there are three point masses: two on the legs (m) and one on the hip (m_h). The end of each leg can be considered as a tip, providing an unactuated point contact on the ground. Note that all angles shown in Figure 1 are measure clockwise with respect to vertical. For each simulation in this paper, the classic CG walker is actuated by a single torque actuator at the hip, and a low-level, PD controller is used to regulate the inter-leg angle, $\alpha = \theta_1 - \theta_2$. Although a passive compass gait walker may have stable limit cycle gaits for some downhill slopes, an actuator is required for successful walking on level ground or rough terrain.
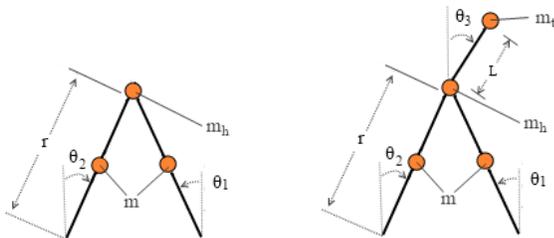
Fig. 1. Left: The classic compass gait biped model. Right: Compass gait with a torso. Simulations parameters: m_h = 15kg, m = 5kg, L = 0.5m, r = 1m, m_t = 10kg (for the compass gait with torso, only) .

The compass gait with a torso, shown on the right side of Figure 1, is based on the compass gait model, but it includes a third rigid link, attached to the hip and ending in a point mass m_t. The CG model with a torso has two actuators at the hip; one provides torque between the swing leg and the torso and the other provides torque between the stance leg and the torso. Although the torso walker has a second actuator, the degree of underactuaction of the model is still one, the same as the classic CG, due to the passive ankle joint.

The walking dynamics can be separated into two phases: a continuous single support phase in which only the stance leg contacts the ground, and an instantaneous double support phase in which the preceding swing leg becomes the stance leg for the next step. The low-level controllers we use to regulate torso and swing leg angle are both PD controllers, which operate only during the single support phase. As for the ground collision, it is assumed to be perfect inelastic and instantaneous, so the equation of motion at impact can be derived using the principle of conservation of angular momentum. Also, we assume that the swing leg will retract automatically when taking a step forward and extend automatically before the impact happens. Equations of motion for both systems can be found in [9].

B. Single Perturbation Experiments

Toward verifying our hypothesis that the stability of the compass gait on rough terrain can be improved by adding a torso, we simulated the compass gait models from Figure 1 both with and without a torso on a long stretch flat terrain with only a single step change in height, Δh.   In these experiments, we use the parameters in the caption of Fig. 1 and the proportional-derivative (PD) gains below in Table 1.

TABLE I
COEFFICIENTS USED FOR LOW-LEVEL PD CONTROL OF JOINTS

PD Controller	CG	CG + torso: torso and swing leg	CG + torso: torso and stance leg
Kp	100	100	18
Kd	10	10	9

The proportional and derivative gains, Kp and Kd, that regulate the (relative) interleg angle of the classic compass gait model are the same values used in regulation of the absolute swing leg angle in the CG model with a torso. We use a desired interleg angle of 20° for the compass gait model, which we determined to be the optimal choice for maximizing the uphill perturbation height this classic model can handle. To compare a similar step-size for the torso walker, we selected a desired absolute swing leg angle of -10° for the swing leg angle (θ_2, positive clockwise) for the CG with a torso. We note briefly that we found the performance of the classic CG to be much better when relative, not absolute, swing leg angle is regulated, while the torso walker performs better when commanding absolute swing angle. In this initial benchmarking, we therefore chose to regulate swing leg angle as best fit each dynamic model: interleg for the classic CG, and absolute for the torso walker. We tested both an "upright" torso angle of 4° for the torso CG, which is optimal when θ_{2des} is constrained to be -10°, and an optimal θ_{2des} and θ_{3des} combination, found as described further on, in Section III.

For each trial in our single perturbation experiments, we start each walker at its fixed point for level-ground walking. We perform successive trials on both models for a range of positive step changes, Δh_+, and negative step change Δh_-, increasing the magnitude by 0.001 [m] until the robot can no longer converge afterward to the fixed point – and falls down.

C. Results of Single-Perturbation Testing

As shown in Table II, the ability of the compass gait to negotiate an unseen step height change in terrain was improved by adding a torso. Performance of the torso depends highly on the combination of desired swing and torso angles that are commanded. In Table II, we show results for both an "upright" and "bent over" set of angles for the torso walker. Recall that PD set points for the upright walker were selected to match the step length of the optimal classic CG model for level-ground walking. For the upright torso model, the *maximum* positive step change in height that could be negotiated, or max(Δh_+), was about 2.5 times that of the two-link CG, and the maximum negative step change in height, max(Δh_-), increases in magnitude by a factor of 9. Once an optimal set for θ_{2des} and θ_{3des} is used, the largest single perturbation the torso walker can handle is over 12 times larger than it is for the classic compass gait. The optimal values for θ_{2des} and θ_{3des} were found using a hill-climbing algorithm described ahead in Section III.

