KAGGLEYEs ACADEMY E-PRINT REGISTRY
ARCHIVAL PREPRINT / OPEN SCHOLARLY RECORD

KAGG RECORD
KAGG:2610.00002

Range-Aligned Control Allocation for Agile Quadruped Navigation

Jonas Meier • **Diego Navarro** • **Clara Fischer**

Robot Dynamics Group, ETH Zurich, Zurich, Switzerland

PUBLISHED	REGISTERED	RECORD TYPE
October 3, 2026	October 6, 2026	Preprint

ABSTRACT

Whole-body controllers for quadrupedal robots typically distribute contact forces without knowledge of where the next obstacle will be, so stability reserves and swing clearance are allocated uniformly rather than where the upcoming motion needs them. This paper presents range-aligned control allocation, which uses a predicted future range field, computed in the style of FutureRay along the control-aligned directions of the robot, to set the weights of a whole-body quadratic program. A softmin over the field locates the predicted closest approach; stance-force weights and a linear load-shift term then bias stabilizing support toward the side opposite that approach, and a barrier-type constraint enforces swing-foot clearance only inside the predicted path corridor. All range-dependent terms scale with a field-contrast statistic that vanishes exactly when the field is uniform, and we prove that the allocator then reduces to the nominal controller, with solution deviation otherwise bounded linearly in the contrast. In an evaluation on ANYmal comprising Isaac Lab simulation across ten scenes with raised obstacles of 15 to 60 cm at 1.0 to 3.0 m/s and hardware runs beside raised obstacles, the method lowered energy per kilometer from 1334 to 1238 J/km relative to range-agnostic allocation, reduced the 95th-percentile solve time from 2.9 to 2.6 ms, and matched the unweighted controller in 99.6 percent of uniform-field evaluations. Its mean clearance margin, however, was lower (20.8 versus 24.0 cm), so the results indicate a tighter, less conservative allocation rather than uniformly larger clearance.

INDEX TERMS legged robots, whole-body control, control allocation, quadratic programming, future range prediction, quadruped navigation, ANYmal

KAGG NOTE Allocation weights that depend on the predicted future range only through a contrast statistic reduce exactly to the nominal controller when the range field is uniform, so any change in behavior can be attributed to information carried by the field.

01 Introduction

Quadrupedal robots such as ANYmal [1] now traverse unstructured terrain at speeds of one to several meters per second, using model-based whole-body controllers [2, 3, 4], learned locomotion policies [5, 6, 7], or combinations of both [8]. Agile navigation among raised obstacles, such as benches, curbs, and low shelves, couples two problems that are usually solved in separate layers. A navigation layer decides where the body should go, and a whole-body controller decides how the contact forces and joint torques that realize that motion are distributed among the legs.

The second decision is made, in most deployed systems, without any knowledge of where the next obstacle is. A whole-body quadratic program (QP) typically regularizes contact forces with fixed weights, tracks swing-foot references generated from a nominal gait, and enforces friction cones and torque limits uniformly over all legs [9, 10, 11]. When an obstacle approaches on one side, such an allocator still spreads its stability reserve symmetrically. It may load the legs adjacent to the obstacle as heavily as the legs that could push the body away from it, and it may spend swing-height margin on feet whose paths never intersect the obstacle while giving no special treatment to the foot that will pass next to a shelf edge. In effect, the allocator spends its stability budget without regard to where it is needed.

The perception side has recently made a forward-looking query of this kind practical. FutureRay [12] predicts the future free range along the robot's own predicted, control-aligned directions, so that obstacle queries follow motions the controller can actually execute rather than a fixed sensor-frame grid. Such a field already contains the information an allocator needs: where along the upcoming motion the closest approach to an obstacle is predicted, on which side of the body it lies, and which portion of the path corridor will be occupied by raised geometry.

This paper studies whether this information can be used inside the allocation QP itself without sacrificing the properties of the nominal controller. We propose *range-aligned control allocation*. A softmin over the predicted range field locates the predicted closest approach. Stance-force weights and a linear load-shift term bias stabilizing support toward the side of the support polygon opposite that approach. A barrier-type acceleration constraint enforces swing-foot clearance only for feet inside the predicted path corridor. Every range-dependent term is scaled by a contrast statistic that is zero exactly when the field is uniform.

The contributions are as follows.

1. A formulation of whole-body force allocation as a strictly convex QP whose weights and constraints are functions of a FutureRay-style predicted future range field, combining a stance-side bias with a swing-clearance corridor constraint.

KAGGLEYEs ACADEMY • E-PRINT REGISTRY • KAGG:2610.00002

2. A field-contrast statistic with a short proof that it lies in $[0, 1)$ and vanishes if and only if the field is uniform. This yields exact recovery of the nominal allocator for an uninformative field and a bound on the solution deviation that is linear in the contrast.
3. A real-time implementation on ANYmal with an explicit deadline and a logged solution-reuse fallback.
4. An evaluation in Isaac Lab and in hardware trials beside raised obstacles, comparing against range-agnostic, present-range, and stance-only allocation, with an ablation and a discussion of recorded failure cases. Range-aligned allocation used less energy per kilometer at a lower mean clearance margin than every baseline.

