TAMOLS: Terrain-Aware Motion Optimization for Legged Systems

Fabian Jenelten, Ruben Grandia, Farbod Farshidian, and Marco Hutter

***Abstract*—Terrain geometry is, in general, non-smooth, non-linear, non-convex, and, if perceived through a robot-centric visual unit, appears partially occluded and noisy. This work presents the complete control pipeline capable of handling the aforementioned problems in real-time. We formulate a trajectory optimization problem that jointly optimizes over the base pose and footholds, subject to a heightmap. To avoid converging into undesirable local optima, we deploy a graduated optimization technique. We embed a compact, contact-force free stability criterion that is compatible with the non-flat ground formulation. Our experiments demonstrate stair climbing, walking on stepping stones, and over gaps, utilizing various dynamic gaits.**

***Paper Type* – Recent Work [1], under review**

I. Introduction

LEGGED locomotion has been studied and designed for the last couple of decades. Recent advances in both software and hardware have triggered the transition from experimental platforms used under laboratory conditions to (semi-) autonomous machines deployed in real-world scenarios, e.g., on industrial sites for inspection [2] or in underground mines for exploration and mapping [3]. Yet, assumptions made in "classical" control approaches limit applications to flat or mildly rough ground [4]–[7] or restrict locomotion to static stability [8]. However, the true potential of legged locomotion is undeniably rooted in a combination of rough environments and dynamic agility. More recent control approaches have eliminated both the flat-ground and static-gait restrictions but struggle to match the computational overhead with the onboard compute budget [4], [9], [10].

When it comes to rough terrain locomotion, we can isolate three major challenges: 1) Low computation time and generalizability over various terrains are contradicting requirements. Hence, it is essential to strike a balance between simplicity and complexity. 2) When embedding the true terrain geometry, many local optima might appear, rendering a motion optimization problem sensitive to the initial guess. 3) Sensors used to generate perceptive data are mounted onboard, and large parts of the field of view might appear occluded.

The main contributions of this work can be split into three independent parts, each addressing one of these challenges. 1) First, we introduce a contact-force free and fully differentiable dynamic stability metric. Compared to Single Rigid Body Dynamics (SRBD) [9], [12]–[14], our model eliminates the contact forces and has thus a smaller problem size; opposed to contact wrench cone (CWC) models [15]–[19], we do not depend on algorithmic construction of the feasible set; and compared to Zero-Moment Point (ZMP) constraints [20]–[22], our model generalizes to uneven ground with vertical contact planes. 2) We integrate this model into a trajectory optimization (TO) framework, hereinafter referred to as TAMOLS (terrain-aware motion optimization for legged systems). It optimizes simultaneously footholds and base pose at frequencies larger than 100 Hz. 3) Finally, we complete our work by integrating a map processing pipeline.

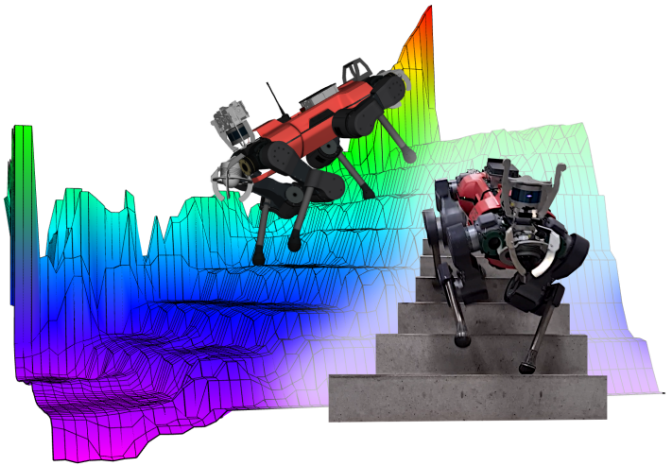

Fig. 1. Right: ANYmal [11] is walking down a stair case using an ambling gait. Left: Corresponding gazebo model visualization.

