IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 14 April 2026, accepted 21 May 2026, date of publication 25 May 2026, date of current version 29 May 2026.

Digital Object Identifier 10.1109/ACCESS.2026.3696363

RESEARCH ARTICLE

Reactive and Decoupled Control of Mobile Manipulators via Spatially Modulated Velocity Fields

TADEJ PETRIČ, (Member, IEEE), AND LEON ŽLAJPAH, (Member, IEEE)
Department of Automatics, Biocybernetics and Robotics, Jožef Stefan Institute, 1000 Ljubljana, Slovenia

Corresponding author: Tadej Petrič (tadej.petric@ijs.si)

This work was supported in part by Slovenian Research Agency (ARIS) under Grant P2-0076; and in part by the Ministry of Higher Education, Science and Innovation of Slovenia, Slovenian Research and Innovation Agency, and European Union—NextGenerationEU, through Digital Transformation of Robot-Supported Factories of the Future (DIGITOP), under Grant TN-06-0106.

ABSTRACT Mobile manipulators enable robotic systems to operate flexibly in extended, dynamic, and unstructured environments. A key challenge in such systems is achieving coordinated and adaptive behavior between the mobile base and the manipulator without relying on global localization or predefined plans. This paper presents a reactive and fully local control strategy for mobile base based on spatially modulated velocity fields defined in the mobile base's frame. The method enables real-time end-effector tracking across the full workspace by exploiting only the internal kinematics of the arm. Using distance- and angle-based activation thresholds, the controller generates smooth linear and angular velocity commands for mobile base that support context-aware transitions between stationary, extension, and retreat modes. Importantly, the control framework is decoupled from the manipulator, allowing compatibility with autonomous planning, teleoperation, or physical guidance. The proposed architecture is particularly well suited for dynamic and collaborative manufacturing scenarios, where mobile robots must operate safely, robustly, and without dependence on global pose estimates. Experiments with a real mobile manipulator demonstrate stable and precise behavior across diverse trajectories, confirming the suitability of the method for adaptive, robust, and autonomous robotic systems.

INDEX TERMS Mobile manipulators, nonholonomic control, reactive control, velocity vector fields, workspace adaptation.

I. INTRODUCTION

Mobile manipulators integrate the dexterity of articulated arms with the mobility of wheeled or tracked bases, enabling interaction across extended and unstructured workspaces typical of modern manufacturing, construction, logistics, and service environments. As deployments move from structured cells toward semi-structured and human-shared spaces, systems must coordinate the whole-body motion under nonholonomic constraints, dynamic coupling, contact, and sensing limitations, all while remaining computationally lightweight and certifiably safe. Recent assessments emphasize that operational success hinges less on isolated improvements in planning or perception and more on tightly integrated stacks that preserve task feasibility while reacting robustly to layout changes, occlusions, and task re-specification in real time [1], [2].

This context has motivated a spectrum of strategies spanning global, planning-centric coordination and local, feedback-centric control. Planning methods enhance reachability and task consistency through base–arm co-design of trajectories and postures, whereas feedback control prioritizes real-time safety, constraint handling, and disturbance rejection. Across both families, practical systems increasingly combine geometric insight (e.g., capability maps, nonholonomic alignment) with optimization or robust control primitives, and augment them with learning and calibration to mitigate modeling gaps. The remainder of this section surveys these directions and positions our approach relative to recent advances.

The associate editor coordinating the review of this manuscript and approving it for publication was Nasim Ullah.

A. STATE OF THE ART

Modern planning methods address the fundamental challenge of coordinating the mobile base with the end-effector under nonholonomic constraints and workspace limits by making the geometry of feasible motions explicit. Within this perspective, task–space formulations align the commanded end effector trajectory with base kinematics that the platform can reliably realize, thereby reducing errors that arise from unreachable or ill-conditioned poses [3]. Capability map–based approaches extend this idea by guiding the choice of base configurations so that manipulability is improved and the reachable workspace expands along the task, allowing the arm to avoid singularities or unstable postures as the base moves [4]. In application-driven domains, planning often integrates the mechanics and task logic of the object itself; for instance, in dual–arm mobile manipulation for installing long deformable components, feasibility requires respecting deformation and contact constraints at the same time as maintaining valid base motions [5]. Other work emphasizes redundancy-aware strategies that blend heuristic search with analytic solvers, showing how genetic exploration, pose interpolation, and inverse kinematics can be combined to preserve feasibility even in confined spaces with multiple constraints [6]. Collectively, these contributions illustrate how planning not only enlarges the set of achievable tasks but also promotes coherent whole-body motion. Yet they also expose a recurring limitation: by depending on accurate maps, robust localization, and substantial precomputation, many planning frameworks struggle to adapt when layouts shift or objectives evolve during execution.

Feedback–centric control takes a complementary stance: rather than depending on precomputed plans, it treats feasibility, safety, and tracking as problems that must be resolved continuously under sensing and modeling uncertainty. This perspective has motivated controllers that embed geometric safety margins directly into their optimization layer, so the robot avoids self–collision without compromising the task intent [7], and others that fuse base and arm into a unified reactive system, making it possible to retarget goals on the fly as perception updates arrive [8]. As scenes grow more cluttered, researchers have shown how online trajectory generation can respect nonholonomic constraints while keeping motions smooth near singularities [9], and how robustness techniques, ranging from disturbance rejection grounded in differential flatness [10] to sliding mode control enhanced with learned compensation [11], can sustain tracking even when the model is uncertain. More recent work emphasizes whole–body safety and dynamic consistency: local optimization strategies balance obstacle avoidance with fidelity to the commanded path [12], while decomposition methods capture the coupled dynamics of wheels, joints, and flexible structures to refine accuracy [13]. From a deployment perspective, modular control architectures have proven valuable because they remain tunable and reliable as sensing conditions or task objectives shift during runtime [14]. Taken together with studies that deliberately perform manipulation while the platform is in motion and controllers that anticipate the demands of contact–rich tasks [15], [16], [17], the literature underscores a consistent message: in mobile manipulation, robustness and feasibility must be maintained not only at the planning stage but throughout execution, moment by moment, as the robot moves and interacts with its environment.

Learning and perception provide the adaptability that modeling alone cannot. End-to-end policies for pick-and-place on mobile manipulators show that deep reinforcement learning can absorb the variability typical of logistics scenarios and still succeed at the task [18]. At the autonomy level, advances in active and visual SLAM emphasize that when mapping and localization are uncertain, control performance depends critically on belief-space planning, next-best-view selection, and robust state estimation [19], [20]. Precision in contact, by contrast, often hinges on calibration and accurate models of base–arm coupling: kinematic calibration methods for collaborative robots on mobile platforms and closed-chain calibration both improve geometric fidelity [21], [22], while experimentally validated coupled dynamics create a safer basis for close interaction with humans [23]. Evidence from field deployments reinforces this point. Localized, feedback-driven pipelines have shown that robots can achieve reliable grasping and placement even in semi-structured environments with variable layouts and frequent occlusions, such as construction sites [24]. Beyond these case studies, work at very different scales has pointed to the same underlying principle: effective mobile manipulation depends on tightly integrated locomotion and manipulation. Examples such as holonomic micromanipulation guided by visual servoing and independently actuated bimanual soft manipulators in confined spaces highlight control motifs that transfer directly to ground robots, where reach, safety, and precision must be carefully balanced in constrained workspaces [25], [26].

B. OUR CONTRIBUTION

The state of the art reveals two strong but complementary directions. Planning–centered methods expand the feasible task set by co–designing base placement and end effector motion, relying on detailed models and maps to align task–space goals with nonholonomic feasibility and manipulability [3], [4], [5]. Feedback–centered controllers, in contrast, deliver whole–body behavior in real time by enforcing safety and constraints through optimization or robust control, yet they often tie base motion directly to the commanded end effector orientation or to multi–layer coordination schemes that assume stable localization and accurate dynamics [7], [8], [9], [10], [11], [12], [13], [14]. What remains missing is a lightweight and deployable policy that achieves coherent whole–body motion using only local sensing, that does not explicitly depend on end-effector orientation at the base control interface, and that can work seamlessly with existing arm controllers without the overhead of global planning or online optimization.

We address this gap with a local velocity field controller for nonholonomic mobile manipulators, where the control space is defined solely by the translational part of the end–effector position relative to the base. By excluding orientation from the policy input, the controller becomes orientation–agnostic at the base control interface, enabling the base to maintain coverage of the reachable translational workspace regardless of the desired tool orientation. The velocity field is expressed in the base frame and governed by a small set of spatial activation thresholds that organize extension, alignment, and retreat, while explicitly respecting nonholonomic motion through smooth gating and rate limiting. Control of the base and arm is decoupled at a task–space interface: the base consumes only a reduced kinematic estimate of the end–effector translation and issues constant–time linear and angular velocity commands. This design eliminates the need for global pose estimates or map–based references and remains compatible with kinesthetic, teleoperated, and autonomous arm modes. Relative to our preliminary study [27], which coupled base control to the full 6–DoF end–effector pose, the present work introduces a field–driven, task–space policy based only on translation and thus orientation–agnostic at the base control interface by construction. This reparameterization enables a clearer and more rigorous characterization of the reachable workspace. Building on this formulation, we examine how spatial activation thresholds, retreat and alignment mechanisms, and velocity smoothing and rate limiting shape the controller's behavior, and we demonstrate their influence on smoothness, stability near reach limits, and workspace coverage in both visualization and experimental validation.

The remainder of the paper is organized as follows. Section II introduces the proposed velocity–field control strategy and its spatial modulation mechanism. Section III visualizes the resulting control behavior across the three-dimensional workspace. Section IV describes the experimental setup and presents validation results on a real mobile manipulator. Finally, Section V discusses broader implications and outlines directions for future research, and Section VI concludes the paper and sketches avenues for future work.

II. CONTROL CONCEPT

This section introduces the proposed control concept, which is centered on a velocity vector field strategy tailored for nonholonomic mobile platforms. The key idea is to govern the mobile platform motion solely based on the relative position of the manipulator's translational end-effector (TEE), computed from a reduced kinematic model. To ensure computational efficiency and numerical stability, the model includes only the primary arm joints responsible for major positioning, omitting the distal joints such as those in the wrist that primarily influence orientation. This yields a representation that renders the base policy orientation-agnostic at the control interface and removes the need for full pose estimation or external tracking.

In this work, the term "velocity field" denotes a spatial mapping from the translational end–effector position expressed in the base frame to admissible base velocity commands. The mapping is constructed through smooth activation functions, proportional scaling, and explicit saturation, yielding a continuous and bounded velocity policy over the operational workspace. Although the modulation is structured through spatial activation regions, the resulting control law defines a piecewise smooth spatial mapping rather than a navigation function derived from a global scalar potential.