TABLE II
SINGLE-PERTURBATION CAPABILITIES OF CG WALKING MODELS

max(Δh)	Compass Gait (no torso)	Compass gait with a torso	
		Upright Torso	Bent Torso (Optimal)
max(Δh_+)	0.011[m]	0.026[m]	0.139[m]
max(Δh_-)	-0.07[m]	-0.63[m]	-0.33[m]

Maximum step change in height (Δh) that the compass gait with and without a torso is capable of negotiating. The desired interleg angle is 20° for the compass gait model, and the desired torso and desired swing leg angles are 4° and -10°, respectively, for the "upright torso" walker, and are 70° and -24° for the optimal "bent torso" walker.

III. HILL-CLIMBING SEARCH FOR OPTIMAL TORSO CONTROL

A hill-climbing search [19] was used to find the optimal value of optimal combination of θ_{2des} and θ_{3des}. Pseudo code for the algorithm is given below. Our goal is to maximize the overall step height, uphill or downhill, that might be randomly encountered during blind walking. In practice, as shown in Table II, the limiting perturbation size going uphill, max(Δh_+), was always smaller in magnitude than max(Δh_-). Thus, our algorithm effectively seeks only to maximize max(Δh_+). To initiate the algorithm, we fixed θ_{2des} at -10° and searched using with 1° step intervals for the value of torso angle, θ_{3des}, that maximizes max(Δh_+). The resulting initial pair of PD set point angles is 4° for θ_{3des}, and -10° for θ_{2des}.

In each trial, the initial conditions used for the walker dynamics are the fixed point on flat terrain for the corresponding θ_{2des} and θ_{3des} of the controller. In the pseudo code, $\overline{\theta}_{2des}$ and $\overline{\theta}_{3des}$ are the optimal θ_{2des} and θ_{3des} to have been found thus far. The search is illustrated in Figure 2.

HILL-CLIMBING PSEUDO CODE

function: find optimal θ_{2des} and θ_{3des} to maximize max(Δh_+)
initialization: $\overline{\theta}_{2des} = -10°$, $\overline{\theta}_{3des} = 4°$, and

$\overline{\theta}_{2des} = \overline{\theta}_{3des} = 0°$, max($\Delta h_+$) = max($\Delta h_+$)$_{saved}$ = 0.026 [m]

while ($\overline{\theta}_{2des} \neq \theta_{2last}$) or ($\overline{\theta}_{3des} \neq \theta_{3last}$)

for $\theta_{2des} = \overline{\theta}_{2des} + [-1°, 0, 1°]$

  for $\theta_{3des} = \overline{\theta}_{3des} + [-1°, 0, 1°]$

    find new max(Δh_+) for these set point angles

    if max(Δh_+) >= max(Δh_+)$_{saved}$

      max(Δh_+)$_{saved}$ ← max(Δh_+)

      $\theta_{2last} \leftarrow \overline{\theta}_{2des}$; $\theta_{3last} \leftarrow \overline{\theta}_{3des}$

      $\overline{\theta}_{2des} \leftarrow \theta_{2des}$; $\overline{\theta}_{3des} \leftarrow \theta_{3des}$

  end

  end

end

returns locally-optimal $\overline{\theta}_{2des}$ and $\overline{\theta}_{3des}$.
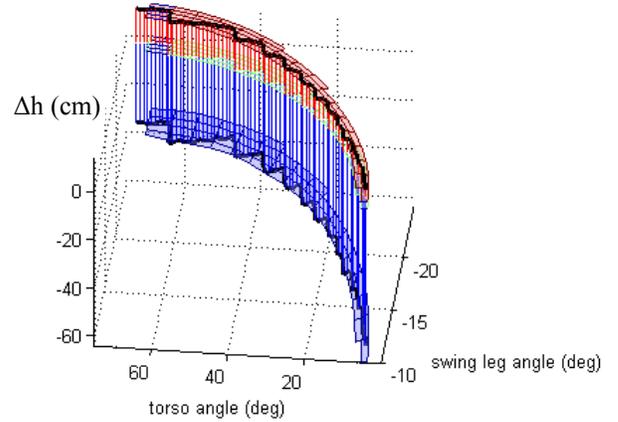
Fig. 2. Path taken during hill-climbing algorithm for optimal torso control. The algorithm starts with θ_{3des} = 4° and θ_{2des} =-10°, and follows the local gradient to a locally-optimal solution θ_{3des} = 70°, θ_{2des} =-24°. We have included patches of the local surface along the paths for max(Δh_+) (top surface, in red) and max(Δh_-) (in blue) to demonstrate that the max uphill and downhill single-step perturbations vary smoothly throughout the search.