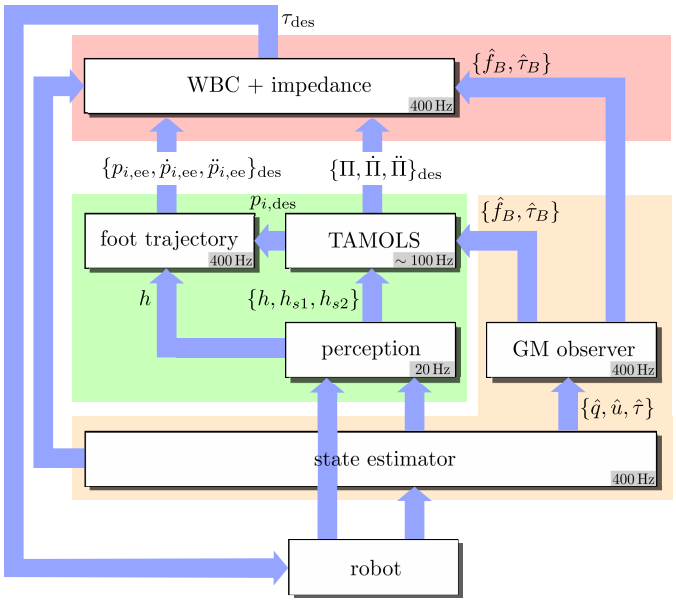

Fig. 2. Overview of the control structure, consisting of trajectory generation (green), tracking (red), and state/disturbance observer (orange). Foot positions are denoted by $\boldsymbol{p}_{\text{ee}}$, base pose by $\boldsymbol{\Pi}$, and elevation maps by h.

II. Method

A. Control Structure

Generalized positions $\hat{\boldsymbol{q}}$, generalized velocities $\hat{\boldsymbol{u}}$ and torques $\hat{\boldsymbol{\tau}}$ are estimated/measured at 400 Hz. The point clouds

gathered from two on-board LiDARs are processed to different elevation maps $\{h,\, h_{s1},\, h_{s2}\}$ at a rate of 20 Hz. The three elevation layers are tightly embedded in the TO problem and the swing trajectory generation. These two modules output desired reference signals for the base pose and swinging end-efectors, which are tracked by a whole-body controller (WBC). External disturbances of the base $\{\hat{\boldsymbol{f}}_B,\, \hat{\boldsymbol{\tau}}_B\}$, estimated by a generalized momentum (GM) observer, are compensated in the tracking controller and the motion optimizer. The control architecture used in this work is outlined in Fig. 2.

B. The GIAC model

We start by introducing the following set of assumptions:

Assumption 1. *1) The limbs have zero mass, 2) the rate of change of the angular momentum has a negligible effect on the contact forces, 3) contact forces can only push on the ground, 4) friction coefficients are constant, 5) contacts are established on horizontal planes, and 6) the base position is located above all the grounded feet.*

Under the first assumption of 1, a legged robot can be modeled as SRBD with frictional constraints, i.e.

$$m\begin{bmatrix}\ddot{\boldsymbol{p}}_B - \boldsymbol{g}\\ \boldsymbol{0}\end{bmatrix} + \begin{bmatrix}\boldsymbol{0}\\ \dot{\boldsymbol{L}}_B\end{bmatrix} = \sum_{i=1}^{N}\begin{bmatrix}\boldsymbol{f}_i\\ (\boldsymbol{p}_i - \boldsymbol{p}_B)\times \boldsymbol{f}_i\end{bmatrix} \tag{1a}$$

$$\boldsymbol{f}_i \in \mathcal{F}_i, \quad \forall i \in \{1,\dots,N\}, \tag{1b}$$

with m the mass, $\boldsymbol{p}_B$ the base position, $\boldsymbol{f}_i$ the contact force at the ith contact location $\boldsymbol{p}_i$, $\mathcal{F}_i$ the friction cone, $\boldsymbol{L}_B$ the angular momentum of the base, and N the number of grounded feet.