We represent the mobile base state vector as $\mathbf{q}_b = [x_b, y_b, \theta]^\top$, describing planar position and heading, and the manipulator joint vector as $\mathbf{q}_m = [q_1, \ldots, q_n]^\top$, with n corresponding to the number of joints used in the reduced model. The TEE position is computed directly from joint angles and expressed in the robot base frame $\mathcal{F}_\mathrm{b}$, enabling fully local control that is independent of global localization systems. This structure supports reactive behavior and reduces infrastructure requirements.

The control inputs to the mobile base are the linear and angular velocities, denoted as v and ω. The platform is modeled as a nonholonomic system with kinematics given by

$$\dot{x}_b = v\cos\theta, \tag{1}$$
$$\dot{y}_b = v\sin\theta, \tag{2}$$
$$\dot{\theta} = \omega, \tag{3}$$

where (x_b, y_b) denotes the base position and θ its heading. The control laws defined in the following specify the inputs v and ω as functions of the relative position of the TEE.

The proposed control law is defined entirely in the base-local frame and depends only on the spatial relationship between the base and its own manipulator. This makes the method particularly suitable for use cases such as kinesthetic teaching, teleoperation, or human-guided demonstrations, where global localization is unavailable or unnecessary. The control scheme is robust to minor localization errors, generalizes well across platforms, and ensures safe coordination based on internal sensing and feedback alone.

In our experimental setup, the TEE position is calculated using the reduced forward kinematics of a Franka Emika arm mounted on a PAL Tiago nonholonomic mobile base (see Fig. 1). The resulting position, denoted as $\boldsymbol{p} = [x, y, z]^\top$, is the sole input to the velocity field controller. Four geometric quantities are derived from $\boldsymbol{p}$: the total distance from the base to the TEE, $d_{\boldsymbol{p}} = \|\boldsymbol{p}\|$; the planar angular deviation $\varphi = \arctan 2(y, x)$, representing the direction of the TEE relative to the base heading; the signed distance to the outer workspace boundary, d_f; and the signed distance to the inner retreat region, d_r. These serve as the foundation for a threshold-based control strategy that modulates linear and angular velocity commands according to the spatial configuration of the end-effector. An illustration of these quantities, along with the coordinate frames and spatial

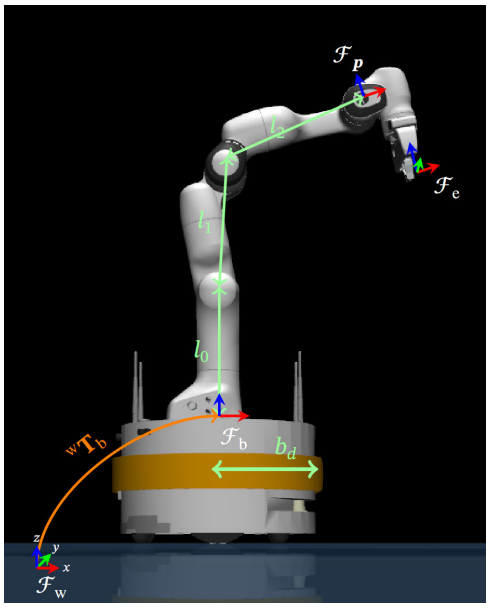

FIGURE 1. Kinematic structure and reference frames of the mobile manipulation system. The figure illustrates the Franka Emika manipulator mounted on a PAL Tiago nonholonomic mobile base. Coordinate frames are shown for the world frame $\mathcal{F}_w$, the base frame $\mathcal{F}_b$, the reduced translational end-effector (TEE) frame $\mathcal{F}_p$, and the true tool center point (TCP) frame $\mathcal{F}_e$. The point p is implicitly defined by the origin of the corresponding reference frame. Control-relevant geometric quantities include the segment lengths l_0, l_1, and l_2, and the mobile base front distance b_d.

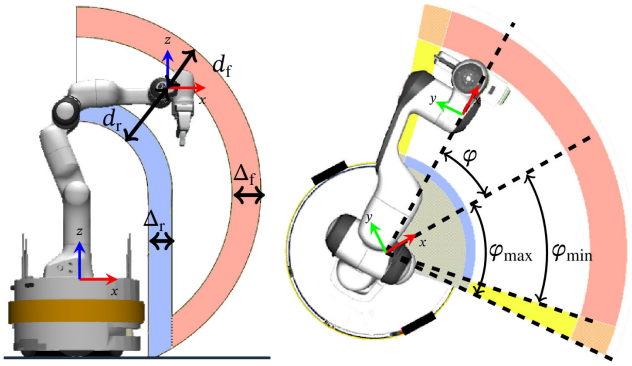

FIGURE 2. Side view (left) and top view (right) of the robot operational space, showing threshold boundaries and transition regions between stationary and active base behavior. Color coding indicates the type of motion response: blue regions trigger backward motion (retreat), red regions trigger forward motion (extension), and yellow regions trigger rotational alignment. Combinations of red or blue with yellow represent coordinated linear and angular motion — either retreat with rotation (blue + yellow) or forward motion with rotation (red + yellow).

arrangement of the system, is shown in Figs. 1 and 2. Unless stated otherwise, all quantities are expressed in the base frame $\mathcal{F}_b$.

The motion behavior of the mobile platform is governed by spatial activation thresholds that modulate transitions between stationary and active states. This partitioning is based on evaluating the position of the translational end-effector (TEE), defined as the position of the origin of frame $\mathcal{F}_p$ with respect to the robot base frame $\mathcal{F}_b$, denoted as $\boldsymbol{p} = [x, y, z]^\top$. Two signed geometric distance measures are used to assess the spatial relationship between the TEE and the base.

The first quantity, d_f, is the signed distance to a spherical workspace boundary centered at height l_0 and with radius $r_f = l_1 + l_2$, defined as

$$d_f = r_f - \|\boldsymbol{p} - \mathbf{c}\|, \tag{4}$$

$$\mathbf{c} = \begin{bmatrix} 0 \\ 0 \\ l_0 \end{bmatrix}. \tag{5}$$

The second quantity, d_r, represents the signed distance to a virtual safety region that approximates the footprint of the mobile base and its immediate vicinity. This region is modeled as a vertically aligned capsule centered at the origin of $\mathcal{F}_b$, with radius b_d, corresponding to the horizontal distance from the base center to its front boundary, and height l_0.

For positions below l_0, the constraint is evaluated in the horizontal plane using the projection $\boldsymbol{p}_\perp = [x, y]^\top$, reflecting the planar footprint of the base. For positions above l_0, the distance is computed relative to the point $\mathbf{c}$ to ensure a smooth transition to the upper workspace constraint. The signed distance is defined as

$$d_r = b_d - \begin{cases} \|\boldsymbol{p}_\perp\| & \text{if } z \le l_0, \\ \|\boldsymbol{p} - \mathbf{c}\| & \text{if } z > l_0, \end{cases} \tag{6}$$

The signed distances d_f and d_r form the basis of a spatial modulation scheme that governs the motion behavior of the mobile platform. Although the underlying control is implemented through smooth velocity fields, the system response can be intuitively interpreted through zone-based activation. These spatial activation zones define continuous transitions between stationary, forward, and retreat behaviors, depending on the relative pose of the end-effector. As illustrated in Fig. 2, the control behavior is shaped by two threshold values: Δ_f, which denotes proximity to the outer workspace boundary and triggers forward motion, and Δ_r, which defines the inner retreat region and activates backward motion when violated.

The platform remains stationary when the end-effector is positioned well within the safe and comfortable manipulation range, i.e., when $d_f \ge \Delta_f$ and $d_r > \Delta_r$. As soon as either condition is violated — when the TEE approaches the outer limit ($d_f < \Delta_f$) or encroaches too close to the base ($d_r \le \Delta_r$) — the platform becomes active. In such cases, it moves forward to extend reach or backward to preserve a safe configuration, depending on the spatial deviation.

Rather than relying on binary switching or explicitly defined "forbidden zones," the system ensures smooth transitions between behavioral modes through continuous interpolation functions. These functions generate stable, gradual responses that naturally discourage configurations beyond the robot's physical or safe operational limits. As a result, the spatial modulation strategy supports robust,

context-aware behavior while avoiding abrupt motion commands or hard-coded constraints.

The linear velocity command v governs the forward or backward motion of the mobile platform based on the spatial configuration of the manipulator's TEE. As established in the zone-based framework, two signed distances determine the activation of translational motion: d_r, measuring proximity to the inner retreat zone, and d_f, indicating closeness to the outer workspace boundary. The corresponding control law is defined as

$$v = \begin{cases} -K_r \cdot (\Delta_\mathrm{r} - d_\mathrm{r}) & \text{if } d_\mathrm{r} < \Delta_\mathrm{r}, \\ K_f \cdot (\Delta_\mathrm{f} - d_\mathrm{f}) & \text{if } d_\mathrm{f} < \Delta_\mathrm{f} \text{ and } d_\mathrm{r} \geq \Delta_\mathrm{r}, \\ 0 & \text{otherwise}, \end{cases} \tag{7}$$

where K_r and K_f are proportional gains for retreat and forward extension, respectively. When the TEE enters the retreat zone ($d_\mathrm{r} < \Delta_\mathrm{r}$), the platform moves backward, with a velocity proportional to the depth of penetration. Conversely, when the forward threshold is reached ($d_\mathrm{f} < \Delta_\mathrm{f}$), the platform advances to maintain reachability.

This formulation enables the base to react continuously and predictably to changes in the manipulator's position, extending the effective workspace when needed while preserving safety near the platform. The system remains stationary when neither threshold is violated, thereby avoiding unnecessary movement and ensuring energy-efficient operation.

The angular velocity command ω controls the heading of the mobile platform to align with the direction of the manipulator's TEE. It is based on the planar angular deviation $\varphi = \arctan 2(y, x)$, where x and y denote the coordinates of the TEE expressed in the base frame $\mathcal{F}_\mathrm{b}$. A smooth angular control policy ensures stability and prevents abrupt or oscillatory motions, particularly near alignment. The nominal control law is defined as

$$\omega_n = K_\omega \cdot \mathrm{sign}(\varphi) \cdot \gamma(|\varphi|), \tag{8}$$

where K_ω is the angular gain and $\gamma(|\varphi|) \in [0, 1]$ is a smooth activation function:

$$\gamma(|\varphi|) = \begin{cases} 0 & \text{if } |\varphi| \leq \varphi_\mathrm{min}, \\ 3s^2 - 2s^3 & \text{if } \varphi_\mathrm{min} < |\varphi| < \varphi_\mathrm{max}, \\ 1 & \text{if } |\varphi| \geq \varphi_\mathrm{max}, \end{cases}$$
$$s = \frac{|\varphi| - \varphi_\mathrm{min}}{\varphi_\mathrm{max} - \varphi_\mathrm{min}}. \tag{9}$$

This smoothstep interpolation enables gradual activation between thresholds φ_min and φ_max, as illustrated in Fig. 2 (right), ensuring precise heading correction without jitter.

