

Optimization-Based Control for Dynamic Legged Robots

Patrick M. Wensing¹, Michael Posa², Yue Hu³, Adrien Escande⁴, Nicolas Mansard⁵, Andrea Del Prete⁶

Abstract—In a world designed for legs, quadrupeds, bipeds, and humanoids have the opportunity to impact emerging robotics applications from logistics, to agriculture, to home assistance. The goal of this survey is to cover the recent progress toward these applications that has been driven by model-based optimization for the real-time generation and control of movement. The majority of the research community has converged on the idea of generating locomotion control laws by solving an optimal control problem (OCP) in either a model-based or data-driven manner. However, solving the most general of these problems online remains intractable due to complexities from intermittent unidirectional contacts with the environment, and from the many degrees of freedom of legged robots. This survey covers methods that have been pursued to make these OCPs computationally tractable, with specific focus on how environmental contacts are treated, how the model can be simplified, and how these choices affect the numerical solution methods employed. The survey focuses on model-based optimization while paving its way for broader combination with learning-based formulations to accelerate progress in this growing field.

I. INTRODUCTION

Over the past decade, we have witnessed rapid growth in the capabilities of legged robots, transitioning from a state of the art where only a few research groups had access to capable platforms, to one where robust locomotion is now common in industry and academic laboratories. This rapid progress has been enabled by a combination of advances across design and control, with optimization-based control strategies playing a central role in many of the milestone demonstrations during this period. Through these advances, current bipeds, quadrupeds, and humanoids can now walk reliably in nominal environments, and these improved capabilities have led to the first practical deployments of legged systems (e.g., the robot Spot from Boston Dynamics). In a world built for legs, the growing scope of these deployments offers a broad opportunity for impact on applications spanning logistics (e.g., delivery), agriculture, and home assistance, among many others.

¹University of Notre Dame: pwensing@nd.edu

²University of Pennsylvania: posa@seas.upenn.edu

³University of Waterloo: yue.hu@uwaterloo.ca

⁴Inria centre at the University Grenoble Alpes: adrien.escande@inria.fr

⁵LAAS-CNRS: nmansard@laas.fr

⁶University of Trento: andrea.delprete@unitn.it

This work was supported by NSF award CMMI-2220924 (Wensing), NSF award CMMI-1830218 (Posa), Natural Sciences and Engineering Research Council of Canada (NSERC) funding reference number RGPIN-2022-03857 (Hu), the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101066915 (Escande), and PRIN Project DOCEAT under Grant CUP E63C22000410001 (Del Prete).

© 2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

DOI 10.1109/TRO.2023.3324580




Fig. 1. Optimization-based control has been a common enabler for legged systems to handle complex environments [3], [10], [34], [58]–[60].

Across these applications, optimization-based control strategies offer many advantages as a pathway to (or component of) capable autonomy. Two prevailing types of optimization-based control have been pursued: predictive and reactive. Predictive controllers consider an explicit system model to “reason” (in a strictly mathematical sense) about the consequences of their actions, iteratively devising and improving motion plans in response to the situation at hand. Reactive controllers, by contrast, only consider their actions for the current instant. The execution of these controllers may depend on the system model, or may be model-free (e.g., the result of offline policy optimization). The benefits of the model-based approach are that it allows one to naturally consider safety constraints, which will be critical for many of the envisioned applications for legged robots. An additional motivation is that first-principles models generalize well to unforeseen situations, which will be an asset for flexible autonomous deployments. There is also a growing interest in extending model-based methods by augmenting or complementing them with learning-based approaches, e.g., learning hard-to-model aspects of the physics, accounting for perception, or addressing higher-level aspects of decision making that are not amendable to physics-based reasoning. With this big picture in mind, a central goal of the review is to synthesize the past decade of model-based advances so that they can be further improved in the future and also incorporated with learning-based strategies—ultimately accelerating the practical deployment of legged robots.

A. Historical Context for the Survey

Much of the recent progress in optimization-based control for robots rests on decades of progress from the mathematical programming community and its early translation to our field. The idea of using optimization for the specification and control of robot movement traces back to decades before its current

popularity (e.g., [61]). The 1990s saw the use of quadratic programming for torque control of redundant manipulators [62], with the 2000s seeing generalizations of these ideas for the control of humanoids in simulation [63]–[65]. The advent of polynomial-time solutions to important classes of convex optimization problems in the mid-1990s [66] and subsequent commercial solvers (e.g., [67]) has contributed to the rise in maturity and accessibility of optimization solutions for control. Building on these advances, work in the late 2000s and in the wave of research before the 2015 DARPA Robotics Challenge (DRC) focused on convex optimization for maintaining balance via reactive control (e.g., [43], [68], [69]). The role of optimization in these accomplishments has since led to recent work on pushing the boundaries of predictive control for quadrupeds and humanoids executing broader dynamic locomotion in more challenging environments.

B. Survey Goals

The overarching goal of this survey is to provide a synthesized entry point into this recent body of work. As such, our intention is for the survey to serve as the most effective resource for early-stage graduate students, while also providing new perspectives and targeted follow-on reading for established experts. Work since the DRC has led to many new methods for optimization through contacts, for real-time optimization using new simplified models, and through advances to numerical methods for handling ever more complicated robot models. These advances are the primary focus of the review.

Beyond covering this past work, a secondary motivation for this survey is to provide the background for the next steps in legged locomotion research, enabling these systems to move beyond nominal environments and unlock mobility in unstructured terrains. A great deal of current research interest focuses on how to combine machine learning strategies with previously developed model-based ones. In this sense, the survey will provide valuable background for machine learning practitioners to come up to speed on existing model-based approaches. Other research continues to push forward numerical methods and algorithms for legged robots to make informed model-based decisions. The survey should serve as a valuable overview of recent work for these groups as well.

II. PROBLEM STATEMENT AND OVERVIEW

During the last decade, most of the research community working on legged robots (humanoids and quadrupeds in particular) has converged to the idea of generating motion based on the formulation of an Optimal Control Problem (OCP). In a high-level form, such an OCP can be written as:

$$\underset{x(\cdot), u(\cdot), \lambda(\cdot)}{\text{minimize}} \quad \text{Cost}(x(\cdot), u(\cdot), \lambda(\cdot)) \quad (1a)$$

$$\text{subject to} \quad M(q)\dot{\nu} + C(q, \nu)\nu + \tau_g(q) = S^T\tau + J_c(q)^T\lambda \quad (1b)$$

$$\text{ContactConstraints}(x(t), \lambda(t), u(t), \text{Env}) \quad (1c)$$

$$\text{KinematicsConstraints}(x(t)) \quad (1d)$$

$$\text{InputConstraints}(u(t), x(t)) \quad (1e)$$

$$\text{TaskConstraints}(x(t), u(t), \lambda(t)) \quad \forall t. \quad (1f)$$

The (infinite-dimensional) decision variables of the OCP are the *trajectories* of state $x = (q, \nu)$, control $u \triangleq \tau$, and contact forces λ exchanged between the robot and the environment Env. The robot configuration is represented as q , whereas its velocity as ν . Eq. (1a) is a user-defined cost function representing a metric to minimize, such as energy consumption or distance to a desired target. Eq. (1b) represents the nonlinear whole-body robot dynamics [70], where $M(q)$ is the mass matrix, $C(q, \nu)\nu$ accounts for Coriolis and centrifugal forces, $\tau_g(q)$ contains the gravity forces, S is typically a matrix that selects the actuated degrees of freedom, and $J_c(q)$ is the contact Jacobian. Eq. (1c) encodes the contact-related constraints, such as non-penetration of rigid objects and friction-cone force constraints. Eq. (1d) includes state constraints resulting from the kinematics, such as joint position and velocity limits. Eq. (1e) typically includes the motor torque limits, but it could in general represent any input constraint. Finally, (1f) can be used to include any task-specific constraint in the OCP, such as fixed initial/final state, or field-of-view limits for a camera.

Clearly, this OCP is very general and therefore allows one to encode a wide variety of movements. For instance, point-to-point locomotion can be generated by specifying the initial and final robot states [29]. An object manipulation behavior could instead be computed by including the state of the object in x and specifying its desired value with an appropriate cost function [32], [71], [72]. While appreciating its versatility, we should acknowledge that some intrinsic limitations do exist in the OCP approach. For instance, it is hard to ensure stability [29] or robustness for the computed trajectories [39]. However, this review focuses on another limitation and how to address it: the computational complexity of the problem. Indeed, problem (1) hides several challenges, the main ones being: non-smoothness/stiffness of the dynamics, nonconvexity, and dimensionality. The rest of this section acts as an executive summary of the rest of the paper by briefly discussing different approaches to tackle these challenges. Each aspect is analyzed more thoroughly in Sections III–VI.

A. Contact Models

The first issue (non-smoothness/stiffness), which is discussed in detail in Section III, arises from physical contacts modeled in (1c). Contacts can be modeled as either *rigid* or *visco-elastic*. With appropriate choices, visco-elastic models ensure continuous dynamics, and smoothing techniques can be used to make them differentiable. The downside of visco-elastic models is that large stiffness values are necessary to generate realistic behaviors. This feature leads to stiff differential equations with corresponding numerical challenges for simulation and optimization.

Alternatively, contacts can be modeled as rigid (i.e., no penetration is allowed). When two points make contact, their relative velocity must immediately become zero to avoid penetration. Therefore, the robot dynamics must be described by a mix of continuous-time and discrete-time equations, i.e., as a hybrid dynamic system. The resulting OCP can then be tackled either as a Linear Complementarity Program (LCP) [10], [13], [73] or as a Mixed Integer Program (MIP) [21], both of which

require customized optimization techniques that are typically much less efficient than classic smooth optimization strategies.

A common way around this non-smoothness is to let the user fix the order in which contacts are made and broken. This makes the dynamics time-switched (a special case of hybrid dynamics) and the OCP differentiable, therefore efficient smooth optimization can be used. The obvious downside is that it may be hard to guess the contact phases.

B. Dynamic Models

Besides the contact dynamics, the other major source of complexity is the robot multi-body dynamics (1b), which is high-dimensional and nonlinear. In turn, this makes the resulting optimization problem high-dimensional and non-convex. The high dimension is especially concerning in the context of online optimization, where fast computation times are mandatory. The non-convexity instead is always concerning, as it makes the solver sensitive to the initial guess that is used. As partial or total remedy to these issues, several *simplified models* have been proposed in the literature (also known as reduced-order models, or template models), and they are discussed in Section IV. These models should capture the most important part of the robot dynamics with a reduced state size. For example, for locomotion, the widely-used Linear Inverted Pendulum (LIP) model [47] considers only the contact locations and the Center of Mass (CoM) of the robot, neglecting the details of joint angles and velocities. While simplified models are key enablers for fast online computation, their simplifying assumptions (e.g., constant CoM height), or neglected constraints [74] (e.g., joint position and torque bounds) can severely limit the generated motions.

C. Optimal Control Solution Methods

After the OCP is mathematically formulated, there are a range of design choices that remain for solving it. General optimal control methods can be categorized as direct methods, indirect methods, or those based on dynamic programming. Within robotics, direct optimization methods prevail almost universally. However, even after choosing to use a direct method, the optimization formulation needs to be approximated as a finite-size nonlinear program to be solved by a numerical solver. This approximation process is called *transcription* and it crucially affects the results in terms of accuracy, numerical stability, and computational complexity.

Two families of transcription methods have been used in robotics (shooting and collocation) and we discuss them in Section V. A shooting method known as Differential Dynamic Programming (DDP) has been a topic of recent interest in the community due to its favorable computational properties and provision of a locally optimal feedback policy. More generally, but specific to direct optimization for locomotion, we examine how contact modeling choices affect the choice of transcription tools and algorithms available, which becomes most critical when optimization is left to freely choose contact sequences.

Table I serves to categorize the choices that state-of-the-art papers have used in adopting different approaches for modeling the contacts and their sequencing, for simplifying

TABLE I
OVERVIEW OF THE STATE OF THE ART.

Papers	Contact model	Gait	Dynamics	Transcription
[1]–[5]	Soft/Smoothed	Optimized	Full	DDP
[6]	Soft/Smoothed	Optimized	Full	Collocation
[7]	Soft/Smoothed	Optimized	Simplified	Collocation
[8], [9]	Rigid	Optimized	Full	DDP
[10]–[15]	Rigid	Optimized	Full	Collocation
[16]–[24]	Rigid	Optimized	Simplified	Collocation
[25]–[28]	Rigid	Fixed	Full	DDP
[29]–[31]	Rigid	Fixed	Full	Multiple Shooting
[32]–[34]	Rigid	Fixed	Full	Collocation
[26], [35]	Rigid	Fixed	Simplified	Multiple Shooting
[36]–[38]	Rigid	Fixed	Simplified	DDP
[39]–[46]	Rigid	Fixed	Simplified	Collocation
[47]–[57]	Rigid	Fixed	Simplified	Single Shooting

the dynamics, and, ultimately, for numerically solving variants of (1). In practice, the highlighted approaches play a key role in breaking problem (1) into many smaller subproblems that focus on building an overall solution in a hierarchical fashion (e.g., solving for footsteps first, then optimizing motions with fixed footsteps).

D. Trajectory Stabilization via Whole-Body Control

Depending on the complexity of the formulated OCP, the computation time may be too large for it to be solved inside a fast control loop for Model Predictive Control (MPC). In these cases, a reactive stabilizing controller is required to execute the computed motion on real hardware, or to provide control of aspects of the system that were ignored during trajectory optimization (e.g., due to modeling simplifications). In the last decade, the legged robotics community has converged on a certain class of reactive optimization-based whole-body control techniques, which mainly rely on the fast solution of small convex Quadratic Programs (QPs) to compute motor commands as a function of state feedback.

These QP techniques have represented an evolution of previous operational-space control paradigms for the control of manipulators. However, for legged systems, new problems related to contact constraints, impacts, self-collisions, etc., have motivated a broader perspective on this classical problem. The most general extensions have focused on treating constraints via task-space inequalities. In spite of many extensions, the formulation of reactive control as an instantaneous control strategy has enabled the method to retain a convex formulation that is lost in more general transcription strategies. The convex-optimization perspective on reactive control has also enabled other paradigms [43], [75]–[77], conventionally separate from operational-space control, to be incorporated as well. We describe these techniques in Section VI, while also reflecting on recent advances in whole-body MPC.

III. CONTACT

Computationally tractable treatment of physical, frictional contact with the environment, which leads to stiff and/or discontinuous equations of motion, is a fundamental challenge in optimization-based control for robotics and a primary distinction between (1) and standard optimal control problems. As a result, design decisions for (i) how to model the effects of contact and (ii) the sequencing or scheduling of contact events

play a significant role in distinguishing different approaches to solving (1). Specifically, the details in this section center on how the choice of contact model affects parameterizations of the force λ and the form of the contact constraints in (1c).

This section is comprised of two parts. In Section III-A, we outline common techniques and associated numerical challenges for numerical models of contact dynamics. In Section III-B, we describe how the need to sequence contact events affects the algorithmic and modeling choices.

A. Modeling contact

Highly accurate mechanics models define contact as *visco-elastic*, representing local deformations of surfaces and the corresponding stress-strain relationships. Some approaches in robotics, particularly in soft robotics (e.g., [78]), do attempt to accurately represent deformation, which prevents interpenetration between bodies. Rigid-body approximations, by contrast, do not consider deformation but do permit some penetration between contacting bodies. The resulting contact forces are modeled as a function of penetration depth, relative normal and tangential velocities, and material properties. Rigid-body models are most commonly used and will be the focus of this review. Noting critical definitions here, we write $\phi(q)$ to be the vector-valued signed distance function, for all possible contacts. $J_n(q)$ and $J_t(q)$ are the normal and tangential Jacobians, such that $J_n(q)v$ and $J_t(q)v$ are the contact frame velocities. Normal and tangential forces are similarly decomposed into λ_n and λ_t , and, where convenient, we use $J_c(q)$ and λ to represent stacked contact Jacobians and forces. The contact force, therefore, can be expressed as

$$\lambda = F_{\text{contact}}(\phi, J_c v).$$

Common methods treat this function as a nonlinear spring and damper (e.g., [79], [80]), where unilateral constraints determine when the spring releases; typically, contact forces cannot “pull”, constraining $\lambda_n \geq 0$, and so separation occurs if $\lambda_n < 0$. However, to accurately represent the interaction between rigid bodies, F_{contact} is numerically stiff [81], inheriting the mechanical stiffness of the material properties in the robot. When making and breaking contact ($\phi = 0$), F_{contact} may also be non-differentiable. Frictional effects exhibit similar issues: in dry Coulomb friction, λ_t depends on the sign of the tangential velocity, which is discontinuous.

