

Robust Bipedal Locomotion: Leveraging Saltation Matrices for Gait Optimization

Maegan Tucker¹, Noel Csomay-Shanklin², and Aaron D. Ames^{1,2}

Abstract—The ability to generate robust walking gaits on bipedal robots is key to their successful realization on hardware. To this end, this work extends the method of Hybrid Zero Dynamics (HZD) – which traditionally only accounts for locomotive stability via periodicity constraints under perfect impact events – through the inclusion of the saltation matrix with a view toward synthesizing robust walking gaits. By jointly minimizing the norm of the extended saltation matrix and the torque of the robot directly in the gait generation process, we demonstrate that the synthesized gaits are more robust than gaits generated with either term alone; these results are shown in simulation and on hardware for the AMBER-3M planar biped and the Atalante lower-body exoskeleton (both with and without a human subject). The end result is experimental validation that combining saltation matrices with HZD methods produces more robust bipedal walking in practice.

I. INTRODUCTION

Achieving stable and robust locomotion on legged systems is a challenging control task due to underactuation, power limitations, and ground impacts. Two main approaches that have proven successful towards mitigating these challenges in the real world include: 1) generating stable reference trajectories [1]–[4] and modifying these behaviors online using regulators (such as modifying the swing foot location based on lateral velocity [5], [6]); and 2) determining the desired behavior of the robot in real time using online planning via model predictive control [7]–[10] or reinforcement learning [11]–[14]. In this work, we aim to improve these existing approaches by synthesizing *robust* reference trajectories. This is motivated by previous work, which has shown that optimizing the robustness of nominal trajectories improves overall performance regardless of the chosen method of online stabilization [15], [16], and that online planning strategies can have unpredictable behavior without the use of a reference trajectory [17], [18].

It is important to note that there exists previous work towards generating robust limit cycles [16]. However, these existing methods can be computationally expensive and do not scale easily to high-dimensional systems. Thus, the goal of this work is to develop a method of generating robust limit cycles in a way that is scalable to high-dimensional systems (such as the 18 degree-of-freedom (DOF) Atalante lower-body exoskeleton shown in Fig. 1).

This research was supported by NSF Graduate Research Fellowship No. DGE-1745301 and the Zeitlin Family Fund. Research involving human subjects was conducted under IRB No. 21-0693

¹ Authors are with the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 91125.

² Authors are with the Department of Control and Dynamical Systems, California Institute of Technology, Pasadena, CA 91125.

mtucker@caltech.edu




Fig. 1. This work improves the robustness of nominal reference trajectories (gaits) by evaluating the extended saltation matrix directly in the gait generation framework. The approach is demonstrated in simulation and on hardware for both the 7-DOF AMBER-3M planar biped (left) and the 18-DOF Atalante lower-body exoskeleton (right).

Our approach for generating robust walking gaits builds upon the Hybrid Zero Dynamics (HZD) method [19], [20]. Importantly, this approach has been successfully deployed on several high-dimensional real-world legged robots [3], [21]–[25]. The HZD framework encodes walking via nontrivial periodic orbits and directly accounts for discrete impact events. One advantage of this approach is that the stability of the walking can be analyzed using the method of Poincaré sections for systems with impulse effects [26]. However, in practice, it can be difficult to synthesize and control these periodic orbits since 1) it can be computationally expensive to evaluate the Poincaré section for high-dimensional systems and 2) the timing of impact events for real-world systems is noisy, throwing the system off the nominal periodic orbit. For these reasons, the existing approach towards generating robust gaits necessitates extensive user tuning [27].

This paper proposes a robust HZD framework that leverages saltation matrices – first-order approximations of a system’s sensitivity to discrete events – to directly synthesize robust nominal walking trajectories. These matrices, originally used in the field of non-smooth analysis, have been receiving growing attention and have been recently demonstrated towards state estimation for hybrid systems [28], [29] and hybrid event shaping [30]. Inspired by this recent research, our work similarly utilizes saltation matrices to generate stable periodic walking gaits that lead to robust behavior in the real world, even for high-dimensional systems; specifically, we propose including the induced norm of the extended saltation matrix in the HZD optimization cost function. As previewed in Fig. 1, we demonstrate the proposed framework in simulation and on hardware for the 7-DOF AMBER-3M planar biped and the 18-DOF Atalante lower-body exoskeleton with and without a human subject.

II. PRELIMINARIES ON GAIT GENERATION VIA THE HYBRID ZERO DYNAMICS METHOD

As walking is comprised of alternating sequences of continuous dynamics followed by intermittent discrete impact events, it is naturally modeled as a hybrid system [31]–[33]. Consider a robotic system with coordinates $q \in \mathcal{Q} \subset \mathbb{R}^n$ and system state $x = (q, \dot{q}) \in \mathcal{T}\mathcal{Q}$, where $\mathcal{T}\mathcal{Q}$ denotes the tangent bundle of the configuration manifold \mathcal{Q} . To define the hybrid dynamics, we begin by introducing the admissible domain $\mathcal{D} \subset \mathcal{T}\mathcal{Q}$ on which the system evolves, i.e. the set of states with swing foot above the floor, which respecting joint limits, etc. During the continuous domain, the dynamics can be derived from the Euler-Lagrange equations as:

$$\dot{x} = \underbrace{\begin{bmatrix} \dot{q} \\ -D(q)^{-1}H(q, \dot{q}) \end{bmatrix}}_{f(x)} + \underbrace{\begin{bmatrix} 0 \\ D(q)^{-1}B \end{bmatrix}}_{g(x)} u. \quad (1)$$

where $D : \mathcal{Q} \rightarrow \mathbb{R}^{n \times n}$ is the mass-inertia matrix, $H : \mathcal{T}\mathcal{Q} \rightarrow \mathbb{R}^n$ contains the Coriolis and gravity terms, $B \in \mathbb{R}^{n \times m}$ is the actuation matrix, and $u \in \mathbb{R}^m$ is the control input.