When the end-effector enters the retreat zone — i.e., comes into close proximity to the base — the nominal angular control is overridden, and a predefined retreat rotation is initiated. The platform aligns with a fixed target heading φ_t, typically oriented away from the direction of approach. The retreat control law is

$$\omega_r = -\mathrm{sign}(\varphi_\mathrm{t} - \varphi) \cdot \omega_\mathrm{max} \cdot \eta(d_r), \tag{10}$$

where ω_max is the maximum angular velocity, and $\eta(d_r)$ is an activation function based on proximity to the retreat boundary. The target heading φ_t is computed once upon entering retreat mode and remains constant until the system exits that mode:

$$\varphi_\mathrm{t} = -\mathrm{sign}(\varphi) \cdot \varphi_\mathrm{max}, \tag{11}$$

where φ is the angular deviation at entry, and φ_max is the maximum heading deviation. This ensures consistent and predictable rotation away from the manipulator.

The activation function $\eta(d_r)$ modulates the retreat rotation according to proximity:

$$\eta(d_r) = \begin{cases} 0 & \text{if } d_r \geq \Delta_\mathrm{r}, \\ 1 & \text{if } d_r \leq \Delta_\mathrm{r} - \epsilon, \\ 3s^2 - 2s^3 & \text{otherwise}, \end{cases}$$
$$s = \frac{d_r - (\Delta_\mathrm{r} - \epsilon)}{\epsilon}. \tag{12}$$

This cubic interpolation provides a smooth transition, avoiding control discontinuities and improving stability near the retreat boundary. The soft activation margin ϵ ensures that retreat behavior begins with backward motion, while rotational alignment is progressively introduced as the end-effector penetrates deeper into the retreat zone.

The selected angular velocity command ω before smoothing and rate limiting is determined according to the retreat condition: when the end-effector enters the retreat zone ($d_r < \Delta_\mathrm{r}$), the platform executes the retreat-specific rotation ω_r; otherwise, it follows the nominal heading correction ω_n. It is given by

$$\omega = \begin{cases} \omega_r & \text{if } d_r < \Delta_\mathrm{r}, \\ \omega_n & \text{otherwise}. \end{cases} \tag{13}$$

To ensure smooth transitions and suppress high-frequency noise in the control signals, exponential smoothing is applied to both linear and angular velocity commands. Let k denote the current discrete-time control step. The filtered velocities $\hat{v}_k$ and $\hat{\omega}_k$ are computed as:

$$\hat{v}_k = (1 - \alpha) v_{k-1} + \alpha v_k, \tag{14}$$
$$\hat{\omega}_k = (1 - \alpha) \omega_{k-1} + \alpha \omega_k, \tag{15}$$

where $\alpha \in (0, 1)$ is a smoothing coefficient, and v_k, ω_k are the raw commands generated at the current step, while v_{k-1} and ω_{k-1} represent the commands applied in the previous cycle. This low-pass filter attenuates rapid fluctuations in velocity, contributing to improved stability and control smoothness during motion execution.

To ensure adherence to both velocity and acceleration constraints of the mobile platform, the smoothed velocity commands are passed through a two-stage rate limiter. First, the difference between the current and previous command is saturated to respect the maximum allowable rate of change

(acceleration). Then, the resulting value is clipped again to respect the platform's maximum velocity limits:

$$\tilde{v}_k = \psi_{v_{\max}}\left(v_{k-1} + \psi_{\Delta v_{\max}}(\hat{v}_k - v_{k-1})\right), \tag{16}$$

$$\tilde{\omega}_k = \psi_{\omega_{\max}}\left(\omega_{k-1} + \psi_{\Delta\omega_{\max}}(\hat{\omega}_k - \omega_{k-1})\right), \tag{17}$$

where $v_{\max}$ and $\omega_{\max}$ denote the maximum linear and angular velocities, and $\Delta v_{\max}$, $\Delta\omega_{\max}$ the corresponding velocity change limits.

The limiting function $\psi_r(x)$ is defined as:

$$\psi_r(x) = \begin{cases} -r & \text{if } x < -r, \\ x & \text{if } -r \le x \le r, \\ r & \text{if } x > r. \end{cases} \tag{18}$$

This saturation function ensures that the commanded velocities change smoothly and remain within safe dynamic bounds, preventing sudden accelerations that could destabilize the platform or violate actuator limits.

The final implemented control input to the mobile platform, corresponding to the kinematic inputs v and ω after smoothing and rate limiting, is thus:

$$\mathbf{u}_b = \begin{bmatrix} \tilde{v}_k \\ \tilde{\omega}_k \end{bmatrix}. \tag{19}$$

The use of a reduced forward kinematic model, excluding the wrist joints, focuses the control logic on the dominant contributors to translational reach, thereby improving numerical stability and simplifying both workspace representation and implementation. Despite this simplification, the resulting velocity control law remains expressive and reactive, governing meaningful platform behavior without explicit dependence on the end-effector's orientation.

For clarity, the control pipeline proceeds as follows. First, the TEE position $\boldsymbol{p}$ is computed from the reduced forward kinematics. From this, the geometric quantities $d_{\boldsymbol{p}} = \|\boldsymbol{p}\|$, φ, d_f, and d_r are evaluated. These quantities are then used to generate the linear and angular velocity commands v and ω according to the control laws defined in Eqs. (7)–(13), including both nominal and retreat behaviors with their associated activation functions. The resulting commands are subsequently smoothed using Eqs. (14)–(15) and passed through rate limiters in Eqs. (16)–(17), yielding the final control input $\mathbf{u}_b$ in Eq. (19).

Motion decisions are driven by threshold-based spatial modulation using signed distance measures between the platform and the translational end-effector. Smooth activation functions ensure continuous transitions between stationary, extension, and retreat behaviors. This spatially driven strategy enables the base to remain passive when the manipulator operates within a comfortable workspace region and to activate only when necessary—either to extend reach or to prevent unsafe proximity to the platform. As such, the system maintains safe and predictable operation without relying on binary switching or global logic.

A. CLOSED-LOOP PROPERTIES AND BOUNDEDNESS

The proposed controller operates at the kinematic coordination level and defines a mapping $\mathbf{u}_b(\mathbf{x}) : \mathbb{R}^3 \to \mathbb{R}^2$ from the translational end-effector position $\mathbf{x}$ to bounded base velocity commands. Throughout the analysis, we assume that: (i) the manipulator is governed by a stable position-level controller, and (ii) the mobile platform implements a stable low-level velocity tracking loop. Under these standard architectural assumptions, the present controller acts as a higher-level kinematic policy that modulates feasible velocity references.

Although the control law is spatially modulated and piecewise defined, its structure ensures regular and well-behaved closed-loop motion.

1) CONTINUITY AND REGULARITY

The activation functions $\gamma(|\varphi|)$ and $\eta(d_r)$ are constructed using cubic smoothstep polynomials and are therefore continuously differentiable within their transition regions. The saturation operator $\psi_r(\cdot)$ is piecewise linear and exhibits a bounded slope over its entire domain. As a result, the composed velocity mapping remains continuous throughout the workspace and exhibits bounded gradients across activation regions. In discrete-time implementation, exponential smoothing and rate limiting further eliminate abrupt transitions and suppress high-frequency switching near activation boundaries.

2) RETREAT-REGION INVARIANCE

When the end-effector enters the retreat region ($d_r < \Delta_r$), the linear velocity command is given by

$$v = -K_r(\Delta_r - d_r), \qquad K_r > 0. \tag{20}$$

Let the penetration depth into the retreat zone be defined as $V_r = \Delta_r - d_r$. For $V_r > 0$, it follows that $v < 0$, which induces motion that increases d_r and consequently decreases V_r. Therefore, the boundary $d_r = \Delta_r$ is locally repelling, and trajectories inside the retreat region are driven toward the admissible workspace. Under the assumed kinematic base model, this establishes forward invariance of the safe region $d_r \ge \Delta_r$.

3) BOUNDEDNESS OF MOTION

All commanded velocities are explicitly saturated,

$$|v| \le v_{\max}, \qquad |\omega| \le \omega_{\max}, \tag{21}$$

with additional rate limits constraining the admissible velocity increments. Therefore, the base input remains uniformly bounded for all time. For bounded manipulator motion (bounded $\mathbf{x}$), the induced base motion remains bounded as well under the assumed stable low-level velocity tracking loop. The rate-limiting stage further ensures bounded acceleration and eliminates impulsive behavior or chattering near activation surfaces.

Taken together, these structural properties characterize the controller as a bounded, locally invariant, and smooth reactive

kinematic policy. While a global asymptotic stability proof is beyond the scope of the present work, the formulation guarantees well-defined, regular, and corrective closed-loop behavior within the admissible workspace under the stated architectural assumptions.

III. SPATIAL EVALUATION OF THE LOCAL VELOCITY VECTOR FIELD

To validate and illustrate the behavior of the proposed velocity control law, we evaluate the resulting vector field over a discretized spatial domain. The field is generated based on the translational end-effector position $\boldsymbol{p} = [x, y, z]^\top$ expressed in the base frame, and the corresponding linear velocity command v is computed at each point according to the signed distances d_f and d_r as described in the previous section.

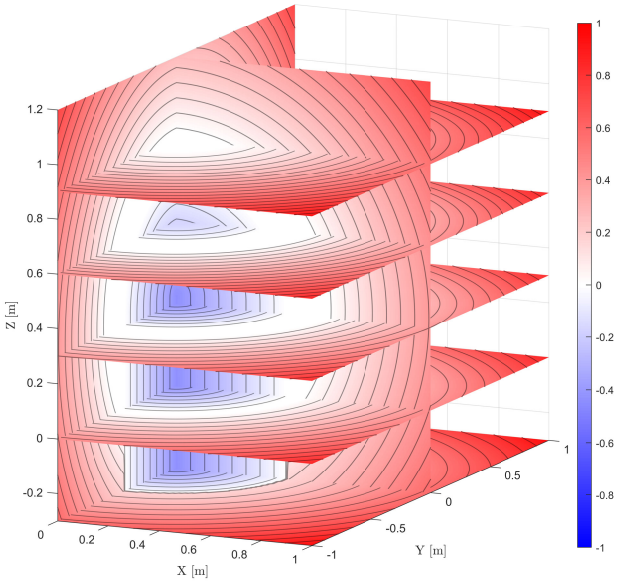

FIGURE 3. Three-dimensional contour plot of the velocity field defined over the robot's workspace. The color encodes the commanded linear velocity v as a function of the translational end-effector position p. Blue regions indicate backward motion (retreat), red regions indicate forward motion (extension), and white areas represent passive behavior near the center of the reachable space. Each horizontal slice corresponds to a fixed z level.

