COMPA: Compensated Holonomic Off-Road Mobile Platform with Attitude Control

Giuseppe Sutera[1], Francesco Cancellier[1], Aniket Datar[2], Cedric Thomas Sebastien Hollande[3], Kartik Virmani[3], Anuj Pokhrel[2], Dario Calogero Guastella[1], Giovanni Muscato[1], Mark Yim[3], and Xuesu Xiao[1]

***Abstract*— We present COMPA (Compensated holonomic Off-road Mobile Platform with Attitude control), an off-road differential-drive with offset turret platform that delivers holonomic control of a stabilized turret on uneven terrain. COMPA uses simple drive wheels on a rocker suspension to keep both wheels in contact with the ground, while an actively stabilized gimbal with three degrees of freedom, i.e., roll, pitch, and yaw, holds the turret at a commanded pose with respect to the world frame. The result is lateral, longitudinal, yaw, roll, and pitch motions that can be commanded independently at the turret, with the robot chassis free to absorb the terrain irregularities. We evaluate the physical COMPA and also in simulation on test courses and quantify its holonomicity using mobility ellipsoid. Experiments demonstrate holonomic turret motion and arbitrary trajectory execution without stop-to-reorient behavior as seen in robots with non-holonomic bases.**

I. Introduction

Off-road robots must maneuver through uneven ground, surface discontinuities, deformable terrain, and varying friction. Their ability to navigate through these difficulties has facilitated their use in agriculture, disaster response, environmental monitoring, and beyond. Because of the ease of design and implementation, many platforms inherit on-road kinematics, such as Ackermann-steering or differential-drive with wheeled or tracked vehicles. Such non-holonomic bases cannot change direction without re-orientation and struggle to follow arbitrary paths which limits their mobility.

Holonomic mobile robots do not have directional kinematic constraints. They command lateral, longitudinal, and yaw motions independently, allowing instantaneous direction changes along an arbitrary trajectory. Holonomic robots shorten local plans, avoid stop-to-reorient maneuvers, and can keep the payload oriented at a constant heading. But they usually require sophisticated wheel design or extra actuators to enable holonomic motion.

Differential-DRive with Offset Turret (DDROT) [1] platforms have no holonomic constraints with a combination of simple drive and caster wheels. DDROT includes a rotating platform which is actuated to provide independent yaw motion. Using the combination of just three actuators, the turret can be fixed in any orientation while the base of the robot rotates to orient itself in any direction. However, it cannot be deployed on off-road terrain due to surface irregularities and large roll and pitch motions.

To this end, we present COMPA (Fig. 1), Compensated holonomic Off-road Mobile Platform with Attitude control.

[1]University of Catania [2]George Mason University [3]University of Pennsylvania

Fig. 1: COMPA demonstrating the ability to keep the turret horizontal (white plate) and directed (blue arrow), independent of the base pose, during holonomic off-road navigation.

It is an open-source, off-road, holonomic DDROT platform, with a stabilized turret mounted on an actuated 3-axis gimbal that holds independently commanded roll, pitch, and yaw angles, while the chassis negotiates through the terrain. COMPA's two adjacent motorized wheels are connected by a rocker suspension, allowing them to conform to the underlying terrain and maintain ground contact for maximum traction. Furthermore, two adjacent caster wheels support the base and rapidly align with the direction of motion.

Measured as the extent to which a vehicle can move equally in all global Degrees of Freedom (DoFs) [1], the holonomicity of COMPA enabled by active turret yaw control along with its ability to compensate for roll and pitch opens up the possibility of applications ranging from remote sensing to outdoor filming, for which robot payload needs to stay stable and directed while the robot base can execute arbitrary trajectories on uneven terrain. We analyze the measured holonomicity using mobility ellipsoid [1] and demonstrate the mobility of the robot in Gazebo simulation and in the real world on uneven terrain by performing comprehensive maneuvers. Our contributions include:

- An open-source design of an off-road DDROT platform with a rocker suspension and an actively stabilized turret that maintains world-fixed attitude on uneven terrain;
- Holonomicity analysis of the agility of the proposed

design using the mobility ellipsoid;
- Simulation and real-world experiments quantifying mobility using measured holonomicity.

II. Related Work

We present related work in terms of off-road mobility and robot chassis design.