Next, we define the *guard*, the set of states where a discrete impact event will occur. Letting $p_{sw}^z : \mathcal{Q} \rightarrow \mathbb{R}$ return the vertical position of the swing foot, the guard for the systems investigated in this work (assuming a known and constant ground height of zero) is defined as:

$$\mathcal{S} = \{x \in \mathcal{D} \mid p_{sw}^z(q) = 0, \dot{p}_{sw}^z(x) < 0\}, \quad (2)$$

or the set of states when the swing foot strikes the ground with a negative velocity. As the system flows into the guard, $n_c \in \mathbb{N}$ holonomic constraints $c : \mathcal{Q} \rightarrow \mathbb{R}^{n_c}$ are enforced in the continuous domain succeeding the impact event, and the system undergoes a discrete jump in the state as captured by the momentum transfer equation [34]:

$$D(q^-)(\dot{q}^+ - \dot{q}^-) = J_c(q^-)^\top \delta F \quad (3)$$

where $q^-, q^+ \in \mathcal{Q}$, and $\dot{q}^-, \dot{q}^+ \in \mathbb{R}^n$ represent the robot configuration and velocity just before and after impact, respectively, $\delta F \in \mathbb{R}^{n_c}$ is the impulse force of the impact event, and $J_c : \mathcal{Q} \rightarrow \mathbb{R}^{n_c \times n}$ is the Jacobian of the holonomic constraints.

Enforcing the holonomic constraints through impacts can equivalently be represented as $J_c(q^-)\dot{q}^+ = 0$. With this, we can rewrite the impact equation [35] as:

$$\begin{bmatrix} D(q^-) & -J_c(q^-)^\top \\ J_c(q^-) & 0 \end{bmatrix} \begin{bmatrix} \dot{q}^+ \\ \delta F \end{bmatrix} = \begin{bmatrix} D(q^-)\dot{q}^- \\ 0 \end{bmatrix}. \quad (4)$$

Noting that the configuration is continuous through impact, solving for \dot{q}^+ in (4) allows us to define the *reset map* $\Delta : \mathcal{S} \rightarrow \mathcal{D}$, i.e. the map describing the discrete event at footstrike [36], as:

$$x^+ = \Delta(x^-) := \begin{bmatrix} Rq^- \\ R(-D^{-1}J_c^\top(J_c D^{-1}J_c^\top)^{-1}J_c + I)\dot{q}^- \end{bmatrix} \quad (5)$$

where the dependence of D and J_c on q^- is suppressed and $R \in \mathbb{R}^{n \times n}$ denotes the relabeling matrix which is used to

maintain state consistency between domains. We are now fully equipped to define the hybrid system of walking as:

$$\mathcal{HC} = \begin{cases} \dot{x} = f(x) + g(x)u, & x \in \mathcal{D} \setminus \mathcal{S}, \\ x^+ = \Delta(x^-), & x \in \mathcal{S}. \end{cases} \quad (6)$$

Note that hybrid systems with a more diverse collection of contact sequences can be modeled via the same framework [3], [37], [38], at the cost of introducing a directed graph describing how the continuous and discrete domains are related – this is omitted in this work for the sake of simplicity.

Trajectory Optimization. The hybrid zero dynamics (HZD) method of gait generation leverages trajectory optimization in the context of the aforementioned hybrid control system to synthesize provably stable walking trajectories (gaits), encoded as nontrivial limit cycles [20].

We begin by defining a collection of $k \in \mathbb{N}$ *outputs* or *virtual constraints* $y_\alpha : \mathcal{Q} \rightarrow \mathbb{R}^k$ which we would like to converge to zero. These virtual constraints encode the desired behavior of the system via:

$$y_\alpha(q) = y^\alpha(q) - y_\alpha^d(\tau(q)), \quad (7)$$

where $y^\alpha : \mathcal{Q} \rightarrow \mathbb{R}^k$ represents the actual (measured) outputs of the system, $y_\alpha^d : \mathbb{R} \rightarrow \mathbb{R}^k$ represents the desired outputs commonly parameterized via a p^{th} -order Bézier polynomial with Bézier coefficients $\alpha \in \mathbb{R}^{k \times p+1}$, and $\tau : \mathcal{Q} \rightarrow [0, 1]$ is a monotonically increasing variable over the gait cycle, termed a *phasing variable*. The HZD framework reduces the stability analysis of the system \mathcal{HC} to a lower-dimensional manifold, the *zero dynamics surface*:

$$\mathcal{Z}_\alpha := \{x \in \mathcal{D} \mid y_\alpha(q) = 0, \dot{y}_\alpha(x) = 0\}, \quad (8)$$

which can be rendered impact-invariant by enforcing the *HZD condition*:

$$\Delta(\mathcal{S} \cap \mathcal{Z}_\alpha) \subset \mathcal{Z}_\alpha. \quad (9)$$

Driving the outputs to zero using a stabilizing controller $u^* : \mathcal{D} \rightarrow \mathbb{R}^m$, for example a feedback linearizing or control Lyapunov function based controller, results in a closed loop dynamical system: $\dot{x} = f_{cl}(x) := f(x) + g(x)u^*(x)$. This stabilizing controller, paired with the HZD condition and an orbit which is stable on the zero dynamics surface renders the closed loop hybrid dynamical system stable. Importantly, the stability of the zero dynamics on \mathcal{Z}_α can be shaped through the choice of outputs and Bézier coefficients α . Therefore, an optimization problem is constructed to synthesize trajectories with desired outputs such that (9) is enforced. This optimization problem takes the form:

$$\{\alpha^*, X^*\} = \underset{\alpha, X}{\operatorname{argmin}} \Phi(X) \quad (10)$$

$$\text{s.t. } \dot{x} = f_{cl}(x) \quad (\text{Closed-loop Dynamics})$$

$$\Delta(\mathcal{S} \cap \mathcal{Z}_\alpha) \subset \mathcal{Z}_\alpha \quad (\text{HZD Condition})$$

$$X_{\min} \preceq X \preceq X_{\max} \quad (\text{Decision Variables})$$

$$a_{\min} \preceq a(X) \preceq a_{\max} \quad (\text{Physical Constraints})$$

In this NLP, $X = (x_0, \dots, x_N, T) \in \mathcal{X}$ is the collection of all decision variables with $x_i \in \mathbb{R}^{2n}$ being the state at

the i^{th} collocation point and $T \in \mathbb{R}$ the total duration, $\Phi : \mathcal{X} \rightarrow \mathbb{R}$ denotes the cost function (such as torque-squared or mechanical cost of transport [3]), and $a(X)$ is the set of physical constraints which includes holonomic constraints, workspace limits, power limits, etc. The solution of (10) is a limit cycle which encodes stable walking, described by some static set of Bézier coefficients $\alpha^* \in \mathbb{R}^{k \times p+1}$.