Fig. 3 provides a 3D contour visualization of the linear velocity field across the X-Y-Z workspace. The color encoding represents the magnitude and direction of the forward velocity component v, with red indicating forward motion (positive velocity) and blue indicating retreat behavior (negative velocity). Each horizontal slice corresponds to a different height z, visualizing how the field varies along the vertical axis of the workspace.

For a more intuitive spatial interpretation, Fig. 4 shows 2D planar slices at five selected heights ($z = -0.3$ m to $z = 0.9$ m). The background color map again encodes v, while arrows specifically represent the direction and magnitude of the angular velocity component ω. These arrows illustrate the local rotational adjustments made by the mobile platform, highlighting how it gradually aligns its orientation with the position of the manipulator's end-effector.

The smooth variation in vector orientations confirms the continuity of the angular control law, particularly near the alignment threshold φ_min and φ_max. This ensures a graceful transition between stationary and turning behavior, avoiding discontinuities that might otherwise cause jitter or instability. Together, these visualizations confirm the spatial consistency, responsiveness, and smoothness of the computed velocity field, validating its suitability for safe and reactive mobile base control in dynamic manipulation scenarios.

IV. EXPERIMENTAL VALIDATION

The experimental system comprises a Franka Emika manipulator mounted on a PAL Robotics Tiago nonholonomic mobile platform. To accurately determine the global position of the mobile base during validation, we employ a motion capture system, consistent with the methodology described in [21]. However, it is important to note that this global positioning is used solely for evaluation purposes; the proposed vector field controller does not depend on it.

Unless stated otherwise, the controller parameters used in the experiments were as follows. For the mobile-base controller, the forward, retreat, and angular gains were set to $K_f = 2.0$, $K_r = 2.0$, and $K_\omega = 1.0$, respectively, with spatial activation thresholds $\Delta_f = 0.15$ and $\Delta_r = 0.15$, and angular limits $\varphi_\mathrm{min} = 30°$ and $\varphi_\mathrm{max} = 45°$. The retreat target angle φ_t was defined with a magnitude of 30°, the smoothing coefficient to $\alpha = 0.25$, and the maximum angular velocity to $\omega_\mathrm{max} = 1.0$, while the linear and angular rate limits were set to $\Delta v_\mathrm{max} = 1.0$ and $\Delta\omega_\mathrm{max} = 1.0$, respectively. The reduced workspace geometry was defined by $l_0 = 0.333$, $l_1 \approx 0.328$, and $l_2 \approx 0.394$. For the manipulator controller, the task-space gain $\boldsymbol{K}_p$ was chosen to ensure stiff Cartesian tracking (typically on the order of 10^3 for collaborative manipulators), while the null-space velocity $\dot{\boldsymbol{q}}_n$ was defined to promote joint configurations close to the center of the joint workspace.

One of the key strengths of our approach is that the control of the mobile platform is entirely decoupled from the control strategy of the manipulator. This independence allows the use of arbitrary motion controllers for the arm, including kinesthetic guidance, teleoperation, or autonomous planning, without requiring modifications to the base control logic. The platform's behavior is driven solely by the position of the end effector relative to the base, making the system highly modular and adaptable.

A. HUMAN-GUIDED MOTION AND BASE ADAPTATION

In this scenario, we demonstrate how the proposed vector field-based control framework enables intuitive, direct teaching of mobile base behavior through physical manipulation of the robotic arm. The operator interacts with the Franka Emika

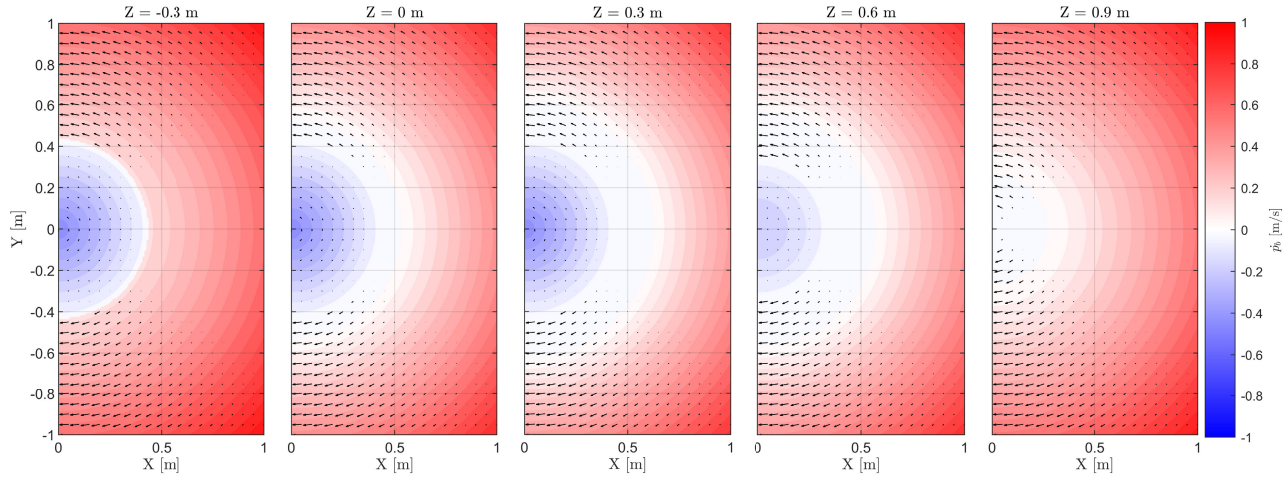

FIGURE 4. Planar slices of the velocity field in the *X*-*Y* plane at multiple *z* heights. The color map represents the linear velocity component *v*, while arrows show the direction and magnitude of the overall velocity vector. These slices illustrate how the platform transitions smoothly between passive, extension, and retreat regions across different workspace heights.

FIGURE 5. Sequence of kinesthetic guidance. As the operator physically moves the robot arm, the mobile base responds automatically by following the end-effector to maintain manipulability.This sequence corresponds to the experimental demonstration shown in Supplementary Multimedia 1.

manipulator by physically guiding it through kinesthetic motion. As the end-effector is moved, the control system continuously evaluates its position in the base frame and generates corresponding velocity commands to move the mobile platform.

Crucially, the base control is entirely driven by the relative position of the TEE with respect to the mobile base and does not rely on external references, global pose estimation, or explicit commands to the platform. As a result, the operator can reposition the entire robot by manipulating only the arm, while the mobile base responds reactively to maintain reachability and avoid undesirable configurations near its body. Similar kinesthetic interaction schemes have also been used to calibrate redundant collaborative manipulators using closed kinematic chains without relying on external motion capture systems [22]. Fig. 5 illustrates a representative sequence of kinesthetic guidance, showing how the base automatically follows the end-effector's motion to maintain a safe and reachable posture.

This use case highlights the local and reactive nature of the controller. Since the logic operates independently of how the arm is controlled, the system supports a wide range of applications, including human-in-the-loop teaching, programming by demonstration, or shared autonomy. Furthermore, because the base control requires only joint state feedback, it is inherently robust to sensor dropouts or infrastructure failures.

B. FULL-WORKSPACE TRAJECTORY TRACKING

The motion of the manipulator is governed by a velocity-based controller structured around a task-priority hierarchy. The primary objective is to follow a desired trajectory of the end-effector (TCP), while secondary tasks, such as joint

posture regulation, are handled in the null space of the Jacobian. All reference signals and errors are expressed in the world base frame.

The manipulator control operates in parallel with the mobile base controller described in Section II. The coupling between the two subsystems is established through the translational end-effector (TEE), whose position is determined by the manipulator configuration and used as the input to the mobile base control law. In this way, the manipulator influences the base motion through the TEE, while the inverse differential kinematics ensures consistent execution of the desired end-effector behavior.

The desired joint velocities are computed as:

$$\dot{\boldsymbol{q}}_m = \mathbf{J}_m^{\#} \left(\boldsymbol{K}_p \mathbf{e} + \dot{\mathbf{x}}_d \right) + \mathcal{N}_m \dot{\boldsymbol{q}}_n, \tag{22}$$

where $\mathbf{J}_m^{\#}$ represents the generalized inverse of the Jacobian matrix for the manipulator, $\boldsymbol{K}_p$ is a task-specific gain matrix constant, $\boldsymbol{e}_r$ is the position error between the desired and actual end-effector positions, and $\dot{\mathbf{x}}_d = {}^{\mathrm{b}}\mathbf{T}'_{\mathrm{w}}\, {}^{\mathrm{w}}\dot{\mathbf{x}}_d$ denotes the feed-forward Cartesian velocity expressed in the robot base frame $\mathcal{F}_{\mathrm{b}}$. The null space projection $\mathcal{N}_m$ and null space velocity $\dot{\boldsymbol{q}}_n$ support prioritization of secondary tasks, where $\dot{\boldsymbol{q}}_n$ denotes a secondary joint-space velocity policy (e.g., for redundancy resolution or posture control) that lies in the null space of the primary task.

The position error $\boldsymbol{e}_r$ in the robot base frame $\mathcal{F}_{\mathrm{b}}$ is defined as:

$$\boldsymbol{e} = \mathbf{x}_d - \mathbf{x}, \tag{23}$$

$$\mathbf{x}_d = {}^{\mathrm{b}}\mathbf{T}'_{\mathrm{w}}\, {}^{\mathrm{w}}\mathbf{x}_d, \tag{24}$$

where ${}^{\mathrm{b}}\mathbf{T}'_{\mathrm{w}}$ represents the transformation from the world frame to the robot base frame, ${}^{\mathrm{w}}\mathbf{x}_d$ is the desired position in the world frame $\mathcal{F}_{\mathrm{w}}$, and $\mathbf{x}$ is the current end-effector position in the robot base frame $\mathcal{F}_{\mathrm{b}}$. Note that $\mathbf{x}$ is equivalent to ${}^{b}\mathbf{x}$; unless explicitly stated otherwise, the reference frame notation b has been omitted throughout the remainder of the text for clarity. The generalized Jacobian inverse $\mathbf{J}_m^{\#}$ is computed using a damped least-squares approach:

$$\mathbf{J}_m^{\#} = \mathbf{J}_m^T \left(\mathbf{J}_m \mathbf{J}_m^T + \lambda \mathbf{I} \right)^{-1}, \tag{25}$$

where λ is a regularization parameter (e.g., $\lambda = 0.05$) that stabilizes the inversion near singularities. The null space projector $\mathcal{N}_m$, which allows the prioritization of secondary tasks without impacting the primary task, is defined as:

$$\mathcal{N}_m = \mathbf{I} - \mathbf{J}_m^{\#} \mathbf{J}_m. \tag{26}$$

enables secondary objectives to be pursued without interfering with the primary task. Importantly, the motion of the mobile base is controlled independently, according to the vector field strategy described in Section II. This decoupled structure allows the base to respond continuously to the manipulator's pose, ensuring coverage of the reachable translational workspace while maintaining smooth and coordinated behavior.