By exploiting the remaining assumptions, we can eliminate the contact forces from (1a) and (1b) and and obtain an analytic expression for the friction cone,

$$\mu \boldsymbol{e}_z^T \cdot \boldsymbol{a}_B \geq ||(\mathbb{I}_{3\times 3} - \boldsymbol{e}_z\boldsymbol{e}_z^T)\cdot \boldsymbol{a}_B|| \qquad N > 0 \tag{2a}$$

$$m\det(\boldsymbol{p}_{ij}, \boldsymbol{p}_B - \boldsymbol{p}_i, \boldsymbol{a}_B) \leq \boldsymbol{p}_{ij}^T \cdot \dot{\boldsymbol{L}}_B \qquad N \geq 3 \tag{2b}$$

$$m\det(\boldsymbol{p}_{12}, \boldsymbol{p}_B - \boldsymbol{p}_1, \boldsymbol{a}_B) = \boldsymbol{p}_{12}^T \cdot \dot{\boldsymbol{L}}_B \qquad N = 2 \tag{2c}$$

$$\det(\boldsymbol{e}_z, \boldsymbol{p}_{12}, \boldsymbol{M}_1) \geq 0 \qquad N = 2 \tag{2d}$$

$$m(\boldsymbol{p}_B - \boldsymbol{p}_i)\times \boldsymbol{a}_B - \dot{\boldsymbol{L}}_B = \boldsymbol{0} \qquad N = 1 \tag{2e}$$

$$\boldsymbol{a}_B = 0 \quad \dot{\boldsymbol{L}}_B = \boldsymbol{0} \qquad N = 0, \tag{2f}$$

where $\boldsymbol{a}_B = \boldsymbol{g} - \ddot{\boldsymbol{p}}_B$ is the gravito-inertia acceleration (GIA) vector, μ the friction coefficient, $\boldsymbol{p}_{ij} = \boldsymbol{p}_j - \boldsymbol{p}_i$, $\boldsymbol{M}_1 = (\boldsymbol{p}_B - \boldsymbol{p}_1)\times(\boldsymbol{g} - \ddot{\boldsymbol{p}}_B) - \dot{\boldsymbol{L}}_B/m$, and where we have used the determinant notation $\det(\boldsymbol{a},\boldsymbol{b},\boldsymbol{c}) = \boldsymbol{a}\cdot(\boldsymbol{b}\times\boldsymbol{c})$.

We can summarize the most important properties, which are partially visualized in Fig. 3, as follows: 1) The constraints (2b) to (2f) can be interpreted as the largest inscribing convex cone $\tilde{\mathcal{C}}$ of the *gravito-inertia acceleration cone (GIAC)* $\mathcal{C}$, which is defined as the convex hull of rays connecting the base position with the footholds. Rate of change of the angular momentum changes shape and size of these cones. 2) On flat ground, (2b) to (2f) simplify to the well known ZMP stability criterion. 3) Under assumption 1, the constraints (2) can be shown to be *weak contact stable*, a property shared with the closely related CWC models.

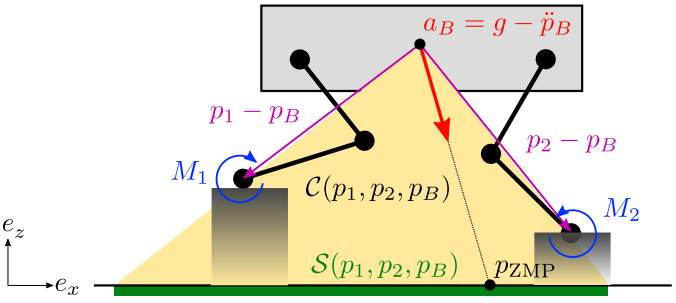

Fig. 3. The GIA vector $\boldsymbol{a}_B = \boldsymbol{g} - \ddot{\boldsymbol{p}}_B$ induces the moment $\boldsymbol{M}_{1,2}$ about the contact position $\boldsymbol{p}_{1,2}$. The cone $\mathcal{C}$, spanned by the base $\boldsymbol{p}_B$ and the two footholds, defines the set of all admissible $\boldsymbol{a}_B$. The ZMP equals the intersection of the ray $\boldsymbol{p}_B + \gamma \boldsymbol{a}_B$ with the ground. The associated support polygon $\mathcal{S}$ is obtained by projecting the footholds onto the ground.