Simply introducing these stiff or discontinuous dynamics to (1) leads to a poorly conditioned optimization problem and is generally avoided. Algorithmic approaches, therefore, tend to focus on one of two choices: (i) *infinitely* stiff representations, as hybrid models, or (ii) smoothing or softening F_{contact} to improve numerical performance. An illustrative example of a single contact is shown in Fig. 2, demonstrating the resulting state trajectory and, critical for optimization-based control, the sensitivity of the solution with respect to initial conditions.

1) *Hybrid dynamics*: In the limit of infinite stiffness, when contact is initiated, the resulting forces become *impulsive*; impacts, therefore, cause an instantaneous jump in velocity. While a full description of hybrid systems is outside the scope of this review (see [84], [85] for an overview), we briefly note





Fig. 2. In an illustrative conceptual example, a particle heads toward a wall, with inelastic impact when $q = 0$. Three common contact models are compared, showing (top-to-bottom) position, velocity, and sensitivity of velocity to initial conditions. (Blue) A Hunt-Crossley spring model [82] activates when contact occurs at $q(t) = 0$, allowing some penetration. The sensitivity $\partial v(t)/\partial q_0$ is well-defined, though not smooth (C^∞) at the impact. (Orange) a relaxed (softened) complementarity model, like MuJoCo, is smooth, with contact forces occurring prior to contact [83]. (Green) a hybrid model with inelastic impact captures rigid behavior with an instantaneous jump in velocity, noting that $\partial v(t)/\partial q_0$ is undefined at impact. All models have well-defined gradients *away* from the impact event, enabling effective differentiation in these regimes. For the hybrid and spring model, these gradients are zero prior to contact; for the relaxed model, as it is smooth, these gradients are non-zero everywhere.

the principles here. Hybrid systems are defined by *modes*, *guards*, and *resets*. In this context, the mode is the contact state: identifying which objects are touching, and whether they are sticking or sliding. Guards determine when mode transitions occur, for instance making or breaking contact, and the reset map R specifies the result of the transition (e.g., impact) $x^+ = R(x^-)$. Within a given mode, the dynamics are well-defined and differentiable; thus, this formulation compresses the complexity of contact to the guards and transition events. Referring to the root optimal control problem (1), a hybrid formulation typically explicitly splits the decision variables $x(\cdot), u(\cdot), \lambda(\cdot)$ by mode, introducing boundary constraints at the resets when transcribing the dynamics in Section V.

2) *Complementarity models*: A mathematically equivalent formulation of hybrid contact dynamics relies upon *complementarity constraints* [86], [87], where a relationship between λ and x implicitly defines the hybrid dynamics. For example, a non-penetration condition $\phi(q) \geq 0$, combined with $\lambda_n \geq 0$, and the property that forces can only be non-zero when in contact leads to the complementarity constraint

$$0 \leq \lambda_n \perp \phi(q) \geq 0,$$

where $\lambda_n \perp \phi(q)$ means that the vectors λ_n and $\phi(q)$ must be orthogonal, which is equivalent to $\lambda_n^T \phi(q) = 0$. This formulation can be expressed in time-stepping [73] or

continuous [88] models, with a similar set of complementarity constraints used to capture Coulomb friction.

3) *Pathologies*: Hybrid formulations of multi-contact robotics, while effective for simulation and control, can exhibit certain pathologies where solutions do not exist or where infinite solutions are possible [86], [89]. These challenges are the limiting case of the high sensitivity to initial conditions seen in the stiff differential equations. From the perspective of optimization-based control, stiff dynamics leads directly to poorly conditioned optimization problems. Alternatively, for instance when using hybrid or complementarity models, planning algorithms that assume uniqueness therefore implicitly (perhaps even unbeknownst to the algorithm designer) select from the set of possible solutions (a property, for instance, of [10]). In many scenarios, this selection is benign; in others, it indicates that the resulting trajectory is practically impossible to track, due to the non-uniqueness. Non-uniqueness can occur in many scenarios; most commonly, if $J_c(q)$ is rank deficient. Think, for example, of static indeterminacy with a four-legged table, where the normal force distribution is indeterminate.

4) *Hybrid differentiability*: In the hybrid formulation, while both the resulting trajectories $x(t)$ and the equations of motion may be discontinuous, under certain circumstances it is possible to generate well-conditioned derivatives of $x(t)$ with respect to initial conditions and control actions [90], [91] (see Fig. 2 for an example). Typically, these derivatives exist if the sequence of modes is constant; the event times themselves might change, but the order does not and all transitions are inevitable. When deviations might change the mode sequence, resulting trajectories are often *not* differentiable with respect to initial conditions or parameterizations, but may admit more general sensitivity characterizations [90]. Practically, if a trajectory $x(t)$ does not make contact, then no amount of local differentiation provides insight into the effects of initiating contact; put more explicitly, solutions $x(t)$ are non-analytic in their initial conditions and control inputs. This property poses natural challenges for optimizers to discover new contacts through local information alone, and this motivates many relaxations or other tailored strategies for scheduling contacts.

To mitigate this issue, certain simulators ensure global differentiability or even smoothness (termed *differentiable simulation*, e.g., MuJoCo [83] or TDS [92]). While these methods do consistently generate local gradients, this inevitably ties the accuracy (stiffness) of the underlying model to stiffness in the resulting optimization problem. The precise nature of this trade-off and the relevance of the different modeling inaccuracies (see, e.g., [93], [94]) remain unknown.

B. Scheduling contact

As the discussion of differentiability above implies, there are clear distinctions in computational tractability between versions of (1) with a fixed, known mode sequence and variations where the optimization problem must also determine the ordering. In many robotics problems, this sequence may be clear; for example, bipedal walking over flat terrain typically follows a “left foot, right foot” ordering with minimal deviation. In others, for instance, movements using whole-body contact

or locomotion over varied surfaces with multiple potential foothold locations, the challenge of finding an optimal ordering may dominate the control problem.

1) *Known modes sequences*: In this setting, once discretized (see Section V), the hybrid optimal control problem is differentiable with respect to the transcribed decision variables. In some settings, particularly when deploying simplified models or when seeking stability proofs, it is common to use *minimal coordinates* to represent the configuration of a robot in contact (e.g., [95], [96]). Contact points are transformed into pin joints, removing degrees of freedom and simplifying the resulting optimal control problem. When contact modes change, however, the hybrid jump must also capture the change in dimension of the state space. The use of minimal coordinates is computationally efficient, although it is difficult to impose constraints (e.g., friction) on λ , as inverse dynamics are necessary to reconstruct the constraint forces. Furthermore, in multi-contact settings, where the contacts generate a closed chain, global definitions of minimal coordinates may not exist.

The more general formulation utilizes the same excess, or *floating-base coordinates* for all contact modes. The constrained dynamics then enforce that active contacts remain touching [29], [97]. Some approaches directly compute the force, e.g., $\lambda(x, u)$ [97]. This approach has the advantage of relative simplicity, though heuristics are necessary when the rigid-body force is non-unique. Alternate methods include $\lambda(t)$ as a decision variable (see (1)), along with corresponding constraints to ensure physical accuracy. Imposing constraints *differentially* in this fashion comes at the risk of constraint drift due to integration error, typically addressed via implicit integration schemes [33], Baumgarte stabilization [27], or a mixture of approaches [98].

2) *Hybrid sequence optimization*: As an immediate extension of the known-sequence scenario, one could jointly optimize over the discrete states of the hybrid sequence and the corresponding robot motion. This can naturally be expressed via *mixed-integer optimization* [17], [21], or as a bilevel optimization problem [16]. Mixed-integer programs (MIPs) have been used to compute contact sequences accounting for obstacle avoidance and step-to-step reachability [99] and approximated as computationally efficient L1-norm minimization (SLIM) [100]. Other strategies ensure “quasi-static” feasibility by limiting the search to “quasi-flat” contact surfaces (i.e., surfaces where the friction cone contains the gravity direction) [100]. In the worst case, these methods must explore every possible mode sequence, and so are most effective when there are relatively few potential sequences or when effective heuristics are available to guide the search.

Sampling-based techniques are an alternative to MIPs. For instance, rapidly-exploring random trees (RRTs) have been used to plan footsteps on flat ground, avoiding obstacles [101]. Similarly, probabilistic road maps (PRMs) have been used to plan a collision-free path for the robot’s base, keeping it close enough to the environment to allow for contact creation [102].

3) *Contact-implicit planning*: Contact-implicit (alternatively, contact-invariant) methods are variations on hybrid optimization that embed the relationship between state and force into a nonlinear program, implicitly representing the hybrid




Fig. 3. The centroidal dynamics is a common modeling simplification that focuses on the CoM and net momentum. It is a special modeling simplification in that it exactly projects the whole-body dynamics, while relaxing joint-space constraints. Restrictions placed on this model then lead to a tree of other models.

mode without discrete variables. Approaches here are most often based on either the complementarity formulation [10], [103] or on smooth approximations of the contact dynamics [3], [7]. Prior to convergence, these methods typically violate strict complementarity, thus simultaneously exploring multiple contact sequences. Contact-implicit optimization, however, can suffer from poor numerical conditioning and can require high-quality initial guesses. Contact-implicit methods have found applications in many robotic domains, including MPC [1] and gait optimization for microrobots [14]. Recent work has focused on improving numerical performance, for example by using simplified models [18], [20] or improving the accuracy of the numerical integration schemes [12], [13]. While these methods are typically limited to offline computation, progress has been made in real-time scheduling of contact [15], [23].

These previous contact-implicit approaches explicitly account for the hybrid nature of planning through contact. Alternatively, gradient-based methods, like iLQR or DDP (see Section V-C), can be applied. This is most common when coupled with a differentiable contact model, enabling the optimization method to “discover” new modes (when it happens to bump into them, perhaps guided by the smoothing in the differentiable model) [1], [5]. Other methods maintain a rigid model, and consider gradient analyses around the contact sequence found during forward simulation (e.g., [8], [9]), leveraging the almost everywhere differentiability discussed above. Broadly speaking, these methods are typically highly sensitive to their initializations and struggle to discover contact sequences that vary dramatically from that of the initial guess.

C. Summary

When the sequence of environmental contacts can be known *a priori*, then the choice of contact model should focus purely on physical realism, as nearly any transcription method can be applied without disruption. It is substantially more difficult, however, to simultaneously plan motion and contact schedule. While significant progress has been made over the last decade, contact planning remains the main challenge for generating arbitrary locomotion behaviors in complex environments. Ultimately, this challenge arises from our reliance on gradient-based optimization, which despite being the workhorse behind

the progress discussed in this survey, is fundamentally unsuited for non-smooth contact-implicit problems.

1) *Relationship to learning*: While this review focuses on model-based optimization, there is an implicit requirement that such a model must first be identified or learned. Methods for learning and identification of contact align closely with the modeling choices detailed in this section. When contact modes are known or can be easily identified for the training data, hybrid approaches are commonly used [93]. Otherwise, machine learning can be directly applied to a smoothed or differentiable contact model (e.g., [104]). Alternatively, contact-implicit methods have led to data-efficient mechanisms for contact model learning [105]. Aside from identifying *models*, machine learning has also shown promise in identifying (or providing high-quality guesses for) potential mode sequences for control [106], [107], where an effective pre-solve for the mode schedule substantially reduces the computational difficulty of the subsequent trajectory optimization. We note, as well, that *model-free* reinforcement learning (RL), which has been widely applied to similar problems in control (e.g., legged locomotion [108]), also commonly leverages smoothed (or stochastic [109]) contact models.

IV. SIMPLIFIED MODELS

A second source of complexity for the efficient solution of (1) arises from the whole-body dynamics model (1b). As noted earlier, the dynamics of legged systems are high-dimensional and nonlinear, and these features correspondingly make the optimization problem high-dimensional and nonconvex. These challenges motivate simplified models in place of (1b) that capture its most salient features in a reduced set of differential equations. A central question is then how to select such a simplified model. We use the evolution in Fig. 3 to walk through the most common modeling simplifications adopted.

The role that contacts play within the whole-body dynamics (1b) has motivated the majority of simplified models employed to date. For a floating-base system, it is common to partition the generalized velocity as $\nu = (\nu_b, \nu_j)$ where $\nu_b = (\omega_b, v_b) \in \mathbb{R}^6$ gives the angular and linear velocity of

the floating base and ν_j the generalized velocity of the joints. With this partitioning in ν , the dynamics are partitioned as

$$\begin{bmatrix} M_{bb} & M_{bj} \\ M_{jb} & M_{jj} \end{bmatrix} \begin{bmatrix} \dot{\nu}_b \\ \dot{\nu}_j \end{bmatrix} + C(q, \nu)\nu + \tau_g(q) = \begin{bmatrix} 0 \\ \tau \end{bmatrix} + J_c(q)^T \lambda. \quad (2)$$

Suppose now that the robot is given a desired motion $(q_d(t), \nu_d(t))$. Considering (2), the left-hand side of the equation is fixed along the desired motion, and so the motion is possible only when the equation can be satisfied through proper choice of joint torques τ and contact forces λ . If the robot has sufficiently strong actuators, then τ can be chosen to satisfy the bottom set of equations. However, since joint torques do not affect the first six rows, the main dynamic limitations on movement are determined by the contact force constraints (i.e., friction and unilaterality).

A. Centroidal Dynamics - Modeling and Optimization

It can be shown that the first six rows of (2) describe the evolution of the net linear and angular momentum of the system as a whole [110], [111]. It is common to consider the angular momentum about the CoM denoted $k_G \in \mathbb{R}^3$, as this quantity is both conserved during free flight and empirically remains close to zero during human walking. Denoting the total linear momentum as $l_G \in \mathbb{R}^3$, these momenta compose the *centroidal momentum* $h_G = (k_G, l_G) \in \mathbb{R}^6$ [112]. This quantity is related to the generalized velocities via $h_G = A_G(q)\nu$, where $A_G(q)$ is called the centroidal momentum matrix [112]. Consider a case with n_c contact points at locations $\{p_i\}_{i=1}^{n_c}$. It can be shown that the top rows of (2) are equivalent to:

$$\dot{h}_G = \begin{bmatrix} \dot{k}_G \\ \dot{l}_G \end{bmatrix} = \begin{bmatrix} 0 \\ -Ma_g \end{bmatrix} + \sum_{i=1}^{n_c} \begin{bmatrix} (p_i - p_{\text{CoM}}) \times \lambda_i \\ \lambda_i \end{bmatrix}, \quad (3)$$

where $p_{\text{CoM}} = [c_x, c_y, c_z]^T$ gives the CoM position, M the total mass, and a_g the gravitational acceleration [110], [111]. It is important to note that $M \frac{d}{dt} p_{\text{CoM}} = l_G$. These *centroidal momentum dynamics* (3) (also called the *centroidal dynamics*) motivate solving a reduced problem posed over trajectories for the CoM p_{CoM} , angular momentum k_G , contact locations p_i , and contact forces λ_i . Such a reduction loses detail about contacting limbs, with forces no longer restricted according to the contact models of the previous section.

While the reduction to a centroidal dynamics model addresses the high-dimensionality of the original problem, it remains non-trivial to ensure that centroidal solutions are whole-body feasible, since they neglect the geometric constraints and actuation bounds. This omission in the model has motivated planning with whole-body kinematics and centroidal dynamics [18], or by alternating between centroidal dynamics and whole-body solves [26], [55]. Despite these advances, heuristic constraints on the CoM and footholds (e.g., to ensure reachability) remain prevalent, and simple bounds on the angular momentum (or fixing it to zero) are often used to simplify the problem.