III. PRELIMINARIES ON SALTATION MATRICES

In preparation for introducing the proposed *robust* HZD method, we will first present some preliminary information on saltation matrices. The saltation matrix is a standard tool used in the field of non-smooth analysis that describes a systems sensitivity to discontinuities (otherwise called ‘saltations’ or ‘jumps’) [39]. Typically, the saltation matrix $S(t_i, x(t_i)) \in \mathbb{R}^{2n \times 2n}$ is defined at time $t_i \in \mathbb{R}^+$ for states in the guard $x(t_i) \in \mathcal{S}$, by the relationship:

$$\delta x(t_{i+1}) = S(t_i, x(t_i))\delta x(t_i), \quad (11)$$

where $x(t_{i+1}) \in \Delta(\mathcal{S})$ denotes the post-transition state. Since we are interested in evaluating the saltation matrix for pre-computed gaits with a known impact time $T_I \in \mathbb{R}^+$, we will specifically derive the saltation matrix $S := S(T_I, x^-)$ for the relationship:

$$\delta x^+ = S\delta x^-, \quad (12)$$

where $\delta x^-, \delta x^+ \in \mathbb{R}^{2n}$ define the variation in the pre- and post- impact states, respectively, at the time of impact T_I . A visualization of these variations are provided in Fig. 2. Explicitly, these state variations are defined as:

$$\delta x^+ := \tilde{x}^+ - x^+, \quad \delta x^- := \tilde{x}^- - x^-, \quad (13)$$

where $\tilde{x}^-, \tilde{x}^+ \in \mathbb{R}^{2n}$ denote the perturbed states. Let $F^-, F^+ \in \mathbb{R}^{2n}$ respectively denote the vector fields of the linearized closed-loop dynamics evaluated at the pre-impact state x^- and post-impact state x^+ , defined as:

$$F^- := f_{cl}(x^-), \quad F^+ := f_{cl}(x^+). \quad (14)$$

As explained in [39], the derivation of the saltation matrix approximates the perturbed states by flowing the system along F^-, F^+ for some duration of time $\delta t \in \mathbb{R}$. Specifically, without loss of generality, the post-impact state can be represented as:

$$x^+ = \Delta(x^-) + F^+\delta t, \quad (15)$$

and the perturbed states as:

$$\tilde{x}^+ := \Delta(\tilde{x}^-), \quad (16)$$

$$\tilde{x}^- := x^- + \delta x^- + F^-\delta t. \quad (17)$$

By substituting these approximations into (13) and taking the first-order Taylor series expansion, we obtain:

$$\delta x^+ = \Delta(x^- + \delta x^- + F^-\delta t) - \Delta(x^-) - F^+\delta t, \quad (18)$$

$$\approx J_\Delta \delta x^- + J_\Delta F^-\delta t - F^+\delta t. \quad (19)$$

Here, $J_\Delta := \frac{\partial}{\partial x} \Delta(x^-) \in \mathbb{R}^{2n \times 2n}$ denotes the Jacobian of the reset map evaluated at the pre-impact state x^- . Note that




Fig. 2. Illustration of a perturbed flow (blue) with uncertain guard conditions (illustrated by the grey region) compared to a nominal periodic orbit (red) which assumes a known guard (illustrated by the black vertical line). The perturbed initial condition $\tilde{x}_0 \in \Delta(\mathcal{S})$ results in pre-impact state error δx^- and post-impact state error δx^+ . In general, saltation matrices capture the relationship between these errors: $\delta x^+ = S\delta x^-$.

because this work exclusively utilizes the reset map shown in (5), we continue the derivation using the assumption that the reset map has no dependence on time, i.e. $\frac{\partial \Delta}{\partial t} = 0_{2n \times 2n}$.

Lastly, we can represent δt in terms of the pre-impact state error δx^- by observing the perturbed guard condition. For generality, we will denote the guard condition as $h : \mathbb{R}^{2n} \rightarrow \mathbb{R}$, but note that in our work we specifically define $h(x) := p_{sw}^z(q)$. Applying a first-order Taylor series expansion to the perturbed guard condition, we obtain:

$$0 = h(\tilde{x}^-), \quad (20)$$

$$\approx \underbrace{h(x^-)}_{=0} + J_h^\top (\delta x^- + F^-\delta t). \quad (21)$$

Here, $J_h := \frac{\partial}{\partial x} h(x^-) \in \mathbb{R}^{2n}$ denotes the Jacobian of the guard condition with respect to x^- . Similar to before, we assume that the guard condition has no dependence on time, i.e. $\frac{\partial h}{\partial t} = 0$; such reset maps are termed *autonomous switching boundary functions*. For information on deriving the saltation matrix for non-autonomous switching boundaries (i.e. $h(x, t)$) refer to [39]. Through the manipulation of (21), we arrive at our expression of δt in terms of δx^- :

$$\delta t = \frac{-J_h^\top \delta x^-}{J_h^\top F^-}. \quad (22)$$

Using this relationship, we can substitute (22) into (19) to obtain the ‘‘traditional saltation matrix’’, $S \in \mathbb{R}^{2n \times 2n}$:

$$\delta x^+ := \underbrace{\left(J_\Delta + \frac{(F^+ - J_\Delta F^-) J_h^\top}{J_h^\top F^-} \right)}_S \delta x^-. \quad (23)$$