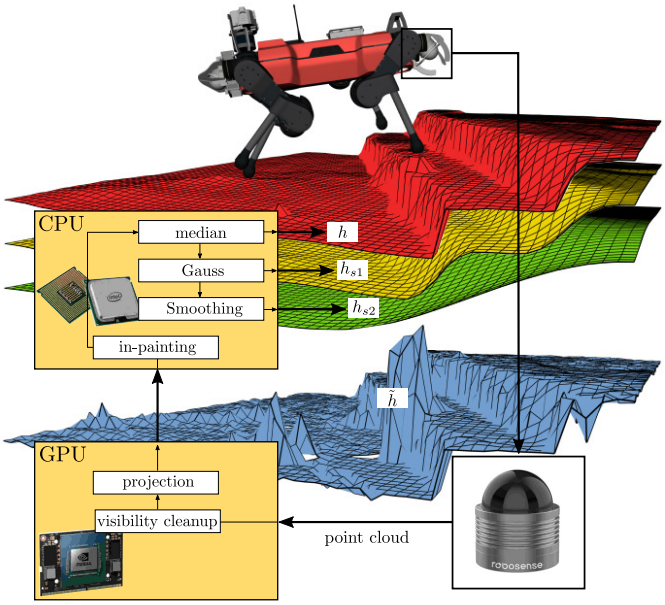

Fig. 4. Illustration of the mapping pipeline using data recorded in a real-world experiment. The raw map $\tilde{h}$ (blue) is generated on GPU and updated at 20 Hz. A filter chain, implemented on CPU, removes noise and artifacts.

C. Perception and Mapping

The perception pipeline is illustrated in Fig. 4. Two onboard LiDARs are used to generate a detailed *height map* [23], abbreviated by $\tilde{h}$. For performance reasons, projection and visibility clean-up are performed on an onboard GPU [6]. Further map processing algorithms run on CPU and generate three height layers of different smoothness, h, h_{s1}, and h_{s2}.

We deploy a filter chain consisting of in-painting, outlier rejection, and smoothing. First, we iterate over all grid cells in the raw map $\tilde{h}$, replacing empty cells with the minimum found across the occlusion border. Outliers are removed by a sequentially repeated median filter [24] on the inpainted map. Given the de-noised map h, we compute two additional layers, h_{s1} and h_{s2}. The former is a slightly Gaussian filtered version of the original map h and is used to compute gradients of edges. The third layer h_{s2} represents a "virtual floor" and is used to guide the base pose trajectory, as exemplified in Fig. 5.

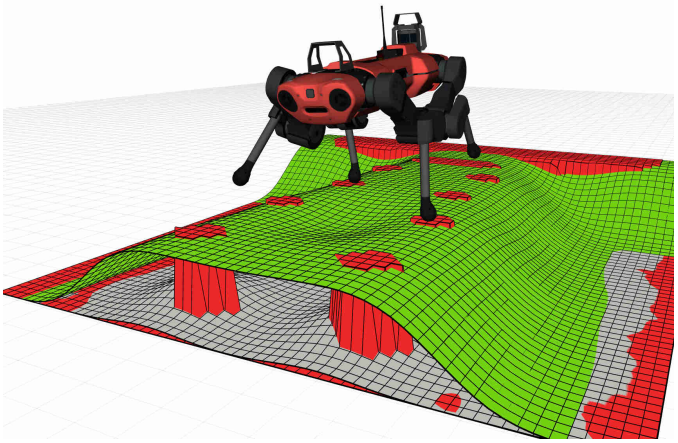

Fig. 5. The red map h, representing geometric features of the terrain, is used for foot placement constraints, while the green map h_{s2} serves as a base pose reference. In contrast to the gray layer, which would be obtained from h by pure Gaussian smoothing, h_{s2} spawns above the stepping stones. The robot tries to align its torso with h_{s2}, eventually encouraging a motion suggested by the cylindrical pillars.

D. Motion Optimization

We parametrize the base pose trajectory as a 6D splines of fixed order five. The first three dimensions capture position and the last three dimensions describe orientation using Euler angles. To increase the feasible space, we allow the base acceleration to evolve discontinuous at contact transitions [20], [21]. We do so by breaking the trajectory into different spline segments connected with each other at the transition times.