02 Related Work

2.1 Whole-body control and force allocation

Optimization-based whole-body control computes generalized accelerations, contact forces, and torques by solving one or more QPs at every control cycle. Herzog et al. formulated momentum control as a hierarchy of inverse-dynamics QPs on a torque-controlled humanoid [10], and Escande et al. developed efficient hierarchical QP solvers that make strict task priorities tractable online [11]. Kuindersma et al. showed that a single weighted QP with carefully chosen costs can stabilize a full humanoid in the field [13]. On quadrupeds, Bellicoso et al. used hierarchical whole-body control on ANYmal for perception-less terrain adaptation [9] and later for dynamic gaits driven by online motion optimization [2]. Gehring et al. described the full optimization-based control stack on StarlETH and ANYmal [14]. Kim et al. combined whole-body impulse control with a convex MPC on the MIT Mini Cheetah [3], building on convex MPC for force planning [15]. Contact stability in these controllers is expressed through friction pyramids or contact wrench cones [16], and the QPs are solved by active-set or operator-splitting methods [17]. In all of these controllers, force-regularization weights and friction margins are either constant or scheduled by gait phase. Our allocator keeps this structure, including the dynamics, contact, and friction constraints, and changes only how weights and a clearance constraint depend on the predicted environment.

2.2 Model-predictive and learned locomotion

Perceptive nonlinear MPC plans base and foot motions over a horizon using a terrain representation and hands the result to a whole-body tracking controller [4]. TAMOLS optimizes footholds and base trajectory jointly with terrain-aware costs [18], and Crocoddyl provides a differential dynamic programming framework for contact-rich motions [19]. Learned policies trained in massively parallel simulation have reached robust blind and perceptive locomotion [5, 6, 7, 20], high-speed running [21], rapid online adaptation [22], and parkour-style agility [23, 24, 25, 26]. DTC combines a learned tracking policy with model-based footholds [8]. These systems reason about terrain mainly through foothold selection or end-to-end learned features. They do not expose an explicit, inspectable mechanism that shifts the stability reserve of the allocation toward the side away from an obstacle. Our method is complementary: it can sit below an MPC or a learned base-motion policy, and its effect is confined to how a commanded motion is distributed among legs.

2.3 Perception for legged navigation and future range

Elevation mapping [27] and learned traversability [28] provide terrain representations for legged navigation, and ArtPlanner uses such representations for robust field navigation on ANYmal [29]. Systems from the DARPA Subterranean Challenge demonstrated the full perception and navigation stack at scale [30]. Learned local navigation policies take height scans or latent state representations as input [31, 32]. Configuration-driven local planners [33, 34] and adaptive safety-margin critics [35] make body geometry and clearance budgets explicit. Classical reactive planners such as the dynamic window approach [36] and gradient-based planners such as CHOMP [37] and EGO-Planner [38] query distance information along candidate motions.

FutureRay [12] is the closest perception counterpart to this work. Instead of querying free space along fixed sensor rays or a body-centered grid, it predicts the future free range along the directions the robot is expected to move under its own control, over a short horizon. The resulting field is control-aligned: each range value answers the question of how far the robot can travel along a motion it can actually execute before reaching an obstacle. FutureRay uses this field for navigation decisions. We use the same kind of field one layer lower. The field sets allocation weights and activates swing-clearance constraints, and it is not used to choose a velocity command. We do not modify the FutureRay predictor and treat its output as an interface. A direct consequence is that the quality of our allocation is bounded by the quality of the predicted field, which we return to in Section 6.

2.4 Safety constraints in control QPs

Control barrier functions (CBFs) encode set invariance as linear constraints on the control input of a QP [39, 40]. High-order CBFs extend this to constraints of higher relative degree [41], and Grandia et al. layered CBFs with MPC to enforce foothold and collision constraints on a quadruped [42]. Our swing-clearance corridor uses a second-order barrier constraint of this type on swing-foot height. It differs in two respects: the constraint is instantiated only for swing feet inside the predicted corridor, and its activation is driven by a contrast statistic of the predicted field, which is what makes exact nominal recovery possible.

2.5 Simulation for legged robots

GPU-parallel simulators [43, 20] and the Orbit framework [44], on which Isaac Lab is built, allow thousands of trials across scenes, seeds, and speeds. We use Isaac Lab for the simulated portion of the evaluation.

03 Problem Formulation

3.1 Robot model and nominal allocation

Let $q \in \mathbb{R}^{n_q}$ denote the generalized coordinates of the floating-base robot, $\nu \in \mathbb{R}^{n_v}$ its generalized velocity ($n_v = 18$ for ANYmal: six base and twelve joint coordinates), and $\tau \in \mathbb{R}^{12}$ the joint torques. Let $\mathcal{C}_t \subseteq \{1, 2, 3, 4\}$ be the set of legs in stance at time t and $\mathcal{S}_t$ its complement (swing legs). The equations of motion are

$$M(q)\dot{\nu} + h(q, \nu) = S^\top \tau + \sum_{i \in \mathcal{C}_t} J_i(q)^\top f_i, \qquad (1)$$

where M is the mass matrix, h collects Coriolis, centrifugal, and gravity terms, S selects the actuated coordinates, J_i is the

KAGGLEYEs ACADEMY • E-PRINT REGISTRY • KAGG:2610.00002

translational Jacobian of foot i, and $f_i \in \mathbb{R}^3$ is its contact force. We write $f_i = (f_{i,x}, f_{i,y}, f_{i,n})$ in the local contact frame, with $f_{i,n}$ the normal component.