To demonstrate the effectiveness of the proposed vector field control framework, we conducted a series of experiments in which the manipulator's end-effector was guided throughout its entire reachable workspace. These experiments explore the platform's motion behavior under different spatial configurations, illustrating how the control law enables smooth and adaptive tracking while maintaining safe operation. Each scenario highlights a specific aspect of the system's design, such as forward and backward extension, angular realignment, or coordinated motion across height and lateral deviation. The following figures present visual comparisons and motion profiles recorded during these trials. It should be noted that the spatial zones defined in Section II are expressed in the base frame and depend on the instantaneous TEE position. As such, they are not fixed in the global frame and are not directly visualized in the trajectory plots presented below.

Fig. 6 presents a detailed comparison between two system configurations during a forward end-effector tracking task: one with the mobile base fixed (left column) and one with base motion enabled (right column). The commanded trajectory moves along the longitudinal x-axis, challenging the manipulator to maintain tracking as the task progresses. The first row shows the actual and reference positions of the end-effector (TCP), where the deviation from the desired trajectory is clearly visible in the fixed-base scenario due to limited reach. This is further quantified in the second row, which plots the position tracking error and reveals significantly higher deviations when the base remains stationary. The third row illustrates the translational velocity of the end-effector in the world frame, highlighting how the manipulator compensates for the missing base motion with more aggressive joint actuation. The fourth row expresses the same end-effector velocity in the robot base frame, which forms the input for base control in the vector field strategy. Note that oscillations that can be observed when mobile platform is moving is due to suspension of the mobile platform. The fifth row shows the linear velocity of the mobile base: in the fixed case it remains zero, while in the enabled case it responds adaptively to support the manipulator. Finally, the sixth row depicts the position of the mobile base, confirming forward displacement when motion is allowed. Together, these plots underscore the importance of coordinated base behavior for achieving full-workspace tracking. When the base remains stationary, the manipulator alone cannot maintain accuracy due to kinematic constraints; however, when base motion is enabled, the system preserves alignment with the reference trajectory and minimizes tracking error through smooth, cooperative movement.

Fig. 7 illustrates the retreat behavior of the mobile platform at two different vertical positions of the end-effector. The left column corresponds to a higher elevation, while the right column shows the behavior when the end-effector operates closer to the ground. Each plot contains four rows capturing key aspects of the system's motion. The first

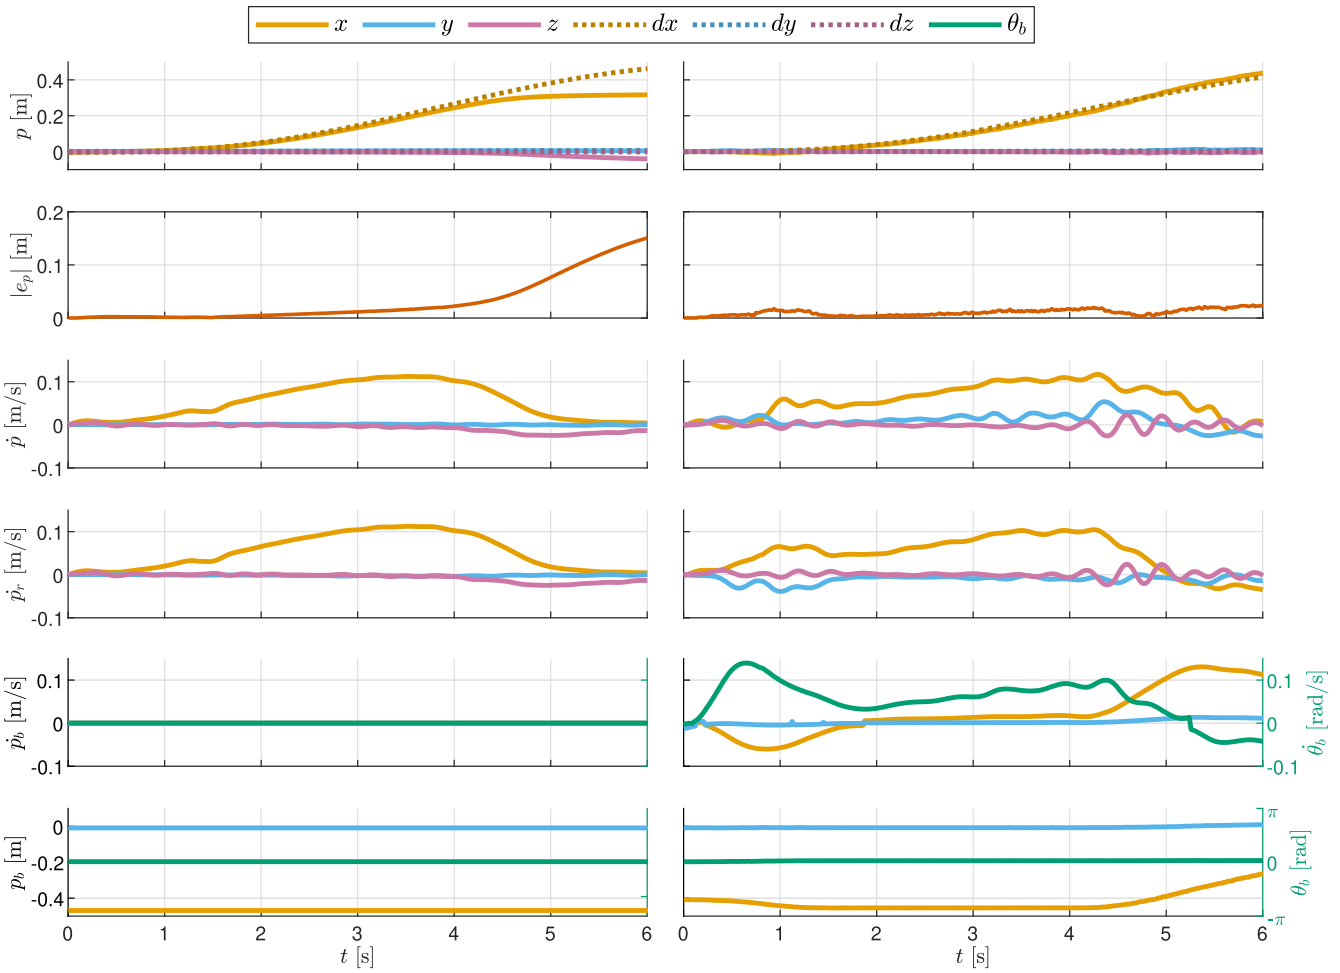

FIGURE 6. Comparison of forward trajectory tracking with mobile base motion disabled (left column) and enabled (right column). From top to bottom, the rows show: (1) end-effector (TCP) position and reference along the *x*-axis, (2) position tracking error, (3) TCP linear velocity in the world frame, (4) TCP linear velocity in the robot base frame, (5) linear and rotational velocity of the mobile base, and (6) mobile base position. The results correspond to the experiment shown in Supplementary Multimedia 2.

row shows the full 3D position of the end-effector (TCP), illustrating the spatial trajectory over time. The second row presents the position tracking error, indicating how accurately the manipulator follows the reference path under each condition. The third row shows the linear velocity of the mobile base, which remains low at higher elevations — where only minor rotation is required—but increases significantly during low-height operation, triggering a full retreat. The fourth row displays the 2D position and orientation of the mobile base, capturing both translational displacement and heading adjustments. In the low-height scenario, the platform not only moves backward but also rotates to align away from the manipulator, consistent with the retreat strategy. These results confirm that the controller adapts its response based on vertical end-effector position, modulating the retreat behavior to ensure safety across the full manipulation volume.

Fig. 8 illustrates complex, multi-axis coordination between the manipulator and the mobile platform during lateral and vertical end-effector movement. In the left column, the end-effector is guided laterally in the x–y plane, prompting the platform to first execute a slight backward motion before rotating to align with the new direction of reach. This behavior demonstrates the controller's prioritization of retreat for safety before attempting reorientation, avoiding unstable or abrupt turning. In the right column, the end-effector follows a more dynamic 3D trajectory: it first moves upward, causing the base to drive forward to maintain reachability, and then descends, triggering a retreat to prevent proximity violations. Throughout this motion, the platform dynamically modulates both linear and angular velocity to remain aligned with the arm's direction. This experiment highlights the layered spatial behavior of the control strategy, combining vertical sensitivity with smooth rotational alignment and confirming the system's ability to coordinate motion across all axes in response to real-time changes in end-effector pose.

Fig. 9 investigates the system's response to forward tracking under sinusoidal trajectories with different lateral amplitudes. In the left column, the end-effector follows a smooth sinusoidal path with a relatively small amplitude, resulting in predominantly forward motion of the platform.