Even with these simplifications, the centroidal dynamics (3) contains nonlinearities in the angular momentum equation due to bilinear terms from the cross product $(p_i - p_{\text{CoM}}) \times \lambda_i$.

Tailored methods to address these terms have been a topic of great focus. Dai et al. [41] consider fixing a-priori a polytopic or ellipsoidal region for the CoM, enabling optimization of an upper bound on the magnitude of the angular momentum. Valenzuela [113] considers relaxing the bilinear equalities with McCormick envelopes and employing mixed-integer convex optimization (see also, [21], [22]). Fernbach et al. [46] presented a development where contact locations are fixed and the CoM trajectory is parameterized via a Bézier curve with one free knot point, which avoids nonlinear effects. Overall, many choices exist for optimizing centroidal trajectories, with tradeoffs between accuracy and computational efficiency.

It is possible to further simplify the optimization of contact forces by considering their net effect. Given frictional limitations, each individual contact force λ_i is constrained to a friction cone \mathcal{C}_i . Then, the net effect of these forces is a 6D cone known as the *Contact Wrench Cone* (CWC)

$$\text{CWC} = \left\{ \sum_{i=1}^{n_c} \begin{bmatrix} p_{i/\text{CoM}} \times \lambda_i \\ \lambda_i \end{bmatrix} \mid \lambda_i \in \mathcal{C}_i, \forall i \in \{1 \dots, n_c\} \right\}.$$

When each friction cone is replaced with a polygonal approximation (e.g., friction pyramid), then the CWC is a polyhedral convex cone, and many tools from computational geometry are available to support its use in planning (c.f., [114]). The CWC plays an integral role in checking for static stability in multicontact scenarios [115], [116] and more recently has seen applicability for multicontact CoM motion generation [42]. Feasibility constraints based on the CWC are natural generalizations of the zero-moment point (ZMP) criteria [117] to multicontact. Since the computation of the CWC is itself costly, it is most useful when the CWC can be precomputed offline (e.g., when contact locations are known in advance).

B. Other Simplified Models

The centroidal dynamics (3) are a special modeling simplification in that they represent a projection of the equations of motion without adding artificial restrictions on possible movements (i.e., any whole-body feasible motion can be projected to a feasible solution of the centroidal dynamics). Many other simple models follow by adopting artificial motion restrictions (or approximations) to further simplify analysis and optimization. The suitability of these approximations/restrictions is often dictated by the robot morphology being modeled.

A common simple model for quadruped locomotion is the single-rigid-body (SRB) model (e.g., [57], [118], [119]). This simplification is motivated by the light leg designs of many quadruped robots, which makes the total rotational inertia of the system I_G about the CoM approximately invariant with configuration [120]. Thus, the SRB can be viewed as a restricted version of the centroidal momentum model in the case when the mass distribution takes a constant shape. Variations on the model accommodate inertia shaping [119], [121], however, there is an important subtlety that must be kept in mind. While the linear momentum of a system is related to its CoM velocity (i.e., the rate of change in some average position), the angular momentum cannot, in general, be equated to the rate of change in any meaningful average

orientation of the system [122]. This property is intimately related to the fact that the conservation of angular momentum, in general, defines a non-holonomic constraint [110].

The most common simple model for humanoid gait planning considers the centroidal dynamics with $k_G = 0$ (i.e., thus removing consideration of orientation dynamics) and the additional restriction that the CoM height c_z takes a fixed value denoted h . Rather than considering a collection of contact points, when walking on level ground, the overall Center of Pressure (CoP) position $p_{\text{CoP}} = [p_x, p_y, p_z]^T$ can be considered with $p_z = 0$ an assumed ground height. These restrictions lead to the LIP Model [47] with dynamics: $\ddot{c}_{x,y} = \omega^2(c_{x,y} - p_{x,y})$ where $\omega = \sqrt{g/h}$ represents its natural frequency. The beauty and power of the LIP reside in that these dynamics are linear, enabling convex optimization for planning CoM and CoP trajectories [49] or linear systems tools (e.g., LQR) for trajectory tracking [43].

More recently, work has focused on discharging assumptions on the CoM height for added flexibility [123]. One approach is to model height variations as a perturbation to the LIP [52] along with constraint tightening to ensure robust feasibility. Rather than treating height variations as a disturbance, they can be introduced into the dynamics model [36], [123]. A common simple form is $\ddot{p}_{\text{CoM}} = a_g + k(p_{\text{CoM}} - p_{\text{CoP}})$ where k represents a variable stiffness-like parameter, constrained to be positive. Strategies for running (e.g., based upon the spring-loaded inverted pendulum model) [124], [125] can be seen as placing additional restrictions on the form of k to mimic a Hookean spring. Overall, these point mass models remain most applicable to bipeds and humanoids where orientation dynamics can be minimized by ensuring all GRFs pass close to the CoM. This is also the case with a subset of quadruped gaits (e.g., trot walking), while others (e.g., bounding) cannot satisfy $k_G = 0$, and are more amenable to SRB models.

C. Summary

Modeling simplifications are commonly employed to reduce the computational burden for dynamic planning in legged robots. The most successful of these models (e.g., the centroidal model, SRB model, and LIP model) focus their modeling detail on addressing motion limitations imposed by contact interactions. The centroidal model does not impose any restrictions on motions (i.e., it represents a projection of the whole-body dynamics), while the other simplified models result from adding motion restrictions. These simplifications can unlock new structure (e.g., linearity) that can be leveraged for motion optimization problems over them (e.g., using the methods in the next section). In all cases, many whole-body details remain to be planned (e.g., swing foot trajectories). This omission represents a main downside to simple model planning, motivating more advanced numerical methods (Section V) for the efficient optimization of whole-body plans.

1) *Relationship with learning*: There are many ways emerging learning strategies have been used to advance the application of simplified models for locomotion planning. With regard to the models themselves, this review has covered existing strategies based on considerations of physics and

expert intuition. Moving forward, the automatic discovery of simple models remains an important open problem (c.f., [126]). Other work has looked at how to close the gap between simplified models and their whole-body counterparts, for example, to learn constraints on simplified models that address kinematics constraints at the whole-body level [74] or that address the gap via learning robust MPC strategies [127]. Yet other strategies have wrapped full RL pipelines around low-level control based on simple models [128] to increase the sample efficiency of learning. It should be acknowledged that even though RL strategies are often architected (e.g., via autoencoders) to learn low-dimensional representations within their layers (c.f., [129]), these latent representations remain difficult to interpret. It is an open problem whether the existing simplified models may be used to increase the interpretability of learned controllers as well.

V. NUMERICAL METHODS FOR SOLVING OCPs

While the decomposition of our general problem formulation in Section II (e.g., via modeling simplifications) remains a bit of an art, the technical details of solving a trajectory optimization problem represent a field unto its own. This field has benefited from contributions across engineering applications ranging from chemical process control [130], to flight planning [131], and others. We begin this section by making a few simplifications to our general problem (1), so that we may review the most common methods for solving OCPs at large, which provides context for the most common approaches taken for robotics in particular. We then gradually discharge our original simplifications, and explain nuances that arise in OCPs for locomotion. The treatment of these nuances has recently fueled a rapid increase in whole-body optimization, where the need for online re-planning has led to the development of fast structure-exploiting solvers tailored for robotics.

To begin, let us consider (1) in the case of a single contact phase (e.g., optimizing a trajectory for a humanoid doing an in-place dance on two feet). For better alignment with the dynamic optimization literature, we rewrite the dynamics (1b) as a system of first-order ordinary differential equations (ODEs) $\dot{x}(t) = f(x(t), u(t), p(t))$, with contact forces $\lambda(t)$ treated (mathematically) as time-varying parameters $p(t) = \lambda(t)$. The formulation can also be readily extended to consider time-invariant parameters. In the presence of contacts, the contact points must not move relative to the ground, which can be written as an algebraic constraint $g(t, x(t), u(t), p(t)) = 0$. These ODEs, combined with the algebraic constraints, lead to a system of differential-algebraic equations (DAEs) [130]. A generic nonlinear OCP for such a system can then be considered as:

$$\underset{x(\cdot), u(\cdot), p(\cdot)}{\text{minimize}} \int_{t_0}^{t_f} \ell(x(t), u(t), p(t)) dt + L(x(t_f)) \quad (4a)$$

$$\text{subject to } \dot{x}(t) = f(x(t), u(t), p(t)) \quad (4b)$$

$$0 = g(x(t), u(t), p(t)) \quad \forall t \in [t_0, t_f], \quad (4c)$$

where the objective function, with running cost ℓ and terminal cost L , is minimized over the time interval $[t_0, t_f]$ subject to the system dynamics defined by the DAEs. While the variables

	<i>Supporting Theory</i>		<i>Computation</i>
Global	Hamilton-Jacobi-Bellman (HJB)		Numerical PDE Solvers or Dynamic Programming
Local	<i>Direct</i> (discretize then optimize)	Numerical Integration	Direct Shooting/Collocation & NLP
	<i>Indirect</i> (optimize then discretize)	Pontryagin Maximum Principle (PMP)	Indirect Shooting/Collocation & Root Find

TABLE II

HIGH-LEVEL RELATIONSHIP BETWEEN OPTIMAL CONTROL SOLUTION METHODS. THIS SURVEY FOCUSES MOSTLY ON DIRECT METHODS FOR CONTINUOUS-TIME PROBLEMS DUE TO THEIR PROMINENCE IN ROBOTICS APPLICATIONS.

above match those in our original formulation (1), the problem could equally apply to a whole-body or simplified model.

For locomotion problems, the formulation (4) requires extension to consider how the dynamics change in different contact modes. Here, we assume that the mode sequence is fixed a-priori, and consider the multi-phase problem:

$$\begin{aligned} & \underset{\substack{x(\cdot), u(\cdot), p(\cdot), \\ s_1, \dots, s_{n_{\text{ph}}}}}{\text{minimize}} && \sum_{j=1}^{n_{\text{ph}}} \left[\int_{s_{j-1}}^{s_j} \ell_j(x(t), u(t), \lambda(t)) dt + L_j(x(s_j)) \right] \end{aligned} \quad (5a)$$

$$\text{subject to } \dot{x}(t) = f_j(x(t), u(t), \lambda(t)) \quad (5b)$$

$$g_j(x(t), u(t), \lambda(t)) = 0 \quad (5c)$$

$$h_j(x(s_j^-)) = 0 \quad (5d)$$

$$x(s_j^+) = R_j(x(s_j^-)) \quad (5e)$$

$$\forall t \in [s_{j-1}, s_j], j = 1, \dots, n_{\text{ph}} \quad (5f)$$

$$s_0 = t_0, s_{n_{\text{ph}}} = t_f, \quad (5g)$$

where $j = 1, \dots, n_{\text{ph}}$ denotes the phase index, with the time interval of the j -th phase given by $[s_{j-1}, s_j]$, and n_{ph} the number of phases. At the mode switch, a guard constraint (5d) is typically used to ensure the active contact points are in contact with the environment. Depending on the contact model, the reset maps (5e) may correspond to the impact dynamics or to switching between coordinate representations.

Generally speaking, many methods for solving an OCP only apply directly to (4) in the ODE case, but often admit extensions to the DAE and multi-phase settings. These methods can be categorized in three groups: global methods based on the Hamilton-Jacobi-Bellman (HJB) equation (or Dynamic Programming in discrete-time), local indirect methods, and local direct methods (see Table II).

Methods based on Dynamic Programming exploit Bellman's principle of optimality [132] to solve a discrete-time version of the OCP (4). They do so via finding the optimal cost-to-go $V(t_0, x_0)$ (also called the *Value function*). In continuous time, the Value function must satisfy a Partial Differential Equation (PDE) known as the HJB. Its solution can be approximated by discretizing time and space and applying dynamic programming. However, it is well-known that the complexity of this strategy increases exponentially with the number of states and controls. Therefore, it is not directly applicable to most legged robots, requiring clever approximate decomposition strategies for application to high DoF models [133], [134].

Indirect methods transform the original OCP into a Boundary Value Problem by using Pontryagin's Maximum Principle [135] to formulate the so-called co-state equations. This approach enables pre-optimizing the control as a function of state and co-state if the dynamics are control affine—which




Fig. 4. Depiction of multiple shooting. Each shooting segment is associated with a numerical simulation (e.g., via a numerical integration method). Optimization variables include the states x_i at the beginning of each segment and variables describing the controls applied in each segment. Constraints enforce continuity, i.e., that the defect between segments is zero.

is the case for most legged robots. Indirect methods turn the OCP into a root-finding problem, which can provide fast and accurate solutions. It can be challenging to initialize the co-state for these methods, and additional mechanisms (e.g., homotopy) are needed to enforce state inequalities [136].

Direct methods have been more broadly adopted by the legged robotics community, and are generally compatible with DAEs. In direct methods, the OCP is transcribed into a finite-dimensional Nonlinear Program (NLP) by discretizing the controls and states with respect to time. Therefore, direct methods directly find the minimum of (1a), where the NLP can be solved with well-established optimization techniques, e.g., Sequential Quadratic Programming (SQP) [137]. The differences between various direct methods lie in how the discretization is carried out. In the following, we focus exclusively on the three prominent direct methods in the community, namely *direct multiple shooting*, *direct collocation*, and *Differential Dynamic Programming (DDP)*. A more comprehensive review can be found in [138].

A. Multiple shooting

In direct shooting methods [139], the controls $u(t)$ are discretized over the time horizon, while the state trajectories are obtained via forward integration. In contrast to *single shooting*, where a single integration is performed over the whole time horizon, in multiple shooting, the time horizon is discretized, then integration is performed on each of the segments (Fig. 4). The initial state of each segment is added to the decision variables, and *continuity constraints* are introduced to guarantee that the final state of each segment i (computed with integration) matches the initial state of the next segment $i + 1$. In essence, as in Fig. 4, the initial optimization problem is divided into smaller optimization problems over the discretized grid, whose initial conditions




Fig. 5. Depiction of collocation. Each finite element is associated with a polynomial that approximates the trajectory. Optimization variables include the states x_i at each endpoint, and states/controls at each collocation point. Constraints enforce that the endpoint defect is zero, and that the slope of the approximating polynomial matches the evaluation of the dynamics function f at each collocation point.

can be set separately. By parameterizing the time horizon with a series of Initial Value Problems (IVP), the time interval over which integration is performed is shortened, reducing the high sensitivity issue of single shooting and preventing states from diverging. This makes multiple shooting deliver rather robust solutions, as shown in [140], which motivated many researchers to use it for motion generation with complex robots.

In particular, the use of multiple shooting has been popular for generating motions considering whole-body dynamics where contact modes are set a priori, i.e., predefined contact sequences. This is the case of [29] and [31], where contacts are assumed to be rigid and impacts are instantaneous and inelastic. The resulting OCP is then a multi-phase problem with discontinuous phase transitions, as in (5). In [30], the authors used multiple shooting to optimize virtual constraints as part of gait design for a Hybrid Zero Dynamics (HZD) controller with a compliant bipedal robot.

B. Collocation

In direct collocation (also known as direct transcription [138]) the original OCP is transformed into an NLP having control and state trajectories as decision variables. Controls and states are discretized over a time grid, where the intervals are called *finite elements* (see Fig. 5). Controls are approximated in each finite element by a finite-dimensional representation, and the states are approximated by polynomials. The system dynamics are then enforced by imposing *continuity constraints* to eliminate endpoint defects at the end of each element and through internal *collocation constraints* imposed at each of the *collocation points* (see Fig. 5). These collocation constraints enforce that the slope of the approximating polynomial matches the system dynamics $f(x(t), u(t), p(t))$ at each collocation point. The choice of the number and location of these collocation points has been thoroughly investigated [141] because it affects the accuracy of the method. More details can be found in [142], [143].