Accounting for Guard Uncertainty. Recently, Payne et al. [29] extended the traditional saltation matrix to also account for guard uncertainty. This is accomplished by adapting (20) to also account for perturbations in the guard location along the normal direction, denoted as $\delta_h \in \mathbb{R}$ (shown in Fig. 2):

$$\delta_h = h(\tilde{x}^-) \quad (24)$$

$$\delta_h \approx J_h^\top (\delta x^- + F^-\delta t) \quad (25)$$

$$\delta t = \frac{-J_h^\top \delta x^- + \delta_h}{J_h^\top F^-}. \quad (26)$$

Then, substituting (26) into (19), we arrive at the expression:

$$\delta x^+ = S\delta x^- + \underbrace{\left(\frac{J_\Delta F^- - F^+}{J_h^\top F^-}\right)}_{S_g} \delta h, \quad (27)$$

where S is the same as in (23), and $S_g \in \mathbb{R}^{2n \times 1}$ is termed the *guard saltation matrix*. Together, these matrices can be combined together to obtain the *extended saltation matrix*, $S_e \in \mathbb{R}^{2n+1 \times 2n+1}$, which is defined by the expression:

$$\begin{bmatrix} \delta x^+ \\ \delta h \end{bmatrix} = \underbrace{\begin{bmatrix} S & S_g \\ 0 & 1 \end{bmatrix}}_{S_e} \begin{bmatrix} \delta x^- \\ \delta h \end{bmatrix}. \quad (28)$$

IV. ROBUST GAIT GENERATION

Now that we have presented the preliminaries on the Hybrid Zero Dynamics (HZD) method and saltation matrices, we will discuss how and why we incorporate the saltation matrix evaluation in the HZD framework. First, to address the question of *why* evaluating the saltation matrix improves the robustness of generated gaits, we refer to the field of contraction theory. As first noted by Lohmiller in 1988, discussing stability alone does not capture the behavior of a system relative to a nominal motion [40]. Instead, Lohmiller proposed a new field of analysis – contraction analysis – which explores how trajectories evolve relative to nearby trajectories. Specifically, a system is defined as *contractive* if all trajectories converge to some nominal trajectory. In this work, we leverage the notion of contractivity to define *robust* walking behaviors as those that are more contractive. In other words, when disturbed, a robust gait will converge to a nontrivial periodic orbit faster than a non-robust gait.

Moreover, Burden et al. [41] recently leveraged saltation matrices to induce contractivity of a hybrid system through discrete events. Motivated by this, we propose optimizing the contractivity of discrete events by minimizing the induced matrix norm of the extended saltation matrix; the induced matrix norm is equivalent to the largest singular value of the extended saltation matrix (i.e. $\|S_e\|_2 := \sqrt{\lambda_{\max}(S_e^\top S_e)} \in \mathbb{R}$). Lastly, since the HZD gait generation framework relies heavily on the cost function, we propose adding the induced norm of the extended saltation matrix directly to the cost.

To investigate the influence of including the saltation matrix in the optimization problem, the remainder of this paper compares gaits generated with different weightings of the commonly used cost function torque-squared and the induced matrix norm of the extended saltation matrix:

$$\Phi(X) = w_1 \|U\|_2^2 + w_2 \|S_e\|_2^2. \quad (29)$$

Here, $w_1, w_2 \in \mathbb{R}_{\geq 0}$ denote constant weighting terms, $U \in \mathbb{R}^{m \times N}$ denotes the vectorized torques throughout the nominal gait (assuming a decision variable with $N \in \mathbb{R}$ discretizations) and $S_e \in \mathbb{R}^{2n+1 \times 2n+1}$ denotes the extended saltation matrix evaluated at the pre-impact state of the generated nominal gait. The extended saltation matrix is

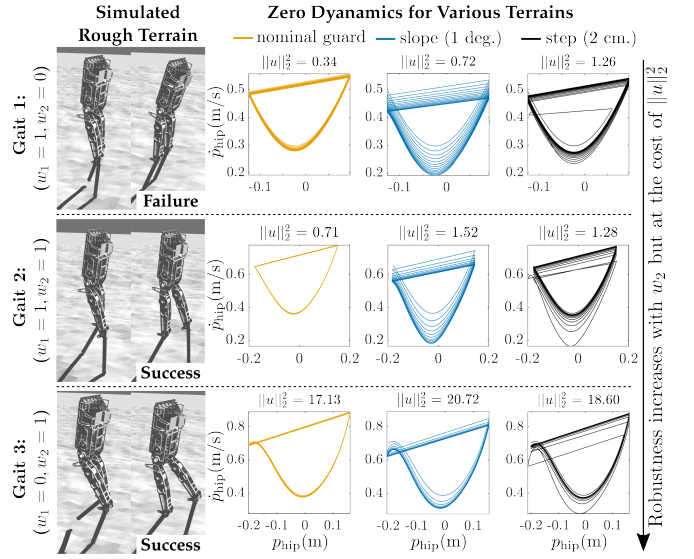


Fig. 3. This figure illustrates the simulated behavior of three gaits for AMBER-3M, where each gait was generated with different weighting terms w_1 and w_2 . The behavior is illustrated using the zero dynamics coordinates of the linearized hip position and velocity ($p_{\text{hip}}, \dot{p}_{\text{hip}} \in \mathbb{R}$), across three different environment conditions: flat ground as captured by the nominal guard (left); 1 degree slope (middle); and 2cm step height (right). The phase portraits show that the robustness of the walking behavior increases as w_2 increases relative to w_1 , but the gait with $w_1 = 0$ results in significantly increased torque.

again computed as in (28), with the traditional saltation matrix and guard saltation matrix explicitly computed as:

$$S = J_\Delta + \frac{(\dot{x}^+ - J_\Delta \dot{x}^-) J_h^\top}{J_h^\top \dot{x}^-}, \quad S_g = \frac{J_\Delta \dot{x}^- - \dot{x}^+}{J_h^\top \dot{x}^-}.$$

To preview the results presented in Sec. V, it was found that increasing w_2 in (29) relative to w_1 improves robustness of the generated gaits, but also results in increased torque (this is illustrated in Fig. 3). This trade-off between performance (characterized by successful implementation on hardware) and robustness (characterized by a systems ability to return to nominal periodic orbits in the presence of disturbances) is further explored in the experimental results.

