2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 25-29, 2020, Las Vegas, NV, USA (Virtual)

Waste Not, Want Not: Lessons in Rapid Quadrupedal Gait Termination from Thousands of Suboptimal Solutions

Stacey Shield[1]*, and Amir Patel[1],

***Abstract*— Elaborate trajectory optimization models with many degrees of freedom can be a useful locomotion-planning tool, as they provide rich solutions that take advantage of the robot's specific morphology. They are, however, prone to falling into local minima. Depending on the seed that initializes the solver, the trajectories themselves and the extent to which they minimize the cost function can vary widely, making it impossible to judge the quality of any solution without generating many more. In this paper, we argue that this perceived drawback can actually be a powerful advantage in exploratory studies, since the resulting set of diverse motions can reveal which features tend to be associated with good performance, and therefore aid in the formulation of strategies for executing challenging maneuvers. We selected rapid gait termination from a high-speed gallop as our case study - a dangerous and scarcely-researched movement. By analyzing a set of over 3000 monopedal and quadrupedal trajectories, we were able to extract conclusions about how braking and sliding should be performed to reduce the stopping distance, and identify a hindlimb action that creates large braking forces.**

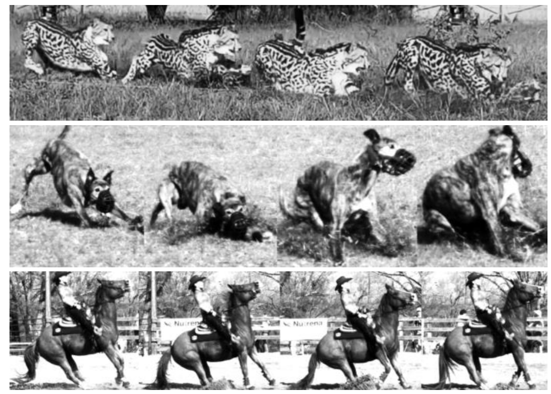

Fig. 1. Cheetah, greyhound and horse stopping. Footage courtesy of Dr. Robert Gilette and the Canine Performance Sciences Program at the Auburn University College of Veterinary Medicine.

I. INTRODUCTION

Fast, dynamic gait is certainly not a solved problem in legged robotics, but compared to high-speed maneuverability, it is a well-defined one. We have a vocabulary of words like ”canter”, ”gallop” and ”trot” to describe steady-state locomotion, so even if robots cannot yet execute these motions with the same power, efficiency and robustness as the fastest animals, we can at least specify which foot should go forward next. Bringing the robot to a sudden standstill from one of these fast gaits is a more mysterious problem.

If some quadrupedal robot of the future is barrelling along in the 20 m/s rotary gallop of a racing greyhound, and detects that an unsuspecting obstacle has wandered into its path, what should it do next? This question still has to be answered at the planning level, let alone the level of a control policy.

We could look for inspiration in the results of the 3.5 billion-year-old evolutionary algorithm that has already been applied to legged locomotion, but perhaps because high-speed gait termination is reserved for situations so desperate that the danger inherent in performing the maneuver is only surpassed by the danger of *not* performing it, few studies of human [1][2] or animal [3] subjects have been conducted. Instances of rapid deceleration in footage of dogs, cheetahs and horses show the maneuver being performed in various ways (Fig. 1), but many more examples would be required to determine whether these differences should be attributed to speed, morphology, or external conditions.

Trajectory optimization has the potential to provide a useful alternative in the form of synthetic data. The simulated subjects might not capture the full dynamic complexity of their real-world counterparts, but they have the advantages of uniform, known, and infinitely-adjustable parameters, and no fear of harm.

While optimization-based approaches have an established record of effective application in legged locomotion studies concerning steady-state gait [4][5][6][7][8][9], and even agile maneuvers [10][11][12][13][14], the models used are often highly simplified. When more complicated whole-body models are used, one of the many difficulties encountered is the problem of local minima. With more degrees of freedom available, there may be myriad ways to achieve similar costs, and initiating the solver from different seeds can yield solutions of widely-varying quality. This makes it impossible to tell how good a solution is without the context of many others seeded from diverse points.

This is a hindrance if the goal is to find the *One True Solution* to be implemented on a specific robot, but it can be a strong advantage when the aim is more exploratory. When you have nothing to work with but an approximate idea of how a motion should start and end, the power of trajectory optimization is less in its ability to find a path that minimizes some cost than its ability to find a path at all. If its purpose is thought of foremost as trajectory *generation*, the suboptimal solutions go from discards to data: by analyzing them in aggregate, it is possible to identify recurring motion features that correlate with better performance, and thereby extract lessons about the successful execution of a maneuver that could guide its eventual implementation on a robot.

In this paper, we apply this approach to a preliminary study into rapid, high-speed deceleration in quadrupeds. Using both a monopedal model and a whole-body quadrupedal model, we generate families of incrementally-improving solutions through trajectory optimization, resulting in a set of over 3000 stopping motions. By finding the features that separate the better solutions from the worse ones, we were able to infer how braking and sliding should be performed to improve the stopping distance, and identify a hindlimb action that can momentarily increase the braking force.

Especially considering how dangerous a poorly-executed stop could be at the speeds considered, these insights into a sparsely-researched movement would be difficult to obtain soley through experiments with animals or robots, so they are a contribution in themselves. The main objective of the paper, however, is to demonstrate the potential for individually-useless suboptimal solutions to become useful collectively. As such, we believe the way trajectory optimization is used in this study could be applicable to many other motion-planning problems throughout robotics.