Collocation has been a popular method to compute walking motions for legged robots using whole-body models. In [34], a periodic motion was generated for a whole-body bipedal robot, where the convergence of the periodic cycle was guaranteed through the implementation of HZD virtual constraints. In

[144], the authors used collocation to generate motions for a quadruped robot, where the contact sequences were predefined. Collocation constraints were modified to address contact constraints in [33], providing third-order integration accuracy with demonstrations on a full-body humanoid.

Most robotics methods to date have considered collocation where the system dynamics are expressed as a system of first-order ODEs. Vanilla collocation methods apply separate approximating polynomials for each component of the state (e.g., separate polynomials for q and \dot{q}), with their consistency only enforced at collocation points [145]. Recent work [145], [146] has shown the benefit of using a unified interpolating polynomial for the components of q and \dot{q} , which ensures their consistency across the trajectory, and improves accuracy [145].

C. Differential Dynamic Programming (DDP)

In recent years a high interest has grown for DDP and its variants. DDP [147] is an age-old method for solving discrete-time unconstrained OCPs. The method is based on Dynamic Programming, but it overcomes the curse of dimensionality by working with a local quadratic estimate of the Value function. This makes DDP different from Dynamic Programming, because the method does not attempt to exhaustively explore the state space of nonconvex problems. DDP can be roughly seen as an efficient iterative algorithm for solving the banded system of linear equations associated with the KKT conditions of an unconstrained OCP transcribed with collocation. Indeed, DDP closely resembles Netwon's method [148], but it is not completely equivalent to it. The main difference is that to solve the KKT system, DDP expresses the control inputs as linear functions of the state, but then, instead of computing control perturbations by using linearized dynamics from the KKT system, it forward simulates with the nonlinear dynamics. This makes DDP a direct single shooting strategy (since the state is computed by integrating the dynamics), but it is better suited to handle unstable dynamics compared to vanilla shooting since its use of state feedback helps prevent divergence.

DDP requires the second derivatives of the dynamics to achieve quadratic convergence. While this is a desirable property, it may be challenging to compute these terms for complex systems. Variants of DDP that make use of Hessian approximations have been more popular, such as iLQR and iLQG [1], [149], [150], with other recent work focused on efficiently computing the full Hessian [151], [152].

Many extensions to the original DDP algorithm have been considered to address constraints beyond those from the system dynamics. Box DDP [153] accounts for box constraints on the control inputs. Hierarchical DDP [150] allows for the optimization of a hierarchy of cost functions in lexicographic order. DDP extensions that can treat arbitrary nonlinear inequality state-control constraints have been proposed based either on interior-point techniques [154], active-set methods [155], Augmented Lagrangian [156], [157], or relaxed barrier [38] approaches. Recently DDP has also been extended to handle implicit dynamics [5], multi-phase dynamics [27], [97], mode switching constraints [158], and switch timing [28], [158], all of which are useful for handling contacts. Even

though DDP was born as a single-shooting method, it has been extended to multiple shooting [159]–[161] to further improve its handling of sensitive dynamics.

D. Contact Implicit Considerations

As a departure from the main assumptions of this section, contact-implicit (or contact-invariant) formulations (e.g., as first discussed in Section III-B2) instead seek to optimize the mode sequence. The choice of transcription strategy is highly coupled with the contact model adopted. When using shooting methods, the key need is for gradients that relate changes in states/controls to trajectory outcomes. For example, when adopting a relaxed contact model, the outcomes of simulation are always differentiable, enabling shooting solvers and/or DDP [1], [5]. When adopting an event-driven hybrid model, these methods can still be used when paired with suitable hybrid sensitivity analysis, as in [9]. Likewise, discrete-time shooting can be applied with time-stepping LCP solvers [8], although both this and the hybrid case exhibit pathologies that prevent differentiability in corner cases (c.f., Section III-A4).

When using direct collocation methods, existing contact-implicit methods rely on complementary constraints in the optimization. First-order methods [10] have been extended to high-order schemes, assuming the mode is not changing during an element [11]–[13]. These methods either require impacts at the element boundaries [13], [162] or sacrifice accuracy to allow impacts during the element [12], [13]. There is a trade-off, since enforcing impacts at the element boundaries requires additional complementarity constraints. The alternative is to consider relaxations of the LCP constraints [23], which again present trade-offs for accuracy and computational complexity.

E. Summary

Within robotics, direct methods remain the most attractive numerical approaches for trajectory optimization. While both direct shooting and direct collocation can be used in cases when the mode sequence is fixed a priori, contact-implicit strategies require more careful consideration regarding the contact modeling choices adopted and how those affect transcription. Recent years have seen progress in accelerating the solution of shooting problems through the use of DDP, with recent sparse QP solvers (e.g., [163]) opening the door for similar accelerations to closely related collocation formulations.

1) *Relationship to learning*: There have been many motivating examples of learning being used to support trajectory optimization. A common strategy is to employ learned networks to warm start optimization (see, e.g., [164]). Another common way for learning to support trajectory optimization is by learning the value function, which enables online trajectory optimization over shorter horizons [165]. Other strategies simplify the optimization by learning footholds [166], e.g., akin to a higher-level contact planner.

VI. ONLINE WHOLE-BODY PLANNING AND CONTROL

There are many ways in which state-of-the-art locomotion architectures combine planning with optimal control solutions




Fig. 6. Different potential strategies for breaking Problem 1 into smaller pieces to ease computational challenges.

for deployment online, depending on computational hardware, robot complexity, and task complexity, as depicted in Fig. 6. The most hierarchical strategy (left column) decomposes locomotion control into footstep planning, CoM planning (via MPC with a simplified model), and whole-body reactive control (e.g., [43]). Recent trends have been flattening this hierarchy, in particular, as optimization solvers have matured. Several recent demonstrations are relying on whole-body MPC (right column). The current state-of-the-art still breaks the contact planning into a high-level module, as contact implicit schemes do not yet run at the rates necessary for online control. Owing to its connection with the previous section, we review whole-body MPC before returning to the more classical methods of reactive whole-body control.

A. Online Whole-Body MPC

Within the past few years, the deployment of MPC with whole-body models of legged robots has gone from an aspirational target to one with broad experimental support. While we provide a short overview here, the interested reader is referred to an excellent and complementary recent review on the topic [167]. Work by Katayama and colleagues and Mastalli and colleagues have shown how to accelerate DDP methods by leveraging inverse dynamics modeling [28], [168] of full rigid-body dynamics for quadrupeds. As a representative data point, [168] ran MPC at 50 Hz using a 1 s horizon with 100 timesteps of 10 ms. Other quadruped work has shown the viability of using whole-body kinematics and centroidal/single-rigid-body dynamics [169]–[171] for MPC. Contrary to the previous DDP results, [171] employed sparse QP solvers [163] for MPC (100 Hz update, 1 s horizon, 15 ms timesteps). A commonality in these online strategies is the willingness to live with the sub-optimal. Real-time iteration schemes [172] are applied where trajectory optimization only runs a few iterations before re-sampling the initial state and proceeding to solve again [171]. Beyond quadrupeds, a recent milestone [173] demonstrated online MPC with a 22 DoF model of the TALOS humanoid operating at 100 Hz (0.5 s planning horizon, 10 ms timesteps) and employing DDP’s local feedback between updates [38].

Across these demonstrations, a few commonalities emerge. The first is the use of a long planning horizon. The work in [169] nicely characterizes the effects of the horizon on performance and notes that planning for less than two gait




Fig. 7. A classical use of a Whole-Body QP in a pipeline to execute a simple-model plan obtained from the methods of the previous sections. A typical list of tasks and constraints is also depicted.

cycles leads to a loss of stability. Within the literature on MPC theory, a proper design of terminal constraints and costs allows the use of shorter horizons without sacrificing recursive feasibility and stability—however such theoretical certificates remain out of reach for the robot models currently in use. Instead, significant effort has gone into warm starting solvers [173], which might be unneeded if the theoretical problem was more thoroughly addressed. Given the computation challenges with high DoF humanoid robots, and limited applicability of single-rigid-body models, we see the most opportunity for growth in whole-body MPC for humanoids moving forward.

B. Instantaneous Whole-Body Control

In cases when computational hardware or model/task complexity prevents the use of whole-body MPC, hierarchical schemes such as the one on the left of Fig. 6 are necessary. For example, consider the setup in Fig. 7. When performing trajectory optimization over simplified models (e.g., the SRB or a point-mass model), the OCP may be able to be solved quickly, but there are still many details that need to be defined (e.g., swing leg motion). Instantaneous control schemes are widely used in practice as a means to accomplish this goal, or more generally to track desired trajectories. The popularity of these methods owes to their ease of implementation and the fact that, with well-chosen objectives/constraints, the problem can be written as a Quadratic Program (QP), or a close variant, and solved quickly (from hundreds of μs to a few ms).

1) *Formulation*: In instantaneous settings, one common way to define the motion is by specifying quantities e_i that we want to regulate to 0, or keep above 0. For example, let us suppose we want the robot to follow a walking gait that was generated using the LIP model. For the robot CoM to track the LIP trajectory, we want to achieve $e_i(q, t) = 0$ where $e_i \in \mathbb{R}^3$ gives the error between the target position from the LIP and the true CoM position computed from q via forward kinematics. Other tasks, such as collision avoidance, are naturally captured by inequality constraints $e_i(q) \geq 0$. More generally, e_i can be a function of any subset of (q, ν, t) . Such a definition is often called a *task error function* by reference to the task function [174] or operational-space [175] formalism.

A key point is that the dynamics (1b) are linear in $\dot{\nu}$, τ , and λ . This points at writing the regulation at the acceleration level where τ and λ can also be used directly to affect motions.

Task Dynamics: Consider a general task in the case where $\dot{q} = \nu$ and $e_i \in \mathbb{R}^m$. Differentiating $e_i(q, t)$ twice w.r.t. time:

$$\ddot{e}_i(q, \nu, \dot{\nu}, t) = J_i(q)\dot{\nu} + \dot{J}_i(q, \nu)\nu + a_i(q, \nu, t), \quad (6)$$

where J_i is the Jacobian matrix of e_i and $a_i = \partial^2 e_i / \partial t^2$. To bring or keep e_i to a desired value, one then writes $\ddot{e}_i = \ddot{e}_i^d$ or $\ddot{e}_i \geq \ddot{e}_i^{min}$ where $\ddot{e}_i^d / \ddot{e}_i^{min}$ is a function of any subset of (q, ν, t) that indicates how we want to regulate e_i . This gives an equality/inequality that is linear in $\dot{\nu}$:

$$J_i(q)\dot{\nu} \geq \ddot{e}_i^d - \dot{J}_i(q, \nu)\nu - a_i(q, t). \quad (7)$$

If e_i also depends on ν , a single differentiation leads to:

$$\frac{\partial e_i}{\partial \nu} \dot{\nu} \geq \dot{e}_i^d - J_i(q, \nu)\nu - \frac{\partial e_i}{\partial t}(q, \nu, t). \quad (8)$$

Both (7) and (8) can be written as

$$A_i(q, \nu)\dot{\nu} \geq b_i(q, \nu, t). \quad (9)$$

Task expressions: The most common tasks involve geometric features (e.g., body or CoM position) that are direct functions f_i of q . Typical equality tasks have the form $e_i(q, t) = f_i(q) - f_i^d(q, t)$, where f_i^d is the target value. For Lie group tasks (e.g., orientation), this error can be generalized as $e_i(q, t) = f_i(q) \ominus f_i^d(q, t)$, where \ominus denotes a meaningful difference in the Lie group [176]. This general definition can capture tasks for the CoM, body positions, body orientations, joint posture, and others. Other equality tasks include a gaze task [177], visual servoing [178], or any subpart of one of the above [59], [179]. Inequality tasks commonly include joint bounds, collision avoidance, and balance [177]. Few tasks depend directly on ν . This is the case of joint speed limits, but also the centroidal momentum task [180]–[182] $e_i = A_G(q)\nu$.

We can also call tasks, by extension, constraints or motion requirements that are directly written on $\dot{\nu}$, τ , and λ and do not need regulation. These include, for example, limits on $\dot{\nu}$, τ , or friction constraints on the forces (to prevent sliding), that are usually approximated using a linear form $C\lambda \geq 0$ (e.g., as was also considered to simplify the CWC in Section IV-A). Direct references can also be given for these variables. For example, $\lambda - \lambda^d = 0$ can be used for force control.

Regulation of Equality Tasks: The task error e_i is often driven to 0 by first specifying a desired value for the derivative $e_i^{(l)}$ where l is the number of differentiations (1 or 2) such that $e_i^{(l)}$ is an affine function of $\dot{\nu}$, τ or λ . For geometric equality tasks ($l = 2$), the most common approach follows a PD construction: $\ddot{e}_i^d = -K_d \dot{e}_i - K_p e_i$, where K_d and K_p are positive definite matrices, usually taken as diagonal matrices with $K_d = 2K_p^{1/2}$ to obtain critical damping. A common (and essential) equality task relates to maintaining contacts. Most often, hard contacts are assumed and \ddot{e}_i^d is set to 0 or a damping term is used to stabilize any slippage [59].

Within the framework of HZD, a critical task comes from the construction of virtual constraints that specify actuated joints as a function of a state-dependent phase variable [96]. Historically, HZD controllers were first implemented via analytical feedback linearization [96], while later being connected to constrained optimization methods herein [75].

Born out of HZD implementation advances, an alternative for making e_i converge to zero is to design a Control Lyapunov Function (CLF) V_i for e_i . Then, the inequality $\dot{V}_i \leq -\gamma_i V$ can be exploited to pre-solve some variables and constraints before passing the problem to a solver.

Regulation of Inequality Tasks: Inequality tasks at the acceleration level, such as joint torque limits or force friction cones, can be directly imposed. Inequality tasks at the velocity level, such as joint velocity limits, can be regulated by setting \dot{e}_i^{min} such that e_i remains positive after a controller time step Δt : $e_i^+ = e_i + \Delta t \dot{e}_i \geq 0$. This approach works well for small Δt , and implies setting $\dot{e}_i^{min} = -e_i/\Delta t$. Inequality tasks at the position level, such as joint position limits or obstacle avoidance, are much harder to enforce. The most common approach is to compute \ddot{e}_i^{min} so that e_i^+ is positive, but this can easily lead to conflicts with acceleration-level constraints [183]. Alternative approaches range from a trick of increasing Δt in the computation of e_i^+ [177], [183], to bounding velocities based on the distance to the position limit [184], to control barrier functions (CBFs) [76]. However, none of these approaches ensure compatibility between position-level and acceleration-level inequalities. The only methods that can guarantee them are those that account for position-level and acceleration-level inequalities at once [185], [186], albeit under the assumption of constant acceleration bounds. Enforcing position constraints remains precarious with instantaneous control, since it is the evolution of the system over an extended time period that determines whether the constraint can be respected in the future. This observation motivates predictive strategies (e.g., MPC) to address position inequality constraints via lookahead.

2) *Resolution:* Instantaneous control schemes make the assumption that the state x is known and constant during the control period Δt . This is supported by the fact that Δt is small (typically between 1 and 10 ms), and relies on the ability to solve the problem in that time. The contact state (i.e., which bodies are in contact) is assumed fixed at each control cycle.

Under these conditions, and if the friction cones are approximated as pyramids, aggregating all the desired behaviors of the tasks leads to a set of linear constraints:

$$M\dot{\nu} - J_c^T \lambda - S^T \tau = -C\nu - \tau_g \quad (10a)$$

$$J_c \dot{\nu} = a_c \quad (10b)$$

$$A_{\dot{\nu},i} \dot{\nu} + A_{\lambda,i} \lambda + A_{\tau,i} \tau \geq b_i \quad \text{for non-contact tasks } i, \quad (10c)$$

where (10a) is a reorganization of (1b), (10b) gives contact constraints, and (10c) is a generalization of (9) to account for tasks directly written on τ or λ . Problem (10) represents what we would like the robot to achieve. The idea is to solve (10) for $\dot{\nu}$, τ , and λ as best as possible (as discussed below). On position- or velocity-controlled robots, $\dot{\nu}$ is integrated and the resulting configuration [59] or velocity [178] is sent to the robot. For a torque-controlled robot, τ is used as a command.