Although the hill-climbing search finds on a locally optimal solution for θ_{2des} and θ_{3des}, we notice that the values of both max(Δh_+) and max(Δh-) vary smoothly, as illustrated in Figure 2. It therefore seems likely that the locally-optimal solution may also be globally optimal; however, we can only guarantee that the solution gives a conservative bound on the true performance possible by the torso compass gait walker. The motion of the optimized system during steady-state walking on level ground is shown in Figure 3.

To further investigate the potential of the torso walker on generalized rough terrain, in which we might have knowledge of the average upcoming slope over time, we used the same gradient-based search algorithm to find optimal combinations of θ_{2des} and θ_{3des} for a total of five different terrain slopes. Results are summarized in Table III. Note that there is only

one matched swing leg angle for each torso angle to generate the maximum max(Δh_+). Although the optimal commanded (desired) torso angle is 70°, the corresponding actual torso angle swings between 80° and 90°, so that the torso is bent quite close to horizontal, as illustrated in Figure 3, as opposed to the more upright posture depicted in Figure 1. Also, it is interesting to note that the optimal desired swing leg for a bent posture walker is larger in magnitude that for near-upright walking (see Table II).

TABLE III
SINGLE-PERTURBATION CAPABILITIES OF CG WALKING MODELS

Terrain Slope (deg)	Δh [m], wrt steady-state slope value	Torso angle (deg)		Swing leg (deg)	
		des	actual	des	actual
8	[-0.52,0.033]	83+[-2,2]	[91,98]	-22	-28 – -23
4	[-0.43,0.093]	74+ [-3,3]	[82,89]	-23	-29 – -24
0	[-0.33,0.139]	70+[-3,2]	[80,90]	-24	-31 – -24
-4	[-0.28,0.189]	60+ [-3,2]	[67,74]	-23	-29 – -23
-8	[-0.25,0.232]	70+ [-4,5]	[77,84]	-22	-28 – -23

The desired (des) torso angle and swing leg angle that is locally optimal for negotiating one Δh_+. For each desired torso angle listed, the range in [] brackets indicates values for which max(Δh_-) and max(Δh_+) remain nearly identical (flat). The column of actual torso angle provides the range observed during a steady state bobbing motion. See Fig. 3 as an illustration of actual torso motion. The actual range for the swing leg angle indicates the range of overshoot caused by the PD controller.
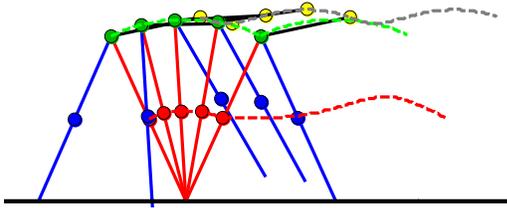
Fig. 3. The torso walker during its optimal gait. The PD controller moves the swing leg forward rapidly during walking, resulting in high ground clearance (i.e., nearly 14% of the leg length), as shown.

In our results, the torso CG model can withstand an isolated perturbation of up to 13.9% of leg length, providing some support to our initial hypothesis that a torso may improve walking stochastically rough terrain. The rest of this paper focuses on analyses of the torso walker on terrain that varies in height at every step. Before continuing, however, we have two brief but important comments. First, the second column in Table III lists heights with respect to the anticipated (steady-state) slope that is listed in column 1, and so the numbers may look a bit misleading. The *absolute* Δh values for uphill step perturbations are given in order below:

0.092 (+8°), 0.128, 0.139, 0.148, 0.143 (-8°) [units in m] These values do not vary quite so significantly as those in the Table. They also show that the walker is more susceptible to uphill perturbations after local periods of walking uphill for some time. This fact leads nicely to our second comment: falls during rough terrain walking may not require a step height as great as the 13.9 cm single-step perturbation, because the local time history of walking may make it more vulnerable to falls. Most particularly, we anticipate more falls after local periods of uphill walking lasting for multiple steps.

IV. QUANTIFICATION OF ROUGH TERRAIN WALKING

Next, we study our controlled torso walkers on stochastic terrain, where the change in height at each step is drawn from a probability distribution. Experiments were performed to analyze the dynamics of both the upright torso configuration, with a commanded torso angle of 4°, and the optimal one, with θ_{3des} of 70°. Our goal is to represent the step-to-step dynamics of rough terrain walking via a compact transition matrix, which can be analyzed to predict the long-term behavior and stability of a particular walker-terrain combination. The details of our simulations, analyses, and results are explained in the sub-sections that follow.

A. One-Step Testing from Non-Equilibrium States

Our first experiments in this Section study the upright torso walker. In initial simulations, the relative change in height of terrain at each step, Δh, is drawn a Gaussian distribution with zero mean and a standard deviation of 0.005 [m], i.e., 0.5% of leg length. This noise level was chosen based roughly on max(Δh_+) from Table II, such that the probability of any single Δh event exceeding max(Δh_+) was quite rare (i.e., 5σ). This was experimentally determined to be an "interesting" level of noise, such that the walker will generally take around 1000 steps or so between failures: long sequences of steps can be recorded, but fall events can also be observed.