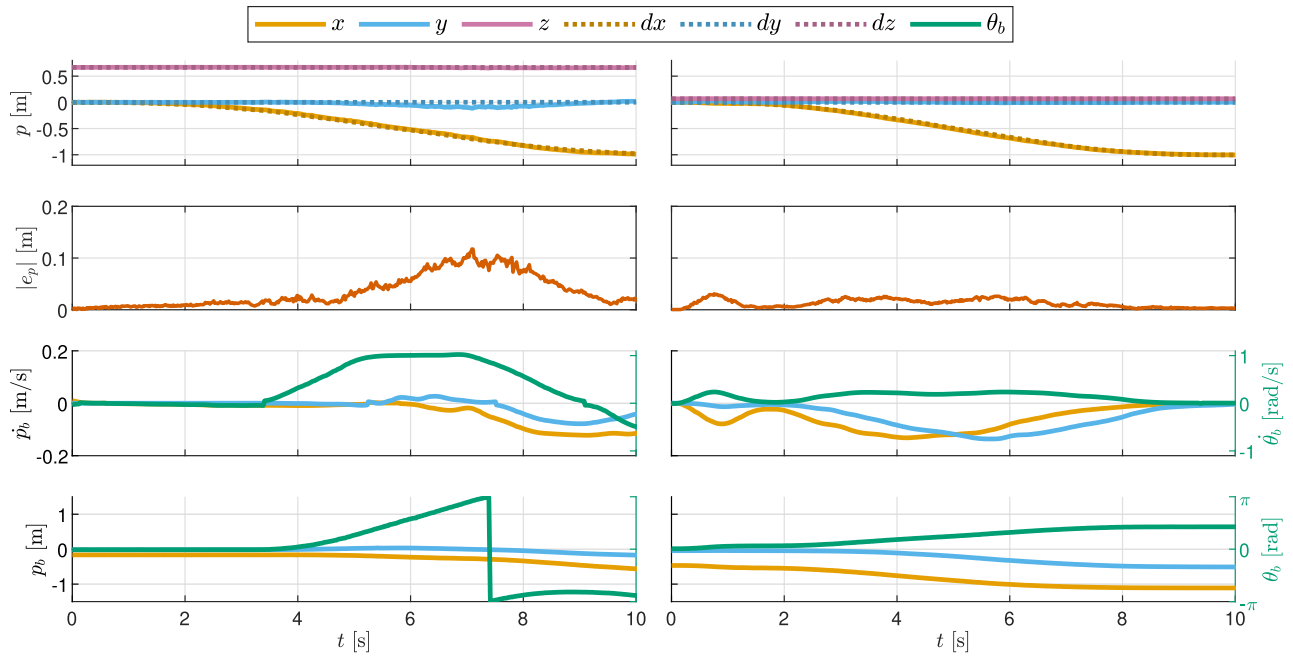

FIGURE 7. Comparison of retreat behavior at two different end-effector heights: higher (left column) and lower (right column). Rows from top to bottom: (1) end-effector (TCP) position in 3D space, (2) position tracking error, (3) linear and rotational velocity of the mobile base, (4) mobile base position and orientation in the plane. At higher elevation, the platform primarily rotates; at lower heights, it retreats and reorients to maintain a safe configuration. The results correspond to the experiment shown in Supplementary Multimedia 3.

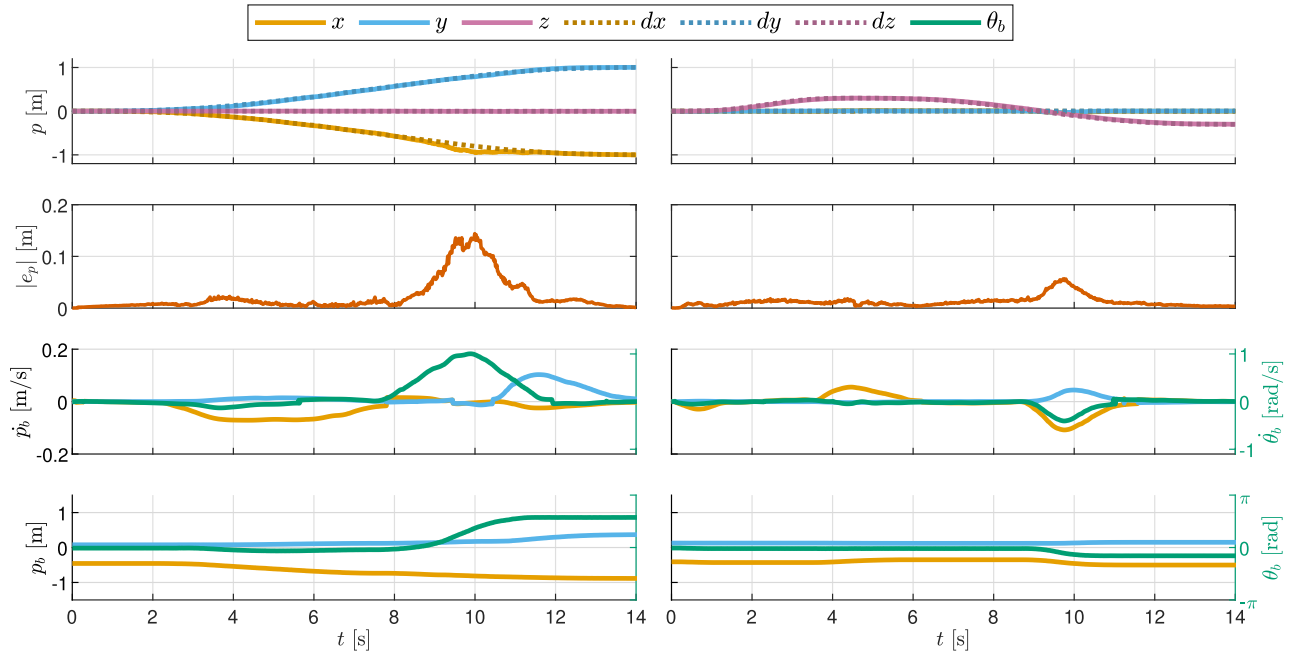

FIGURE 8. Coordinated platform behavior during lateral and vertical end-effector motion. Left column: the end-effector moves sideways in the *x–y* plane, triggering an initial retreat followed by rotation. Right column: the end-effector moves upward, prompting the base to advance, and then downward, triggering retreat and reorientation. Rows from top to bottom: (1) end-effector (TCP) position in 3D space, (2) position tracking error, (3) linear and rotational velocity of the mobile base, (4) mobile base position and orientation in the plane. The results correspond to the experiment shown in Supplementary Multimedia 4.

Because the lateral deviation remains below the angular threshold, the base maintains alignment with minimal heading correction. In the right column, the amplitude of the sinusoidal trajectory is increased, introducing more pronounced side-to-side motion. As a result, the platform must not only move forward but also adjust its orientation dynamically to maintain alignment with the end-effector. This example demonstrates how the angular velocity controller activates proportionally as the end-effector deviates laterally, enabling smooth and coordinated rotation. The system effectively

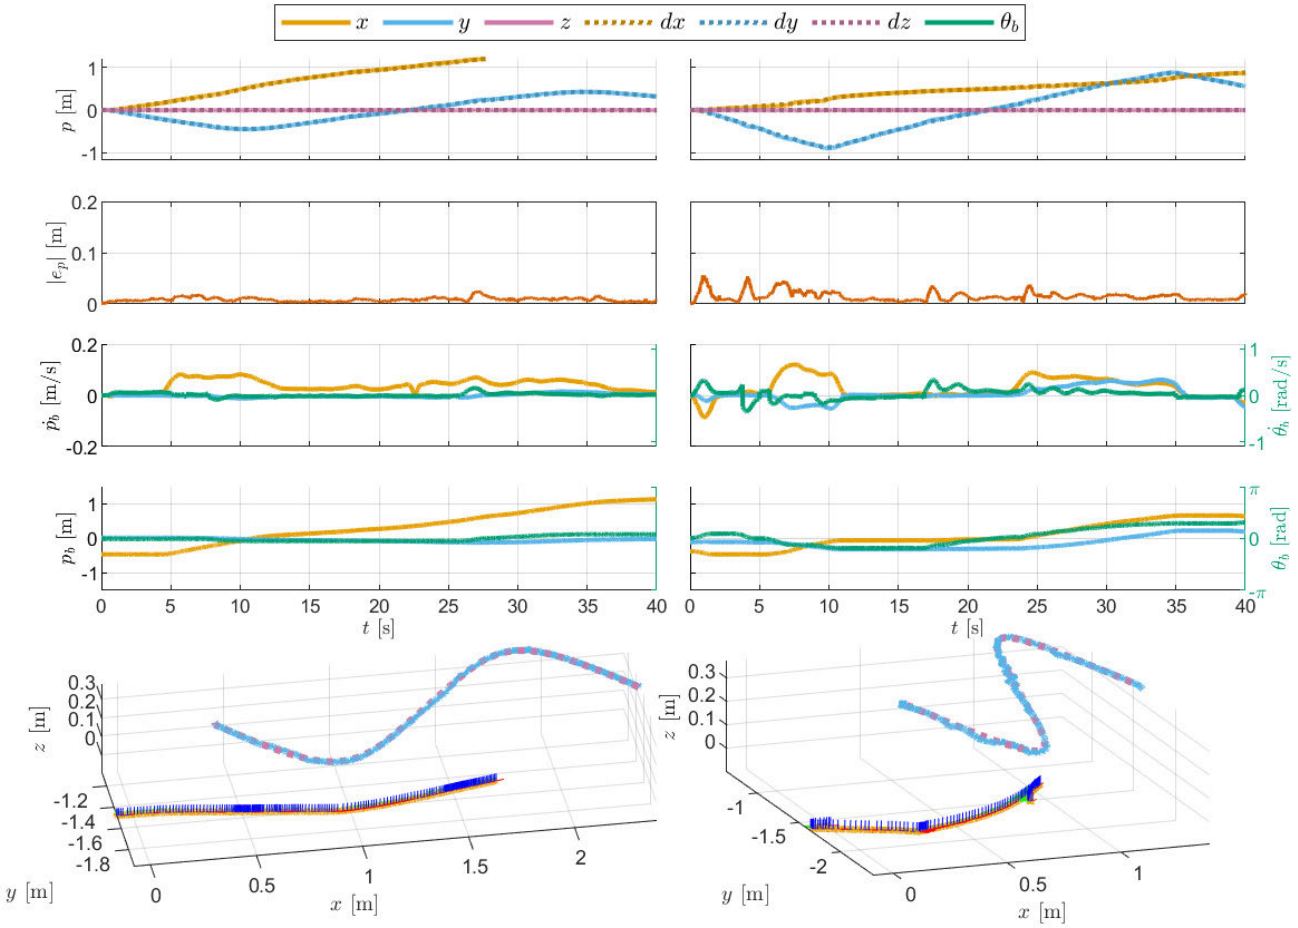

FIGURE 9. Tracking behavior under sinusoidal end-effector trajectories with small (left column) and large (right column) lateral amplitudes. Rows from top to bottom: (1) end-effector (TCP) position in 3D space, (2) position tracking error, (3) linear and rotational velocity of the mobile base, (4) mobile base position and orientation in the plane, (5) 3D path trace of the end-effector and mobile base. For small lateral deviation, the platform advances with minimal rotation; for larger amplitudes, it also rotates to follow the curved path. The results correspond to the experiment shown in Supplementary Multimedia 5.

handles both cases, confirming that the base can fluidly adapt to complex paths and support full-body trajectory tracking through integrated linear and angular motion.