We stack all unknown variables (which are spline coefficients, current/desired footholds and slack variables) together into a state vector $\boldsymbol{x}$. The constraint non-linear program (NLP) can be formulated as

$$\min_{\boldsymbol{x}} \sum_i f_i(\boldsymbol{x}) \quad \text{s.t.} \quad \boldsymbol{c}_{\text{eq}}(\boldsymbol{x}) = \boldsymbol{0},\ \boldsymbol{c}_{\text{ineq}}(\boldsymbol{x}) \leq \boldsymbol{0}. \tag{3}$$

The objective functions f_i are used for foot-height constraints, foot-self collision avoidance, base pose alignment w.r.t. to virtual floor h_{s2}, edge avoidance based on h_{s1}, and base twist tracking. Inequality constraints $\boldsymbol{c}_{\text{ineq}}$ enforce the GIAC constraints and kinematic constraints formulated in task space. The equality constraints $\boldsymbol{c}_{\text{eq}}$ are used to formulate initial conditions and to smoothly connect two adjacent splines.

E. Optimization Method

Instead of finding a sophisticated initial guess, capable of avoiding the common pitfalls of gradient based methods, we deploy a global optimization technique: At each leg touchdown, we solve a sequence of TO problems, starting with a greatly simplified problem and progressively approaching towards the original problem. Such a technique is generally known as *graduated optimization* [25].

In the case of TAMOLS, non-convexity and discontinuity is introduced mostly by the terrain. This observation gives rise to the idea of solving a "nearly convex" TO in the first iteration by replacing the height maps h and h_{s1} with h_{s2}. The robot's measured state suffices already as an initial guess for this simplified NLP to succeed.

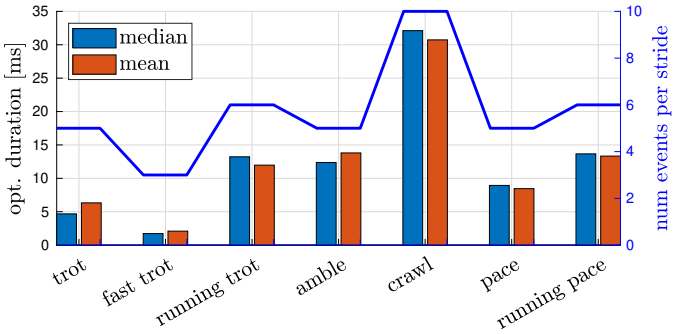

Fig. 6. Optimization duration (elapsed time between initialization to retrieving the solution) for a prediction horizon of one stride. Fast trot eliminates the full stance phase of trot, while the swing phases overlap for running trot. Amble is obtained from crawl by overlapping all swing phases until the triple stance phase vanishes. Pace/running pace has the same gait timing as trot/running trot but has a lateral leg pair alternating.

We could now continue in the sense of graduated optimization and introduce several virtual floors of different smoothness that gradually approach to the original map. For each of these virtual floors, a new TO problem needs be solved. In favor of low computation time, we truncate the sequence after solving the first TO while replacing the remaining sequence with a foothold refinement step [6]. In the last stage, the optimized and refined state is subsequently passed over the actual NLP as a new initial guess.

III. Results

We evaluate the performance of our perceptive control pipeline using the ANYmal platform. Elevation mapping runs at 20 Hz on an onboard GPU while control and state estimation is updated on a separate onboard CPU at 400 Hz. TAMOLS runs asynchronously at the maximum possible rate and optimizes trajectories for a prediction horizon of one gait cycle.

A. Computation Times

TAMOLS is able to produce motions for a wide variety of dynamic gaits. We record optimization durations while ANYmal is walking on flat ground. Mean and median are plotted in Fig. 6. For the most common gait trot, we achieve an average optimization duration of 6.3 ms, which is 48 times faster than the latest state-of-the-art perceptive locomotion controller [9]. If we also eliminate the full stance phase (fast trot), the average computation duration drops to 2.1 ms.

B. Gaps and Stepping Stones

The strength of the proposed TO method lies in the generalization to stairs, gaps, and stepping stones. The last two scenarios are supported by the refinement step with the batch search, which tends to select elevated footholds.

We validate the control performance in a stepping stone experiment as described in Fig. 7. By penalizing gradients of the smooth map h_{s1}, the robot centers the footholds in the middle of the stepping stones. Due to odometry drift, the bricks, as seen by the elevation map, may be translated by a few centimeters from their actual locations. This explains why the robot sometimes steps close to the border.