The allocation decision variable is $x = (\dot\nu, f) \in \mathbb{R}^{n_x}$ with f stacking f_i for $i \in \mathcal{C}_t$. Torques follow from the actuated rows of (1). The *nominal* allocator solves

$$\min_{x \in \mathcal{X}_t} \ \tfrac{1}{2} x^\top H_0 x + g_0^\top x, \tag{2}$$

where the cost collects the base-motion task $\|A_b x - b_b\|^2_{Q_b}$, the swing-foot tracking task $\sum_{j \in \mathcal{S}_t} \|J_j \dot\nu + \dot J_j \nu - \ddot p_j^{\text{ref}}\|^2_{Q_s}$, force regularization $\sum_{i \in \mathcal{C}_t} f_i^\top W_i^0 f_i$ with $W_i^0 \succ 0$, and a small acceleration damping term $\epsilon \|\dot\nu\|^2$ that makes $H_0 \succ m_0 I$ for some $m_0 > 0$. The feasible set $\mathcal{X}_t$ is a polyhedron defined by the unactuated rows of (1), the stance no-slip conditions $J_i \dot\nu + \dot J_i \nu = 0$ for $i \in \mathcal{C}_t$, linearized friction pyramids $|f_{i,x}|, |f_{i,y}| \le \mu f_{i,n}$ with $f_{i,n} \ge f_{\min}$, and torque limits expressed linearly in x. Because H_0 is positive definite and $\mathcal{X}_t$ is convex, (2) has a unique minimizer $x_0^\star$ whenever it is feasible.

3.2 Predicted future range field

A navigation layer issues a planar velocity command and, from it, a predicted base path $\hat p(\tau)$ over a horizon $\tau \in [0, T_p]$, sampled at $\tau_m = m\,\Delta\tau$, $m = 0, \dots, M$. Following FutureRay [12], the perception module predicts, for each of K control-aligned directions $d_k(\tau_m)$ and each of L elevation layers at heights z_ℓ above the local ground, the free range

$$\rho_{k\ell m} \in [\rho_{\min}, \rho_{\max}], \qquad \rho_{\min} > 0, \tag{3}$$

measured from $\hat p(\tau_m)$ along $d_k(\tau_m)$. We call $\mathcal{R} = \{\rho_{k\ell m}\}$ the *predicted future range field*. The lower clamp $\rho_{\min}$ is a numerical safeguard. We partition layers into foot layers $\mathcal{L}_f$, which cover the swing-height band, and body layers $\mathcal{L}_b$, which cover the trunk.

Definition (uniform field). A field is uniform on an index set $\mathcal{I}$ if $\rho_{k\ell m} = c$ for all $(k, \ell, m) \in \mathcal{I}$ and some constant c.

A uniform field carries no directional information. Every predicted direction is equally free, so no side of the body and no portion of the corridor is distinguished.

3.3 Problem statement

Given $\mathcal{R}$, $\hat p$, and the nominal allocation (2), the problem is to construct a range-dependent allocation QP whose solution (i) places stance stability preferentially on the support side opposite the predicted closest approach, (ii) enforces swing-foot clearance over raised obstacles within the predicted path corridor, (iii) remains a strictly convex QP solvable within the allocation period, and (iv) coincides with $x_0^\star$ whenever $\mathcal{R}$ is uniform.

04 Method

4.1 Architecture and timing

The stack has three layers. The navigation layer and the range-field predictor run at 20 Hz, which is a 50 ms period, with a prediction horizon $T_p = 1.0$ s. At the evaluated speeds of 1.0–3.0 m/s this horizon spans one to three meters of travel, which covers several strides. The range-aligned allocator runs at 200 Hz with a solve deadline inside the 5 ms period. Between field updates, the field statistics below are held constant while the robot state (q, ν) is updated every allocation cycle. A gait generator supplies the contact schedule $\mathcal{C}_t$ and swing references $\ddot p_j^{\text{ref}}$. The resulting torques are tracked by the joint-level actuator controllers. In our configuration, $K = 32$ directions and $L = 4$ elevation layers are used.

4.2 Closest approach and field contrast

Over the body layers, define softmin weights with temperature $\sigma > 0$,

$$\alpha_{k\ell m} = \frac{\exp(-\rho_{k\ell m}/\sigma)}{\sum_{(k',\ell',m') \in \mathcal{I}_b} \exp(-\rho_{k'\ell'm'}/\sigma)}, \tag{4}$$

for $(k, \ell, m) \in \mathcal{I}_b$, the index set of body-layer samples. The softmin range and the arithmetic mean are

$$\rho_{\text{s}} = \sum_{\mathcal{I}_b} \alpha_{k\ell m} \rho_{k\ell m}, \qquad \bar\rho = \frac{1}{|\mathcal{I}_b|} \sum_{\mathcal{I}_b} \rho_{k\ell m}. \tag{5}$$

The *field contrast* is

$$\kappa = 1 - \frac{\rho_{\text{s}}}{\bar\rho}. \tag{6}$$

Lemma 1. $0 \le \kappa < 1$, and $\kappa = 0$ if and only if the field is uniform on $\mathcal{I}_b$.