The compass gait with a torso, shown in Figure 1, has six state variables, which we define in a state vector, *X*:

$$[X_1, X_2, X_3, X_4, X_5, X_6]^T = [\theta_1, \theta_2, \theta_3, \dot{\theta}_1, \dot{\theta}_2, \dot{\theta}_3]^T$$

Recall that θ_1 is the angle of the stance leg, θ_2 is the angle of the swing leg, and θ_3 is the angle of the torso.

1) Definition of Poincaré map: When walking on perfectly level ground, the controlled torso walker has a stable gait with a periodic orbit through state space, which can be analyzed using a Poincaré map. We employ a similar step-to-step mapping analysis for rough terrain walking. On flat ground, it is typical to select the "post-impact" state (as a new leg begins a continuous stance phase) as the state of interest with which to build a Poincaré map. In such situations, with the ground always defined to have a known, constant slope value, and with both feet in contact with the ground post-impact, only 5 additional states are necessary and sufficient to fully define the post-impact state. For example, given the interleg angle and a known ground slope, both θ_1 and θ_2 are defined.

In our work, a Poincaré snapshot is instead taken for each time instant when the stance leg passes vertical position, i.e., when $X_1 = 0$. The only cases in which such a state did not exist in our simulations were due to imminent failure (falling backward) of the walker, and, fortunately, such events simply map to a failure state in our simulations, regardless.

2) Low-dimensional state representation: As shown in Figure 4, our Poincaré states all tended to fall on or near a thin manifold within the five-dimensional state space of Poincaré snapshots. Moreover, three of the five states could each be expressed as a function of X_4 and X_5. In other words, only X_4 and X_5, seem to be necessary to uniquely define the entire Poincaré state. As a result, we chose to represent the Poincaré

state using only two states: the angular velocity of the stance leg (X_4) and of the swing leg (X_5)
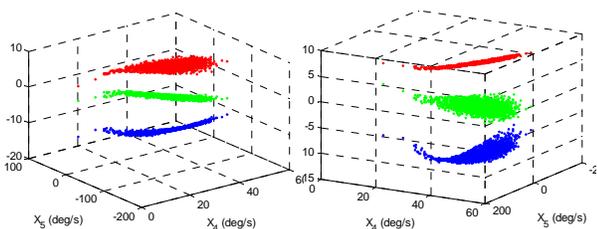
Fig. 4. Poincaré states for the torso walker lie on a thin manifold in a 5D space, and three of the five states can each be represented as function of the other two. To illustrate this, the plots above show X_2 (swing leg angle, in blue), X_3 (torso angle, in red), and X_6 (torso velocity, in green). Note that X_1 is by definition zero for our Poincaré sections. Units on the z-axis are degrees and deg/s, as appropriate.

3) Sparse state selection for a transition matrix: We wish to represent the step-to-step dynamics of rough terrain walking by using an efficient distribution of sparse, discrete points along our 2D manifold of Poincaré states. We hope to demonstrate that such a simplified representation can still be reasonably accurate in describing the system dynamics, providing a promising tool in analyzing walking gaits. To select a sparse set of Poincaré states, we performed 30 Monte Carlo trials in which the torso walker began at the fixed point for zero-slope terrain. Each simulation ran until the walker eventually fell down.

Figure 5 shows the next randomly-selected upcoming step height (z-axis) as a function of Poincaré states (X_4, X_5) for these particular Monte Carlo trials. Successful steps are shown in red, and steps resulting in failure are in blue. The green points show the one-step limit in terrain height change that can be negotiated for any given (X_4, X_5) state visited during walking, vs the $X_4 - X_5$ plane. These limits were obtained empirically, by testing successively larger step heights from each visited state.
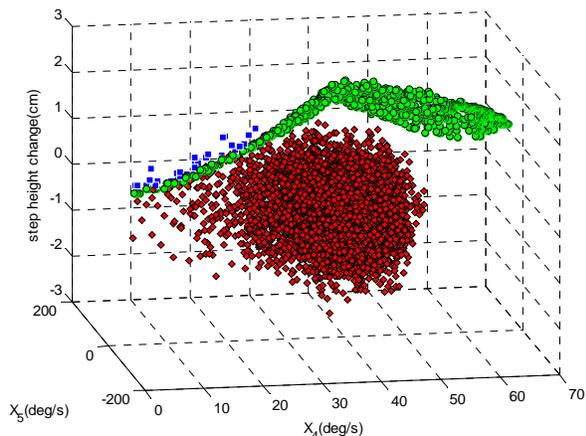
Fig. 5. Poincaré states visited during Monte Carlo trials. Red points show the next, (randomly selected) upcoming step height during Monte Carlo trials. Blue points are the last steps before falling down in each trial. The maximum possible one-step uphill limit, max($\Delta h+$), is shown in green. Note that failures occur from "dangerous" regions of state space, where the green ceiling of points is low.