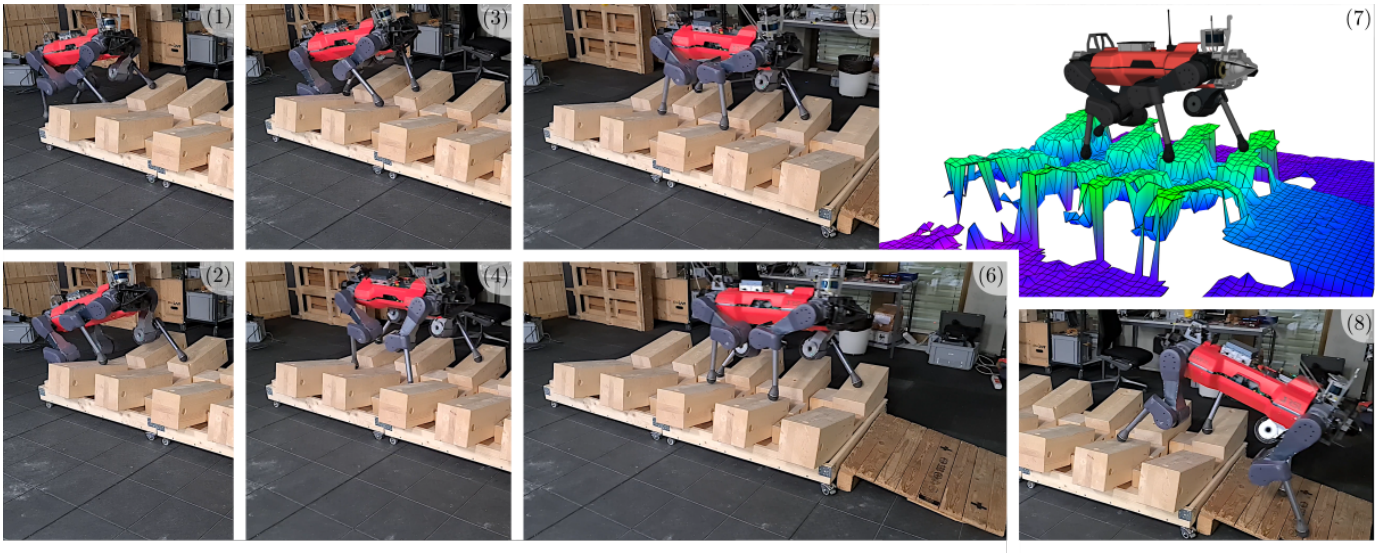

Fig. 7. Anymal traversing a stepping stone parkour made of inclined wooden bricks using a trotting gait. The commanded heading velocity is 0.4 m/s The dimensions of each brick are 20 cm × 20 cm × 50 cm and the gap between any two adjacent bricks measures 20 cm. Due to the inclined surfaces, the model assumptions are not perfectly satisfied in this experiment. Picture 7 shows the visualization of the robot along the raw elevation map.

TABLE I
SUCCESS RATE SCORED WHILE WALING ON STAIRS.

method	batch search	TAMOLS
gait	trot	trot, amble, running trot, pace
up	10/18	18/18
down	14/18	18/18

C. Comparison to Previous Work

In our previous work [6] we have presented a control pipeline separating foothold from base pose optimization, and discussed its limitations. We found that a significant amount of failure cases on stairs were caused by knee joint collisions with edges (upstairs) or by violating reachability constraints (downstairs). The experimental set-up included a simulated staircase of 12 treads, which was passed 18 times in ascending and descending direction using a trotting gait. The success rates are recapitulated in table I.

We repeat the exact same experiment with TAMOLS. The statistical results are presented in table I. While trotting upstairs, the knee joints occasionally collide with the edges, but the robot is always able to keep its balance. Contrary to our previous control framework, we do not encounter issues with leg over-extensions. This can be explained by the kinematic constraints that are present in the prediction and the tracking level. Moreover, we successfully repeat the experiment with three additional dynamic gaits: running trot, amble, and pace.