A. Off-Road Mobility

Off-road terrain presents risks of uneven surfaces and friction variations that can cause a robot to lose stability or become stuck [2]. Terrain traversability estimation [3]–[5] using features such as elevation, roughness, or semantics provides the information needed for safe motion, while planning algorithms [6] use this information to generate trajectories that avoid high-risk regions. Risk-aware or uncertainty-aware planning approaches [7], [8] account for terrain cost and associated stability, producing safe and reliable trajectories. Approaches based on kinodynamics modeling across terrain classes [9]–[15] and platforms [16], [17] ensure maximum traction and stability during navigation.

Despite these advances, algorithmic design alone cannot eliminate mobility limitations such as non-holonomic constraints. COMPA addresses these challenges with a holonomic DDROT platform, an actively stabilized turret, and a rocker suspension. Together, these enable traversal of arbitrary trajectories through uneven terrain, while keeping the payload stable and preserving wheel-ground contact.

B. Robot Chassis Design

Off-road robot navigation has driven a number of chassis innovations like active suspension [18], [19] and differential lock [12]. These features improve traction and stability but do not add to holonomicity. Quadruped robots [20] offer exceptional mobility on rough terrain, yet their movement involves complex and dynamic gaits, as well as foot placement strategies which lack granularity in motion control and stability. These highly articulated robots also tend to carry less payload and cost much more than their wheeled counterparts.

Classic holonomic designs rely on specialized wheels with rollers like omni wheels, Swedish/Mecanum wheels [21], orthogonal wheels [22], and ball wheels [23]. These classic designs work well on flat, high-friction surfaces, but struggle to generate traction off-road as the terrain traversal is limited by the roller radius. Omni-directional off-road robots solve this issue with multiple wheels oriented in certain configurations that maximize mobility [24], but they still have non-holonomic constraints. Additionally, these designs have higher complexity and cost.

The DDROT concept was introduced to achieve holonomic control at the turret using only two simple drive wheels and a turret rotation actuator. HAMR [1] implements this idea to achieve holonomic turret motion. However, its applicability is largely confined to flat, continuous surfaces. On uneven or deformable terrain, the platform is prone to losing traction or getting stuck. Moreover, because the turret is rigidly coupled to the robot body, its roll and pitch change with the base as the robot traverses uneven surfaces. COMPA addresses these limitations by combining a rocker suspension and a turret mounted on an actively controlled gimbal on a DDROT platform. This design maintains turret stability and holonomic behavior on slopes and uneven surfaces, making the platform suitable for off-road deployment.

III. System Overview

In this section, we present a system overview of COMPA's mechanical and electrical design.

A. Mechanical Design

The mechanical structure, as shown in Fig. 2, is organized into three main subsystems: a differential-drive mobile base with a rocker suspension, an actuated gimbal providing two DoFs for roll and pitch stabilization, and an active turret dedicated to yaw heading control. This layered architecture enables progressive disturbance reduction as the suspension conforms to terrain irregularities, the gimbal compensates residual pitch and roll oscillations, and the turret maintains yaw heading. As a result, the platform can operate in off-road environments while keeping the payload in the desired orientation and heading.

1) Mobile Base: The mobile base is a four-wheel platform, with two front driving wheels and two rear caster wheels, resulting in a non-holonomic differential-drive platform. Its mechanical architecture consists of a central frame connected to two rocker links. Each rocker link carries a wheel pair (drive and caster) on its side and is coupled through a lever, thus forming a rocker suspension with an inter-axle differential. The rocker links are interconnected through passive joints, which facilitate the rotation of the axle that is rigidly connected to the rocker link (Fig. 2 left). This design enables the axle and frame to move independently while linking the rockers and ensuring the platform's stability. The suspension system permits vertical travel equivalent to 80% of the wheel diameter (Fig. 2 bottom right). The current COMPA design focuses exclusively on uneven terrain characterized by moderate irregularities, without considering scenarios involving large obstacles.

The central frame of the platform, illustrated in Fig. 3, has overall dimensions of 450×450 mm and is constructed from four 40×40 mm extruded aluminum bars. Inside the frame, a plate of 300 mm radius is mounted, free to move within the gimbal assembly, enabling active payload stabilization. The remaining mechanical elements are either 3D printed or lightweight aluminum tubes, following a hybrid construction strategy that balances structural strength with weight reduction. The complete platform weighs approximately 20 kg.