The sparse set of (X_4, X_5) points we chose to create our transition matrix of the dynamics is shown as a set of green points in Figure 6. We began by iteratively selecting points from our set of visited states (shown in red) one-at-a-time. Points were selected randomly, such that no new point fell within a given distance metric of any previously selected point. Then, we systematically simulated the step-to-step dynamics from each sparse point, for step heights from -2.7 cm to +2.7 cm, at 1 mm spacing. In the course of these tests, we also expanded our sparse mesh whenever necessary, to ensure no mesh point transitioned to a new state that exceeded our distance metric. This explains why the green points in Figure 6 go beyond the range of the red ones from Monte Carlo trials. The new points correspond to multi-step periods of downhill walking. In the case of multi-step periods of significant uphill step heights, the walker would simply fall down, going to an absorbing failure state. Thus, the mesh only grew in one general direction.
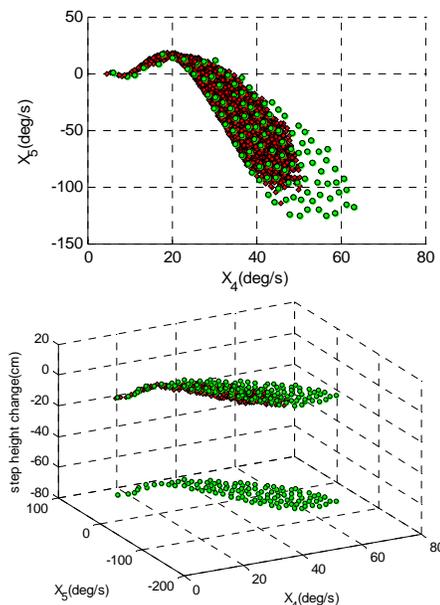
Fig. 6. Performance limits for the upright torso walker. Top: Data from the Monte Carlo trials (red points) and a sparse set of points (green) used to build an approximate transition matrix of the dynamics. Bottom: the max($\Delta h+$) (largest uphill step) and max($\Delta h-$) (downhill) one-step perturbations are shown for each point in the sparse set.

4) Intuition on failure modes: The bottom of Figure 6 shows max(Δh) values for our sparse mesh points. The largest max(Δh_+) value is about 0.026 [m], compared to the max(Δh_-), which is about -0.63 [m]. The fact that max(Δh_+) is so much smaller in magnitude than the max(Δh_-) reemphasizes our earlier comments improved walking intuitively depends on increasing max(Δh_+).

Figure 5 in fact gives a closer view of the upper limit shown in Figure 6. Note that there are two distinct regions within the ceiling of green points shown. For states with higher kinetic energy, i.e., larger magnitude angular velocities, the limit is nearly flat at around 2.6 cm. The failure mode here corresponds to a lack of sufficient swing leg clearance. Unlike the optimal walking gait shown in Figure 3, the upright torso gait simply never gets its swing leg foot high

enough to clear a step change in terrain of more than about 2.7 cm. The walker is "tripping forward" for failures that occur in this region

By contrast, at states with lower kinetic energy, where angular velocities approach zero, the swing leg can clear terrain, but the stance leg never successfully goes past vertical. The failure events here correspond to "falling backward". See the upper diagram in Figure 6 to better visualize this region, in which both leg velocities approach zero. Reaching this region typically occurs because a walker has traveled uphill for several, successive steps.

B. Analyses of walking via state transition matrix

We will examine the dynamics of rough terrain walking by analyzing a state transition matrix. This matrix is built using our sparse set of observed Poincaré states along with data from trials (described earlier) in which we simulate what happens from each state, for discrete heights ranging from -2.7 to 2.7 cm. We will in fact build transition matrices for each of several magnitudes of noise, to calculate the expected number of step to failure as a function of the level of noise.

Let us assume that each new relative change in height is selected from a particular Gaussian distribution with zero mean and some particular standard deviation, σ. Since the probability of the next step height change is drawn from a Gaussian distribution, there is always finite (non-zero) probability any given step will cause failure, and so the walker is guaranteed to eventually fall, with probability one. However, the system may still demonstrate exceptionally long-living periods of continuous walking. Such long-living behaviors before eventual failure are not strictly stable, but they are well-described as metastable [17] dynamics.

For each value of terrain noise, σ, that we wish to examine, we calculate a probability transition matrix, $\hat{T}$, by using equations (1) and (2) below:

$$f_H(h_i) = \int_{h_i - \delta_h/2}^{h_i + \delta_h/2} f_H dh \qquad (1)$$

$$\hat{T} = \sum_{i=1}^{n} T_{\Delta h} \cdot f_H(h) \qquad (2)$$

where f_H is a probability mass function, and f_h is the probability density function of the Gaussian distribution with the desired variance to be examined. $T_{\Delta h}$ is the state transition matrix at the corresponding step height change Δh, with an extra "fallen down" state included along with our sparse (X_4, X_5) points. Once a walker enters the fallen state, it remains fallens, so that falling down is represented as an absorbing failiure state. Hence, the last row of $T_{\Delta h}$ is always $[0,0,\ldots,0,1]$, representing fallen states transition to fallen states with probability one. In our simulations, 110 sparse states were selected and step-to-step transitions were simply mapped to the nearest neighbor node. Note that we also tested more sophisticated weighting strategies, notably barycentric weighting for an irregular mesh of the sparse points, but results were not significantly different in either case. Thus,

we focus on the more simple, nearest-neighbor approach in this presentation.