Proof. The weights α are a strictly decreasing function of ρ. By Chebyshev's sum inequality, the α-weighted mean of ρ is at most the uniform mean, with equality if and only if all ρ are equal, so $\rho_{\text{s}} \le \bar\rho$ and $\kappa \ge 0$ with equality exactly for a uniform field. Since $\rho \ge \rho_{\min} > 0$, $\rho_{\text{s}} > 0$ and hence $\kappa < 1$. □

The predicted closest-approach point is the softmin-weighted ray endpoint

$$p^\star = \sum_{\mathcal{I}_b} \alpha_{k\ell m} \big(\hat p(\tau_m) + \rho_{k\ell m}\, d_k(\tau_m)\big). \tag{7}$$

With P the projection onto the horizontal plane and c the center of mass, the closest-approach direction is $u = P(p^\star - c)/\|P(p^\star - c)\|$, and $u = 0$ if the denominator vanishes. The softmin, instead of a hard minimum, keeps $p^\star$ and u continuous in the field, so the weights do not jump between field updates when two rays have nearly equal range.

4.3 Stance-side bias

For each stance foot i with position p_i, define its signed offset toward the predicted approach and its side score,

$$e_i = u^\top P(p_i - c), \qquad s_i = -\frac{e_i}{\max(\|P(p_i - c)\|,\, r_0)} \in [-1, 1], \tag{8}$$

where $r_0 > 0$ prevents division by zero. A foot on the side away from the predicted closest approach has $s_i > 0$. The range-aligned stance cost is

$$J_{\text{st}}(f; \kappa) = \sum_{i \in \mathcal{C}_t} (1 - \beta \kappa s_i)\, f_i^\top W_i^0 f_i + \gamma \kappa \sum_{i \in \mathcal{C}_t} e_i\, f_{i,n}, \tag{9}$$

with gains $0 \le \beta < 1$ and $\gamma \ge 0$. The first term lowers the regularization cost of legs on the far side, so these legs absorb a larger share of the commanded wrench and of the corrective forces. It also raises the cost on legs next to the obstacle, which reduces their tangential load and preserves their friction-cone margin. The second, linear term shifts the normal-load distribution, and with it the center of pressure, away from

KAGGLEYEs ACADEMY • E-PRINT REGISTRY • KAGG:2610.00002

the predicted approach. The total normal force is fixed by the base-motion task and the unactuated dynamics, so this term redistributes load and does not add any. Because $|s_i| \le 1$ and $\kappa < 1$, every weight satisfies $1 - \beta\kappa s_i \ge 1 - \beta > 0$, so the stance cost stays strictly convex in f.

4.4 Swing-clearance corridor

The predicted path corridor is

$$\mathcal{K} = \Big\{ x \in \mathbb{R}^2 : \min_m \|x - P\hat{p}(\tau_m)\| \le \tfrac{w_b}{2} + c_w \Big\}, \quad (10)$$

with body half-width $w_b/2$ and lateral swing allowance c_w. For swing foot $j \in \mathcal{S}_t$, let $\mathcal{I}_j$ be the foot-layer samples whose ray endpoints lie in $\mathcal{K}$ within one step length of the foot. The corridor contrast κ_j^c is defined as in (6), restricted to $\mathcal{I}_j$, and so by Lemma 1 it vanishes exactly when the field is uniform on $\mathcal{I}_j$. If $\kappa_j^c > \kappa_{\min}$, the foot is marked active and an obstacle-top estimate $\hat{h}_j$ is taken as the highest foot layer z_ℓ whose range in $\mathcal{I}_j$ falls below the foot's distance to the end of its swing. For active feet, define the clearance function $h_j(q) = z_j(q) - \hat{h}_j - h_m$, where z_j is foot height and h_m is a fixed margin, and impose the second-order barrier constraint [39, 41]

$$\ddot{h}_j + k_1 \dot{h}_j + k_0 h_j \ge -\xi_j, \qquad \xi_j \ge 0, \quad (11)$$

where $\ddot{h}_j = e_z^\top (J_j \dot{\nu} + \dot{J}_j \nu)$ is affine in x and $k_0, k_1 > 0$ place the roots of $s^2 + k_1 s + k_0$ in the open left half-plane. The slack ξ_j is penalized by $w_\xi \xi_j^2$ with a large w_ξ, so that infeasibility caused by torque or friction limits degrades clearance gracefully and does not make the QP infeasible. Restricting the constraint to $\mathcal{K}$ avoids spending swing height, and therefore energy, on feet whose paths do not cross raised geometry.

4.5 Range-aligned allocation QP

Let $\mathcal{A}_t = \{ j \in \mathcal{S}_t : \kappa_j^c > \kappa_{\min} \}$ and $\xi = (\xi_j)_{j \in \mathcal{A}_t}$. The range-aligned allocator solves

$$\begin{aligned} \min_{x,\,\xi} \; & \tfrac{1}{2} x^\top H_0 x + g_0^\top x + w_\xi \|\xi\|^2 \\ & + \big[J_{\mathrm{st}}(f;\kappa) - J_{\mathrm{st}}(f;0) \big] \\ \text{s.t. } & x \in \mathcal{X}_t, \quad (11)\ \forall j \in \mathcal{A}_t, \quad \xi \ge 0. \end{aligned} \quad (12)$$

The bracket replaces the nominal force regularization by (9). The problem can be written as $\min \tfrac{1}{2} x^\top H(\kappa) x + g(\kappa)^\top x + w_\xi \|\xi\|^2$, with $H(\kappa) = H_0 + \kappa \Delta H$ and $g(\kappa) = g_0 + \kappa \Delta g$. Here $\Delta H = -2\beta\, \mathrm{blkdiag}(s_i W_i^0)$ acts on the force block and $\Delta g = \gamma\, \bar{e}$, where $\bar{e}$ places e_i on the normal-force entries. Problem (12) is a strictly convex QP of the same size as the nominal one, plus at most two slack variables in a trot.