II. Background

Rapid termination of high-speed gaits is near-unexplored research territory. The few studies that address legged deceleration at all primarily concern a case as far-removed from the sprinting cheetah as it is possible to be: termination of bipedal walking [15][16][17][18][19][20], often in the elderly or pathological cases [21][22]. Still, we can identify some fundamental principles and extrapolate how they might transfer from geriatrics to greyhounds:

A. Mechanics of deceleration

Using the spring-mass conception of legged locomotion rendered into a control scheme by Raibert [23], deceleration is achieved by placing the center of pressure (COP) further ahead of the centre of mass than the *neutral point* - the position that would result in zero net acceleration over the stride [23][15]. This causes the portion of the stride where the leg functions as a damper (compressing and absorbing energy) to be longer than the portion where it functions as a motor (extending and expelling energy) [23]. Two factors restrict the velocity it is possible to reduce to rest in one stride: balance and actuator power [24][3]. If these limits are exceeded, the subject must take another step or they will topple forward [24]. Motion of other appendages such as the arms [17] or a tail [25] can assist with balance by absorbing undesired angular momentum. In quadrupeds, it is suggested that the primary limit shifts from balance to power as velocity increases [3].

Decreasing the stopping time or distance must ultimately come down to applying more opposing external force to the body, but this trade-off between braking and balance opens up two possible approaches: generating greater force magnitudes is the obvious option, but maximizing braking time - that is, reducing the time spent in non-contact states, and using dynamic compensation to extend the safe duration over which the feet can remain in contact without toppling - could have a similar effect.

B. Limitations of previous studies

1) Sliding: Prior studies on gait termination have made the assumption that the stance foot must remain stationary. When low-friction surfaces have been included [17], avoiding slipping is the primary concern. Failure to consider sliding may be the result of failure to consider velocity: all the animals shown in Fig. 1 skidded to a halt, indicating that slipping might be unavoidable at higher speeds. With balance in mind, slipping might even be desirable earlier in the maneuver, as stopping the feet while the body is still moving fast seems likely to end in the subject tumbling posterior over paws.

If sliding does turn out to be an inextricable aspect of fast gait termination, we will need to specify a way to do it safely. A widely-used metric for assessing the dynamic stability of robots is the Zero Rate of Angular Momentum (ZRAM) criterion [26], which states that the subject is stable if no external moments are acting on it. Practically, this is satisfied when the ground reaction force vector passes through the centre of mass. We can adapt this criterion to define a stable sliding condition as follows:

1) The body is positioned such that the angle θ_G of the COP-COM vector corresponds to the angle of kinetic friction, θ_μ
2) There is no relative velocity ${}^G\dot{x}_P$ between the COM and COP.

2) Minimal-DOF models: Another limitation of past work is the widespread use of spring-mass monopedal templates [27]. Although these models can describe changes in velocity [23], including transient deceleration within gait, or even gradual gait termination, a sudden stop from high-speed could depart from the basic form of constant-speed locomotion drastically enough to warrant a model that was not primarily designed for periodic strides. These simple models also miss morphology-specific strategies that could improve performance, such as the use of limbs to assist in balancing, and cannot give any information about foot contact sequences.

While we will be working primarily with a whole-body quadrupedal model to develop a more detailed idea of how this maneuver should be executed, we will also include a monopedal template in the study. One reason for this is to counter the unavoidable generality problem of the whole-body model. Our goal is to draw conclusions for all quadrupeds, but any model that is specifically quadrupedal cannot be truly abstract, because it ultimately requires parameters like limb masses, inertias and segment lengths to be defined. The results will therefore always come with the caveat that they may only hold for quadrupeds that are sufficiently similar to the chosen model, but supporting them with the results of a more abstract model could provide some indication of broader applicability. The monopod will also offer a useful sanity check by making our results more directly comparable with prior work.

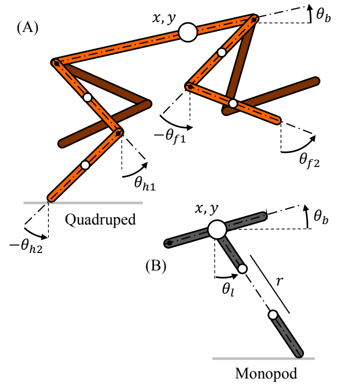

Fig. 2. Planar quadruped (A) and half-quadruped (B) models.

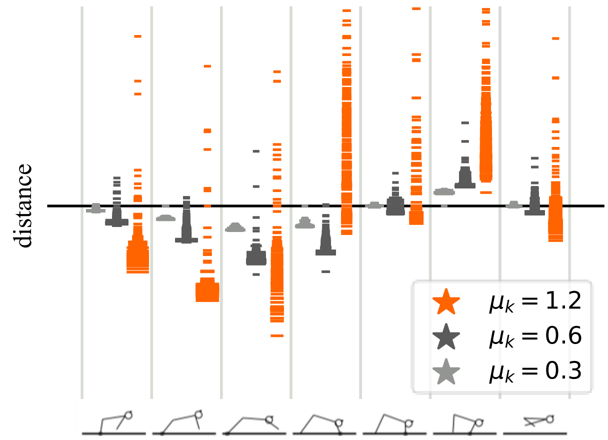

Fig. 3. Normalized stopping distances on surfaces with different coefficients of friction (μ_k) for termination motions initiated from various points in the galloping gait cycle. Adapted from [33]

III. Trajectory Optimization

A. Models

The models used in this paper are shown in Fig. 2. The primary model (Fig. 2A) is an 11-DOF planar quadruped. Each leg is actuated by two revolute joints, with the directions of the second joints corresponding to an 'X' configuration. Rather than using the relative joint angles as generalized coordinates, the angles of all segments are referenced counter-clockwise from the global vertical axis [28]. The leg segments are of equal length, and the fully-extended leg is the same length as the body link, l_b. The mass of the body and leg links are $0.6m$ and $0.05m$, respectively, where m is the total mass of the model. All links have the COM in the middle, except the body link, where it is situated $0.4l_b$ from the shoulder joint.