The probability transition matrix can be used to estimate the probability of being at any neighborhood of interest in state space after n steps of walking on stochastically rough terrain, as shown in equation (3).

$$p[n] = p[0]\hat{T}^n \qquad (3)$$

Here, $p[n]$ is the state distribution vector of the system; it is a vector containing the probability of being at any mesh state at the n^{th} step, and can be calculated given any initial state distribution vector, $p[0]$.

Furthermore, we can decompose any initial state distribution vector, $p[0]$, into a sum of the eigenvectors of $\hat{T}^T$, which is the transpose of the transition matrix, as:

$$p[0] = S_j^0 = V \cdot k, \quad V \in \Re^{(N+1)\times(N+1)} \qquad (4)$$

$$S_j^0 = \begin{cases} 1, & n = j; \\ 0, & \text{otherwise} \end{cases} \qquad (5)$$

where k is the initial weight constant vector, N is the number of the selected states, and V contains all the eigenvectors of $\hat{T}^T$ in columns. S_j^0 is a vector representing a particular, known initial state for the walker, before taking the first step. Note that it is also completely valid to assume a probability distribution for the initial condition, rather than a particular state; the derivation below is identical in either case.

The probability mass function (PMF) for each mesh state after n steps can be found from the state distribution vector, S_j^n, as shown in equation (6),

$$S_j^n = \sum_i^N k_i \lambda_i^n V_i \qquad (6)$$

where λ_i, $i \in \{1,2,3,\cdots,N+1\}, |\lambda_1| \geq |\lambda_2| \geq |\lambda_3| \geq \cdots \geq |\lambda_{N+1}|$, are the eigenvalues of $\hat{T}^T$ and V_i are the corresponding eigenvectors. Note that V_i is the i^{th} column of V, and k_i is the initial weight constant at the i^{th} row. λ_1 will always be 1 since it is the eigenvalue of the absorbing state. In other words, once you are in an absorbing state, you remain there with probability one. For systems that demonstrate long-living behaviors, the magnitude of the second-largest eigenvalue is, correspondingly, only slightly less than 1. If the third and lower eigenvalues are not very close to unity as well, then the expected mean number of steps to failure for the system can be approximated as:

$$M_{fail} \approx 1/(1-\lambda_2) \qquad (7)$$

This theory is discussed in much greater detail in [17].

C. Application to the Upright Torso Walker

For the upright torso walker, walking on terrain where $\sigma=5$mm, we find that $\lambda_2 = 0.9991$ and $\lambda_3 = 0.83$ in our simulations. This corresponds to an expected average of about 1,200 steps before failure, for initial conditions that are well-represented within the second eigenvector (e.g., in a particular neighborhood about the fixed point for walking).

Figure 7 shows how rapidly initial conditions are "forgotten" for this system. After 2 to 5 steps, the PMF of our discretized system has nearly converged a steady-state value.
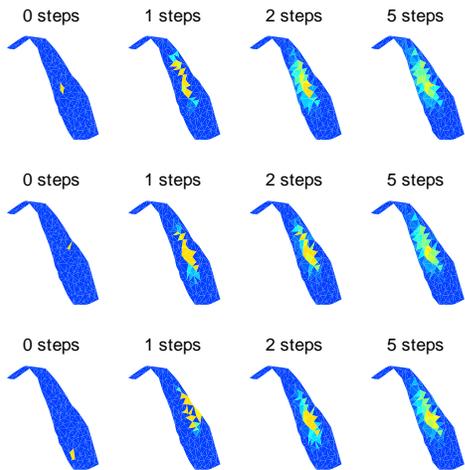
Fig. 7. Starting at each of three, different initial conditions, shown in the leftmost images above, the probability mass function (PMF) is shown after taking 1, 2, and 5 steps on stochastically rough terrain. The similarity toward the right side indicates that initial conditions are rapidly forgotten as the dynamics evolve toward a probabilistic distribution for metastable walking. (Sigma is 5mm in this case.)