2) Pitch and Roll Stabilization: The 3-DoF gimbal is designed to keep the plate horizontal or independently maintain predefined roll, pitch, and yaw angles. Each gimbal axis is actuated by a continuous-rotation servo motor, which inherently lacks direct positional feedback. To achieve precise motion control, a closed-loop strategy is implemented using a nine-axis Witmotion WT901 IMU mounted beneath the

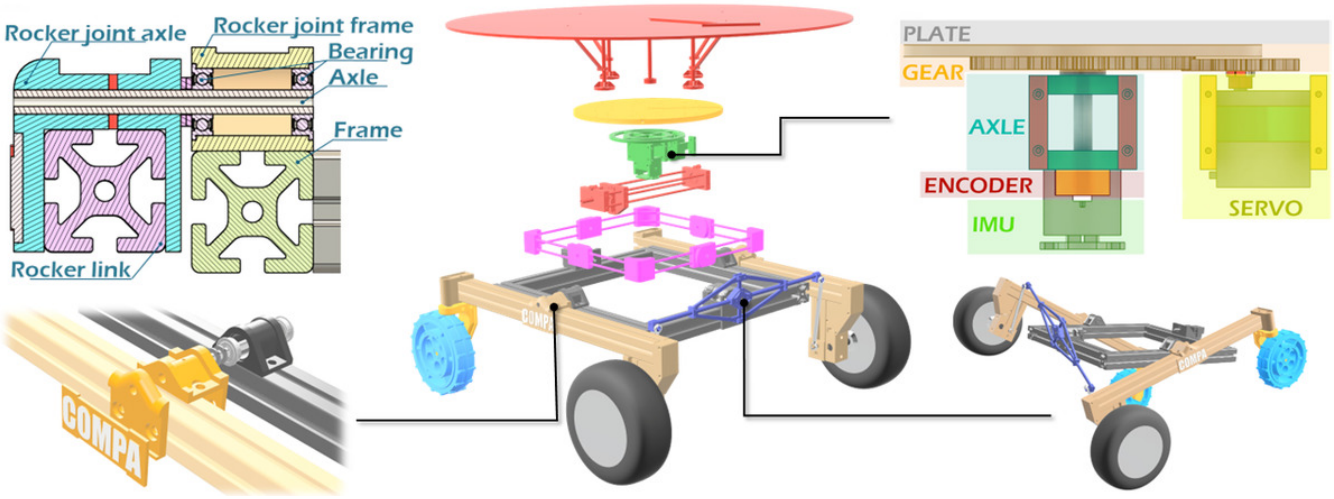

Fig. 2: Composite view of COMPA. Center: the whole assembly, where the two rocker links (beige) are connected through rods by a front-mounted lever (blue), and the main frame (gray) houses the gimbal structure. The gimbal provides three rotational DoFs, with roll (purple), pitch (red), and yaw (green) highlighted. Left: detail of the rocker joint, a sectional view (top), and an exploded 3D view (bottom). Right: the gimbal's layered structure (top) and the range of motion of the rocker suspension system (bottom).

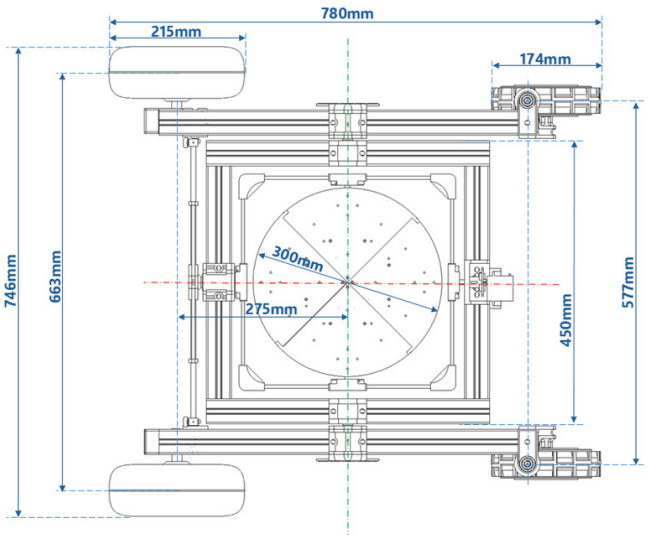

Fig. 3: Platform Dimensions.

gimbal block, which provides real-time orientation data. This architecture allows the gimbal to maintain accurate stabilization within an effective operational range of approximately ±20 degree with a maximum payload of 3 kg.