D. Staircase

We take ANYmal to the real world and walk upstairs using a trot in the building of our lab. Two floors are connected with each other by 20 treads where each platform has the dimensions 29 cm × 17 cm, forming an inclination of 36 deg. The operator commands a heading speed of 0.45 m/s whereas the realized average velocity is 0.37 m/s. The tracking error originates mainly from the reduced feasible space imposed by the stair geometry: The robot prefers to place the footholds s.t. all feet either clear zero, one, or two treads per stride. Depending on the magnitude of the reference velocity, a different optimum is generated, which was found at one tread per step for our experiment.

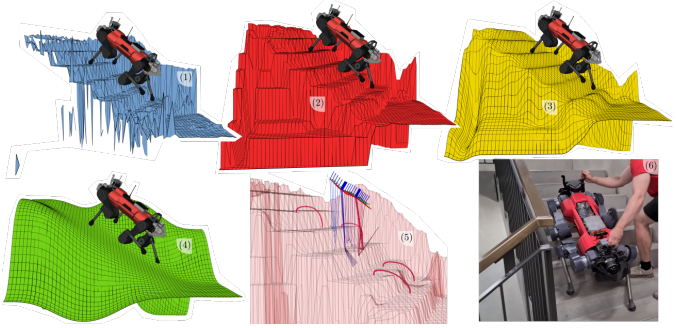

Fig. 8. ANYmal trotting up stairs. 1) Visualization of the robot together with the raw elevation map h. 2) The in-painted and de-noised elevation map h is used for height constraints on the footholds. 3) The Gaussian-smoothed height map h_{s1} is used for edge avoidance. 4) The robot aligns its torso with the virtual floor map h_{s2}. 5) Visualization of the motion plan in task space. 6) Snap-shot of the experiment.

In Fig. 8 we visualize the three height maps used for motion optimization along with the raw heightmap. For a staircase, the computation of the virtual floor h_{s2} equals Gaussian filtering of the heightmap h. The smooth map h_{s1} attains zero gradient in the middle of the tread, thereby pushing footholds away from edges towards the center line.

REFERENCES

[1] F. Jenelten, R. Grandia, F. Farshidian, and M. Hutter, "TAMOLS: Terrain-aware motion optimization for legged systems," *(submitted to) IEEE Transactions on Robotics*, 2021.

[2] D. Bellicoso, M. Bjelonic, L. Wellhausen, K. Holtmann, F. Günther, M. Tranzatto, P. Fankhauser, and M. Hutter, "Advances in real-world applications for legged robots," *Journal of Field Robotics*, 10 2018.

[3] M. Tranzatto, F. Mascarich, L. Bernreiter, C. Godinho, M. Camurri, S. Khattak, T. Dang, V. Reijgwart, J. Loeje, D. Wisth, S. Zimmermann, and K. Alexis, "CERBERUS: autonomous legged and aerial robotic exploration in the tunnel and urban circuits of the DARPA subterranean challenge," *CoRR*, vol. abs/2201.07067, 2022. [Online]. Available: https://arxiv.org/abs/2201.07067

[4] C. Mastalli, I. Havoutis, M. Focchi, D. G. Caldwell, and C. Semini, "Motion planning for quadrupedal locomotion: Coupled planning, terrain mapping, and whole-body control," *IEEE Transactions on Robotics*, vol. 36, no. 6, pp. 1635–1648, 2020.

[5] O. Magaña, V. Barasuol, M. Camurri, L. Franceschi, M. Focchi, M. Pontil, D. Caldwell, and C. Semini, "Fast and continuous foothold adaptation for dynamic locomotion through cnns," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2140–2147, April 2019.

[6] F. Jenelten, T. Miki, A. E. Vijayan, M. Bjelonic, and M. Hutter, "Perceptive locomotion in rough terrain – online foothold optimization," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5370–5376, 2020.

[7] S. Gangapurwala, M. Geisert, R. Orsolino, M. Fallon, and I. Havoutis, "Rloc: Terrain-aware legged locomotion using reinforcement learning and optimal control," 2020.

[8] P. Fankhauser, M. Bjelonic, C. Dario Bellicoso, T. Miki, and M. Hutter, "Robust rough-terrain locomotion with a quadrupedal robot," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 5761–5768.