Pre-solving variants: Problem (10) is written over all 3 variables $\dot{\nu}$, τ and λ . This gives great flexibility to express new

tasks/constraints and offers a clear vision of what we want to achieve. However, the specific structure, of (10a) and (10b) can be exploited to pre-solve some variables and constraints before passing the problem to a solver.

The form of $S = \begin{bmatrix} 0 & I \end{bmatrix}$ for legged robots allows expressing τ as a linear function of $\dot{\nu}$ and λ and eliminating it from the problem. Only the non-actuated rows of (10a) are kept [59], [182], leading to a near-universal decrease in computation time. Alternatively, joint accelerations $\dot{\nu}$ can be eliminated via (10a) with the inversion of M [187], or forces λ can be eliminated via projection on the nullspace of the contact constraints [188], [189]. Finally, both $\dot{\nu}$ and λ can be eliminated via (10a)-(10b), leaving τ as the only variable. Regardless, the optimal pre-solving variant is highly coupled with the numerical solver used.

Arbitrating conflict: For most task sets considered in practice, task requirements will conflict, meaning (10) has no solution. In this case, one can measure violations with an $L2$ norm, and specify how important it is to minimize each violation. Two main approaches exist that can be combined: weighing each violation (also known as *soft priorities*) or defining a hierarchy between them (*strict priorities*).

Numerous works use a QP approach: all inequality constraints and some equality constraints in (10) are kept in the QP, making them the top priority. The weighted $L2$ norms of all other constraint violations form the objective of the QP. For example, let us consider the problem of tracking a LIP-based reference trajectory, while trying to use minimal joint torques within bounds. In this case, a QP could be formulated as:

$$\begin{aligned} & \underset{\dot{\nu}, \tau, \lambda}{\text{minimize}} \quad w_1 \|J_{CoM} \dot{\nu} + \dot{J}_{CoM} \nu - (\ddot{p}_{CoM}^d + \ddot{e}_{CoM}^d)\|^2 \\ & \quad \quad \quad + w_2 \|\tau\|^2 \\ & \text{subject to} \quad M\dot{\nu} + C\nu + \tau_g = S^T \tau + J_c^T \lambda \quad (\text{dynamics}) \\ & \quad \quad \quad J_c \dot{\nu} = a_c \quad (\text{fixed contacts}) \\ & \quad \quad \quad C\lambda \leq 0 \quad (\text{friction cones}) \\ & \quad \quad \quad \dots (\text{other top-priority constraints}) \dots \end{aligned}$$

with J_{CoM} being the CoM Jacobian, \ddot{p}_{CoM}^d the target CoM acceleration, and \ddot{e}_{CoM}^d the desired acceleration of e_{CoM} .

The use of QP formulations for reactive control is now nearly universal, but follows a long historical development. After an early work motivated by the inclusion of unilateral contact forces [190], the QP-based inverse dynamics approach started with [63], [65] and closely related convex formulations [64], continuing with [59], [179], [191], [192] among many others. An advantage of this approach is that it relies on mature off-the-shelf solvers, which are often free (e.g., [193], [194]). Its main limitations are the inability to handle inequalities at lower priority levels, and its need for detailed tuning.

One of the benefits of the QP formalism is its flexibility to incorporate other control paradigms into reactive whole-body control. For example [43], [195] embed long-horizon optimal control for simple-model tasks via cost function terms for descending a task-space value function. Other work has explored the incorporation of passivity-based control approaches that improve robustness to unmodeled effects [77]. To improve robustness through other means, recent QP strategies have

explored methods for controls to be invariant to velocity jumps from impact events [196], which are hard to detect in practice. Alternative strategies that consider soft contact can also be incorporated into whole-body QPs [197]–[199].

In contrast to using a weighting approach for tasks, a hierarchical approach assigns an explicit priority level for each task. Tasks at the same level can be combined with weights, and top-priority constraints need not be feasible. The resulting problem is called *lexicographic least-squares* or *hierarchical QP* (HQP) [177], [180], [182]. Kanoun [200] presented the first solver to tackle inequality constraints at any priority level, followed by computational improvements in [180], [182]. Dedicated solvers [201], [202] have been introduced to solve the problem efficiently in one pass. The hierarchical formulation is a strict superset of the QP-based formulation and is the limit case for when the ratio between task weights goes to infinity [203]. An HQP can be solved faster using a dedicated solver [202], both in theory and practice. However, computational costs and the complexity of handling singularities [204] have limited broad use.

C. Summary

This section has focused on common methods for reactive control that coordinate the selection of joint torques and contact forces at the current instant. These methods complement predictive control (i.e., over a horizon). By focusing on an instantaneous control problem, the formulation inherits a desirable structure (e.g., nonlinear constraints become linear ones). This enables solutions at rapid rates, which are amenable to computation using off-the-shelf QP solvers. There are many different variants to such whole-body problems that consider priorities between tasks, or that strategically pre-solve for some decision variables to accelerate the QP solution. While these methods have represented a natural outgrowth of operational-space control to legged robotics applications, the shift from solving these problems analytically to solving them via QPs has enabled many other control paradigms to be considered, including CLFs, CBFs, passivity-based strategies, and others guided by value functions. Overall, while there remain opportunities to further improve these methods, the current state of the art represents a reliable technology for future applications.

VII. OUTLOOK AND PROSPECTS

We conclude this survey by reflecting on some of the main trends that were identified through the literature review, and also some of the notable gaps that remain to be addressed.

A. Main Trends

Main trends for optimization with contacts: The past five years have seen an increase in contact implicit approaches that avoid specifying gait/contact sequences a-priori. These methods have considered either compliant contact or MIP/LCP strategies. Solution speed remains a challenge for these methods, as does non-convexity of the problem. Strategies that combine planning and/or learning with contact implicit methods appear as a necessary next step. In the near term, fixing

contact sequences/timing remains the most viable option for practical deployment in applications that require online optimization (e.g., in MPC). In the end, while much more recent work has concentrated on addressing rigid contacts, other approaches end up being similar since rigid complementarity constraints are often relaxed and smoothed, leading to similar operations as with compliant models.

Main trends in the dynamic models adopted for model-based optimization: The past five years have observed an increase in model complexity used for MPC (with several examples using the full model). This advance has occurred from an increased understanding of gradient-based numerical methods and the availability of fast open-source software libraries to compute robot dynamics and its derivatives [205]. With Moore’s law ending, subsequent advances may center on parallelization. However, even with increases in computational speed, the nonconvexity of problems is an unavoidable challenge with complex models. This observation suggests that simple models will likely remain pertinent in the near future. The review points out that a combination with learning strategies that include exploration (as in conventional RL) presents promise for optimization-based strategies to avoid poor local minima when adopting complex models.

Main trends for transcription: The past five years have included emphasis on DDP methods as well as extensions of collocation schemes to contact-constrained/implicit settings. Adoption of DDP has been motivated in part by its demonstrated scalability to high-DoF systems. However, recent work has also shown specialized Riccati-like solvers [163], which have the promise to bring many of the advantages of shooting via DDP to other transcription approaches. While the numerical methods are improving, convergence to infeasible solutions is a challenge for complex systems, and this holds with any of the transcription methods. The review points that addressing these challenges is an important open point for future work. Further, simultaneous methods (collocation and multiple shooting) will present additional opportunities for algorithm-level parallelism to take advantage of how computational resources are expected to grow in the coming decades.

Main trends for learning for optimal control: Learning has emerged as a valuable ally for the application of trajectory optimization (TO) on challenging systems such as legged robots. Learning can be combined with TO in different ways. Recent work has proposed learning the robot dynamics model, e.g., using Deep Neural Networks [206] or Gaussian Processes [207]. Supervised learning can be used to exploit motion-capture data, for instance, to obtain human-like walking trajectories [208]. Cost functions, instead of being defined by expert users, can be learned, e.g. using active learning techniques [209]. Another promising approach that has been recently investigated is to learn good initial guesses to speed up the convergence of OCP solvers [164], [210]. Finally, another method to help OCP solvers is to learn an approximation of the Value function, which can then be used as a terminal cost to *guide* the solver [106], [165].

B. Overview of Open Problems

Formal Analysis: Despite a great deal of recent progress, there is a broad lack of formal analysis regarding the properties and operation of existing optimization-based robot control paradigms. Existing theoretical results regarding the stability and recursive feasibility of nonlinear MPC are mostly focused on the simplified problem of regulation to an equilibrium or tracking a feasible reference state trajectory [211], which leaves a theoretical gap for the classes of problems considered in the locomotion community. Another barrier to meaningful analysis is a relative lack of work addressing considerations of robustness to bounded or stochastic uncertainty (e.g., [25], [52], [212]). Such a treatment would, for example, be critical to understanding how our control frameworks interact with state estimation. While there are clear academic open questions, it remains an open question of its own as to how much additional performance such analysis would unlock.

Reducing the Expertise Required: Beyond the technical expertise required for model-based control, there remains a great deal of domain-specific expertise required for crafting cost functions and constraints. It remains a challenge to remove the human as an outer loop in the tuning of problem formulations. A part of the issue is that it is often not apparent how to write down cost functions for even simple tasks. For example, the problem of opening a door has binary success, and such cost structures are largely incompatible with current tools. Whether through automated methods, or through new tools that enable a broader set of cost designs, the best path for reducing required domain-specific expertise remains open.

Given the complexity of the numerical methods involved for online optimal control, an additional consideration is the availability of well-maintained open-source software for newcomers to rapidly deploy OCP solutions. In this regard, the recent review of Carpentier and Wieber [213] provides an excellent resource on various dynamics and optimal control packages available for formulating and solving optimal control problems for robotics. While we encourage engagement with these tools, we likewise call on everyone to contribute to their advancement and maintenance so that we might continue to accelerate the impact of OCP methods on robotics practice.

Optimal Control Versus or With RL?: With the recent progress in RL for legged locomotion (e.g., [129], [214]) much debate has surrounded whether real-time optimization or reinforcement learned policies are the path forward. RL methods have some clear benefits, the main one probably being their generality. Indeed, since gradients of the model dynamics are not required, these methods have no trouble with non-smooth contact-implicit problems, which are instead hard to solve with conventional gradient-based OCP solvers. An interesting explanation of why that is the case has been recently provided in [215], drawing a connection between *Policy Gradient* and the stochastic optimization technique *Randomized Smoothing*. Although our community has grown attached to gradient-based optimization, in order to overcome the challenges discussed in this survey it may be worth exploring gradient-free approaches, with RL being among the most promising. Another positive feature of RL methods (especially

those based on Dynamic Programming [216]) is their ability to avoid local minima, which is fundamental to solve non-convex problems. Indeed, while standard TO algorithms simply follow the gradient, some RL methods can exploit the learned action-Value function to get out of local minima.

Still, there remain exciting opportunities for how model-based optimization can help accelerate learning. Strategies such as the popular guided policy search method [217] and relatives (see, e.g., [218], [219]) can leverage TO to cut down on needed exploration and guide learning. Another important opportunity for future work is to use the efficient constraint-handling nature of TO to help make learning *safe* [220]. Overall, since RL and OCP methods try to solve similar problems [221] (especially in offline simulations where models are unavoidable), the history of their synergy is still quite early.

C. Closing Remarks

Reflecting on the trajectory of the field, we see a great deal of commonality between the transformation that is taking place now and the one that occurred in the mid-2010s. In 2007, convex optimization strategies for reactive control were beginning to appear, with no shortage of concern regarding both (a) their online computation requirements or (b) the abandonment of previously available closed-form control solutions. By 2015, these methods were commonplace.

From where we are today, early demonstrations of whole-body MPC on hardware have greatly benefited from the structural understanding of physics-based models from previous reactive control developments, while also adopting reactive methods for low-level control. Yet, open challenges remain regarding (a) online computation and (b) the abandonment of convexity properties that would give solution assurances.

While we believe these state-of-the-art methods will become technologies in the near future, we also look forward to the way the community will build upon them via combination with the next generation of learning-based controllers, where shared optimal control fundamentals will allow model-based methodologies to inform and accelerate learning. Challenges regarding (a) offline computation requirements and (b) a lack of generalization guarantees with RL seem to both be well addressed by structure exploiting methods from model-based optimization, and we hope this survey will play a small role in helping those solutions come to light. Overall, the evolutionary progress of the field shows a steady maturation toward applications in logistics, agriculture, construction, and yet others. With external commercial funding at unprecedented levels, we look forward to the next decade of progress from legged robots, which promises to be greater than the previous.

REFERENCES

- [1] Y. Tassa, T. Erez, and E. Todorov, "Synthesis and stabilization of complex behaviors through online trajectory optimization," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2012.
- [2] J. Koenemann, A. Del Prete, Y. Tassa, E. Todorov, O. Stasse, M. Bennewitz, and N. Mansard, "Whole-body model-predictive control applied to the HRP-2 humanoid," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2015, pp. 3346–3351.
- [3] M. Neunert, F. Farshidian, A. W. Winkler, and J. Buchli, "Trajectory Optimization Through Contacts and Automatic Gait Discovery for Quadrupeds," *IEEE Robot. and Autom. Lett.*, vol. 2, no. 3, 2017.