3) Yaw Heading: The active yaw turret, which is mounted on the pitch axis, plays a fundamental role in the holonomicity of the platform. While the rocker suspension and the gimbal compensate for ground unevenness, the turret ensures that the payload plate maintains a constant heading with respect to the global reference frame, regardless of the robot's instantaneous trajectory during differential-drive maneuvers. The plate is actuated by a continuous-rotation servo, which drives a 1:2 reduction gear, with the larger gear fixed to the plate. The turret is organized into five functional layers: rotating plate, reduction gear, yaw shaft support, absolute encoder, and IMU (Fig. 2 top right). This layered architecture ensures precise alignment of all components along the rotation axis, thereby minimizing eccentricity and reducing mechanical backlash. An AS5600 magnetic absolute encoder monitors angular position, while the nine-axis Witmotion WT901 IMU provides inertial feedback at 100 Hz.

B. Electrical Design

COMPA is powered by two 8000 mAh 4S LiPo batteries in series, which provide a nominal voltage of 33.6V. Two 300W outrunner brushless DC motors are driven by an ODrive V3.6 motor controller each, allowing position, velocity, and torque control.

The continuous servomotors receive power from an 8.4 V UBEC regulator and are controlled through PWM signals generated by an Arduino MEGA 2560, which controls the direction and rotational speed.

The Arduino MEGA 2560 is the central controller, supplying 5 V power to the IMU, the encoder, and the serial telemetry module. It gathers real-time data from the IMU via Serial protocol and the encoder via I2C, transmitting turret orientation and yaw metrics over a wireless serial link to a ROS node operating on an external computer. This ROS node calculates the necessary velocity commands for the wheel drive system and the turret, factoring in the desired motion and the current orientation of the platform before relaying these commands back to the Arduino for execution.

IV. Control System Modeling

We present COMPA's control system modeling using the concept of mobility ellipsoid extended with attitude control.

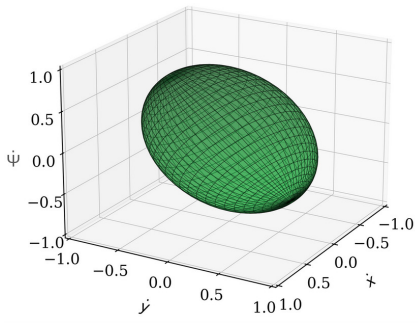

Fig. 4: Projected ME: $\dot{x}$, $\dot{y}$, and $\dot{\psi}$ of COMPA.

A. Mobility Ellipsoid and Holonomicity

Let the actuator rate vector be $\omega \triangleq [\,\omega_r, \omega_\ell, \omega_t\,]^\top$, where ω_r, ω_ℓ are the right and left wheel angular velocities, and ω_t is the turret's relative yaw rate. The commanded twist of the turret reference point is $\mathbf{v} = [\dot{x},\ \dot{y},\ \dot{\psi}]^\top$, with $\dot{x}, \dot{y}$ denoting the translational velocities in the world frame and $\dot{\psi}$ representing the yaw rate. The mapping between actuator space and task space is captured by the Jacobian $\mathbf{J}(\psi) \in \mathbb{R}^{3\times 3}$:

$$\mathbf{v} = \mathbf{J}(\psi)\,\omega.$$

Following Yoshikawa's manipulability ellipsoid for robotic arms [25], HAMR [1] introduced the *mobility ellipsoid* (ME) for wheeled mobile robots as a measure of holonomicity. The idea is simple: if actuator rates are bounded ($\|\omega\| \le 1$), the set of achievable twists forms an ellipsoid in velocity space,

$$\mathbf{v}^\top \left(JJ^\top\right)^{-1} \mathbf{v} \;\le\; 1.$$

The ME is a compact visualization of relative mobility. Shorter ellipse axes indicate harder-to-achieve motion directions. To quantify uniformity the instantaneous *holonomicity* is defined as:

$$\hat{H} \;\triangleq\; \frac{\sigma_{\min}(J)}{\sigma_{\max}(J)}, \tag{1}$$

where $\sigma_{\min}, \sigma_{\max}$ are the minimum and maximum singular values of $\mathbf{J}(\psi)$ [1]. A perfectly spherical ellipsoid yields $\hat{H} = 1$ (isotropic mobility), while values near zero indicate anisotropy or non-holonomic behavior.