The force and power limits for the models were selected to be the minimal values necessary for each model to move at the desired average velocity of 30 body lengths per second - equivalent to the speed of a greyhound [29]. These values were identified through the generation of force-optimal and power-optimal trajectories for two tasks: a symmetrical gait cycle at that speed (constrained to match the characteristics of a rotatory gallop for the quad) and acceleration from rest to that speed. The maximum normal force acting on the feet was constrained to three body weights, the peak value observed in galloping animals [29].

B. Direct Collocation

The time for each trajectory was discretized into $N = 100$ finite elements, each consisting of three collocation points. Constraints describing approximate integration by three-point Radau quadrature linked the state at one node to the state at the next, as described in [30].

C. Contact Model

Hard unilateral contact constraints are imposed between the model's feet and the ground, and as hard end-stops at all joints so the model can hit its range-of-motion (ROM) limits at speed. We applied the complementarity-based implicit contact scheme described in [31] and [30] to allow these contacts to occur without any predefined mode order. When implemented in this way, the foot contacts take the form of inelastic collisions with the possibility of sliding under a Coulomb friction coefficient μ_k. When the foot is stationary, a larger static friction coefficient μ_s is applied. To ensure that the friction coefficient would not make stopping without sliding impossible, we selected high values: $\mu_k = 1.2$ and $\mu_s = 2.4$

To make the associated complementarity constraints more tractable, we solve them using a penalty technique [32]: the ith such constraint is set equal to a penalty variable $p_i \geq 0$, and the sum of these penalties is then added to the objective function to be minimized.

Because the model can only change its contact state at the boundary of a finite element, it is helpful to introduce some flexibility by letting the duration of the ith element h_i vary around a master timestep h_M according to $0.8h_M \leq h_i \leq 1.2h_M$.

D. Boundary value problem

The initial condition for the trajectory was sampled from a galloping gait cycle generated using the quadrupedal model. Selecting a favourable point was crucial: there is a *critical region* of the gait cycle within which gait termination must be initiated, or else the motion will be clumsy and difficult to control [19][18]. In a previous study conducted using a simplified quadrupedal model [33], we determined that the critical region for a rotatory gallop falls between hind stance and foreleg touchdown, whereas initiating the maneuver from contracted positions requires more corrective motion. The effects of initiation point on performance are illustrated in Fig. 3.

Based on this result, we selected the apex of the extended flight phase as the initial state. Likewise, the apex of the flight phase was sampled from a hopping trajectory as the starting point for the monopod.

For the final condition, we considered the gait to be terminated when the body was no longer moving forward, i.e. when the velocity of the body link, $\dot{x}_b$, had been reduced to zero (or less) and all feet were grounded. We did not want the final stabilizing motions to affect the results, so in our analysis, we cropped the trajectories to the point that $\dot{x}_b$ first reaches zero. The stopping time and distance metrics used throughout the paper therefore refer to the change in time and COM x position from the initial state until this moment. The purpose of the grounded foot condition is to avoid trajectories that would lead into dives or falls.

E. Iteratively Minimizing Stopping Distance

Stopping distance is a challenging objective to minimize as it is closely linked to time. Despite the slight variability in the timestep, the simulation time is effectively a constant parameter in these problems, and should remain so, as it would make the problem very difficult to solve otherwise. Technically, there is nothing to stop the model from stopping in much less than the available time, but setting an intentionally over-generous value is not a desirable option, as it effectively reduces the resolution of the solution by decreasing the number of nodes used for the maneuver. An overly-long time also leaves many more possibilities for local minima open. Attempting to minimize distance using a cost function is therefore unlikely to turn up a solution approaching the true minimum unless the minimal stopping time is somehow already known.

To get around this problem, we instead minimize time and distance by iteratively *squeezing* them in an outer loop until a feasible solution can no longer be found:

1) The solver is initialized with a two-step process: first, a procedurally-generated smooth-random *silly walk*[31] is given as a guess to solve a simplified version of the problem, where two of the collocation points are deactivated and first-order integration is used. This solution then initializes the first attempt to solve the full-scale problem.
2) For the first iteration, the simulation time is assigned a random value. If it converges, the master timestep h_m for the next attempt is decreased by 10 percent, and an upper bound is placed on x, restricting it to 0.98 of the previous stopping distance.
3) The previous solution becomes the guess for the next iteration, and the process is repeated until the problem fails to converge, or the complementarity penalties can no longer be minimized to acceptable values.

The only objective applied in each solving iteration is minimizing the complementarity penalties, so in terms of the stopping distance problem, each solution should be regarded as feasible result, rather than even a local minumum, but the overall effect of the itertively-decreasing upper bound on x is to minimize distance. Anecdotally, we did find that the final distances were comparable to those achieved when a distance-minimizing cost function was used, but this process was much less failure-prone. Additionally, this iterative method is well-suited for the type of exploratory study we are conducting, as it produces families of gradually evolving, incrementally improving motions.