Fig. 10 presents two representative examples of full-workspace trajectory tracking. In the left column, the manipulator is guided along a square-shaped path in 3D space, which includes abrupt direction changes at each corner. The mobile base responds by actively repositioning and reorienting itself to maintain alignment with the end-effector, even as the trajectory moves across workspace boundaries. The controller successfully handles the sharp transitions without inducing oscillations or instability. In the right column, a more complex and dynamic trajectory is shown — featuring arbitrary motion in elevation and lateral directions. The mobile base tracks these variations in real time, adapting its motion smoothly and reactively. These results demonstrate the robustness and generality of the proposed control strategy: it enables stable and accurate full-body coordination across highly diverse and unconstrained task-space trajectories.

V. DISCUSSION

The proposed controller proves most effective in precisely those situations where mobile manipulation typically struggles: when global localisation is unreliable, when workcells are frequently reconfigured, or when robots must switch rapidly between autonomous and interactive modes. By driving the base solely from the translational end–effector position relative to the base, the policy is orientation–agnostic at the base control interface, meaning that the base controller does not explicitly depend on end–effector orientation, although orientation may indirectly influence reachability and manipulability through the manipulator kinematics. In this way, the base autonomously restores coverage of the reachable workspace without constraining how the manipulator orientation is controlled.

Because the controller requires only a reduced translational state and generates bounded velocity commands, base motion is decoupled at the coordination layer from the internal arm controller, while remaining compatible with arbitrary manipulator control strategies. This decoupling is grounded

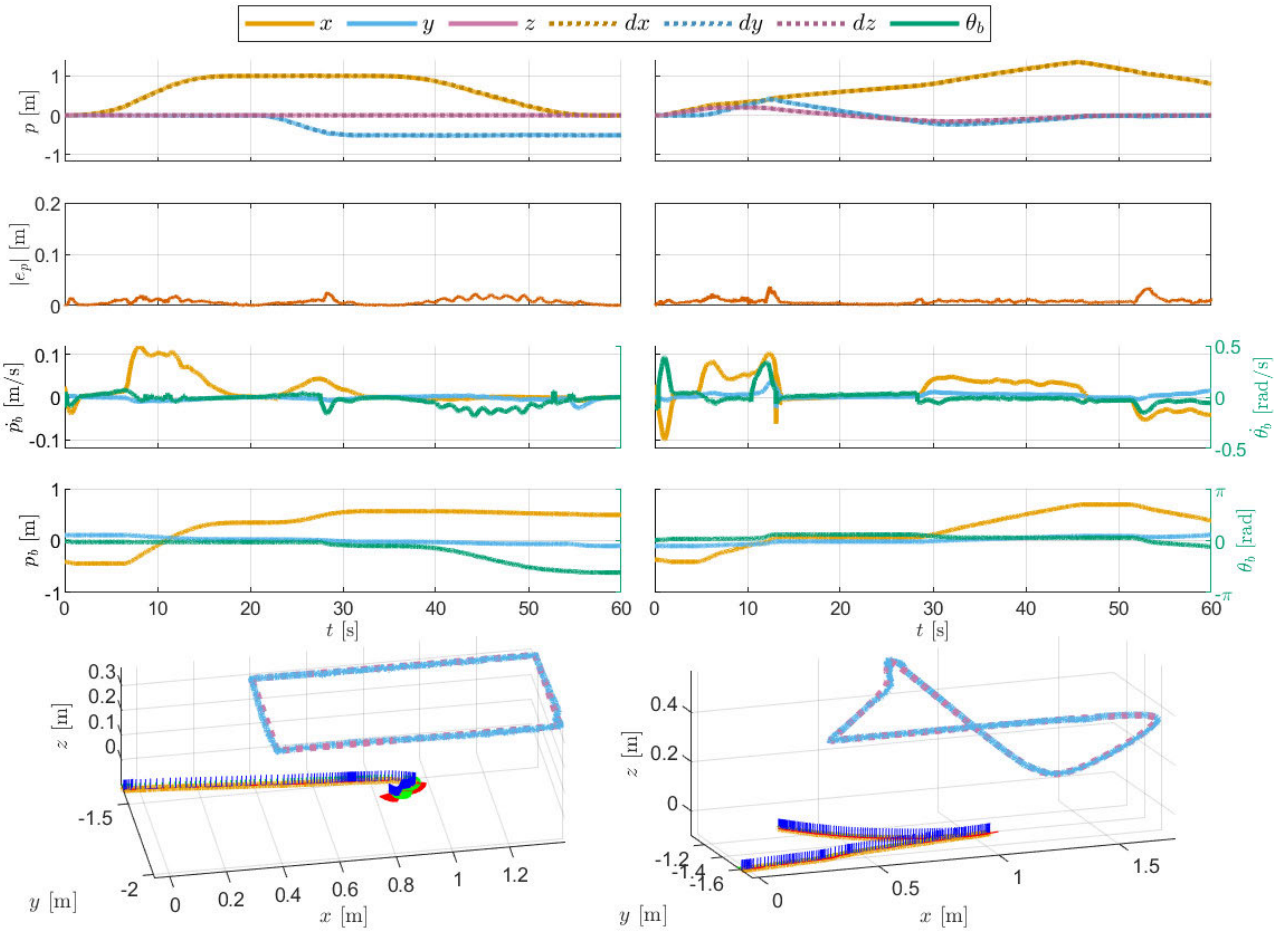

FIGURE 10. Tracking of representative full-workspace trajectories. Left column: a square-shaped 3D trajectory with sharp transitions; right column: a smooth, dynamic trajectory with arbitrary elevation and lateral components. Rows from top to bottom: (1) end-effector (TCP) position in 3D space, (2) position tracking error, (3) linear and rotational velocity of the mobile base, (4) mobile base position and orientation in the plane, (5) 3D path trace of the end-effector and mobile base. The controller enables stable and reactive tracking for both structured and unconstrained motions. The results correspond to the experiment shown in Supplementary Multimedia 6.

in a layered architecture in which the manipulator position controller and the platform velocity controller are assumed to be internally stable. Dynamic coupling effects, including payload variations, are therefore absorbed within these low-level loops. Since the base policy depends exclusively on the measured translational end–effector position, its structure remains unchanged under payload variation; however, degraded end–effector tracking accuracy may indirectly influence base behaviour through the kinematic interface. The present contribution thus intentionally addresses kinematic coordination rather than dynamic whole-body optimal control.

Unlike optimization-based whole-body controllers, the proposed method does not solve a quadratic program, model predictive control problem, or dynamic coupling formulation at each time step. Instead, it defines a closed-form spatial velocity mapping that can be evaluated in constant time, without iterative optimisation or global pose estimation. In contrast to artificial potential field approaches used for navigation, the present formulation is not derived from a global potential function but from task-driven workspace modulation based on end–effector position. Furthermore, while several reactive mobile manipulation strategies explicitly couple base motion to full end–effector pose or dynamic models, the proposed controller operates solely on the translational component at the coordination layer, leaving orientation and dynamics to lower-level controllers.

More broadly, planning–centric methods expand feasibility and manipulability through base–arm co–design [3], [4], [5], but their reliance on accurate models and maps limits adaptability when goals or layouts change online. Similarly, recent reactive controllers achieve impressive whole–body behaviour using online optimisation or detailed coupled models [9], [12], [13], yet they typically tie base motion either to the commanded end–effector orientation or to multi–layer coordination frameworks that presume stable state estimation. Our approach replaces these assumptions with a lightweight, kinematic–level interface that is reactive and deployable.

Relative to geometry–aware safety and self–collision handling in hierarchical QP [7], and to unified whole–body formulations that leverage perception in the loop for agile retargeting [8], our method trades explicit optimisation for a vector field that respects nonholonomic motion by design and structures extension, alignment, and retreat through spatial activations. This simplicity does not preclude stronger guarantees or more advanced posture choices; instead, it provides clear composition points. For instance, barrier–function projections or constraint filters could wrap the field to enforce explicit distance margins when required, robust inner loops could enhance tracking under flexibility and disturbances [10], [11], and capability–map information could bias the field toward regions of higher manipulability [4].

Because the controller relies only on a reduced kinematic estimate of the translational pose, its performance benefits directly from improved calibration and coupled modelling of the manipulator [21], [22], [23]. These properties are consistent with reports from semi–structured deployments, where localized pipelines have sustained precise manipulation despite occlusions and layout variability [24], and with industrial surveys that highlight adaptability and robustness as higher priorities than heavy precomputation [2].

The proposed control principle is not tied to a specific hardware configuration. Because the velocity-field policy depends solely on the translational position of the end–effector expressed in the base frame, it generalizes to manipulators with arbitrary degrees of freedom, provided that forward kinematics are available, as the number of joints does not affect the structure of the base controller. The formulation directly applies to nonholonomic velocity-controlled platforms such as differential-drive systems, while alternative kinematic structures, including omnidirectional or car-like steering bases, can accommodate the same spatial policy through standard kinematic transformations. Furthermore, the framework can be extended to multi-arm systems by defining the coordination variable as a suitable function of multiple end–effector positions, for example through weighted combinations or task-priority selection, although such extensions were beyond the scope of the present study.

The controller gains and spatial thresholds are chosen based on geometric and kinematic considerations rather than numerical optimisation. The distance thresholds Δ_f and Δ_r are directly related to the reachable workspace boundaries, while angular activation limits reflect acceptable heading deviation before corrective rotation is initiated. Linear and angular gains scale the responsiveness of the base motion but do not alter the qualitative structure of the policy. Due to explicit velocity saturation, smooth activation functions, and rate limiting, the closed-loop behaviour remains bounded and structurally stable across a wide range of gain values. In practice, parameter tuning was performed to balance responsiveness and smoothness, and the controller was observed to be insensitive to moderate gain variations.

A. ASSUMPTIONS AND LIMITATIONS

The proposed framework assumes that the translational end–effector position is available with sufficient accuracy from internal joint sensing and forward kinematics. The base controller operates purely in the base-local frame and therefore does not require knowledge of the global pose of the mobile platform for reactive coordination, such as kinesthetic guidance or teleoperation. However, when executing trajectories defined in a global reference frame, knowledge of the base pose becomes necessary for trajectory generation and task specification, though not for the internal velocity-field policy itself. In the experimental validation, global pose measurements were used exclusively for reference tracking and evaluation, not for computing base velocity commands. The method further assumes internally stable position control of the manipulator and stable velocity tracking of the mobile platform. Moderate sensing delays and encoder noise are mitigated by smoothing and rate limiting, whereas large unmodeled disturbances or severe tracking degradation may affect performance. These assumptions delineate the intended operational scope of the proposed framework and clarify the conditions under which its performance can be expected to remain consistent.