Lastly, we would like to note that computing the Jacobian of the reset map, J_Δ , can be computationally expensive for high-dimensional systems because the reset map (5) requires several matrix inversions. Thus, for implementation purposes, we numerically approximate the Jacobian of the reset map. However, future work could more efficiently obtain these terms using autodiff or other tools for computing efficient analytical derivatives [42].

V. RESULTS

We demonstrate the application of saltation matrices towards robust gait generation on two robotic platforms: the AMBER-3M planar biped, and the Atalante exoskeleton. As illustrated in Fig. 3, three gaits were generated for each robotic platform: 1) a nominal gait with the cost function equal to torque squared; 2) a gait with the cost function being a weighting of both torque squared and the induced matrix norm of the extended saltation matrix at impact; and 3) a gait with the cost function only including the

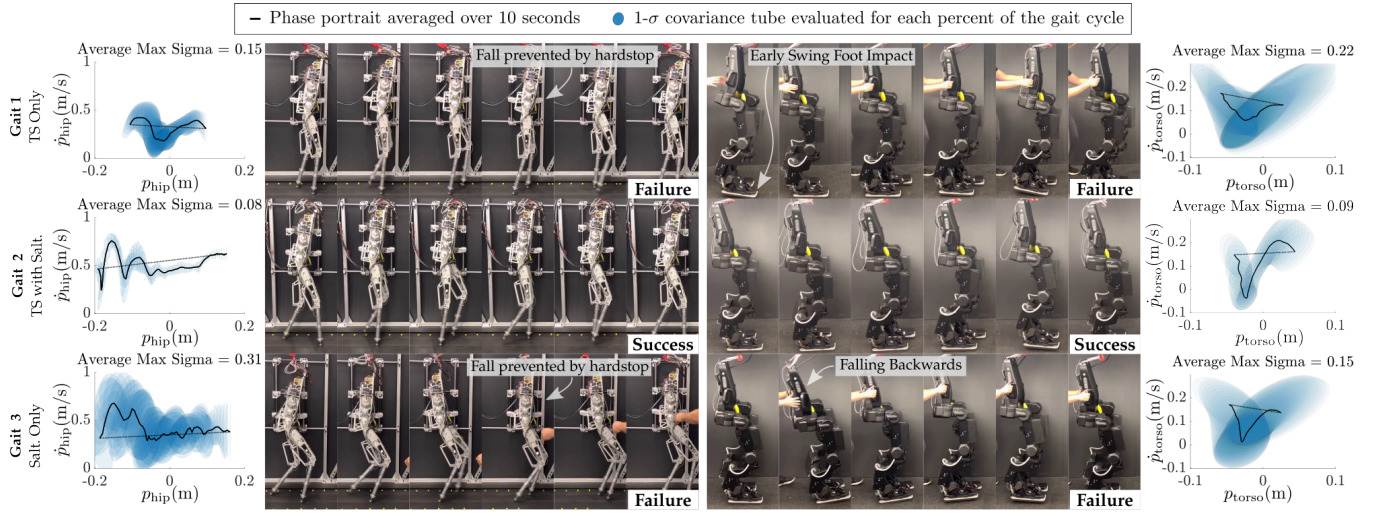


Fig. 4. Gait tiles demonstrating the experimental performance of all three gaits on the AMBER-3M planar biped (left), and the empty Atalante exoskeleton (right). For both platforms, only the gait generated with both torque and the saltation matrix in the cost function ($w_1, w_2 > 0$) was able to sustain stable locomotion. The experimental data is also visualized via phase diagrams of the linearized hip position and velocity ($p_{\text{hip}}, \dot{p}_{\text{hip}} \in \mathbb{R}$) for AMBER-3M and the forward position and velocity of the floating-base frame ($p_{\text{torso}}, \dot{p}_{\text{torso}} \in \mathbb{R}$) relative to the stance foot for the exoskeleton. The black line shows the average zero dynamics across a single step, with the blue region illustrating the 1-sigma tube.

saltation matrix. For AMBER-3M, the three compared gaits were generated with weight values of¹: $w_1 = 1, w_2 = 0$; $w_1 = 1, w_2 = 1$; $w_1 = 0, w_2 = 1$. For Atalante, the weights were selected as: $w_1 = 1e^{-6}, w_2 = 0$; $w_1 = 1e^{-6}, w_2 = 1e^6$; $w_1 = 0, w_2 = 1e^6$. Note that for each robotic platform, the remaining constraints and bounds of the HZD gait generation framework were held constant; only the weights w_1 and w_2 varied. All gaits were generated using the FROST toolbox [43]. The presented experimental results are best demonstrated via the supplemental video [44].

AMBER-3M Results. The AMBER-3M planar biped² is a custom planar robot [45] with four motorized joints (left hip, left knee, right hip, right knee). The measured joint positions are denoted as $q^a \in \mathbb{R}^4$ and are selected as the outputs of the generated gaits, i.e. $y^a(q) := q^a \in \mathbb{R}^4$. The phasing variable $\tau(q) \in \mathbb{R}$ is selected to be the linearized forward hip position, and the desired outputs are described using a 5th-order Bézier polynomial ($\alpha \in \mathbb{R}^{4 \times 6}$). The generated joint-level trajectories are enforced using PD control.

AMBER-3M Simulation Results. The three gaits were first demonstrated in a planar RaiSim [46] simulation environment with randomly generated terrain. As shown in Fig. 3, both gaits generated with the inclusion of the saltation matrix ($w_2 = 1$) were able to walk on rough terrain, while the gait generated with only torque-squared ($w_2 = 0$) failed. The figure also compares three environmental guard conditions: flat ground (as captured by the nominal guard condition), a 1deg slope, and a 2cm step. As illustrated by the phase portraits of the zero dynamics (selected as linearized forward

hip position and velocity), the gaits with $w_2 = 1$ again show improved robustness. However, it is interesting to note that the gait with $w_1 = 0$ suffers from significantly increased torque while the gait with both terms ($w_1, w_2 = 1$) only has a moderate increase in required torque.