F. Solving

The optimization problem was written in Python using the Pyomo library [34][35] and solved using the IPOPT algorithm [36] combined with the MA97 linear solver from the HSL solver library [37].

IV. Results and Discussion

As an overview of the solutions we obtained, the stopping times and distances for both models are plotted in Fig. 4. To facilitate a clearer comparison between the two models, we normalized their performance by comparing it to box benchmarks: the time T_μ and distance x_μ that a rigid mass (or, "box") of equivalent weight would take to stop from the same velocity, sliding under the same kinetic friction conditions.

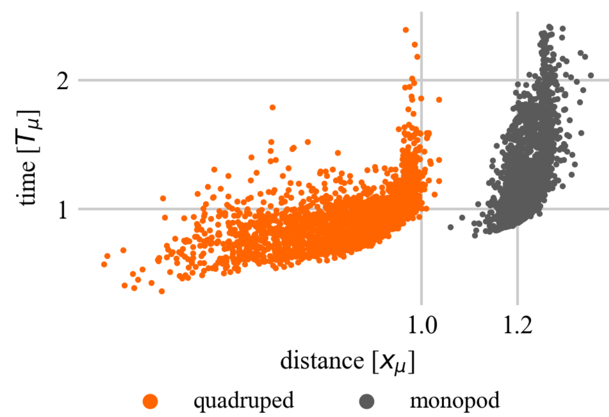

Fig. 4. Stopping distance and time for quadrupedal and monopedal models vs. an equivalent mass sliding on a surface with the same friction coefficient.

We can view this comparison to the box as an indicator of the extent to which a model is able to use articulation to its advantage. Performing better than the box shows that it is using relative motion of body segments or limbs to generate normal forces exceeding its weight. Consistently failing to meet the box benchmark suggests static instability, as a statically stable model should be able to imitate the box by planting its feet and sliding in a fixed pose.

Static stability is one aspect of the quadrupedal configuration that monopedal templates cannot capture. Due to this limitation, the monopod was unable to beat the box's distance, while the worst quadruped solutions at least matched it. The limitations of the template are also clear in the relatively small improvement from its worst results to its best, compared to the much wider range of performance for the model with more degrees of freedom.

Because the monopod must lift and re-position its leg to maintain balance, it cannot apply a consistent braking force throughout the maneuver. This is shown in the upper half of Fig. 5, which plots the portion of the total time that was spent actively braking. While the monopod was forced to spend, at minimum, around 10 percent of its time in the air,

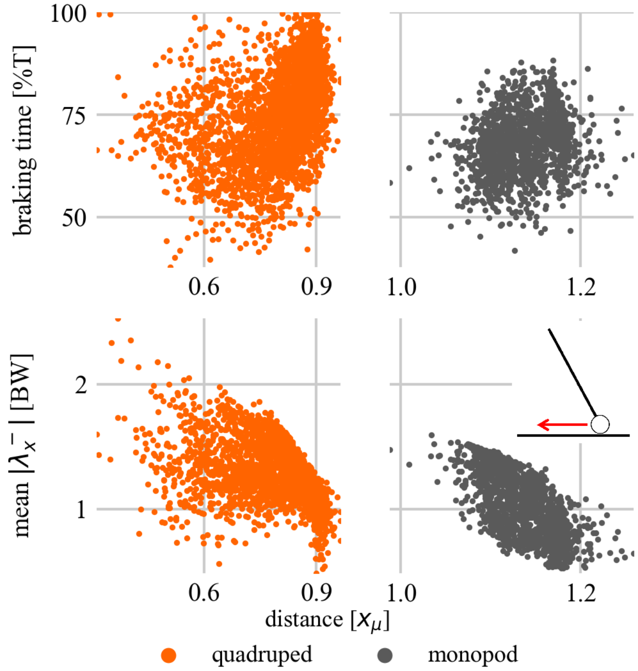

Fig. 5. Percentage of the stopping time spent braking (top) and time-averaged magnitude of the total braking force applied to each model (bottom). Maintaining contact throughout the maneuver does not necessarily decrease the stopping distance, but increasing the braking force does.

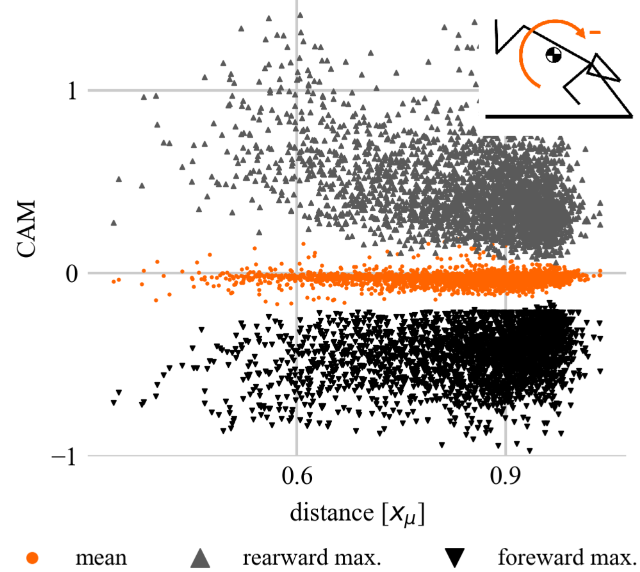

Fig. 6. Centroidal angular momentum of the quadruped. The forward maximum is of particular interest as this represents the dangerous toppling that large, rearward ground reaction forces could cause.

it was possible for the quadruped to maintain contact for the full duration. This does not, however, mean that doing so is necessarily favorable, as prolonged braking does not appear to lead better results for either model. As might be expected, the solutions that do maintain complete contact tend to fall close to the box benchmark. Of course, it is not desirable to spend the majority of the time in flight, either, so the solutions at the extremes - spending either the most or the least time in the air - tend to fall on the less-successful half of the stopping distance spread.