4.6 Nominal recovery and continuity

Proposition 1 (nominal recovery). If $\mathcal{R}$ is uniform on $\mathcal{I}_b$ and on every $\mathcal{I}_j$, then the unique minimizer of (12) equals $x_0^\star$.

Proof. By Lemma 1, $\kappa = 0$ and $\kappa_j^c = 0 \le \kappa_{\min}$ for all j, so $\mathcal{A}_t = \emptyset$ and no slack variables or barrier constraints are instantiated. Then $H(0) = H_0$, $g(0) = g_0$, and the feasible set is $\mathcal{X}_t$, so (12) is identical to (2). Strict convexity gives uniqueness. □

Proposition 2 (linear deviation). Suppose $\mathcal{A}_t = \emptyset$ and both problems are feasible. Then

$$\|x^\star(\kappa) - x_0^\star\| \le \kappa\, \frac{\|\Delta H\, x_0^\star\| + \gamma \|\bar{e}\|}{(1-\beta)\, m_0}. \quad (13)$$

Proof. The force block of $H(\kappa)$ scales each W_i^0 by $1 - \beta\kappa s_i \ge 1 - \beta$, and the remaining blocks are unchanged, so $H(\kappa) \succeq (1-\beta) H_0 \succeq (1-\beta) m_0 I$. The first-order optimality conditions over the common convex set $\mathcal{X}_t$ give $(H_0 x_0^\star + g_0)^\top (x^\star - x_0^\star) \ge 0$ and $(H(\kappa) x^\star + g(\kappa))^\top (x_0^\star - x^\star) \ge 0$. Adding the two inequalities and writing $H(\kappa)x^\star - H_0 x_0^\star = H(\kappa)(x^\star - x_0^\star) + \kappa \Delta H x_0^\star$ yields $(x^\star - x_0^\star)^\top H(\kappa)(x^\star - x_0^\star) \le \kappa \|\Delta H x_0^\star + \Delta g\|\, \|x^\star - x_0^\star\|$, and the triangle inequality gives (13). □

Proposition 1 is the formal content of the claim that the alignment cost vanishes when the field carries no information. Proposition 2 states that a nearly uniform field perturbs the allocation only slightly. This matters in open areas, where small fluctuations of the predicted range should not make the stance distribution chatter. Neither result covers the barrier constraint when it is active. In that case the allocator intentionally departs from the nominal solution.

4.7 Implementation and fallback

The QP (12) is solved by a dense active-set solver warm-started from the previous cycle's active set. If no solution is available by the deadline, the allocator reuses the previous cycle's torque command for one cycle and logs the event. If two consecutive deadlines are missed, it falls back to the nominal allocator (2). The field statistics (4)–(10) are computed once per field update, so their cost does not enter the 200 Hz loop.

05 Experiments

5.1 Setup and platform

Simulation experiments were run in Isaac Lab with the ANYmal model on 10 scenes containing raised obstacles: benches, curbs, and low shelves with heights between 15 and 60 cm. Each scene was run with 5 seeds and 30 trials each, at commanded speeds between 1.0 and 3.0 m/s. Real-world trials used an ANYmal in a laboratory course beside raised obstacles, with 8 runs per obstacle-height class. All methods share the same whole-body QP structure (2), gait generator, and navigation layer. They differ only in how allocation weights and clearance terms are set. Ours uses the range-dependent stance-stability weights and corridor constraint of Section 4, driven by the predicted future range field.

REPORTING NOTE

The tables below reproduce the compiled study results verbatim. The compiled results report one aggregate value per method and metric. They do not include per-seed dispersion or a separate breakdown into simulation and hardware runs, so no confidence intervals are reported and differences should be read with that in mind.

5.2 Baselines and metrics

We compare four allocators. *Range-agnostic allocation* uses uniform (constant) stance and clearance weights that do not depend on the range field. *Present-range weights* applies the same weighting rule to the currently measured range without prediction. *Stance stability only* applies the stance-side bias but no swing-clearance corridor. *Ours* is the full range-aligned allocator (12).

Four metrics are reported. *Clearance margin* is the clearance between the robot and the raised obstacles in centimeters, where

KAGGLEYEs ACADEMY • E-PRINT REGISTRY • KAGG:2610.00002

larger values indicate more clearance. *Energy per km* is the logged locomotion energy normalized by distance traveled. *QP solve time p95* is the 95th percentile of the allocation solve time. *Nominal-recovery parity* is the percentage agreement of each allocator's output with the unweighted nominal controller when it is fed a uniform range field. This is the empirical counterpart of Proposition 1.

5.3 Main comparison

Table 1 summarizes the comparison. Range-aligned allocation had the lowest energy per kilometer, 1238 J/km. Relative to range-agnostic allocation (1334 J/km), this is a reduction of 96 J/km, or 7.2%. Relative to present-range weights (1355 J/km) the reduction is 117 J/km (8.6%), and relative to stance stability only (1389 J/km) it is 151 J/km (10.9%). Its 95th-percentile solve time of 2.6 ms was 0.3 ms below range-agnostic and stance-only allocation (2.9 ms) and 0.2 ms below present-range weights (2.8 ms). All four allocators stayed within the 5 ms allocation period at the 95th percentile. The extra slack variables of (12) therefore did not translate into a higher tail solve time in these runs.