AMBER-3M Experimental Results. Once demonstrated in simulation, the three gaits were also demonstrated on hardware, as shown in Fig. 4. The generated joint-level trajectories are enforced on AMBER-3M using an off-board joint-level PD controller that computes desired torques and sends them to the on-board motor controllers via UDP communication. The motor driver communication and control logic run at approximately 1kHz. As shown in Fig. 4, the gait generated using only torque-squared ($w_2 = 0$), and the gait generated using only the extended saltation matrix ($w_1 = 0$) were unstable, while the gait generated with the inclusion of both torque-squared and the extended saltation matrix ($w_1, w_2 > 0$) in the cost function was independently stable.

To further demonstrate the performance of the gait generated with $w_1, w_2 > 0$, several robustness tests were also performed, as shown in Fig. 5. For these experiments, random wooden objects were placed on the treadmill in front of AMBER-3M. These experiments highlight the robustness of the gait generated with the inclusion of both torque-squared and the extended saltation matrix in the cost function.

Exoskeleton Results. The second robotic platform used in this work is the Atalante lower-body exoskeleton – a 3D bipedal platform capable of realizing crutch-less exoskeleton locomotion for people with complete motor paraplegia [24]. In these experiments, sets of three gaits were generated for the exoskeleton with various human-subject models, as well as for the empty exoskeleton. As before, all constraints and bounds of the gait generation framework were held constant except for w_1 and w_2 , and the generated joint-level

¹The purpose of scaling w_1 and w_2 is to ensure that the platform-specific torque-squared term $\|U\|_2^2$ has a similar magnitude to that of $\|S_e\|_2^2$, thus preventing either term from heavily dominating the NLP objective.

²In this work, we specifically utilize the point-foot configuration of AMBER-3M, termed AMBER3M-PF.




Fig. 5. Gait tiles demonstrating robustness of the gaits generated with both torque and the extended saltation matrix in the cost function ($w_1, w_2 > 0$).

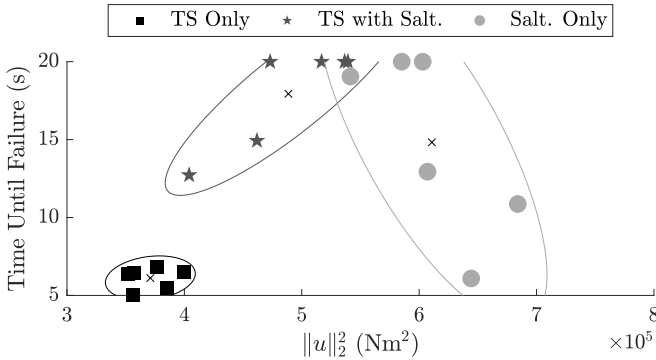


Fig. 6. Simulation results on the Atalante exoskeleton for 6 random exoskeleton subject models. Each marker indicates an individual simulation with the corresponding time until failure (defined as the COM height falling below 0.4 meters) and nominal torque-squared evaluation. Each simulation was limited to 20 seconds total. The results for each gait condition are highlighted by ellipses constructed using 2-sigma fits to the data.

trajectories were enforced using PD control.

The Atalante exoskeleton has 12 motorized joints (3 controlling the spherical position of each hip, 1 for each knee, and 2 in each ankle). As with AMBER, the outputs are again selected as the positions of the motorized joints, i.e. $y^a(q) := q^a \in \mathbb{R}^{12}$. Since the Atalante exoskeleton is fully-actuated when one foot is constrained to remain flat against the ground, gaits are generated using the *Partial Hybrid Zero Dynamics* method [47], an extension of the HZD method. These gaits are described using 7th-order Bézier polynomials ($\alpha \in \mathbb{R}^{12 \times 8}$) and are parameterized using time as the phasing variable. It is important to note that by using time as a phase-variable, the theoretical guarantees of stability are no longer valid. For this reason, many results using periodic orbits on 3D robots also incorporate regulators to stabilize the walking [3]. However, in this work, since we are interested in how *robust* the periodic gaits are independent of regulators, no additional regulators were used (aside from desired output filtering to prevent discontinuities caused by early impacts). The generated joint-level trajectories are enforced on the Atalante exoskeleton using an on-board PD controller which sends current commands to the low-level motor drivers.

Atalante Simulation Results. First, sets of three gaits were generated and deployed for six human models in a 3D simulation environment. As with prior work [48]–[51], a human-exoskeleton model is synthesized by collecting the human’s height, mass, thigh length, and shank length and using this information to approximate the remaining human segment inertias and remaining segment lengths based on anthropomorphic models from [52].

The simulation results, illustrated in Fig. 6, found that the gaits with $w_2 > 0$ resulted in more stable steps being taken before the exoskeleton fell (characterized by the COM vertical height falling below 0.4 meters). However, the gait generated with only the saltation matrix resulted in significantly increased torque. In contrast, using both the saltation matrix and torque-squared in the cost function resulted in increased robustness with only a small increase in torque.

Atalante Experimental Results. Once demonstrated in simulation, a set of three gaits was also generated for the empty Atalante exoskeleton and deployed on hardware. The motivation for conducting experiments with the empty exoskeleton is to isolate the effects of the generated gaits independent of the human-subject’s motion inside of the exoskeleton. The performance of the generated gaits was evaluated by whether or not the exoskeleton could locomote without operator interference for 3 meters. As shown in Fig. 4, only the gait generated with both torque-squared and the saltation matrix ($w_1, w_2 > 0$) was able to successfully and independently walk for the full 3-meter test.

Lastly, a set of gaits was also generated and deployed on the Atalante exoskeleton for a human subject. Again, only the gait generated with the inclusion of both cost function terms ($w_1, w_2 > 0$) was able to successfully complete the 3-meter walk test. These results are shown in Fig. 5 and in the supplemental video [44].