B. Mobility Ellipsoid with Attitude Control of COMPA

COMPA extends the holonomicity of the DDROT HAMR system [1] by adding two additional DoFs, roll ϕ and pitch θ, at the turret. These DoFs are not commanded as part of the Cartesian motion task; instead, an inner attitude controller regulates (ϕ, θ) against gravity and disturbances. In general, the inner loop can track any desired attitude within actuator limits. However, for simplicity, we command the level condition $(\phi, \theta) \approx (0, 0)$, which does not alter the outer-loop formation. As a result, the commanded twist remains

$$\mathbf{v}_d \triangleq [\dot{x}_d,\ \dot{y}_d,\ \dot{\psi}_d\,]^\top,$$

while the roll and pitch rates $[\dot{\phi}_d, \dot{\theta}_d]^\top$ are commanded internally for stabilization.

a) Task prioritization view: Conceptually, COMPA executes two nested tasks. The primary task is to regulate roll and pitch back to zero, thereby preserving stability, while the secondary task is to execute the commanded Cartesian twist $\mathbf{v}_d$. Feasible outer-loop velocities are therefore those consistent with the authority left over after the inner loop. We denote the resulting effective ME as

$$\mathscr{E}_{\text{eff}} \;=\; \underbrace{\{\mathbf{v} : \mathbf{v}^\top (JJ^\top)^{-1}\mathbf{v} \le 1\}}_{\text{kinematic ME from Jacobian } J(\psi)} \;\cap\; \underbrace{\{\mathbf{v} : \mathscr{C}(\mathbf{v}) \le \mathbf{u}_{\max}\}}_{\text{attitude constraints}}, \tag{2}$$

where $\mathscr{C}(\mathbf{v})$ represents the inner-loop control effort required to reject gravitational and inertial disturbances for a given twist $\mathbf{v}$, and $\mathbf{u}_{\max}$ denotes the maximum available authority of the attitude actuators.

b) Flat ground, small angles: On level terrain with small commanded accelerations and a high-bandwidth inner loop, the control effort $\mathscr{C}(\mathbf{v})$ is much smaller than the available authority $\mathbf{u}_{\max}$, and the effective ellipsoid $\mathscr{E}_{\text{eff}}$ reduces to HAMR's kinematic ME. In this regime, the commanded ME is unchanged (Fig. 4); the benefit of the inner loop is improved disturbance rejection and decoupling.

c) Slopes and aggressive maneuvers: On uneven off-road terrain or during aggressive maneuvers, mobility becomes reduced in two ways. A portion of the turret's actuation is dedicated to holding roll and pitch, which reduces the residual authority available for yaw and lateral motions and shrinks the ME along coupled directions. In addition, tilting of the body shifts normal loads across the wheels, tightening slip constraints and causing translational mobility to contract anisotropically. As a result, the effective ellipsoid $\mathscr{E}_{\text{eff}}$ becomes state and terrain dependent.

d) Projected MEs: To make these effects clearer, we separate the projections of the ellipsoid into translational and yaw components. The translational ME is obtained as

$$\mathbf{v}_{\text{xy}}^\top \left(S_{\text{xy}} JJ^\top S_{\text{xy}}^\top\right)^{-1} \mathbf{v}_{\text{xy}} \le 1, \qquad S_{\text{xy}} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix},$$

where S_{xy} projects the commanded twist $\mathbf{v}$ into its translational components $\mathbf{v}_{\text{xy}} = [\dot{x}, \dot{y}]^\top$. The yaw ME is expressed as $|\dot{\psi}| \le \sigma_\psi(J)$, where $\sigma_\psi(J)$ is the principal radius of the ellipsoid along the yaw direction obtained from the singular values of J. These projections highlight three key behaviors: COMPA retains the anisotropy of HAMR, yaw capability contracts first when inner-loop authority is taxed, and on flat ground with ample authority, the projected MEs are identical to those of HAMR.

e) Holonomicity under constraints: The standard holonomicity metric for the commanded task is given by Eqn. (1). When attitude constraints are active, we define an effective holonomicity,

$$\hat{H}_{\text{eff}} \;\triangleq\; \frac{\sigma_{\min}(\mathscr{E}_{\text{eff}})}{\sigma_{\max}(\mathscr{E}_{\text{eff}})},$$

computed from the contracted ellipsoid $\mathscr{E}_{\text{eff}}$. On flat ground, $\hat{H}_{\text{eff}}$ remains close to the theoretical value $\hat{H}$, while on uneven off-road terrain or during aggressive maneuvers, $\hat{H}_{\text{eff}}$ decreases.