Nominal-recovery parity was 99.6% for range-aligned allocation, compared with 95.6% for range-agnostic allocation, 94.3% for stance stability only, and 88.5% for present-range weights. The value for our method is consistent with Proposition 1. The analysis predicts exact agreement for a uniform field, and the remaining 0.4% is not explained by it. Solver tolerances and warm-start differences are plausible contributors, but the compiled logs do not isolate them. The analysis makes no prediction for the baselines, whose uniform clearance terms or present-range rules are not designed to reduce to the unweighted controller. We therefore do not attribute their lower parity to a specific mechanism.

The clearance margin does not favor our method. Range-aligned allocation had the smallest mean clearance margin of the four allocators: 20.8 cm, against 24.0 cm for range-agnostic allocation (3.2 cm less), 22.9 cm for stance stability only (2.1 cm less), and 22.2 cm for present-range weights (1.4 cm less). The original motivation, which was to keep the robot farther from obstacles by placing stability where it is needed, is therefore not supported by the mean clearance figures. The results are more consistent with a less conservative allocation. One interpretation, which the energy figures are compatible with but which the aggregate data cannot confirm, is that uniform clearance weights produce extra margin on approaches where it is not required, and that corridor-restricted clearance removes this excess along with its energy cost. Testing this interpretation would require per-scene margin distributions, which the compiled results do not contain.

5.4 Ablation

Table 2 removes one component at a time. Each removal increased the mean clearance margin. Without the swing-clearance corridor the margin was 22.4 cm (1.6 cm more than the full method), without the stance-side bias it was 22.1 cm (1.3 cm more), and without range prediction, using present range only, it was 22.8 cm (2.0 cm more). Each component therefore contributes to the tighter mean margin of the full method, and replacing the predicted field by the present range had the largest effect.

The ablation and the main comparison overlap conceptually but not numerically. "W/o swing-clearance corridor" (22.4 cm) and "stance stability only" (22.9 cm) describe similar configurations, as do "w/o range prediction" (22.8 cm) and "present-range weights" (22.2 cm). The compiled results list these as separate entries and do not record whether the differences come from retuning, from separate run batches, or from configuration details. We report both as recorded and do not compare values across the two tables.

5.5 Recorded failure cases

Narrow shelf at high speed. In the 30 cm shelf scene at 3.0 m/s, uniform-weights allocation stepped with the wrong swing leg in 6 of 30 trials (20%), with a clearance margin of 4 cm in those trials. Range-aligned allocation held an 11 cm mean margin in the same scene. This is the one recorded condition in which the tail behavior of the baseline is directly visible, and it points the other way from the aggregate means. Two caveats apply. First, the 4 cm figure refers to the failing trials, while the 11 cm figure is a mean, so the comparison is not like-for-like. Second, the contact schedule is generated upstream of the allocator, so the allocator affects swing-leg events only indirectly, through the base motion and load distribution to which the gait generator responds. The compiled logs do not isolate this mechanism.

Latency spike. Under thermal load on the onboard CPU, the QP solve time spiked once to 7.8 ms. This is three times the 2.6 ms p95 and exceeds the 5 ms period. The allocator reused the previous solution for one cycle, as designed, and the event was logged. The nominal fallback was not triggered. A single occurrence does not establish a rate, and longer thermal-stress runs would be needed to characterize it.

Uniform-field recovery. With a uniform range field, parity with the unweighted baseline was 99.6%. The alignment cost vanished when the field carried no information, as Proposition 1 predicts.

5.6 Discussion

The evidence supports three statements. Range-aligned allocation can be added to a whole-body QP without increasing tail solve time. It recovers the nominal controller for an uninformative field to within 0.4% of evaluations. In the aggregate it used less energy per kilometer than the range-agnostic, present-range, and stance-only allocators. The evidence does not support the claim that range-aligned allocation increases mean clearance, because it produced the smallest mean margin in both tables. The single recorded high-speed failure scene suggests that targeting margin can help where the baseline fails, but one scene is not sufficient to establish a general tail-risk benefit. A practitioner choosing between the allocators should therefore treat range-aligned allocation as an energy-efficient, less conservative option whose clearance floor still needs to be checked per deployment.

06 Limitations

First, the compiled results contain one aggregate value per method and metric, without per-seed dispersion, per-scene breakdown, or separation of simulation and hardware runs. The statistical significance of the reported differences, some of which are a few tenths of a millisecond or a few centimeters, cannot be assessed. The hardware portion (8 runs per height class) is small. Second, the mean clearance margin of the proposed method is lower than that of every baseline and ablation. Deployments that require a hard clearance floor should raise h_m or $\kappa_{\min}$ and re-validate. Third, the energy metric is the

KAGGLEYEs ACADEMY • E-PRINT REGISTRY • KAGG:2610.00002

Table 1: Main comparison against range-agnostic and partial allocators. Values are reproduced from the compiled study results.

Method	Clearance margin (cm)	Energy per km (J/km)	QP solve time p95 (ms)	Nominal-recovery parity (uniform field, %)
Range-agnostic allocation (uniform weights)	24.0	1334	2.9	95.6
Present-range weights (no prediction)	22.2	1355	2.8	88.5
Stance stability only (no swing clearance)	22.9	1389	2.9	94.3
Ours	20.8	1238	2.6	99.6

Table 2: Ablation of range-aligned allocation. Values are reproduced from the compiled study results.