In combination with the results shown in the lower part of Fig. 5, which plots the time-averaged magnitude of the decelerating forces, it is clear that a braking harder is a more effective strategy than braking longer. The applicability of this this lesson to a particular robot is obviously contingent on the strength of its limbs and actuators, but it demonstrates the potential usefulness of this sort of sweeping maneuverability study as a preliminary step in the design process. Much like the way that trajectory optimization has been used to speculate about the evolution of cursorial birds' legs to handle the strain induced by specific actions [38], it could inform the mechanical design of robot limbs by indicating where maneuvers of interest are likely to create the most strain in the system.

Another possible concern with this strategy could be that the application of larger braking forces will come with the drawback of decreased stability. We used the centroidal angular momentum (CAM)[], plotted in Fig. 6[39][14], as a metric for the dynamic stability of the model. A large value in either direction is undesirable, with forward rotation being the most critical, as this would indicate the forward toppling that these rearward-directed forces tend to induce. We see that the model was able to stop in a shorter distance without increasing the peak CAM beyond that experienced during less-effective motions, and even the most rapid trajectories still maintained a mean CAM around zero. (They did, however, tend to lead to larger peak values for rearward rotation, for a reason that will be discussed later in this section.)

These results lead to two follow-up questions:

1) How are larger braking forces generated in the superior solutions?
2) How is the model able to maintain stability under the effect of those forces?

There was no single variable that correlated directly with the average braking force or with the forward CAM, but we can identify some contributing factors:

1) Actuator force: Fig. 8 shows the mean actuator force and power exerted by the monopod's prismatic leg for the time window in which the largest deceleration occurred. As would be expected, pushing harder into the ground is an effective way to generate larger braking forces. The results for the quadruped support this, but they do not make for a compelling plot, as both joints hit their torque limits even for trajectories showing only modest improvements over the box distance. So how are the solutions that do significantly better able to exert more force once their actuators have saturated?

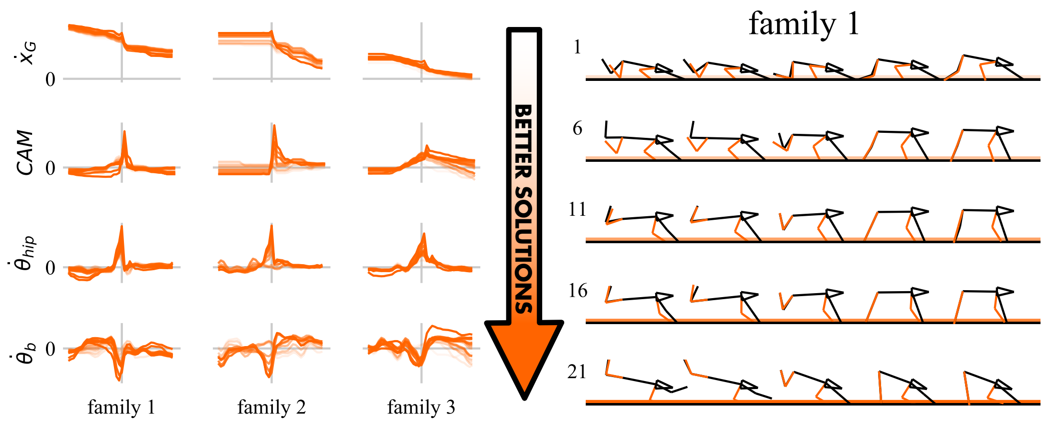

Fig. 7. Examples of the hind leg swing motion that emerged in many of the solution families during maximal acceleration, and its effects on centroidal angular momentum (CAM) and the angular velocity of the body. Family 1 is animated in the supplementary video.

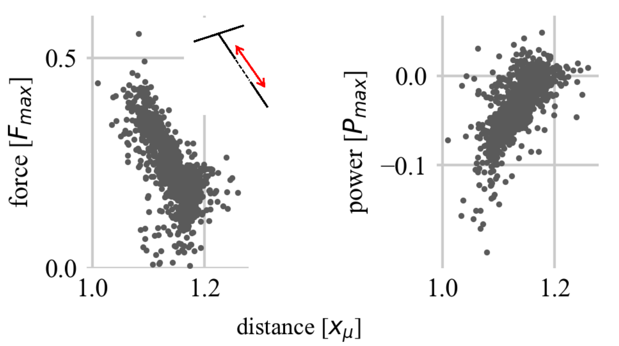

Fig. 8. Mean force and power in the monopod's prismatic joint during maximal acceleration. (The quadruped is excluded as both its joints reached their torque and power limits for nearly all the trajectories.)