- [4] M. Neunert, M. Stauble, M. Gifftthaler, C. D. Bellicoso, J. Carius, C. Gehring, M. Hutter, and J. Buchli, “Whole-Body Nonlinear Model Predictive Control Through Contacts for Quadrupeds,” *IEEE Robot. and Autom. Lett.*, vol. 3, no. 3, pp. 1458–1465, 2018.
- [5] I. ChatziniKolaidis and Z. Li, “Trajectory optimization of contact-rich motions using implicit differential dynamic programming,” *IEEE Robot. and Autom. Lett.*, vol. 6, no. 2, pp. 2626–2633, 2021.
- [6] Y. Zhu, Z. Pan, and K. Hauser, “Contact-Implicit trajectory optimization with learned deformable contacts using bilevel optimization,” in *IEEE Int. Conf. on Robot. and Autom.*, 2021, pp. 9921–9927.
- [7] I. Mordatch, E. Todorov, and Z. Popovic, “Discovery of complex behaviors through contact-invariant optimization,” *ACM Trans. on Graphics*, 2012.
- [8] J. Carius, R. Ranftl, V. Koltun, and M. Hutter, “Trajectory optimization with implicit hard contacts,” *IEEE Robot. and Autom. Lett.*, 2018.
- [9] N. J. Kong, G. Council, and A. M. Johnson, “iLQR for Piecewise-Smooth Hybrid Dynamical Syst.” in *IEEE Conf. on Decision and Control*, 2021, pp. 5374–5381.
- [10] M. Posa, C. Cantu, and R. Tedrake, “A Direct Method for Trajectory Optimization of Rigid Bodies Through Contact,” *Int. J. of Robot. Res.*, vol. 33, no. 1, 2014.
- [11] W. Xi, Y. Yesilevskiy, and C. D. Remy, “Selecting gaits for economical locomotion of legged robots,” *Int. J. of Robot. Res.*, vol. 35, no. 9, pp. 1140–1154, 2016.
- [12] Z. Manchester, N. Doshi, R. J. Wood, and S. Kuindersma, “Contact-implicit trajectory optimization using variational integrators,” *Int. J. of Robot. Res.*, 2019.
- [13] A. Patel, S. L. Shield, S. Kazi, A. M. Johnson, and L. T. Biegler, “Contact-implicit trajectory optimization using orthogonal collocation,” *IEEE Robot. and Autom. Lett.*, 2019.
- [14] N. Doshi, K. Jayaram, B. Goldberg, Z. Manchester, R. J. Wood, and S. Kuindersma, “Contact-Implicit Optimization of Locomotion Trajectories for a Quadrupedal Microrobot,” 2019.
- [15] A. Aydinoglu and M. Posa, “Real-time multi-contact model predictive control via admm,” in *Int. Conf. on Robot. and Autom.* IEEE, 2022, pp. 3414–3421.
- [16] K. Wampler and Z. Popovic, “Optimal gait and form for animal locomotion,” in *ACM Trans. on Graphics*, 2009.
- [17] A. Ibanez, P. Bidaud, and V. Padois, “Emergence of humanoid walking behaviors from mixed-integer model predictive control,” in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2014.
- [18] H. Dai, A. Valenzuela, and R. Tedrake, “Whole-body motion planning with centroidal dynamics and full kinematics,” in *IEEE-RAS Int. Conf. on Humanoid Robots*, Madrid, Spain, 2014, pp. 295–302.
- [19] B. Ponton, A. Herzog, S. Schaal, and L. Righetti, “A convex model of humanoid momentum dynamics for multi-contact motion generation,” in *IEEE-RAS Int. Conf. on Humanoid Robots*, 2016.
- [20] A. W. Winkler, C. D. Bellicoso, M. Hutter, and J. Buchli, “Gait and Trajectory Optimization for Legged Syst. Through Phase-Based End-Effector Parameterization,” *IEEE Robot. and Autom. Lett.*, 2018.
- [21] B. Aceituno-Cabezas and et al., “Simultaneous Contact, Gait, and Motion Planning for Robust Multilegged Locomotion via Mixed-Integer Convex Optimization,” *IEEE Robot. and Autom. Lett.*, 2018.
- [22] Y. Ding, C. Li, and H. W. Park, “Single Leg Dynamic Motion Planning with Mixed-Integer Convex Optimization,” in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2018.
- [23] S. L. Cleach, T. Howell, M. Schwager, and Z. Manchester, “Fast contact-implicit model-predictive control,” *arXiv:2107.05616*, 2023.
- [24] T. A. Howell, S. Le Cleach, S. Singh, P. Florence, Z. Manchester, and V. Sindhvani, “Trajectory optimization with optimization-based dynamics,” *IEEE Robot. and Autom. Lett.*, vol. 7, no. 3, pp. 6750–6757, 2022.
- [25] J. Morimoto and C. G. Atkeson, “Minimax differential dynamic programming: An application to robust biped walking,” *Advances in Neural Information Processing Syst.*, 2003.
- [26] R. Budhiraja, J. Carpentier, and N. Mansard, “Dynamics consensus between centroidal and whole-body models for locomotion of legged robots,” in *IEEE Int. Conf. on Robot. and Autom.*, 2019.
- [27] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, “Crocodyl: An efficient and versatile framework for multi-contact optimal control,” in *IEEE Int. Conf. on Robot. and Autom.*, 2020, pp. 2536–2542.
- [28] S. Katayama and T. Ohtsuka, “Whole-body model predictive control with rigid contacts via online switching time optimization,” in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2022, pp. 8858–8865.
- [29] K. Mombaur, “Using optimization to create self-stable human-like running,” *Robotica*, vol. 27, no. 3, pp. 321–330, 2009.
- [30] A. Hereid, C. M. Hubicki, E. A. Cousineau, J. W. Hurst, and A. D. Ames, “Hybrid zero dynamics based multiple shooting optimization with Appl. to robotic walking,” in *IEEE Int. Conf. on Robot. and Autom.*, 2015, pp. 5734–5740.
- [31] Y. Hu and K. Mombaur, “Bio-inspired optimal control framework to generate walking motions for the humanoid robot iCub using whole body models,” *Applied Sci.s*, vol. 8, no. 2, 2018.
- [32] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, “CHOMP: Gradient optimization techniques for efficient motion planning,” in *IEEE Int. Conf. on Robot. and Autom.*, 2009, pp. 489–494.
- [33] M. Posa, S. Kuindersma, and R. Tedrake, “Optimization and stabilization of trajectories for constrained dynamical Syst.” in *IEEE Int. Conf. on Robot. and Autom.*, 2016.
- [34] A. Hereid, E. A. Cousineau, C. M. Hubicki, and A. D. Ames, “3D dynamic walking with underactuated humanoid robots: A direct collocation framework for optimizing hybrid zero dynamics,” in *IEEE Int. Conf. on Robot. and Autom.*, may 2016, pp. 1447–1454.
- [35] J. Carpentier and N. Mansard, “Multicontact Locomotion of Legged Robots,” *IEEE Trans. on Robot.*, 2018.
- [36] S. Feng, X. Xinjilefu, W. Huang, and C. G. Atkeson, “3D walking based on online optimization,” in *IEEE-RAS Int. Conf. on Humanoid Robots*, 2013.
- [37] F. Farshidian, M. Neunert, A. W. Winkler, G. Rey, and J. Buchli, “An efficient optimal planning and control framework for quadrupedal locomotion,” in *IEEE Int. Conf. on Robot. and Autom.*, 2017, pp. 93–100.
- [38] R. Grandia, F. Farshidian, R. Ranftl, and M. Hutter, “Feedback MPC for Torque-Controlled Legged Robots,” in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2019.
- [39] H. Dai and R. Tedrake, “Optimizing robust limit cycles for legged locomotion on unknown terrain,” in *IEEE Conf. on Decision and Control*, 2012.
- [40] H. Audren, J. Vaillant, A. Kheddar, A. Escande, K. Kaneko, and E. Yoshida, “Model preview control in multi-contact motion-application to a humanoid robot,” in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2014.
- [41] H. Dai and R. Tedrake, “Planning robust walking motion on uneven terrain via convex optimization,” in *IEEE-RAS Int. Conf. on Humanoid Robots*, 2016.
- [42] S. Caron and A. Kheddar, “Multi-contact walking pattern generation based on model preview control of 3D COM accelerations,” in *IEEE-RAS Int. Conf. on Humanoid Robots*, 2016.
- [43] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake, “Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot,” *Auton. Robots*, 2016.
- [44] X. Xiong and A. D. Ames, “Dynamic and Versatile Humanoid Walking via Embedding 3D Actuated SLIP Model with Hybrid LIP Based Stepping,” *IEEE Robot. and Autom. Lett.*, 2020.
- [45] M. Khadiv, A. Herzog, S. A. A. Moosavian, and L. Righetti, “Walking Control Based on Step Timing Adaptation,” in *IEEE Trans. on Robot.*, 2020.
- [46] P. Fernbach, S. Tonneau, O. Stasse, J. Carpentier, and M. Taix, “C-CROC: Continuous and Convex Resolution of Centroidal Dynamic Trajectories for Legged Robots in Multicontact Scenarios,” *IEEE Trans. on Robot.*, 2020.
- [47] S. Kajita and F. Kanehiro, “Biped walking pattern generation by using preview control of zero-moment point,” in *IEEE Int. Conf. on Robot. and Autom.*, 2003.
- [48] P. B. Wieber, “Viability and predictive control for safe locomotion,” in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2008.
- [49] A. Herdt, H. Diedam, P. B. Wieber, D. Dimitrov, K. Mombaur, and M. Diehl, “Online walking motion generation with automatic footstep placement,” *Advanced Robot.*, 2010.
- [50] I. Mordatch, M. D. Lasa, and A. Hertzmann, “Robust physics-based locomotion using low-dimensional planning,” in *ACM SIGGRAPH*, 2010.
- [51] A. Sherikov, D. Dimitrov, and P.-B. Wieber, “Whole body motion controller with long-term balance constraints,” in *IEEE-RAS Int. Conf. on Humanoid Robots*, nov 2014, pp. 444–450.
- [52] C. Brasseur, A. Sherikov, C. Collette, D. Dimitrov, and P. B. Wieber, “A robust linear MPC approach to online generation of 3D biped walking motion,” in *IEEE-RAS Int. Conf. on Humanoid Robots*, 2015.

- [53] M. A. Hopkins, D. W. Hong, and A. Leonessa, "Compliant locomotion using whole-body control and Divergent Component of Motion tracking," in *IEEE Int. Conf. on Robot. and Autom.*, 2015.
- [54] Y. Liu, P. M. Wensing, J. P. Schmiedeler, and D. E. Orin, "Terrain-Blind Humanoid Walking Based on a 3-D Actuated Dual-SLIP Model," *IEEE Robot. and Autom. Lett.*, 2016.
- [55] A. Herzog, S. Schaal, and L. Righetti, "Structured contact force optimization for kino-dynamic motion generation," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2016.
- [56] C. Mastalli, M. Focchi, I. Havoutis, A. Radulescu, S. Calinon, J. Buchli, D. G. Caldwell, and C. Semini, "Trajectory and foothold optimization using low-dimensional models for rough terrain locomotion," in *IEEE Int. Conf. on Robot. and Autom.*, 2017, pp. 1096–1103.
- [57] J. Di Carlo, P. M. Wensing, B. Katz, G. Bledt, and S. Kim, "Dynamic Locomotion in the MIT Cheetah 3 Through Convex Model-Predictive Control," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2018.
- [58] A. Kheddar and et al., "Humanoid robots in aircraft manufacturing: The Airbus use cases," *IEEE Robot. & Autom. Magazine*, vol. 26, no. 4, pp. 30–45, 2019.
- [59] J. Vaillant, A. Kheddar, H. Audren, F. Keith, S. Brossette, A. Escande, K. Bouyarmane, K. Kaneko, M. Morisawa, P. Gergondet, E. Yoshida, S. Kajita, and F. Kanehiro, "Multi-contact vertical ladder climbing with an HRP-2 humanoid," *Auton. Robots*, 2016.
- [60] G. Bledt, M. J. Powell, B. Katz, J. Di Carlo, P. M. Wensing, and S. Kim, "MIT Cheetah 3: Design and control of a robust, dynamic quadruped robot," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2018, pp. 2245–2252.
- [61] J. E. Slotine and M. W. Spong, "Robust robot control with bounded input torques," *J. of Robotic Syst.*, 1985.
- [62] F. T. Cheng, R. J. Sheu, and T. H. Chen, "The Improved Compact QP Method For Resolving Manipulator Redundancy," *IEEE Trans. on Syst., Man, and Cybernetics*, 1995.
- [63] C. Collette, A. Micaelli, C. Andriot, and P. Lemerle, "Dynamic Balance Control of Humanoids for Multiple Grasps and non Coplanar Frictional Contacts," in *Humanoids*, 2007.
- [64] J. Park, J. Haan, and F. C. Park, "Convex optimization algorithms for active balancing of humanoid robots," *IEEE Trans. on Robot.*, 2007.
- [65] Y. Abe, M. Da Silva, and J. Popović, "Multiobjective Control with Frictional Contacts," in *ACM SIGGRAPH/Eurographics*, 2007, pp. 249–258.
- [66] Y. Nesterov and A. Nemirovskii, *Interior-Point Polynomial Algorithms in Convex Programming*. SIAM, 1994.
- [67] E. D. Andersen, C. Roos, and T. Terlaky, "On implementing a primal-dual interior-point method for conic quadratic optimization," *Mathematical Programming, Series B*, 2003.
- [68] S.-H. Hyon, J. G. Hale, and G. Cheng, "Full-body compliant human-humanoid interaction: balancing in the presence of unknown external forces," *IEEE Trans. on Robot.*, vol. 23, no. 5, pp. 884–898, 2007.
- [69] B. J. Stephens and C. G. Atkeson, "Dynamic balance force control for compliant humanoid robots," in *IEEE/RSJ Int. Conf. on Intell. robots and Syst.*, 2010, pp. 1248–1255.
- [70] R. Featherstone, *Rigid body dynamics algorithms*. Springer Berlin, 2008, vol. 49.
- [71] M. Toussaint, K. R. Allen, K. A. Smith, and J. B. Tenenbaum, "Differentiable Physics and Stable Modes for Tool-Use and Manipulation Planning," *Robotics: Sci. and Syst.*, pp. 1–9, 2018.
- [72] F. R. Hogan, E. R. Grau, and A. Rodriguez, "Reactive Planar Manipulation with Convex Hybrid MPC," in *Int. Conf. on Robot. and Autom.*, 2018, pp. 247–253.
- [73] D. Stewart and J. C. Trinkle, "An implicit time-stepping scheme for rigid body dynamics with Coulomb friction," *Int. J. for Numerical Methods in Engineering*, vol. 39, no. 15, pp. 2673–2691, 1996.
- [74] J. Carpentier, R. Budhiraja, and N. Mansard, "Learning Feasibility Constraints for Multicontact Locomotion of Legged Robots," in *Robotics: Sci. and Syst.*, 2017.
- [75] A. D. Ames and M. Powell, "Towards the Unification of Locomotion and Manipulation through Control Lyapunov Functions and Quadratic Programs," in *Lecture Notes in Control and Information Sci.*, 2013.
- [76] Q. Nguyen and K. Sreenath, "Exponential Control Barrier Functions for enforcing high relative-degree safety-critical constraints," in *American Control Conf.*, 2016.
- [77] J. Engelsberger, A. Dietrich, G. Mesesan, G. Garofalo, C. Ott, and A. Albu-Schaeffer, "MPTC – Modular Passive Tracking Controller for stack of tasks based control frameworks," 2020.
- [78] C. Duriez, "Control of elastic soft robots based on real-time finite element method," in *IEEE Int. Conf. on Robot. and Autom.*, 2013.
- [79] D. W. Marhefka and D. E. Orin, "A compliant contact model with nonlinear damping for simulation of robotic Syst." *IEEE Trans. on Syst., Man, and Cybernetics Part A: Syst. and Humans.*, 1999.
- [80] R. Tedrake and the Drake Team Development, "Drake: A planning, control, and analysis toolbox for nonlinear dynamical Syst." 2016.
- [81] U. M. Ascher and L. R. Petzold, *Computer methods for ordinary differential equations and differential-algebraic equations*. SIAM, 1998.
- [82] K. H. Hunt and F. R. E. Crossley, "Coefficient of restitution interpreted as damping in vibroimpact," *J. of Applied Mechanics*, 1975.
- [83] E. Todorov, "Convex and analytically-invertible dynamics with contacts and constraints: Theory and implementation in MuJoCo," in *IEEE Int. Conf. on Robot. and Autom.*, 2014.
- [84] A. J. Van Der Schaft and J. M. Schumacher, *An introduction to hybrid dynamical Syst.* Springer London, 2000, vol. 251.
- [85] R. Goebel, R. G. Sanfelice, and A. R. Teel, *Hybrid Dynamical Syst.: Modeling, Stability, and Robustness*. Princeton University Press, 2012.
- [86] D. E. Stewart, "Rigid-body dynamics with friction and impact," *SIAM Review*, 2000.
- [87] B. Brogliato, *Nonsmooth mechanics : models, dynamics, and control*. Springer, 1999.
- [88] M. Anitescu and F. A. Potra, "Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems," *Nonlinear Dynamics*, vol. 14, no. 3, pp. 231–247, 1997.
- [89] M. Halm and M. Posa, "Modeling and analysis of non-unique behaviors in multiple frictional impacts," 2019.
- [90] A. Saccon, N. Van De Wouw, and H. Nijmeijer, "Sensitivity analysis of hybrid Syst. with state jumps with application to trajectory tracking," in *IEEE Conf. on Decision and Control*, 2014.
- [91] S. A. Burden, S. S. Sastry, D. E. Koditschek, and S. Revzen, "Event-selected vector field discontinuities yield piecewise-differentiable flows," *SIAM J. on Applied Dynamical Syst.*, 2016.
- [92] E. Heiden, D. Millard, E. Coumans, Y. Sheng, and G. S. Sukhatme, "NeuralSim: Augmenting differentiable simulators with neural networks," in *IEEE Int. Conf. on Robot. and Autom.*, 2021, pp. 9474–9481.
- [93] N. Fazeli, S. Zapolsky, E. Drumwright, and A. Rodriguez, "Fundamental Limitations in Performance and Interpretability of Common Planar Rigid-Body Contact Models," in *Int. Symp. of Robotic Res.*, 2017.
- [94] B. Acosta, W. Yang, and M. Posa, "Validating Robot Simulators on Real-World Impacts," *IEEE Robot. and Autom. Lett.*, vol. 7, no. 3, pp. 6471–6478, 2022.
- [95] K. Byl and R. Tedrake, "Approximate optimal control of the compass gait on rough terrain," in *IEEE Int. Conf. on Robot. and Autom.*, 2008.
- [96] E. R. Westervelt, J. W. Grizzle, and D. E. Koditschek, "Hybrid zero dynamics of planar biped walkers," *IEEE Trans. on automatic control*, vol. 48, no. 1, pp. 42–56, 2003.
- [97] R. Budhiraja, J. Carpentier, C. Mastalli, and N. Mansard, "Differential Dynamic Programming for Multi-Phase Rigid Contact Dynamics," in *IEEE-RAS Int. Conf. on Humanoid Robots*, 2019.
- [98] R. Bordalba, T. Schoels, L. Ros, J. M. Porta, and M. Diehl, "Direct Collocation Methods for Trajectory Optimization in Constrained Robotic Syst." *IEEE Trans. on Robot.*, 2022.
- [99] R. Deits and R. Tedrake, "Footstep planning on uneven terrain with mixed-integer convex optimization," in *IEEE-RAS Int. Conf. on Humanoid Robots*, nov 2014, pp. 279–286.
- [100] D. Song, P. Fernbach, T. Flayols, A. D. Prete, N. Mansard, S. Tonneau, and Y. J. Kim, "Solving Footstep Planning as a Feasibility Problem Using L1-Norm Minimization," *IEEE Robot. and Autom. Lett.*, vol. 6, no. 3, pp. 5961–5968, 2021.
- [101] N. Perrin, O. Stasse, L. Baudouin, F. Lamiroux, and E. Yoshida, "Fast humanoid robot collision-free footstep planning using swept volume approximations," *IEEE Trans. on Robot.*, 2012.
- [102] S. Tonneau, A. Del Prete, J. Pette, C. Park, D. Manocha, and N. Mansard, "An efficient acyclic contact planner for multiped robots," *IEEE Trans. on Robot.*, vol. 34, no. 3, 2018.
- [103] S. Dafarra, G. Romualdi, and D. Pucci, "Dynamic complementarity conditions and whole-body trajectory optimization for humanoid robot locomotion," *IEEE Trans. on Robot.*, vol. 38, no. 6, pp. 3414–3433, 2022.
- [104] F. de Avila Belbute-Peres, K. Smith, K. Allen, J. Tenenbaum, and J. Z. Kolter, "End-to-end differentiable physics for learning and control," *Advances in neural information processing Syst.*, vol. 31, 2018.
- [105] S. Pfrommer, M. Halm, and M. Posa, "ContactNets: Learning of Discontinuous Contact Dynamics with Smooth, Implicit Reps." in *Conf. on Robot Learning*, 2020.