Variant	Clearance margin (cm)
Ours (range-aligned allocation)	20.8
w/o swing-clearance corridor	22.4
w/o stance-side bias	22.1
w/o range prediction (present only)	22.8

logged locomotion-energy figure. It is suitable for relative comparison among allocators under identical conditions, but it should not be compared directly with whole-robot cost-of-transport figures computed under other definitions. Fourth, the method inherits the errors of the range predictor. Propositions 1 and 2 concern the allocator given a field and say nothing about the consequences of a confidently wrong prediction, such as a missed shelf edge. Fifth, the analysis does not cover the active barrier constraint, the interaction with the contact-schedule generator, or the solution-reuse fallback. Sixth, the obstacles were static raised geometry of 15–60 cm. Moving obstacles, overhangs taller than the body layers, and deformable terrain were not evaluated.

07 Conclusion

This paper introduced range-aligned control allocation, in which a FutureRay-style predicted future range field sets the stance weights, a load-shift term, and swing-clearance constraints of a whole-body QP for a quadruped. A contrast statistic that vanishes exactly for uniform fields gives exact nominal recovery and a deviation bound linear in the contrast. On ANYmal in Isaac Lab and in hardware trials beside raised obstacles, the method reduced energy per kilometer from 1334 to 1238 J/km relative to range-agnostic allocation, lowered the 95th-percentile solve time from 2.9 to 2.6 ms, and matched the unweighted controller in 99.6% of uniform-field evaluations. It also produced a smaller mean clearance margin (20.8 versus 24.0 cm), while holding 11 cm in the one recorded high-speed shelf scene where the baseline failed in 6 of 30 trials. Future work should report margin distributions and tail statistics per scene, couple the contrast statistic with predictor uncertainty, and let the allocator inform the contact schedule directly.

References

[1] 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. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2016.

[2] 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, 3(3):2261-2268, 2018.

[3] D. Kim, J. Di Carlo, B. Katz, G. Bledt, and S. Kim. Highly Dynamic Quadruped Locomotion via Whole-Body Impulse Control and Model Predictive Control. arXiv:1909.06586, 2019.

[4] R. Grandia, F. Jenelten, S. Yang, F. Farshidian, and M. Hutter. Perceptive Locomotion through Nonlinear Model Predictive Control. IEEE Transactions on Robotics, arXiv:2208.08373, 2023.

[5] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter. Learning Agile and Dynamic Motor Skills for Legged Robots. Science Robotics, arXiv:1901.08652, 2019.

[6] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter. Learning Quadrupedal Locomotion over Challenging Terrain. Science Robotics, arXiv:2010.11251, 2020.

[7] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter. Learning Robust Perceptive Locomotion for Quadrupedal Robots in the Wild. Science Robotics, arXiv:2201.08117, 2022.

[8] F. Jenelten, J. He, F. Farshidian, and M. Hutter. DTC: Deep Tracking Control. Science Robotics, arXiv:2309.15462, 2024.

[9] C. D. Bellicoso, C. Gehring, J. Hwangbo, P. Fankhauser, and M. Hutter. Perception-less Terrain Adaptation through Whole Body Control and Hierarchical Optimization. IEEE-RAS International Conference on Humanoid Robots (Humanoids), 2016.

[10] A. Herzog, N. Rotella, S. Mason, F. Grimminger, S. Schaal, and L. Righetti. Momentum Control with Hierarchical Inverse Dynamics on a Torque-Controlled Humanoid. Autonomous Robots, arXiv:1410.7284, 2016.

[11] A. Escande, N. Mansard, and P.-B. Wieber. Hierarchical Quadratic Programming: Fast Online Humanoid-Robot Motion Generation. The International Journal of Robotics Research, 33(7):1006-1028, 2014.

[12] T. Zang, S. Wang, H. Huang, and S. Cheng. FutureRay: Control-Aligned Future Range for Agile Quadruped Navigation. arXiv:2608.01197, 2026.

[13] 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. Autonomous Robots, 40(3):429-455, 2016.

[14] C. Gehring, S. Coros, M. Hutter, C. D. Bellicoso, H. Heijnen, R. Diethelm, M. Bloesch, P. Fankhauser, J. Hwangbo, M. Hoepflinger, and R. Siegwart. Practice Makes Perfect: An Optimization-Based Approach to Controlling Agile Motions for a Quadruped Robot. IEEE Robotics and Automation Magazine, 23(1):34-43, 2016.

[15] 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. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2018.

[16] S. Caron, Q.-C. Pham, and Y. Nakamura. Stability of Surface Contacts for Humanoid Robots: Closed-Form Formulae of the Contact Wrench Cone for Rectangular Support Areas. IEEE International Conference on Robotics and Automation (ICRA), arXiv:1501.04719, 2015.

KAGGLEYEs ACADEMY • E-PRINT REGISTRY • KAGG:2610.00002

[17] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd. OSQP: An Operator Splitting Solver for Quadratic Programs. Mathematical Programming Computation, arXiv:1711.08013, 2020.

[18] F. Jenelten, R. Grandia, F. Farshidian, and M. Hutter. TAMOLS: Terrain-Aware Motion Optimization for Legged Systems. IEEE Transactions on Robotics, arXiv:2206.14049, 2022.

[19] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard. Crocoddyl: An Efficient and Versatile Framework for Multi-Contact Optimal Control. IEEE International Conference on Robotics and Automation (ICRA), arXiv:1909.04947, 2020.

[20] N. Rudin, D. Hoeller, P. Reist, and M. Hutter. Learning to Walk in Minutes Using Massively Parallel Deep Reinforcement Learning. Conference on Robot Learning (CoRL), arXiv:2109.11978, 2021.