VI. CONCLUSION

We introduced a local velocity–field controller for nonholonomic mobile manipulators that drives the base solely from the translational end–effector position, making the policy orientation–agnostic at the base control interface, lightweight, and fully decoupled from the arm controller. The method organises extension, alignment, and retreat through spatial activation thresholds with smooth gating and rate limiting, requires no global localisation or maps, and integrates seamlessly with kinesthetic, teleoperated, and autonomous arm modes. Experiments on tasks spanning the full manipulator workspace confirmed coherent whole–body motion, stable behaviour near workspace limits, and responsiveness to task changes, while maintaining computational efficiency suitable for embedded deployment. Looking ahead, future work will couple the base field with posture optimisation for redundant arms, introduce soft orientation preferences when applications demand specific approach directions, and add certified safety layers and robustness analyses to extend applicability in contact–rich and uncertain environments.

REFERENCES

[1] T. Sandakalum and M. H. Ang, "Motion planning for mobile manipulators—A systematic review," *Machines*, vol. 10, no. 2, p. 97, Jan. 2022.
[2] N. Ghodsian, K. Benfriha, A. Olabi, V. Gopinath, and A. Arnou, "Mobile manipulators in Industry 4.0: A review of developments for industrial applications," *Sensors*, vol. 23, no. 19, p. 8026, Sep. 2023.
[3] S. Zhou, Y. C. Pradeep, M. Zhu, K. Amezquita-Semprun, and P. Chen, "Motion control of a nonholonomic mobile manipulator in task space," *Asian J. Control*, vol. 20, no. 5, pp. 1745–1754, Sep. 2018.
[4] H. Zhang, Q. Sheng, Y. Sun, X. Sheng, Z. Xiong, and X. Zhu, "A novel coordinated motion planner based on capability map for autonomous mobile manipulator," *Robot. Auto. Syst.*, vol. 129, Jul. 2020, Art. no. 103554, doi: 10.1016/j.robot.2020.103554.

[5] Y. Qin, A. Escande, F. Kanehiro, and E. Yoshida, "Dual-arm mobile manipulation planning of a long deformable object in industrial installation," *IEEE Robot. Autom. Lett.*, vol. 8, no. 5, pp. 3039–3046, May 2023.

[6] K. Vazquez-Santiago, C. F. Goh, and K. Shimada, "Motion planning for kinematically redundant mobile manipulators with genetic algorithm, pose interpolation, and inverse kinematics," in *Proc. IEEE 17th Int. Conf. Autom. Sci. Eng. (CASE)*, Aug. 2021, pp. 1167–1174.

[7] K. Jang, S.-H. Kim, and J. Park, "Reactive self-collision avoidance for a differentially driven mobile manipulator," *Sensors*, vol. 21, no. 3, p. 890, 2021. [Online]. Available: https://www.mdpi.com/1424-8220/21/3/890

[8] J. Haviland, N. Sunderhauf, and P. Corke, "A holistic approach to reactive mobile manipulation," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 3122–3129, Apr. 2022.

[9] G. Wang, H. Ma, H. Wang, P. Ding, H. Bai, W. Xu, W. Wang, and Z. Du, "Reactive mobile manipulation based on dynamic dual-trajectory tracking," *Robot. Auto. Syst.*, vol. 172, Feb. 2024, Art. no. 104589. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921889023002282

[10] D. Feliu-Talegon, H. Sira-Ramirez, and V. Feliu-Batlle, "Trajectory tracking problem for a flexible mobile manipulator: A flatness based approach combined with active disturbance rejection control," *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 6338–6343, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405896323011904

[11] S. Cui, H. Song, T. Zheng, and P. Dai, "Trajectory tracking control of mobile manipulator based on improved sliding mode control algorithm," *Processes*, vol. 12, no. 5, p. 881, Apr. 2024. [Online]. Available: https://www.mdpi.com/2227-9717/12/5/881

[12] C. Zheng, Y. Li, Z. Song, Z. Bi, J. Zhou, B. Zhou, and J. Ma, "Local reactive control for mobile manipulators with whole-body safety in complex environments," *IEEE Robot. Autom. Lett.*, vol. 10, no. 5, pp. 4556–4563, May 2025.

[13] H. Xing, Y. Xu, L. Ding, J. Chen, H. Gao, and M. Tavakoli, "Trajectory tracking control of wheeled mobile manipulators with joint flexibility via virtual decomposition approach," *IEEE Trans. Autom. Sci. Eng.*, vol. 22, pp. 11808–11825, 2025.

[14] T. P. Nguyen, H. T. Nguyen, and H. Q. T. Ngo, "Development of the robotic motion controller for a wheeled manipulator," *Meas. Control*, vol. 58, no. 3, pp. 323–338, 2024.

[15] B. Burgess-Limerick, C. Lehnert, J. Leitner, and P. Corke, "An architecture for reactive mobile manipulation on-the-move," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2023, pp. 1623–1629.

[16] M. V. Minniti, R. Grandia, K. Fäh, F. Farshidian, and M. Hutter, "Model predictive robot-environment interaction control for mobile manipulation tasks," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2021, pp. 1651–1657.

[17] S. Kim, K. Jang, S. Park, Y. Lee, S. Y. Lee, and J. Park, "Whole-body control of non-holonomic mobile manipulator based on hierarchical quadratic programming and continuous task transition," in *Proc. IEEE 4th Int. Conf. Adv. Robot. Mechatronics (ICARM)*, Jul. 2019, pp. 414–419.

[18] A. Iriondo, E. Lazkano, L. Susperregi, J. Urain, A. Fernandez, and J. Molina, "Pick and place operations in logistics using a mobile manipulator controlled with deep reinforcement learning," *Appl. Sci.*, vol. 9, no. 2, p. 348, Jan. 2019. [Online]. Available: https://www.mdpi.com/2076-3417/9/2/348

[19] J. A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, and J. A. Castellanos, "A survey on active simultaneous localization and mapping: State of the art and new frontiers," *IEEE Trans. Robot.*, vol. 39, no. 3, pp. 1686–1705, Jun. 2023.

[20] I. A. Kazerouni, L. Fitzgerald, G. Dooly, and D. Toal, "A survey of state-of-the-art on visual SLAM," *Expert Syst. Appl.*, vol. 205, Nov. 2022, Art. no. 117734. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417422010156

[21] L. Žlajpah and T. Petrič, "Kinematic calibration for collaborative robots on a mobile platform using motion capture system," *Robot. Comput.-Integr. Manuf.*, vol. 79, Oct. 2022, Art. no. 102446. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0736584522001296

[22] T. Petrič and L. Žlajpah, "Kinematic model calibration of a collaborative redundant robot using a closed kinematic chain," *Sci. Rep.*, vol. 13, no. 1, p. 17804, 2023, doi: 10.1038/s41598-023-45156-6.

[23] Z. Zhou, X. Yang, H. Wang, and X. Zhang, "Coupled dynamic modeling and experimental validation of a collaborative industrial mobile manipulator with human–robot interaction," *Mechanism Mach. Theory*, vol. 176, Oct. 2022, Art. no. 105025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0094114X22002737

[24] P. Stibinger, G. Broughton, F. Majer, Z. Rozsypalek, A. Wang, K. Jindal, A. Zhou, D. Thakur, G. Loianno, T. Krajnik, and M. Saska, "Mobile manipulator for autonomous localization, grasping and precise placement of construction material in a semi-structured environment," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 2595–2602, Apr. 2021.

[25] M. Suzuki, Y. Iida, Y. Tsukui, H. Kusama, R. Kinoshita, E. Kusui, Y. Sunohara, R. Minegishi, Y. Sugiyama, Y. Nishimura, C. Sekine, and O. Fuchiwaki, "Automatic holonomic mobile micromanipulator for submillimeter objects inspired by the rhinoceros beetle," *Adv. Intell. Syst.*, vol. 6, no. 4, Apr. 2024, Art. no. 2300517. [Online]. Available: https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/aisy.202300517

[26] Z. Koszowska, M. Brockdorff, T. da Veiga, G. Pittiglio, P. Lloyd, T. Khan-White, R. A. Harris, J. W. Moor, J. H. Chandler, and P. Valdastri, "Independently actuated soft magnetic manipulators for bimanual operations in confined anatomical cavities," *Adv. Intell. Syst.*, vol. 6, no. 2, Feb. 2024, Art. no. 2300062. [Online]. Available: https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/aisy.202300062

[27] T. Petrič, T. Brecelj, and L. Žlajpah, "Integrated control strategy for nonholonomic mobile robots with anthropomorphic manipulators," in *Advances in Service and Industrial Robotics*. Cham, Switzerland: Springer, 2025, pp. 486–494.

TADEJ PETRIČ (Member, IEEE) received the D.Sc. degree in robotics from the Faculty of Electrical Engineering, University of Ljubljana, Slovenia, in 2013. He began his research career at the Jožef Stefan Institute, in 2008. In 2015, he was a Postdoctoral Fellow at the Biorobotics Laboratory, École Polytechnique Fédérale de Lausanne (EPFL). He conducted part of his doctoral research at the German Aerospace Center (DLR), Department of Robotic Systems for Dynamic Control of Legged Humanoid Robots, and was a Visiting Researcher at ATR Computational Neuroscience Laboratories, Japan, in 2013. He is currently a Senior Research Associate with the Department of Automation, Biocybernetics, and Robotics, Jožef Stefan Institute, and an Associate Professor with the Jožef Stefan International Postgraduate School, Slovenia. He is also the Head of the Laboratory for the Advancement of Collaborative Robot Behavior in Physical Human–Robot Interaction (CoBoTaT). His current research interests include biologically inspired robotic controllers that achieve human-like robustness and adaptability in changing environments, particularly in the context of physical human–robot interaction. He is the Co-Chair of the IEEE Technical Committee on Human Movement Understanding.

LEON ŽLAJPAH (Member, IEEE) received the D.Sc. degree in electrical engineering from the University of Ljubljana, Ljubljana, Slovenia, in 1989. He joined as a Researcher at the Jožef Stefan Institute, in 1981, where he has been a Senior Consultant, since 2009. From 2005 to 2014, he was the Head of the Department of Automation, Biocybernetics, and Robotics. His contributions span the control of redundant robots, coherent robot control, sensor systems, robot applications, and, more recently, the control of robots in physical human–robot interaction. He has led several national research projects as well as bilateral applied projects with industry. His research interests include the modeling and simulation of robotic systems and advanced robot control. He has received multiple awards for his research in robotics, including the national award for research excellence in robotics.