- [106] R. Deits, T. Koolen, and R. Tedrake, "LVIS: Learning from value function intervals for contact-aware robot controllers," in *IEEE Int. Conf. on Robot. and Autom.*, vol. 2019-May, 2019, pp. 7762–7768.
- [107] A. Cauligi, P. Culbertson, E. Schmerling, M. Schwager, B. Stellato, and M. Pavone, "CoCo: Online mixed-integer control via supervised learning," *IEEE Robot. and Autom. Lett.*, vol. 7, no. 2, pp. 1447–1454, 2021.
- [108] J. Siekmann, Y. Godse, A. Fern, and J. Hurst, "Sim-to-real learning of all common bipedal gaits via periodic reward composition," in *IEEE Int. Conf. on Robot. and Autom.*, 2021, pp. 7309–7315.
- [109] H. J. T. Suh, T. Pang, and R. Tedrake, "Bundled gradients through contact via randomized smoothing," *IEEE Robot. and Autom. Lett.*, vol. 7, no. 2, pp. 4000–4007, 2022.
- [110] P. B. Wieber, "Holonomy and nonholonomy in the dynamics of articulated motion," in *Lecture Notes in Control and Inf Sci.s*, 2006.
- [111] P. M. Wensing and D. E. Orin, "Improved Computation of the Humanoid Centroidal Dynamics and Application for Whole-Body Control," *Int. J. of Humanoid Robot.*, vol. 13, no. 01, p. 1550039, 2016.
- [112] D. E. Orin, A. Goswami, and S.-H. Lee, "Centroidal dynamics of a humanoid robot," *Auton. Robots*, vol. 35, no. 2, pp. 161–176, 2013.
- [113] A. K. Valenzuela, "Mixed-integer convex optimization for planning aggressive motions of legged robots over rough terrain," Ph.D. dissertation, Massachusetts Institute of Technology, 2016.
- [114] S. Caron, Q. C. Pham, and Y. Nakamura, "Leveraging cone double description for multi-contact stability of humanoids with Appl. to statics and dynamics," in *Robotics: Sci. and Syst.*, 2015.
- [115] T. Bretl and S. Lall, "Testing static equilibrium for legged robots," *IEEE Trans. on Robot.*, 2008.
- [116] H. Audren and A. Kheddar, "3-D Robust Stability Polyhedron in Multicontact," *IEEE Trans. on Robot.*, 2018.
- [117] H. Hirukawa, S. Hattori, K. Harada, S. Kajita, K. Kaneko, F. Kanehiro, K. Fujiwara, and M. Morisawa, "A universal stability criterion of the foot contact of legged robots - Adios ZMP," in *IEEE Int. Conf. on Robot. and Autom.*, 2006.
- [118] Y. Ding, A. Pandala, C. Li, Y.-H. Shin, and H.-W. Park, "Representation-free model predictive control for dynamic motions in quadrupeds," *IEEE Trans. on Robot.*, vol. 37, no. 4, pp. 1154–1171, 2021.
- [119] Z. Zhou, B. Wingo, N. Boyd, S. Hutchinson, and Y. Zhao, "Momentum-Aware Trajectory Optimization and Control for Agile Quadrupedal Locomotion," *arXiv:2203.01548*, 2022.
- [120] Y. Sim and J. Ramos, "Tello Leg: The Study of Design Principles and Metrics for Dynamic Humanoid Robots," *IEEE Robot. and Autom. Lett.*, vol. 7, no. 4, pp. 9318–9325, 2022.
- [121] S.-H. Lee and A. Goswami, "Reaction mass pendulum (RMP): An explicit model for centroidal angular momentum of humanoid robots," in *IEEE RAS Int. Conf. on Robot. and Autom.*, 2007, pp. 4667–4672.
- [122] A. Saccon, S. Traversaro, F. Nori, and H. Nijmeijer, "On centroidal dynamics and integrability of average angular velocity," *IEEE Robot. and Autom. Lett.*, vol. 2, no. 2, pp. 943–950, 2017.
- [123] T. Koolen, M. Posa, and R. Tedrake, "Balance control using center of mass height variation: Limitations imposed by unilateral contact," in *IEEE-RAS Int. Conf. on Humanoid Robots*, 2016.
- [124] P. M. Wensing and D. E. Orin, "High-speed humanoid running through control with a 3D-SLIP model," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2013, pp. 5134–5140.
- [125] G. Piovani and K. Byl, "Approximation and Control of the SLIP Model Dynamics via Partial Feedback Linearization and Two-Element Leg Actuation Strategy," *IEEE Trans. on Robot.*, 2016.
- [126] Y. M. Chen and M. Posa, "Optimal Reduced-order Modeling of Bipedal Locomotion," in *IEEE Int. Conf. on Robot. and Autom.*, 2020.
- [127] A. Pandala, R. T. Fawcett, U. Rosolia, A. D. Ames, and K. A. Hamed, "Robust Predictive Control for Quadrupedal Locomotion: Learning to Close the Gap Between Reduced-and Full-Order Models," *IEEE Robot. and Autom. Lett.*, vol. 7, no. 3, pp. 6622–6629, 2022.
- [128] Z. Xie, X. Da, B. Babich, A. Garg, and M. van de Panne, "Glide: Generalizable quadrupedal locomotion in diverse environments with a centroidal model," *arXiv:2104.09771*, 2021.
- [129] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Sci. Robot.*, vol. 5, no. 47, p. eabc5986, 2020.
- [130] L. T. Biegler and V. M. Zavala, "Large-scale nonlinear programming using IPOPT: An integrating framework for enterprise-wide dynamic optimization," *Comp. & Chem. Eng.*, vol. 33, no. 3, pp. 575–582, 2009.
- [131] D. Malyuta, T. P. Reynolds, M. Szmuk, T. Lew, R. Bonalli, M. Pavone, and B. Acikmese, "Convex Optimization for Trajectory Generation," *IEEE Control Syst. Magazine*, vol. 42, no. 5, pp. 40–113, 2022.
- [132] R. Bellman, "The theory of dynamic programming," *Bulletin of the Amer. Math Society*, vol. 60, no. 6, pp. 503–515, 1954.
- [133] E. C. Whitman and C. G. Atkeson, "Control of instantaneously coupled Syst. applied to humanoid walking," in *IEEE-RAS Int. Conf. on Humanoid Robots*, 2010, pp. 210–217.
- [134] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin, "Decomposition of reachable sets and tubes for a class of nonlinear Syst.," *IEEE Trans. on Automatic Control*, vol. 63, no. 11, pp. 3675–3688, 2018.
- [135] D. Liberzon, *Calculus of variations and optimal control theory: a concise introduction*. Princeton university press, 2011.
- [136] R. Bonalli, B. Hérisse, and E. Trélat, "Optimal control of endoatmospheric launch vehicle Syst.: Geometric and computational issues," *IEEE Trans. Automat. Contr.*, vol. 65, no. 6, pp. 2418–2433, 2019.
- [137] S. Wright and J. Nocedal, "Numerical optimization," *Springer*, 1999.
- [138] J. T. Betts, *Practical methods for optimal control and estimation using nonlinear programming*. SIAM, 2010.
- [139] M. Diehl, H. G. Bock, H. Diedam, and P.-B. Wieber, "Fast direct multiple shooting algorithms for optimal robot control," in *Fast motions in biomechanics and Robot*. Springer, 2006, pp. 65–93.
- [140] H. G. Bock and K.-J. Plitt, "A multiple shooting algorithm for direct solution of optimal control problems," *IFAC Proceedings Volumes*, vol. 17, no. 2, pp. 1603–1608, 1984.
- [141] W. W. Hager, H. Hou, S. Mohapatra, and A. V. Rao, "Convergence Rate for an hp Collocation Method Applied to Unconstrained Optimal Control," *J. of Optimization Theory and Appl.*, vol. 169, no. 3, pp. 801–824, 2016.
- [142] O. Von Stryk and R. Bulirsch, "Direct and indirect methods for trajectory optimization," *Annals of operations Res.*, vol. 37, no. 1, pp. 357–373, 1992.
- [143] M. Kelly, "An introduction to trajectory optimization: How to do your own direct collocation," *SIAM Review*, 2017.
- [144] D. Pardo, M. Neunert, A. W. Winkler, R. Grandia, and J. Buchli, "Hybrid direct collocation and control in the constraint-consistent subspace for dynamic legged robot locomotion," in *Robotics: Sci. and Syst.*, 2017.
- [145] S. Moreno-Martin, L. Ros, and E. Celaya, "Collocation Methods for Second Order System," in *Robotics: Sci. and Syst.*, 2022.
- [146] N. S. Hoang, "Collocation Runge–Kutta–Nyström methods for solving second-order initial value problems," *Int. J. of Computer Mathematics*, vol. 98, no. 12, pp. 2423–2444, 2021.
- [147] D. Mayne, "A second-order gradient method for determining optimal trajectories of non-linear discrete-time Syst.," *Int. J. of Control*, vol. 3, no. 1, pp. 85–95, 1966.
- [148] D. M. Murray and S. J. Yakowitz, "Differential dynamic programming and Newton's method for discrete optimal control problems," *J. of Optimization Theory and Appl.*, vol. 43, no. 3, pp. 395–414, jul 1984.
- [149] E. Todorov and W. Li, "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic Syst.," in *American Control Conf.*, 2005.
- [150] M. Geisert, A. Del Prete, N. Mansard, F. Romano, and F. Nori, "Regularized Hierarchical Differential Dynamic Programming," *IEEE Trans. on Robot.*, vol. 33, no. 4, 2017.
- [151] J. N. Nganga and P. M. Wensing, "Accelerating second-order differential dynamic programming for rigid-body Syst.," *IEEE Robot. and Autom. Lett.*, vol. 6, no. 4, pp. 7659–7666, 2021.
- [152] S. Singh, R. P. Russell, and P. M. Wensing, "Analytical Second-Order Partial Derivatives of Rigid-Body Inverse Dynamics," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2022.
- [153] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," in *IEEE Int. Conf. on Robot. and Autom.*, 2014.
- [154] A. Pavlov, I. Shames, and C. Manzie, "Interior point differential dynamic programming," *IEEE Trans. on Control Syst. Technology*, vol. 29, no. 6, pp. 2720–2727, 2021.
- [155] Z. Xie, C. K. Liu, and K. Hauser, "Differential dynamic programming with nonlinear constraints," *IEEE Int. Conf. on Robot. and Autom.*, pp. 695–702, 2017.
- [156] T. A. Howell, B. E. Jackson, and Z. Manchester, "ALTRO: A Fast Solver for Constrained Trajectory Optimization," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2019.
- [157] W. Jallet, A. Bambade, N. Mansard, and J. Carpentier, "Constrained Differential Dynamic Programming: A primal-dual augmented La-