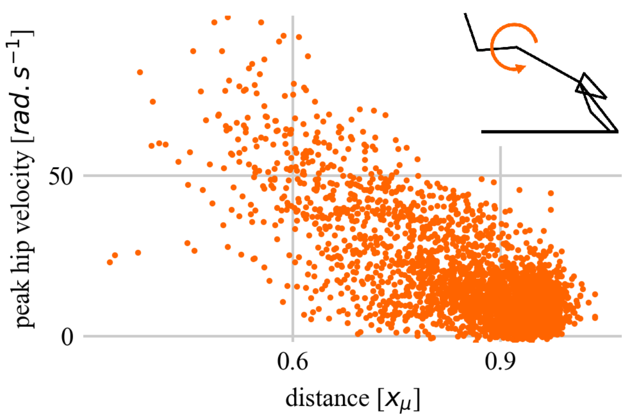

Fig. 9. Peak side-averaged hip velocity during maximal braking for the quadruped. Better-performing trajectories tended to exhibit higher hip velocities, suggesting the forward-swinging action of the hind legs illustrated in Fig. 7.

2) Hind leg swing: An advantage of the iterative way we generated the trajectories is that it allows for the identification of specific features that emerge incrementally in the gait waveforms as the performance improves. A feature that developed in many of the solution families was a rapid forward swing of one or both hind legs occurring at the same time as a sudden, steep deceleration. Often, the instantaneous braking force at this moment was the largest achieved in the trajectory. Representative examples of the leg swing are illustrated in Fig. 7 for three families.

Plotting the peak hip velocity in a window around the largest instantaneous deceleration value (Fig. 9) suggests that this feature is widespread in the data, and correlated with improved stopping distance.

We hypothesize that the leg swing performs two functions: firstly, it is responsible for the large braking force, as the opposing reaction of the front half of the body acts to push the forelimbs down, increasing the normal force and, consequently, the friction. If the feet slam into the ground at the end of the swing, this further contributes to the decelerating force. Secondly, it counteracts the external pitching moment caused by this force to the extent that the centroidal angular momentum is directed rearward during the swinging motion - hence, the tendency for the peak CAM to be larger for the better-performing solutions. Although the plots of the body's angular velocity show that it does experience some forward pitching due to the opposing torque at the hip, it can immediately be corrected following the swing by the now-grounded hind legs.

When the limb loading is considered (Fig. 11), it seems that the hind legs are more useful in this ballast role than as brakes: predictably, the majority of the braking force was exerted by the forelegs, though the model did tend to spend similar amounts of time in double stance (with both a hind- and a foreleg on the ground) and front stance. The model is even able to stop using foreleg braking exclusively, but the solutions which did this were not especially successful. This potentially advocates for the addition of a dedicated ballast

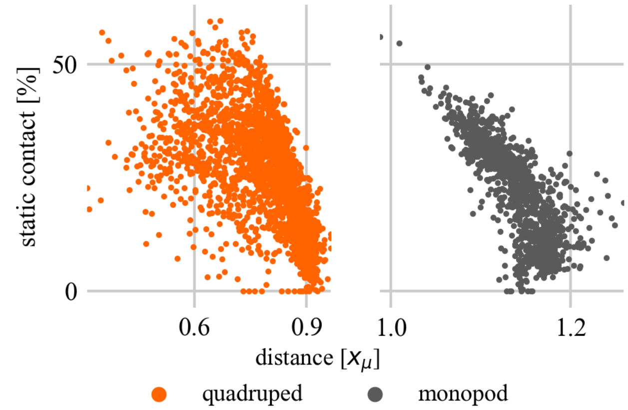

Fig. 10. Percentage of the total braking force exerted in static contact mode.

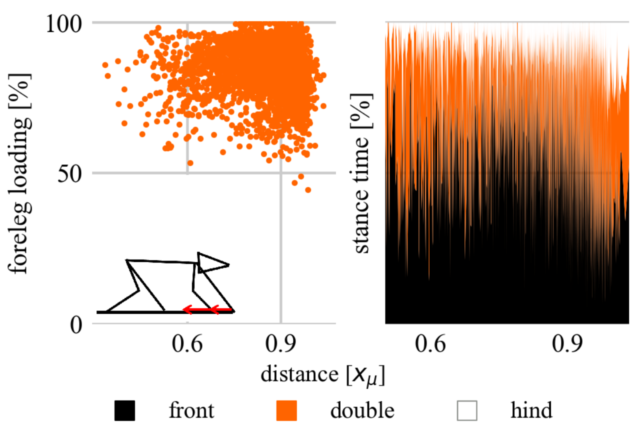

Fig. 11. Percentage of the total braking force exerted by the forelimbs (left) and time spent in different stance configurations (right) for the quadruped.

limb, such as a tail [25].

A. Sliding

Another way that the braking force could be increased was through increased use of static braking. This is shown in Fig. 10, which plots the portion of the total applied braking force that was exerted while the foot was stationary. Despite this trend, and the high coefficients of friction we selected for these experiments, all but a few solutions slid more than they stuck. This indicates that sliding must be incorporated into an effective high-speed stopping strategy, as it is either advantageous to some extent or nearly impossible to avoid.