[21] G. B. Margolis, G. Yang, K. Paigwar, T. Chen, and P. Agrawal. Rapid Locomotion via Reinforcement Learning. Robotics: Science and Systems (RSS), arXiv:2205.02824, 2022.

[22] A. Kumar, Z. Fu, D. Pathak, and J. Malik. RMA: Rapid Motor Adaptation for Legged Robots. Robotics: Science and Systems (RSS), arXiv:2107.04034, 2021.

[23] A. Agarwal, A. Kumar, J. Malik, and D. Pathak. Legged Locomotion in Challenging Terrains using Egocentric Vision. Conference on Robot Learning (CoRL), arXiv:2211.07638, 2022.

[24] Z. Zhuang, Z. Fu, J. Wang, C. Atkeson, S. Schwertfeger, C. Finn, and H. Zhao. Robot Parkour Learning. Conference on Robot Learning (CoRL), arXiv:2309.05665, 2023.

[25] X. Cheng, K. Shi, A. Agarwal, and D. Pathak. Extreme Parkour with Legged Robots. IEEE International Conference on Robotics and Automation (ICRA), arXiv:2309.14341, 2024.

[26] D. Hoeller, N. Rudin, D. Sako, and M. Hutter. ANYmal Parkour: Learning Agile Navigation for Quadrupedal Robots. Science Robotics, arXiv:2306.14874, 2024.

[27] P. Fankhauser, M. Bloesch, and M. Hutter. Probabilistic Terrain Mapping for Mobile Robots with Uncertain Localization. IEEE Robotics and Automation Letters, 3(4):3019-3026, 2018.

[28] J. Frey, M. Mattamala, N. Chebrolu, C. Cadena, M. Fallon, and M. Hutter. Fast Traversability Estimation for Wild Visual Navigation. Robotics: Science and Systems (RSS), arXiv:2305.08510, 2023.

[29] L. Wellhausen and M. Hutter. ArtPlanner: Robust Legged Robot Navigation in the Field. Field Robotics, arXiv:2303.01420, 2023.

[30] M. Tranzatto et al.. CERBERUS in the DARPA Subterranean Challenge. Science Robotics, arXiv:2201.07067, 2022.

[31] D. Hoeller, L. Wellhausen, F. Farshidian, and M. Hutter. Learning a State Representation and Navigation in Cluttered and Dynamic Environments. IEEE Robotics and Automation Letters, arXiv:2103.04351, 2021.

[32] N. Rudin, D. Hoeller, M. Bjelonic, and M. Hutter. Advanced Skills by Learning Locomotion and Local Navigation End-to-End. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), arXiv:2209.12827, 2022.

[33] T. Zang, S. Cheng, H. Huang, S. Wang, and W. Zhang. AgniNav: Configuration-Driven Cross-Embodiment Local Planning for Robot Navigation. arXiv:2606.10903, 2026.

[34] H. Xi, M. Tan, X. Zhang, S. Cheng, S. Wang, Y. Gu, X. Shen, and W. Zhang. CeRLP: A Cross-embodiment Robot Local Planning Framework for Visual Navigation. arXiv:2603.19602, 2026.

[35] J. Hu, S. Yuan, G. C. R. Bethala, A. Tzes, and Y. Fang. Learning Adaptive Safety Margins for Visual Navigation. arXiv:2607.18200, 2026.

[36] D. Fox, W. Burgard, and S. Thrun. The Dynamic Window Approach to Collision Avoidance. IEEE Robotics and Automation Magazine, 4(1):23-33, 1997.

[37] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa. CHOMP: Gradient Optimization Techniques for Efficient Motion Planning. IEEE International Conference on Robotics and Automation (ICRA), 2009.

[38] X. Zhou, Z. Wang, H. Ye, C. Xu, and F. Gao. EGO-Planner: An ESDF-free Gradient-based Local Planner for Quadrotors. IEEE Robotics and Automation Letters, arXiv:2008.08835, 2021.

[39] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada. Control Barrier Function Based Quadratic Programs for Safety Critical Systems. IEEE Transactions on Automatic Control, arXiv:1609.06408, 2017.

[40] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada. Control Barrier Functions: Theory and Applications. European Control Conference (ECC), arXiv:1903.11199, 2019.

[41] W. Xiao and C. Belta. Control Barrier Functions for Systems with High Relative Degree. IEEE Conference on Decision and Control (CDC), arXiv:1903.04706, 2019.

[42] R. Grandia, A. J. Taylor, A. D. Ames, and M. Hutter. Multi-Layered Safety for Legged Robots via Control Barrier Functions and Model Predictive Control. IEEE International Conference on Robotics and Automation (ICRA), arXiv:2011.00032, 2021.

[43] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, and G. State. Isaac Gym: High Performance GPU-Based Physics Simulation for Robot Learning. arXiv:2108.10470, 2021.

[44] M. Mittal, C. Yu, Q. Yu, J. Liu, N. Rudin, D. Hoeller, J. L. Yuan, R. Singh, Y. Guo, H. Mazhar, A. Mandlekar, B. Babich, G. State, M. Hutter, and A. Garg. Orbit: A Unified Simulation Framework for Interactive Robot Learning Environments. IEEE Robotics and Automation Letters, arXiv:2301.04195, 2023.

KAGGLEYEs ACADEMY • E-PRINT REGISTRY • KAGG:2610.00002