Figure 8 illustrates that most falls for the walker are expected to occur from low-energy regions of state space, where the walker (as previously discussed) cannot successfully pivot forward on its passive ankle to complete a step. The average step height causing failure in these trials (in which σ=5mm) was around 1.7cm, considerably lower than the single-perturbation limit of 2.6cm, listed in Table. As we anticipated, the mixing effects of rough terrain walking result in failures due to multi-step events – most particularly, to multiple uphill perturbations in a row.
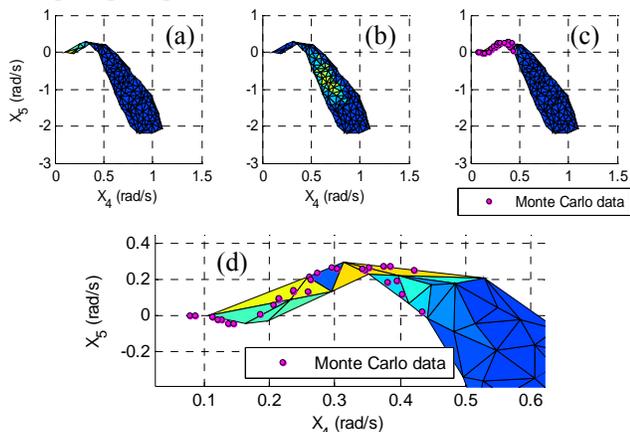
Fig. 8. (a):The conditional probability of failing on the next step, as a function of X_4 and X_5. (b): The probability of being at any given state, after initial conditions (IC) are forgotten. (c) and (d): The total probability of a failure occurring from each state on any one step, after IC are forgotten. (d) magnifies (c) around the "tail" area. Magenta points show actual failures in our Monte Carlo simulations; they are grouped at the tail, as expected.

D. Application to the Optimal Torso Walker

We also created a family of transition matrices for the optimal PD-controlled torso walker for various different levels of noise. Our methods were similar to those described in modeling the upright walker, except that we skipped the painstaking step of simulating each, particular step height for each, particular element in the mesh and instead only post-processed data from initial Monte Carlo trials (with sigma of 4cm) to create each matrix.
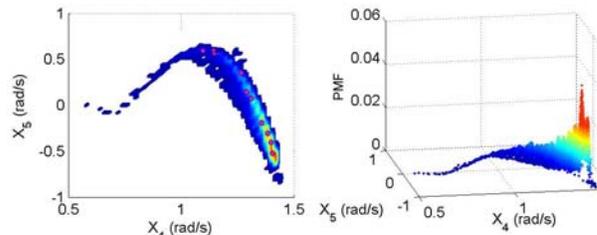
Fig. 9. The probability mass distribution (PMF) for the optimal torso walker, shown from above (left) and in a 3D viewpoint (right).

E. Experimental Results from Monte Carlo Trials

Figure 10 compares steps-to-failure with results from Monte Carlo trials. It is not practical to use Monte Carlo trials to verify all cases here, but the agreement we find in cases we did test is very encouraging, since it is not obvious that our sparse set of points, representing (X4,X5) only, could accurately capture the step-to-step walking dynamics here.
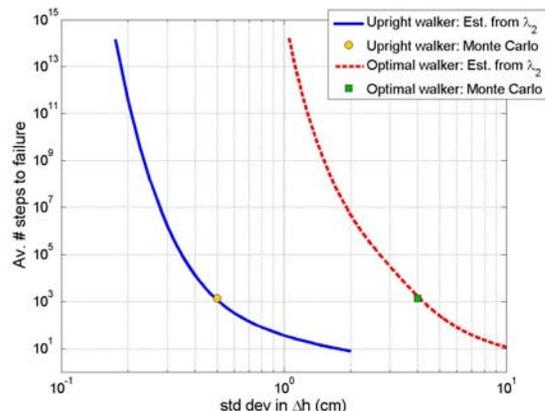
Fig. 10. Expected number of steps between falls, as a function of terrain variability. Solid and dashed lines show estimates based on transition matrix eigenvalues for the upright ($\theta_{3des} = 4°$) and optimal ($\theta_{3des} = 70°$)

V. CONCLUSIONS

In this paper, we identify and analyze a simple but effective model for biped locomotion on stochastically rough terrain. Our contributions are summarized below.

First, we compare the relative effectiveness of a compass gait model with versus without a torso. Using heuristically-tuned controllers for both systems, we find the torso walker can handle a single-step perturbation that is five times greater than our standard CG model. As our second contribution, we optimize low-level controllers for torso angle and swing leg motion, improving performance significantly over our heuristic initial controller.

Previous research has documented that the largest uphill step an optimized CG walker can handle is at most on the order of 2% of leg length [10,15]. Our optimized torso CG can handle a single step perturbation of about 14 cm for a

1-meter leg. On stochastically rough terrain, where terrain height varies at each step and the walker is in a stochastic limit cycle, failures can occur due to lower step changes; however, it can still handle uphill or downhill steps of at least 10% of leg length. As in the case of the heuristic controller, the increase in terrain roughness for the walker with a torso is on the order of a factor of five.

Finally, we develop meshing techniques to capture the step-the-step dynamics of a six-dimensional system. We use an eigenanalysis of this transition matrix to examine the stochastic stability of controlled torso walkers.

VI. DISCUSSION AND FUTURE WORK

Our discretized, low-dimensional modeling of walker dynamics is in part developed as a step toward facilitating the use of the use of dynamic programming in studying and further optimizing control of robot walking. We have presented some promising steps toward the analysis of higher-order dynamic systems, but there are still several directions we plan to explore in future work. Of particular interest, we seek improved methods for using limited and/or real-world data from a robot to develop both transition matrix approximations and to estimate the terrain disturbance which results in failure across the regions of state space that are visited.