The proposed condition for dynamically stable sliding could provide a target for controlling the sliding motion, as behaviour that loosely adheres to this standard emerged in both models. With the exception of some of the low-quality quadruped motions, the COM angle converges around the angle of friction, indicating that these trajectories tended to adhere to the ZRAM criterion [26] by keeping the COM in

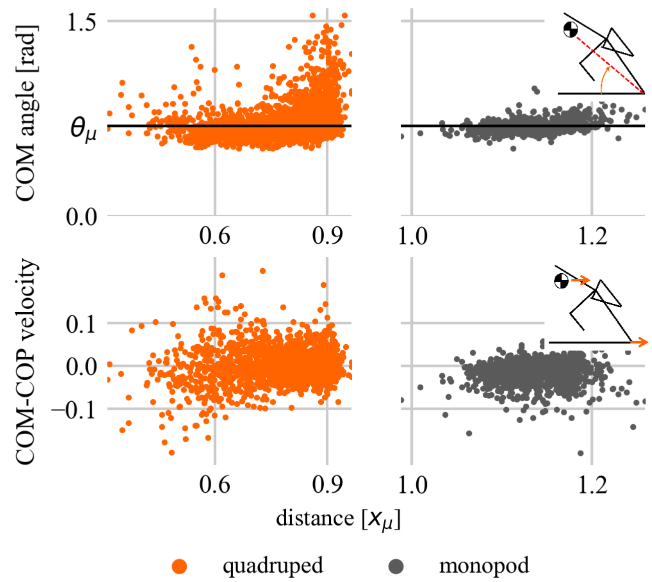

Fig. 12. Median angle of the COM vector, relative to the ground, compared to the angle of friction (top), and median velocity of the COM relative to the COP (bottom). The relative velocities are scaled to the COM velocity: for each point in the trajectory, the difference between the COM x velocity and COP x velocity was divided by the COM x velocity, and the median of these values was plotted.

line with the ground reaction force vector. They also avoided large differences between the velocities of the COM and COP, with these discrepancies typically falling within 10 percent of the COM velocity's magnitude. These quantities are plotted in Fig. 12.

V. Conclusions

We were able to extract the following lessons in rapid gait termination from a dataset of suboptimal solutions:

- Maximizing the magnitude of braking forces is more effective than maximizing the duration of contact.
- Once the forelegs are pushing into the ground at the maximum capacity of their actuators, the normal force can be increased by rapidly swinging the hindlegs forward. The forelegs perform most of the braking function, so a control scheme could conceivably prioritize keeping the hindlegs free to use as a ballast.
- More friction can be generated if the feet stick rather than slide, and though this should be taken advantage of, sliding might be impossible to avoid altogether. To maintain dynamic stability during slipping, the body should be positioned so the COM angle matches the angle of friction, and the relative velocity between the feet and COM should be minimized.

Besides providing the first steps towards executing this challenging maneuver on a quadrupedal robot, the purpose of this paper was to demonstrate that suboptimal solutions to trajectory optimization problems can still provide useful insights into motion when analyzed together. Synthetic datasets that combine quantity with varying quality are potentially a useful exploratory tool that both the robotics and biomechanics fields can use to elucidate complicated locomotion problems - especially those that would be difficult or dangerous to investigate in reality.

REFERENCES

[1] G. M. Cesar and S. M. Sigward, "Dynamic stability during running gait termination: Differences in strategies between children and adults to control forward momentum," *Human movement science*, vol. 43, pp. 138–145, 2015.

[2] G. M. Cesar and S. M. Sigward, "Dynamic stability during running gait termination: Predictors for successful control of forward momentum in children and adults," *Human movement science*, vol. 48, pp. 37–43, 2016.

[3] S. B. Williams, H. Tan, J. R. Usherwood, and A. M. Wilson, "Pitch then power: limitations to acceleration in quadrupeds," *Biology letters*, vol. 5, no. 5, pp. 610–613, 2009.

[4] K. Mombaur, "Using optimization to create self-stable human-like running," *Robotica*, vol. 27, no. 3, pp. 321–330, 2009.

[5] G. Schultz and K. Mombaur, "Modeling and optimal control of human-like running," *IEEE/ASME Transactions on mechatronics*, vol. 15, no. 5, pp. 783–792, 2009.

[6] W. Xi and C. D. Remy, "Optimal gaits and motions for legged robots," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3259–3265, IEEE, 2014.

[7] Z. Gan, Z. Jiao, and C. D. Remy, "On the dynamic similarity between bipeds and quadrupeds: a case study on bounding," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3614–3621, 2018.

[8] C. Gehring, S. Coros, M. Hutter, M. Bloesch, P. Fankhauser, M. A. Hoepflinger, and R. Siegwart, "Towards automatic discovery of agile gaits for quadrupedal robots," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4243–4248, IEEE, 2014.

[9] M. Srinivasan, "Fifteen observations on the structure of energy-minimizing gaits in many simple biped models," *Journal of The Royal Society Interface*, vol. 8, no. 54, pp. 74–98, 2010.

[10] C. Hubicki, M. Jones, M. Daley, and J. Hurst, "Do limit cycles matter in the long run? stable orbits and sliding-mass dynamics emerge in task-optimal locomotion," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5113–5120, IEEE, 2015.

[11] C. Fisher, C. Hubicki, and A. Patel, "Do intermediate gaits matter when rapidly accelerating?," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3418–3424, 2019.

[12] C. Gehring, S. Coros, M. Hutler, C. D. Bellicoso, H. Heijnen, R. Diethelm, M. Bloesch, P. Fankhauser, J. Hwangbo, M. Hoepflinger, *et al.*, "Practice makes perfect: An optimization-based approach to controlling agile motions for a quadruped robot," *IEEE Robotics & Automation Magazine*, vol. 23, no. 1, pp. 34–43, 2016.

[13] A. Blom and A. Patel, "Investigation of a bipedal platform for rapid acceleration and braking manoeuvres," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 426–432, IEEE, 2018.