[9] O. Melon, R. Orsolino, D. Surovik, M. Geisert, I. Havoutis, and M. F. Fallon, "Receding-horizon perceptive trajectory optimization for dynamic legged locomotion with learned initialization," *CoRR*, vol. abs/2104.09078, 2021. [Online]. Available: https://arxiv.org/abs/2104.09078

[10] R. J. Griffin, G. Wiedebach, S. McCrory, S. Bertrand, I. Lee, and J. Pratt, "Footstep planning for autonomous walking over rough terrain," in *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*, 2019, pp. 9–16.

[11] M. Hutter, C. Gehring, D. Jud, A. Lauber, C. D. Bellicoso, V. Tsounis, J. Hwangbo, K. Bodie, P. Fankhauser, M. Bloesch, R. Diethelm, S. Bachmann, A. Melzer, and M. Hoepflinger, "Anymal - a highly mobile and dynamic quadrupedal robot," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2016, pp. 38–44.

[12] R. Grandia, F. Farshidian, R. Ranftl, and M. Hutter, "Feedback mpc for torque-controlled legged robots," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4730–4737.

[13] A. W. Winkler, C. D. Bellicoso, M. Hutter, and J. Buchli, "Gait and trajectory optimization for legged systems through phase-based end-effector parameterization," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1560–1567, July 2018.

[14] O. Melon, M. Geisert, D. Surovik, I. Havoutis, and M. Fallon, "Reliable trajectories for dynamic quadrupeds using analytical costs and learned initializations," 2020.

[15] H. Dai and R. Tedrake, "Planning robust walking motion on uneven terrain via convex optimization," in *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, 2016, pp. 579–586.

[16] S. Caron, Q.-C. Pham, and Y. Nakamura, "Leveraging cone double description for multi-contact stability of humanoids with applications to statics and dynamics," in *Robotics: Science and System*, Jul. 2015.

[17] R. Orsolino, M. Focchi, C. Mastalli, H. Dai, D. G. Caldwell, and C. Semini, "Application of wrench-based feasibility analysis to the online trajectory optimization of legged robots," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3363–3370, Oct 2018.

[18] P. Fernbach, S. Tonneau, O. Stasse, J. Carpentier, and M. Taïx, "C-croc: Continuous and convex resolution of centroidal dynamic trajectories for legged robots in multicontact scenarios," *IEEE Transactions on Robotics*, pp. 1–16, 2020.

[19] P. Fernbach, S. Tonneau, and M. Taïx, "Croc: Convex resolution of centroidal dynamics trajectories to provide a feasibility criterion for the multi contact planning problem," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 1–9.

[20] M. Kalakrishnan, J. Buchli, P. Pastor, M. Mistry, and S. Schaal, "Fast, robust quadruped locomotion over challenging terrain," in *2010 IEEE International Conference on Robotics and Automation*, May 2010, pp. 2665–2670.

[21] C. D. Bellicoso, F. Jenelten, P. Fankhauser, C. Gehring, J. Hwangbo, and M. Hutter, "Dynamic locomotion and whole-body control for quadrupedal robots," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017, pp. 3359–3365.

[22] C. D. Bellicoso, F. Jenelten, C. Gehring, and M. Hutter, "Dynamic locomotion through online nonlinear motion optimization for quadrupedal robots," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2261–2268, July 2018.

[23] P. Fankhauser, M. Bloesch, C. Gehring, M. Hutter, and R. Siegwart, "Robot-centric elevation mapping with uncertainty estimates," in *International Conference on Climbing and Walking Robots*, 2014, pp. 433–440.

[24] E. Arias-Castro and D. Donoho, "Does median filtering truly preserve edges better than linear filtering?" *The Annals of Statistics*, vol. 37, 12 2006.

[25] H. Mobahi and J. W. Fisher, "On the link between gaussian homotopy continuation and convex envelopes," in *Energy Minimization Methods in Computer Vision and Pattern Recognition*, X.-C. Tai, E. Bae, T. F. Chan, and M. Lysaker, Eds. Cham: Springer International Publishing, 2015, pp. 43–56.