C. Platform Geometry and Notation

For our differential-drive base with right/left wheel angular velocities $\omega_r,\ \omega_\ell$ and wheel radius r_w, the wheel contact points are separated by $2a$ (i.e., a is the half-track). The yawing turret is mounted on the base at a body-frame offset $\mathbf{p}_b = [b,\ 0]^\top$ from the axle center and can rotate relative to the base with rate ω_t. For world-frame base yaw ψ, denote $c \triangleq \cos\psi$ and $s \triangleq \sin\psi$. As defined before, the commanded Cartesian velocity of the turret reference point expressed in the world frame is $\mathbf{v}_d \triangleq [\dot{x}_d,\ \dot{y}_d,\ \dot{\psi}_d]^\top$. In the base frame, standard differential-drive kinematics give the forward speed and yaw rate as

$$v_b \;=\; \frac{r_w}{2}\,(\omega_r+\omega_\ell), \qquad \dot{\psi}_{\text{base}} \;=\; \frac{r_w}{2a}\,(\omega_r-\omega_\ell). \tag{3}$$

Eqn. (3) represent the standard differential-drive kinematics under the assumptions of flat terrain and no slip. On uneven terrain, these relations remain approximately valid as long as roll and pitch are small and wheel-ground contact is preserved.

D. Velocity of the Turret Reference Point

The linear velocity at the turret reference point combines the base translational velocity with the contribution from the base's angular velocity about the axle center:

$$\mathbf{v}^{(b)}_{\text{turret}} \;=\; \begin{bmatrix} v_b \\ 0 \end{bmatrix} \;+\; \dot{\psi}_{\text{base}} \begin{bmatrix} 0 \\ b \end{bmatrix} \;=\; \begin{bmatrix} v_b \\ b\,\dot{\psi}_{\text{base}} \end{bmatrix}, \tag{4}$$

where $(\cdot)^{(b)}$ denotes the base (body) frame. Transforming to the world frame using $R(\psi) = \begin{bmatrix} c & -s \\ s & c \end{bmatrix}$ yields

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} \;=\; R(\psi) \begin{bmatrix} v_b \\ b\,\dot{\psi}_{\text{base}} \end{bmatrix} \;=\; \begin{bmatrix} c\,v_b - s\,b\,\dot{\psi}_{\text{base}} \\ s\,v_b + c\,b\,\dot{\psi}_{\text{base}} \end{bmatrix}. \tag{5}$$

The *world* yaw rate of the turret frame is the sum of the base yaw rate and the turret's relative yaw rate:

$$\dot{\psi} \;=\; \dot{\psi}_{\text{base}} + \omega_t. \tag{6}$$

E. Kinematic Jacobian

Substituting Eqn. (3) into Eqns. (5)–(6) and collecting terms in $(\omega_r,\ \omega_\ell,\ \omega_t)$, we obtain the square Jacobian $\mathbf{J}(\psi) \in \mathbb{R}^{3\times 3}$ that maps joint rates to the commanded turret-world twist:

$$\underbrace{\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\psi} \end{bmatrix}}_{\mathbf{v}} \;=\; \underbrace{\begin{bmatrix} \frac{r_w}{2}\left(c - s\,\frac{b}{a}\right) & \frac{r_w}{2}\left(c + s\,\frac{b}{a}\right) & 0 \\ \frac{r_w}{2}\left(s + c\,\frac{b}{a}\right) & \frac{r_w}{2}\left(s - c\,\frac{b}{a}\right) & 0 \\ \frac{r_w}{2a} & -\frac{r_w}{2a} & 1 \end{bmatrix}}_{\mathbf{J}(\psi)} \underbrace{\begin{bmatrix} \omega_r \\ \omega_\ell \\ \omega_t \end{bmatrix}}_{\omega}. \tag{7}$$

Hence, for a feasible commanded twist $\mathbf{v}_d$ away from singular geometries ($a \neq 0$), the joint rates are obtained by solving

$$\omega \;=\; \mathbf{J}(\psi)^{-1}\,\mathbf{v}_d, \tag{8}$$

which is implemented via a linear solve. The b/a terms capture the additional translational velocity at the turret due to the base yaw rate acting on the offset $\mathbf{p}_b = [b,0]^\top$. When $b = 0$ (turret at the axle center), Eqn. (7) reduces to the familiar mapping for a differential-drive point at the axle center, while the third row remains $\dot{\psi} = \frac{r_w}{2a}(\omega_r - \omega_\ell) + \omega_t$.