[14] H. Dai, A. Valenzuela, and R. Tedrake, "Whole-body motion planning with centroidal dynamics and full kinematics," in *2014 IEEE-RAS International Conference on Humanoid Robots*, pp. 295–302, IEEE, 2014.

[15] Y. Jian, D. A. Winter, M. G. Ishac, and L. Gilchrist, "Trajectory of the body cog and cop during initiation and termination of gait," *Gait & Posture*, vol. 1, no. 1, pp. 9–22, 1993.

[16] K. Hase and R. Stein, "Analysis of rapid stopping during human walking," *Journal of neurophysiology*, vol. 80, no. 1, pp. 255–261, 1998.

[17] A. Oates, A. Patla, J. Frank, and M. Greig, "Control of dynamic stability during gait termination on a slippery surface," *Journal of Neurophysiology*, vol. 93, no. 1, pp. 64–70, 2005.

[18] M. Bishop, D. Brunt, N. Pathare, and B. Patel, "The effect of velocity on the strategies used during gait termination," *Gait & posture*, vol. 20, no. 2, pp. 134–139, 2004.

[19] P. Vanitchatchavan, "Termination of human gait," in *2009 IEEE International Conference on Systems, Man and Cybernetics*, pp. 3169–3174, IEEE, 2009.

[20] M. Qiao and D. L. Jindrich, "Leg joint function during walking acceleration and deceleration," *Journal of biomechanics*, vol. 49, no. 1, pp. 66–72, 2016.

[21] F. W. O'Kane, C. A. McGibbon, and D. E. Krebs, "Kinetic analysis of planned gait termination in healthy subjects and patients with balance disorders," *Gait & posture*, vol. 17, no. 2, pp. 170–179, 2003.

[22] W. Sparrow and O. Tirosh, "Gait termination: a review of experimental methods and the effects of ageing and gait pathologies," *Gait & posture*, vol. 22, no. 4, pp. 362–371, 2005.

[23] M. H. Raibert, *Legged robots that balance*. MIT press, 1986.

[24] Y.-C. Pai and J. Patton, "Center of mass velocity-position predictions for balance control," *Journal of biomechanics*, vol. 30, no. 4, pp. 347–354, 1997.

[25] A. Patel and M. Braae, "Rapid acceleration and braking: Inspirations from the cheetah's tail," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 793–799, IEEE, 2014.

[26] A. Goswami and V. Kallem, "Rate of change of angular momentum and balance maintenance of biped robots," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, vol. 4, pp. 3785–3790, IEEE, 2004.

[27] R. J. Full and D. E. Koditschek, "Templates and anchors: neuromechanical hypotheses of legged locomotion on land," *Journal of experimental biology*, vol. 202, no. 23, pp. 3325–3332, 1999.

[28] A. Knemeyer, S. L. Shield, A. Patel, A. Del Prete, and A. Kheddar, "Minor change, major gains: The effect of orientation formulation on solving time for multi-body trajectory optimization," *IEEE Robotics and Automation Letters*, 2020.

[29] P. E. Hudson, S. A. Corr, and A. M. Wilson, "High speed galloping in the cheetah (acinonyx jubatus) and the racing greyhound (canis familiaris): spatio-temporal and kinetic characteristics," *Journal of Experimental Biology*, vol. 215, no. 14, pp. 2425–2434, 2012.

[30] A. Patel, S. L. Shield, S. Kazi, A. M. Johnson, and L. T. Biegler, "Contact-implicit trajectory optimization using orthogonal collocation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2242–2249, 2019.

[31] A. P. Stacey Shield, "On the effectiveness of silly walks as initial guesses for optimal legged locomotion problems," in *2020 Southern African Universities Power Engineering Conference/Robotics and Mechatronics/Pattern Recognition Association of South Africa (SAUPEC/RobMech/PRASA)*, pp. 211–216, IEEE, 2020.

[32] B. Baumrucker and L. Biegler, "Mpec strategies for optimization of a class of hybrid dynamic systems," *Journal of Process Control*, vol. 19, no. 8, pp. 1248–1256, 2009.

[33] S. Shield and A. Patel, "Investigating rapid gait termination with synthetic data." presented at the Symposium on Comparative Biomechanics across Organizational Scales (Tissues to Whole Body Dynamics) at ISB/ASB 2019, 2019.

[34] W. E. Hart, C. D. Laird, J.-P. Watson, D. L. Woodruff, G. A. Hackebeil, B. L. Nicholson, and J. D. Siirola, *Pyomo–optimization modeling in python*, vol. 67. Springer Science & Business Media, second ed., 2017.

[35] W. E. Hart, J.-P. Watson, and D. L. Woodruff, "Pyomo: modeling and solving mathematical programs in python," *Mathematical Programming Computation*, vol. 3, no. 3, pp. 219–260, 2011.

[36] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, no. 1, pp. 25–57, 2006.

[37] HSL, "A collection of fortran codes for large-scale scientific computation," *See http://www.hsl.rl.ac.uk*.

[38] A. V. Birn-Jeffery, C. M. Hubicki, Y. Blum, D. Renjewski, J. W. Hurst, and M. A. Daley, "Don't break a leg: running birds from quail to ostrich prioritise leg safety and economy on uneven terrain," *Journal of Experimental Biology*, vol. 217, no. 21, pp. 3786–3796, 2014.

[39] D. E. Orin, A. Goswami, and S.-H. Lee, "Centroidal dynamics of a humanoid robot," *Autonomous robots*, vol. 35, no. 2-3, pp. 161–176, 2013.