- grangian approach,” in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2022, pp. 13371–13378.
- [158] H. Li and P. M. Wensing, “Hybrid Syst. differential dynamic programming for whole-body motion planning of legged robots,” *IEEE Robot. and Autom. Lett.*, vol. 5, no. 4, pp. 5448–5455, 2020.
- [159] E. Pellegrini and R. P. Russell, “A Multiple-Shooting Differential Dynamic Programming Algorithm,” *AAS/AIAA Space Flight Mechanics Meeting*, no. March, pp. 759–778, 2017.
- [160] B. Plancher and S. Kuindersma, “A performance analysis of parallel differential dynamic programming on a GPU,” in *Int. Workshop on the Algorithmic Foundations of Robot.* Springer, 2018, pp. 656–672.
- [161] M. Gifftthaler, M. Neunert, M. Stäuble, J. Buchli, and M. Diehl, “A family of iterative gauss-newton shooting methods for nonlinear optimal control,” in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2018, pp. 1–9.
- [162] S. Shield, A. M. Johnson, and A. Patel, “Contact-Implicit Direct Collocation With a Discontinuous Velocity State,” *IEEE Robot. and Autom. Lett.*, vol. 7, no. 2, pp. 5779–5786, 2022.
- [163] G. Frison and M. Diehl, “HPIPM: a high-performance quadratic programming framework for model predictive control,” *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 6563–6569, 2020.
- [164] N. Mansard, A. Del Prete, M. Geisert, S. Tonneau, and O. Stasse, “Using a Memory of Motion to Efficiently Warm-Start a Nonlinear Predictive Controller,” in *IEEE Int. Conf. on Robot. and Autom.*, 2018, pp. 2986–2993.
- [165] K. Lowrey, A. Rajeswaran, S. Kakade, E. Todorov, and I. Mordatch, “Plan Online, Learn Offline: Efficient Learning and Exploration via Model-Based Control,” in *Int. Conf. on Learning Reps.*, 2018.
- [166] S. Gangapurwala, M. Geisert, R. Orsolino, M. Fallon, and I. Havoutis, “Rloc: Terrain-aware legged locomotion using reinforcement learning and optimal control,” *IEEE Trans. on Robot.*, 2022.
- [167] S. Katayama, M. Murooka, and Y. Tazaki, “Model predictive control of legged and humanoid robots: models and algorithms,” *Advanced Robot.*, vol. 37, no. 5, pp. 298–315, 2023.
- [168] C. Mastalli, S. P. Chhatoi, T. Corbères, S. Tonneau, and S. Vijayakumar, “Inverse-dynamics mpc via nullspace resolution,” *IEEE Trans. on Robot.*, 2023.
- [169] A. Meduri, P. Shah, J. Viereck, M. Khadiv, I. Havoutis, and L. Righetti, “Biconmp: A nonlinear model predictive control framework for whole body motion planning,” *IEEE Trans. on Robot.*, 2023.
- [170] J.-P. Sleiman, F. Farshidian, M. V. Minniti, and M. Hutter, “A unified mpc framework for whole-body dynamic locomotion and manipulation,” *IEEE Robot. and Autom. Lett.*, vol. 6, no. 3, pp. 4688–4695, 2021.
- [171] R. Grandia, F. Jenelten, S. Yang, F. Farshidian, and M. Hutter, “Perceptive locomotion through nonlinear model-predictive control,” *IEEE Trans. on Robot.*, 2023.
- [172] M. Diehl, R. Findeisen, F. Allgöwer, H. G. Bock, and J. P. Schlöder, “Nominal stability of real-time iteration scheme for nonlinear model predictive control,” *IEE Proc.-Control Theory and Appl.*, vol. 152, no. 3, pp. 296–308, 2005.
- [173] E. Dantec and et al., “Whole body model predictive control with a memory of motion: Experiments on a torque-controlled talos,” in *IEEE Int. Conf. on Robot. and Autom.*, 2021, pp. 8202–8208.
- [174] C. Samson, B. Espiau, and M. L. Borgne, *Robot control: the task function approach*. Oxford University Press, Inc., 1991.
- [175] O. Khatib, “A unified approach for motion and force control of robot manipulators: The operational space formulation,” *IEEE J. of Robot. and Autom.*, vol. 3, no. 1, pp. 43–53, 1987.
- [176] J. Sola, J. Deray, and D. Atchuthan, “A micro Lie theory for state estimation in Robot,” *arXiv:1812.01537*, 2018.
- [177] L. Saab, O. E. Ramos, F. Keith, N. Mansard, P. Soueres, and J. Y. Fourquet, “Dynamic whole-body motion generation under rigid contacts and other unilateral constraints,” *IEEE Trans. on Robot.*, 2013.
- [178] D. J. Agravante, G. Claudio, F. Spindler, and F. Chaumette, “Visual Servoing in an Optimization Framework for the Whole-Body Control of Humanoid Robots,” *IEEE Robot. and Autom. Lett.*, 2017.
- [179] S. Feng, E. Whitman, X. Xinjilefu, and C. G. Atkeson, “Optimization-based full body control for the DARPA Robot. Challenge,” *J. of Field Robot.*, 2015.
- [180] M. de Lasa, I. Mordatch, and A. Hertzmann, “Feature-based locomotion controllers,” *ACM Trans. on Graphics*, 2010.
- [181] P. M. Wensing and D. E. Orin, “Generation of dynamic humanoid behaviors through task-space control with conic optimization,” in *IEEE Int. Conf. on Robot. and Autom.*, 2013.
- [182] A. Herzog, N. Rotella, S. Mason, F. Grimminger, S. Schaal, and L. Righetti, “Momentum control with hierarchical inverse dynamics on a torque-controlled humanoid,” *Auton. Robots*, vol. 40, 2016.
- [183] K. C. Park, P. H. Chang, and S. H. Kim, “The Enhanced Compact QP Method for Redundant Manipulators Using Practical Inequality Constraints,” in *IEEE Int. Conf. on Robot. and Autom.*, 1998.
- [184] F. Kanehiro, M. Morisawa, W. Suleiman, K. Kaneko, and E. Yoshida, “Integrating geometric constraints into reactive leg motion generation,” in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2010.
- [185] W. Decré, R. Smits, H. Bruyninckx, and J. De Schutter, “Extending iTaSC to support inequality constraints and non-instantaneous task specification,” in *IEEE Int. Conf. on Robot. and Autom.*, 2009.
- [186] A. Del Prete, “Joint Position and Velocity Bounds in Discrete-Time Acceleration / Torque Control of Robot Manipulators,” *IEEE Robot. and Autom. Lett.*, vol. 3, no. 1, 2018.
- [187] N. Mansard, “A dedicated solver for fast operational-space inverse dynamics,” in *IEEE Int. Conf. on Robot. and Autom.*, 2012.
- [188] J. Park and O. Khatib, “Contact consistent control framework for humanoid robots,” in *IEEE Int. Conf. on Robot. and Autom.*, 2006.
- [189] L. Sentis, J. Park, and O. Khatib, “Compliant control of multicontact and center-of-mass behaviors in humanoid robots,” *IEEE Trans. on Robot.*, 2010.
- [190] Fujimoto Y., Obata S., and Kawamura A., “Robust Biped Walking with Active Interaction Control between Foot and Ground,” in *IEEE Int. Conf. on Robot. and Autom.*, 1998.
- [191] K. Bouyarmane and A. Kheddar, “Using a multi-objective controller to synthesize simulated humanoid robot motion with changing contact configurations,” in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2011.
- [192] T. Koolen and et al., “Summary of Team IHMC’s Virtual Robot. Challenge Entry,” *IEEE-RAS Int. Conf. on Humanoid Robots*, 2013.
- [193] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl, “qpOASES: A parametric active-set algorithm for quadratic programming,” *Mathematical Programming Computation*, vol. 6, no. 4, pp. 327–363, 2014.
- [194] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, “OSQP: An operator splitting solver for quadratic programs,” *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020.
- [195] M. Chignoli and P. M. Wensing, “Variational-based optimal control of underactuated balancing for dynamic quadrupeds,” *IEEE Access*, vol. 8, pp. 49 785–49 797, 2020.
- [196] W. Yang and M. Posa, “Impact invariant control with Appl. to bipedal locomotion,” in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2021, pp. 5151–5158.
- [197] S. Fahmi, M. Focchi, A. Radulescu, G. Fink, V. Barasuol, and C. Semini, “Stance: Locomotion adaptation over soft terrain,” *IEEE Trans. on Robot.*, vol. 36, no. 2, pp. 443–457, 2020.
- [198] G. Romualdi, S. Dafarra, and D. Pucci, “Modeling of visco-elastic environments for humanoid robot motion control,” *IEEE Robot. and Autom. Lett.*, vol. 6, no. 3, pp. 4289–4296, 2021.
- [199] T. Flayols, A. Del Prete, M. Khadiv, N. Mansard, and L. Righetti, “Reactive Balance Control for Legged Robots under Visco-Elastic Contacts,” *Applied Sci.*, vol. 11, no. 353, 2021.
- [200] O. Kanoun, F. Lamiroux, P.-B. Wieber, F. Kanehiro, E. Yoshida, and J.-P. Laumond, “Prioritizing linear equality and inequality Syst.: Application to local motion planning for redundant robots,” in *IEEE Int. Conf. on Robot. and Autom.*, 2009.
- [201] A. Escande, N. Mansard, and P. B. Wieber, “Hierarchical quadratic programming: Fast online humanoid-robot motion generation,” *Int. J. of Robot. Res.*, 2014.
- [202] D. Dimitrov, A. Sherikov, and P.-B. Wieber, “Efficient resolution of potentially conflicting linear constraints in Robot.” aug 2015.
- [203] C. Van Loan, “On the Method of Weighting for Equality-Constrained Least-Squares Problems,” *SIAM J. on Numerical Analysis*, 1985.
- [204] K. Pfeiffer, A. Escande, and A. Kheddar, “Singularity Resolution in Equality and Inequality Constrained Hierarchical Task-Space Control by Adaptive Nonlinear Least Squares,” *IEEE Robot. and Autom. Lett.*, 2018.
- [205] J. Carpentier, G. Saurel, G. Buondonno, J. Mirabel, F. Lamiroux, O. Stasse, and N. Mansard, “The Pinocchio C++ library: A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives,” *IEEE/SICE Int. Symp. on System Integration*, pp. 614–619, 2019.
- [206] M. Lutter, C. Ritter, and J. Peters, “Deep Lagrangian networks: Using physics as model prior for deep learning,” *Int. Conf. on Learning Reps.*, pp. 1–17, 2019.

- [207] R. Calandra, S. Ivaldi, M. P. Deisenroth, E. Rueckert, and J. Peters, "Learning inverse dynamics models with contacts," *IEEE Int. Conf. on Robot. and Autom.*, vol. 2015-June, no. June, pp. 3186–3191, 2015.
- [208] Viceconte, P. M. et al., "ADHERENT: Learning Human-like Trajectory Generators for Whole-body Control of Humanoid Robots," *IEEE Robot. and Autom. Lett.*, vol. 7, no. 2, pp. 2779–2786, 2021.
- [209] C. Daniel, M. Viering, J. Metz, O. Kroemer, and J. Peters, "Active Reward Learning," *Robotics: Sci. and Syst.*, 2014.
- [210] W. X. Merkt, V. Ivan, T. Dinev, I. Havoutis, and S. Vijayakumar, "Memory clustering using persistent homology for multimodality- And discontinuity-sensitive learning of optimal control warm-starts," *IEEE Trans. on Robot.*, vol. 37, no. 5, pp. 1649–1660, 2021.
- [211] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model predictive control: theory, computation, and design*. Nob Hill Publishing, 2017.
- [212] A. Gazar, M. Khadiv, A. Del Prete, and L. Righetti, "Stochastic and robust mpc for bipedal locomotion: A comparative study on robustness and performance," in *IEEE-RAS Int. Conf. on Humanoid Robots*, 2021, pp. 61–68.
- [213] J. Carpentier and P.-B. Wieber, "Recent progress in legged robots locomotion control," *Current Robot. Reports*, vol. 2, no. 3, pp. 231–238, 2021.
- [214] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Sci. Robot.*, vol. 7, no. 62, p. eabk2822, 2022.
- [215] Q. L. Lidec, L. Montaut, C. Schmid, I. Laptev, and J. Carpentier, "Augmenting differentiable physics with randomized smoothing," *arXiv:2206.11884*, 2022.
- [216] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," in *Int. Conf. on Learning Reps.*, 2016.
- [217] S. Levine and V. Koltun, "Guided policy search," in *Int. Conf. on Machine Learning*. PMLR, 2013, pp. 1–9.
- [218] Z. Zhao, S. Zuo, T. Zhao, and Y. Zhao, "Adversarially regularized policy learning guided by trajectory optimization," in *Learning for Dynamics and Control Conf.* PMLR, 2022, pp. 844–857.
- [219] K. Chatzilygeroudis, V. Vassiliades, F. Stulp, S. Calinon, and J.-B. Mouret, "A survey on policy search algorithms for learning robot controllers in a handful of trials," *IEEE Trans. on Robot.*, vol. 36, no. 2, pp. 328–347, 2019.
- [220] M. Zanon and S. Gros, "Safe reinforcement learning using robust MPC," *IEEE Trans. on Automatic Control*, vol. 66, no. 8, pp. 3638–3652, 2020.
- [221] B. Recht, "A tour of reinforcement learning: The view from continuous control," *Annual Review of Control, Robotics, and Autonomous Syst.*, vol. 2, pp. 253–279, 2019.



Patrick M. Wensing (Senior Member, IEEE) received the B.S., M.S., and Ph.D. degrees in electrical and computer engineering from The Ohio State University, Columbus, OH, USA, in 2009, 2013, and 2014, respectively.

He is currently an Associate Professor with the Department of Aerospace and Mechanical Engineering, University of Notre Dame, where he directs the Robotics, Optimization, and Assistive Mobility (ROAM) Laboratory. Before joining Notre Dame, he was a Postdoctoral Associate with MIT, working on

control system design for the MIT cheetah robots. His current research interest includes aspects of dynamics, optimization, and control toward advancing the mobility of legged robots and assistive devices. He received the NSF CAREER Award in 2020. He serves as the Co-Chair for the IEEE RAS Technical Committee on Model-Based Optimization for Robotics and as an Associate Editor for the IEEE ROBOTICS AND AUTOMATION LETTERS.



Michael Posa (Member, IEEE) received the B.S. and M.S. degrees in mechanical engineering from Stanford University in 2007 and 2008 respectively. He received Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2017.

He is currently an Assistant Professor of mechanical engineering and applied mechanics with the University of Pennsylvania, where he is a member of the General Robotics, Automation, Sensing and Perception (GRASP) Lab. He holds secondary appointments in electrical and systems engineering and in computer and information science. He leads to Dynamic Autonomy and Intelligent Robotics (DAIR) Lab, which focuses on developing computationally tractable algorithms to enable robots to operate both dynamically and safely as they maneuver through and interact with their environments, with applications including legged locomotion and manipulation. He has received multiple awards, including the NSF CAREER Award, the RSS Early Career Spotlight Award, and best paper awards. He is an Associate Editor for the IEEE TRANSACTIONS ON ROBOTICS.



Yue Hu (Member, IEEE) received the Doctorate degree in computer science (Dr. rer. nat.) from Heidelberg University, Germany in 2017. She is currently an Assistant Professor at the Department of Mechanical and Mechatronics Engineering at the University of Waterloo, Canada, and is head of the Active & Interactive Robotics Lab. Prior to joining Waterloo, she was a JSPS Fellow at the CNRS-AIST Joint Robotics Laboratory (JRL) in Japan between 2018 and 2020 and an Assistant Professor at the Tokyo University of Agriculture and Technology in

2021. She also serves as a Co-Chair for the IEEE RAS Technical Committee on Model-Based Optimization for Robotics and is an Associate Editor for IEEE ROBOTICS AND AUTOMATION LETTERS. Her research interests are human-robot interaction, collaborative robots, human factors, optimal control, and humanoid robots.



Adrien Escande received the M.S. degree in science and executive engineering from École des Mines de Paris, Paris, France, in 2005, and the Ph.D. degree in robotics from Université d'Évry Val-d'Essonne, Essonne, France, in 2008.

He worked at CEA-LIST in Fontenay-aux-Roses, Paris, France, and at the CNRS-AIST Joint Robotics Laboratory (JRL), in Tsukuba, Japan. He is now researcher at the Inria centre at the University Grenoble Alpes, France. His research interests include whole-body planning and control for multi-body robots, numerical optimization for robotics and inexact computations to reduce control energy consumption. He currently serves as the Co-Chair for the IEEE RAS Technical Committee on Model-Based Optimization for Robotics and was Associate Editor for the IEEE ROBOTICS AND AUTOMATION LETTERS.



Nicolas Mansard received the MSc in computer science of Univ. Grenoble in 2003 and the PhD in robotics of Univ. Rennes in 2006. He was a post-doctoral researcher in Stanford Univ. with O. Khatib in 2007 and in JRL-Japan with A. Kheddar in 2008. He was an invited researcher in Univ. Washington with E. Todorov in 2014. He is currently a CNRS researcher ("chargé de recherche CNRS") since 2009.

He received the CNRS Bronze Medal in 2015 (one medal is awarded in France in automatic/robotic/signal-processing every year). His main research interests include the motion generation, planning and control of complex robots, with a special regard in humanoid robotics. His expertise covers sensor-based (vision and force) control, numerical mathematics for control, bipedal locomotion and locomotion planning. He published more than 70 papers in international journals and conferences and supervised 10 PhD thesis. He was previously an associate editor of the IEEE TRANSACTIONS ON ROBOTICS.



Andrea Del Prete (Member, IEEE) received the degree (Hons.) in computer engineering from the University of Bologna (Italy) in 2009, and the Ph.D. degree in robotics from the Cognitive Humanoids Laboratory (Italian Institute of Technology, Genova, Italy) in 2013. From 2014 to 2017 he was a Post-doctoral Researcher with the Laboratoire d'Analyse et d'Architecture des Systemes (CNRS, Toulouse, France). In 2018, he had worked as Senior Researcher at the Max Planck Institute for Intelligent Systems (Tubingen, Germany). He is currently an

Associate Professor in the Industrial Engineering Department at the University of Trento (Italy).

His research focuses on the use of optimization and learning algorithms for control, planning and estimation of autonomous robots, with a special focus on legged locomotion. He has published more than 50 papers in international journals and conferences and he serves as Associate Editor for the IEEE ROBOTICS AND AUTOMATION LETTERS.