Tentatively, we believe this might potentially be done with two small-scale sets of experiments. One set would involve mildly rough stochastic terrain; this would be used to generate a sparse set of discrete points in state space and corresponding transition matrix describing step-to-step dynamics. Our plots of probability mass distribution in Section IV use this basic approach, except that the experiments here are performed in simulation, rather than on an actual robot.

A second set would involve steady state walking on a slope, followed by a particular step change in height. Here, the goal would be to approximate the green "ceiling" in possible next Δh, as shown in Figure 4, for other systems. Intuitively, regions in state space corresponding to lower kinetic energy (i.e., velocities closer to zero) have a lower "ceiling", and we can travel toward these regions by walking uphill for several steps. Conversely, walking downhill tends to pump energy into the system, as one would expect. Exact details of such a method remain as open problem.

Finally, we note that it is interesting that the optimal torso angle for blind walking on rough terrain is nearly horizontal. Humans walk with an upright posture, while our walker looks more reminiscent of a bird such as an ostrich or emu. It would be interesting to investigate the relative stability of models approximating the mass distribution and joint trajectories from data captured from both humans and large birds using our eigenanalysis techniques.

REFERENCES

[1] R. Bellman, Dynamic Programming. Princeton University Press, 1957.
[2] T. McGeer, "Passive dynamic walking," International Journal of Robotics Research, vol. 9, pp. 62–82, April 1990.
[3] M. J. Coleman and A. Ruina, "An uncontrolled toy that can walk but cannot stand still," Phys. Rev. Letters, vol. 80, pp. 3658–3661, 1998.
[4] S. H. Collins, A. Ruina, R. Tedrake, and M.Wisse, "Efficient bipedal robots based on passive-dynamic walkers," Science, vol. 307, pp. 1082–1085, February 18 2005.
[5] A. D. Kuo, "Energetics of actively powered locomotion using the simplest walking model," Journal of Biomechanical Engineering, vol. 124, pp. 113–120, 2002.
[6] S. H. Collins and A. Ruina, "A bipedal walking robot with efficient and human-like gait," Proc. IEEE International Conference on Robotics and Automation, Barcelona, Spain, Apr 2005.
[7] A. Goswami, B. Espiau, and A. Keramane, "Limit cycles and their stability in a passive bipedal gait," in Proc. of IEEE International Conference on Robotics and Automation (ICRA), pp. 246–251, 1996.
[8] A. Goswami, B. Thuilot, and B. Espiau, "A study of the passive gait of a compass-like biped robot: symmetry and chaos," International Journal of Robotics Research, vol. 17, no. 12, 1998.
[9] E. R. Westervelt, J. W. Grizzle, C. Chevallereau, J. H. Choi, and B. Morris, Feedback Control of Dynamic Bipedal Robot Locomotion. CRC Press, Boca Raton, FL, 2007.
[10] K. Byl and R. Tedrake, "Approximate optimal control of the compass gait on rough terrain," in Proc. IEEE International Conference on Robotics and Automation (ICRA), 2008.
[11] J. L.-S. Su and J. B. Dingwell, "Dynamic stability of passive dynamic walking on an irregular surface," Trans. of the ASME, vol. 129, no. 6, pp. 802–810, 2007.
[12] T. Erez and W. D. Smart, "Bipedal walking on rough terrain using manifold control," Proc. International Conference on Intelligent Robots and Systems (IROS), pp. 1539–1544, 2007.
[13] C. M. Hubicki, "Energy-economical heuristically based control of compass gait walking on stochastically varying terrain," Master's thesis, Bucknell Universitiy, 2011.
[14] F. Iida and R. Tedrake, "Minimalistic control of a compass gait robot in rough terrain," in Proceedings of the IEEE/RAS International Conference on Robotics and Automation (ICRA), IEEE/RAS, 2009.
[15] F. Iida and R. Tedrake, "Minimalistic control of biped walking in rough terrain," Autonomous Robots, vol. 28, no. 3, pp. 355–368, January 7 2010.
[16] I. R. Manchester, U. Mettin, F. Iida, and R. Tedrake, "Stable dynamic walking over uneven terrain," The International Journal of Robotics Research (IJRR), vol. 30, March 2011.
[17] K. Byl and R. Tedrake, "Metastable walking machines," International Journal of Robotics Research, vol. 28, pp. 1040–1064, August 1 2009.
[18] E. C. Whitman and C. G. Atkeson, "Control of a walking biped using a combination of simple policies," in Proceedings of the IEEE International Conference on Humanoid Robots, (New Orleans, LA), pp. 520–527, Dec 2009.
[19] S. J. Russell and P. Norvig, Artificial Intelligence: A Modern Approach. Prentice Hall, 2nd ed., 2003.