VI. CONCLUSION

This work explored the application of saltation matrices for the generation of robust periodic walking gaits. The results found that by jointly minimizing the norm of the extended saltation matrix and torque, nominal gaits generated using the HZD framework successfully yielded robust locomotion. It is interesting to note that on both platforms (AMBER-3M and the Atalante exoskeleton), gaits generated using only the saltation matrix performed well in simulation, but failed on hardware. This is most likely due to the fact that the saltation only gait result in much higher commanded torques, which leads to undesirable performance on hardware. Thus, future work includes exploring how to systematically balance this performance-robustness trade-off. Ultimately, the results of our work further expands our understanding of the metrics underlying *robust* robotic walking such that in the future, stable and robust nominal gaits can be successfully generated without the need for extensive constraint tuning.

ACKNOWLEDGMENTS

The authors would like to thank the entire Wandercraft team for their continued guidance and technical support with Atalante.

REFERENCES

- [1] V. Paredes and A. Hereid, "Dynamic locomotion of a lower-limb exoskeleton through virtual constraints based ZMP regulation," in *Dynamic Systems and Control Conference*, vol. 84270. American Society of Mechanical Engineers, 2020, p. V001T14A001.
- [2] G. A. Castillo, B. Weng, T. C. Stewart, W. Zhang, and A. Hereid, "Velocity regulation of 3D bipedal walking robots with uncertain dynamics through adaptive neural network controller," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 7703–7709.
- [3] J. P. Reher, A. Hereid, S. Kolathaya, C. M. Hubicki, and A. D. Ames, "Algorithmic foundations of realizing multi-contact locomotion on the humanoid robot DURUS," in *Algorithmic Foundations of Robotics XII*. Springer, 2020, pp. 400–415.
- [4] V. C. Paredes and A. Hereid, "Resolved motion control for 3D underactuated bipedal walking using linear inverted pendulum dynamics and neural adaptation," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 6761–6767.
- [5] M. H. Raibert, "Hopping in legged systems—modeling and simulation for the two-dimensional one-legged case," *IEEE Transactions on Systems, Man, and Cybernetics*, no. 3, pp. 451–463, 1984.
- [6] J. Reher and A. D. Ames, "Dynamic walking: Toward agile and efficient bipedal robots," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, pp. 535–572, 2021.
- [7] D. Kim, J. Di Carlo, B. Katz, G. Bleedt, and S. Kim, "Highly dynamic quadruped locomotion via whole-body impulse control and model predictive control," *arXiv preprint arXiv:1909.06586*, 2019.
- [8] R. Grandia, F. Farshidian, R. Ranftl, and M. Hutter, "Feedback mpc for torque-controlled legged robots," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 4730–4737.
- [9] J.-P. Sleiman, F. Farshidian, M. V. Minniti, and M. Hutter, "A unified mpc framework for whole-body dynamic locomotion and manipulation," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4688–4695, 2021.
- [10] G. García, R. Griffin, and J. Pratt, "Mpc-based locomotion control of bipedal robots with line-feet contact using centroidal dynamics," in *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2021, pp. 276–282.
- [11] P. Fankhauser, M. Hutter, C. Gehring, M. Bloesch, M. A. Hoepflinger, and R. Siegwart, "Reinforcement learning of single legged locomotion," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 188–193.
- [12] Z. Li, X. Cheng, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, "Reinforcement learning for robust parameterized locomotion control of bipedal robots," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 2811–2817.
- [13] J. Siekmann, K. Green, J. Warila, A. Fern, and J. Hurst, "Blind bipedal stair traversal via sim-to-real reinforcement learning," *arXiv preprint arXiv:2105.08328*, 2021.
- [14] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to walk in minutes using massively parallel deep reinforcement learning," in *Conference on Robot Learning*. PMLR, 2022, pp. 91–100.
- [15] K. D. Mombaur, "Stability optimization of open-loop controlled walking robots," Ph.D. dissertation, 2001.
- [16] H. Dai and R. Tedrake, "Optimizing robust limit cycles for legged locomotion on unknown terrain," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*. IEEE, 2012, pp. 1207–1213.
- [17] C. Chevallereau, D. Djoudi, and J. W. Grizzle, "Stable bipedal walking with foot rotation through direct regulation of the zero moment point," *IEEE Transactions on Robotics*, vol. 24, no. 2, pp. 390–401, 2008.
- [18] Z. Xie, G. Berseth, P. Clary, J. Hurst, and M. van de Panne, "Feedback control for cassie with deep reinforcement learning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1241–1246.
- [19] E. R. Westervelt, J. W. Grizzle, and D. E. Koditschek, "Hybrid zero dynamics of planar biped walkers," *IEEE transactions on automatic control*, vol. 48, no. 1, pp. 42–56, 2003.
- [20] J. W. Grizzle, C. Chevallereau, A. D. Ames, and R. W. Sinnet, "3D bipedal robotic walking: models, feedback control, and open problems," *IFAC Proceedings Volumes*, vol. 43, no. 14, pp. 505–532, 2010.
- [21] K. Sreenath, H.-W. Park, I. Poulakakis, and J. W. Grizzle, "A compliant hybrid zero dynamics controller for stable, efficient and fast bipedal walking on mabel," *The International Journal of Robotics Research*, vol. 30, no. 9, pp. 1170–1193, 2011.
- [22] W.-L. Ma, S. Kolathaya, E. R. Ambrose, C. M. Hubicki, and A. D. Ames, "Bipedal robotic running with DURUS-2D: Bridging the gap between theory and experiment," in *Proceedings of the 20th international conference on hybrid systems: computation and control*, 2017, pp. 265–274.
- [23] A. Hereid, C. M. Hubicki, E. A. Cousineau, and A. D. Ames, "Dynamic humanoid locomotion: A scalable formulation for HZD gait optimization," *IEEE Transactions on Robotics*, vol. 34, no. 2, pp. 370–387, 2018.
- [24] O. Harib, A. Hereid, A. Agrawal, T. Gurriet, S. Finet, G. Boeris, A. Duburcq, M. E. Mungai, M. Masselin, A. D. Ames *et al.*, "Feedback control of an exoskeleton for paraplegics: Toward robustly stable, hands-free dynamic walking," *IEEE Control Systems Magazine*, vol. 38, no. 6, pp. 61–87, 2018.
- [25] W.-L. Ma, K. A. Hamed, and A. D. Ames, "First steps towards full model based motion planning and control of quadrupeds: A hybrid zero dynamics approach," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 5498–5503.
- [26] J. W. Grizzle, G. Abba, and F. Plestan, "Asymptotically stable walking for biped robots: Analysis via systems with impulse effects," *IEEE Transactions on automatic control*, vol. 46, no. 1, pp. 51–64, 2001.
- [27] M. Tucker, N. Csomay-Shanklin, W.-L. Ma, and A. D. Ames, "Preference-based learning for user-guided HZD gait generation on bipedal walking robots," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 2804–2810.
- [28] N. J. Kong, J. J. Payne, G. Council, and A. M. Johnson, "The salted kalman filter: Kalman filtering on hybrid dynamical systems," *Automatica*, vol. 131, p. 109752, 2021.
- [29] J. J. Payne, N. J. Kong, and A. M. Johnson, "The uncertainty aware salted kalman filter: State estimation for hybrid systems with uncertain guards," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 8821–8828.
- [30] J. Zhu, N. J. Kong, G. Council, and A. M. Johnson, "Hybrid event shaping to stabilize periodic hybrid orbits," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 01–07.
- [31] J. Lygeros, K. H. Johansson, S. Sastry, and M. Egerstedt, "On the existence of executions of hybrid automata," in *Proceedings of the 38th IEEE Conference on Decision and Control (Cat. No. 99CH36304)*, vol. 3. IEEE, 1999, pp. 2249–2254.
- [32] J. Lygeros, K. H. Johansson, S. N. Simic, J. Zhang, and S. S. Sastry, "Dynamical properties of hybrid automata," *IEEE Transactions on automatic control*, vol. 48, no. 1, pp. 2–17, 2003.
- [33] A. D. Ames, R. W. Sinnet, and E. D. Wendel, "Three-dimensional kneed bipedal walking: A hybrid geometric approach," in *International Workshop on Hybrid Systems: Computation and Control*. Springer, 2009, pp. 16–30.
- [34] Y. Hurmuzlu and D. B. Marghitu, "Rigid body collisions of planar kinematic chains with multiple contact points," *The international journal of robotics research*, vol. 13, no. 1, pp. 82–92, 1994.
- [35] C. Glocker and F. Pfeiffer, "Dynamical systems with unilateral contacts," *Nonlinear Dynamics*, vol. 3, no. 4, pp. 245–259, 1992.
- [36] J. W. Grizzle, C. Chevallereau, R. W. Sinnet, and A. D. Ames, "Models, feedback control, and open problems of 3D bipedal robotic walking," *Automatica*, vol. 50, no. 8, pp. 1955–1988, 2014.
- [37] R. W. Sinnet and A. D. Ames, "2d bipedal walking with knees and feet: A hybrid control approach," in *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*. IEEE, 2009, pp. 3200–3207.
- [38] H. Zhao, A. Hereid, W.-l. Ma, and A. D. Ames, "Multi-contact bipedal robotic locomotion," *Robotica*, vol. 35, no. 5, pp. 1072–1106, 2017.
- [39] R. I. Leine and H. Nijmeijer, *Dynamics and bifurcations of non-smooth mechanical systems*. Springer Science & Business Media, 2013, vol. 18.
- [40] W. Lohmiller and J.-J. E. Slotine, "On contraction analysis for nonlinear systems," *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.
- [41] S. A. Burden, T. Libby, and S. D. Coogan, "On contraction analysis for hybrid systems," *arXiv preprint arXiv:1811.03956*, 2018.
- [42] S. Singh, R. P. Russell, and P. M. Wensing, "Efficient analytical derivatives of rigid-body dynamics using spatial vector algebra," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1776–1783, 2022.