F. Outer-loop Velocity Control and Implementation

A Cartesian PID regulator computes the desired turret-world twist $\mathbf{v}_d = [\dot{x}_d,\ \dot{y}_d,\ \dot{\psi}_d]^\top$ from pose errors $\mathbf{e} = [e_x,\ e_y,\ e_\psi]^\top$, where e_x, e_y are position errors in the world x–y directions and e_ψ is the yaw error wrapped to $(-\pi, \pi]$. Optional feedforward terms $\mathbf{v}_{\text{ff}} = [\dot{x}_{\text{ff}},\ \dot{y}_{\text{ff}},\ \dot{\psi}_{\text{ff}}]^\top$ can be added. Using proportional, integral, and derivative gains collected in diagonal matrices K_P, K_I, K_D, the control law is expressed compactly as

$$\mathbf{v}_d \;=\; \mathbf{v}_{\text{ff}} + K_P\,\mathbf{e} + K_I \int \mathbf{e}\,dt + K_D\,\tilde{\dot{\mathbf{e}}}, \tag{9}$$

where $\tilde{\dot{\mathbf{e}}}$ denotes a first-order low-pass filtered derivative. The integral terms are clamped to prevent windup.

The commanded linear speed $\sqrt{\dot{x}_d^2 + \dot{y}_d^2}$ and angular speed $\dot{\psi}_d$ are saturated to hardware limits before inversion via Eqn. (8). Pose errors are computed from odometry, turret-to-base quaternion, and desired setpoints. The Jacobian (7) is evaluated at the current base yaw ψ and solved each cycle at 100 Hz to obtain actuator rates $(\omega_r, \omega_\ell, \omega_t)$, which are then sent to the corresponding motors. This same control framework is applied in both physical and simulated experiments, allowing us to evaluate holonomicity consistently across hardware and its digital twin.

V. Experiments

To evaluate holonomicity, a suite of tests was conducted in both simulated and real-world environments. The tests involved the robot following three distinct paths: a square, a triangle, and a circle. These paths were chosen because they effectively represent the robot's mobility performance: the sharp turns required for the square and triangle paths are particularly useful for evaluating its holonomic behavior; the curvature of the circle path also causes challenge to keep constant heading. In addition to these experiments, the platform was also tested in a complex maze-like environment in simulation and was qualitatively assessed in an outdoor, real-world setting. The same Jacobian-based controller (Sec. IV) was deployed in both simulation and the real world.

A. Simulation Experiments

To evaluate holonomicity in simulation, we constructed a digital twin of COMPA in Gazebo. A URDF was built from the CAD model, and 2.5D heightmaps were generated from grayscale images to replicate uneven terrain in a reproducible environment.

The three geometrical paths (square, triangle, and circle trajectories) were commanded on generated uneven terrains (Fig. 5). COMPA successfully tracked all paths while maintaining effective holonomicity: lateral and rotational motions remained feasible in every direction.

To assess holonomicity in more challenging scenarios, we integrated PRM [26] and A* planners to execute maze-like trajectories across uneven terrain (Fig. 6). With both