- [43] A. Hereid and A. D. Ames, "FROST: Fast robot optimization and simulation toolkit," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 719–726.
- [44] Supplementary video, <https://youtu.be/BZu-9UStG2E>.
- [45] E. Ambrose, W.-L. Ma, C. Hubicki, and A. D. Ames, "Toward benchmarking locomotion economy across design configurations on the modular robot: AMBER-3M," in *2017 IEEE Conference on Control Technology and Applications (CCTA)*. IEEE, 2017, pp. 1270–1276.
- [46] J. Hwangbo, J. Lee, and M. Hutter, "Per-contact iteration method for solving contact dynamics," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 895–902, 2018. [Online]. Available: www.raisim.com
- [47] A. D. Ames, "Human-inspired control of bipedal walking robots," *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1115–1130, 2014.
- [48] T. Gurriet, S. Finet, G. Boeris, A. Duburcq, A. Hereid, O. Harib, M. Masselin, J. Grizzle, and A. D. Ames, "Towards restoring locomotion for paraplegics: Realizing dynamically stable walking on exoskeletons," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 2804–2811.
- [49] T. Gurriet, M. Tucker, A. Duburcq, G. Boeris, and A. D. Ames, "Towards variable assistance for lower body exoskeletons," *IEEE Robotics and Automation Letters*, vol. 5, no. 1, pp. 266–273, 2019.
- [50] M. Tucker, E. Novoseller, C. Kann, Y. Sui, Y. Yue, J. W. Burdick, and A. D. Ames, "Preference-based learning for exoskeleton gait optimization," in *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 2351–2357.
- [51] M. Tucker, M. Cheng, E. Novoseller, R. Cheng, Y. Yue, J. W. Burdick, and A. D. Ames, "Human preference-based learning for high-dimensional optimization of exoskeleton walking gaits," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 3423–3430.
- [52] D. A. Winter, *Biomechanics and motor control of human movement*. John Wiley & Sons, 2009.