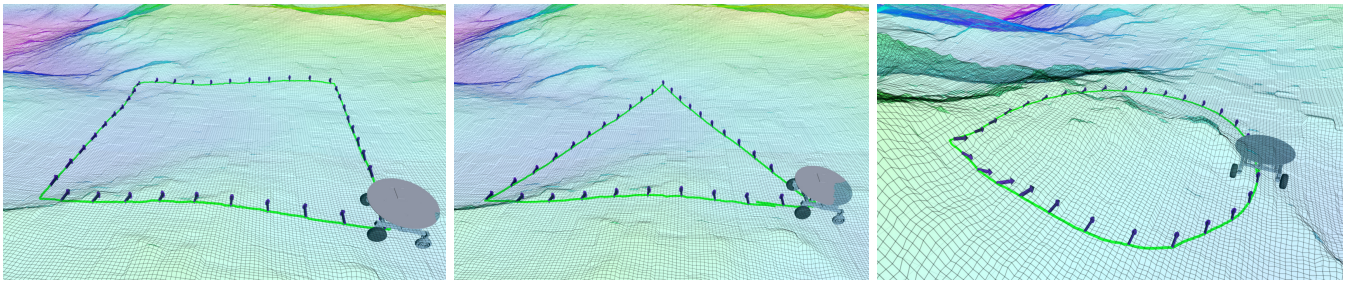

Fig. 5: Canonical holonomicity tests in simulation: square, triangle, and circle trajectories. Green lines denote the executed trajectories, while blue arrows show the heading angle.

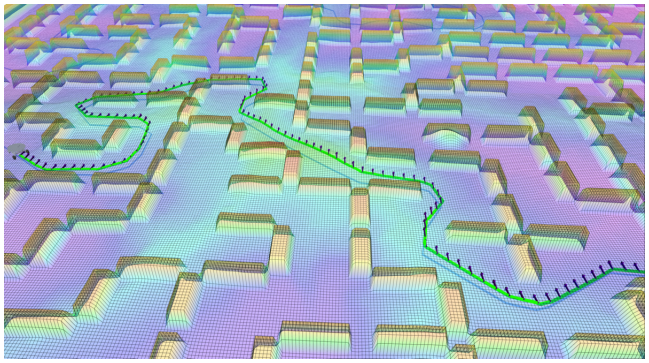

Fig. 6: Maze navigation in simulation with PRM/A*.

planners, COMPA preserved mobility in all commanded directions.

TABLE I: Simulation Experiments Results.

Path	**Path RMSE** (*m*)
Square	0.040
Triangle	0.041
Circle	0.022
Maze	**0.018**

The performance in simulation was evaluated with the Root Mean Square Error (RMSE), comparing the reference path to the robot's true path. Between the canonical paths and the maze, COMPA achieved very low path-following error in simulation (Table I). The best results were seen in the maze tests, where the robot was able to track its path perfectly due to the long, straight sections. Roll, pitch, and yaw deviations were omitted for simulation since they remain negligible in our digital twin. Across all scenarios of off-road terrain, COMPA maintained holonomicity consistent with the mobility ellipse framework. In maze navigation, holonomicity was preserved despite frequent movement direction changes.

B. Physical Experiments

For real-world validation, the tests were conducted in a controlled environment where a motion capture system was used to provide ground truth for the robot's position and orientation. To test the robot's ability to navigate uneven terrain, two sets of modular floor tiles were laid out along the paths. The first set formed a tall but mostly smooth surface, while the second set was made of stepped tiles that the robot had to climb over.

As shown in Fig. 7, the robot's trajectory is plotted against the reference paths. While RMSE is used to evaluate path-following performance, to assess the robot's stability, the standard deviation of the error for the roll, pitch, and yaw angles was evaluated. This approach is preferred because, while the average error for these angles is negligible, the standard deviation effectively measures the extent of the turret's motion, providing a more meaningful representation of stability.

Based on the results reported in Table II, the robot's performance is highly dependent on the path's shape. The circle path demonstrates the best performance, as indicated by the lowest errors in both path-following and yaw. This suggests that the robot handles smooth, continuous movement with high accuracy and stability. In contrast, the more complex square and triangle paths lead to a noticeable increase in errors for both path-following and yaw, with the triangle path showing the largest deviations. This indicates that sharp turns are the primary factor affecting the robot's heading and accuracy. Despite this, the consistently low errors in roll and pitch across tests confirm that the robot maintains stability and balance, regardless of the path. The RMSE values are higher than the simulation results (Table I), reflecting near-perfect execution in simulation where actuator imperfections and slip are absent.

TABLE II: Physical Experiments Results.

Path	**Path RMSE** (*m*)	**Roll Std.** (*rad*)	**Pitch Std.** (*rad*)	**Yaw Std.** (*rad*)
Square	0.1776	**0.0439**	**0.0404**	0.1067
Triangle	0.1973	0.0447	0.0441	0.1597
Circle	**0.1005**	0.0457	0.0438	**0.0501**

C. Outdoor Demonstration

COMPA was evaluated in a qualitative, outdoor setting to demonstrate its real-world performance on uneven terrain. The goal was to supplement the quantitative data from

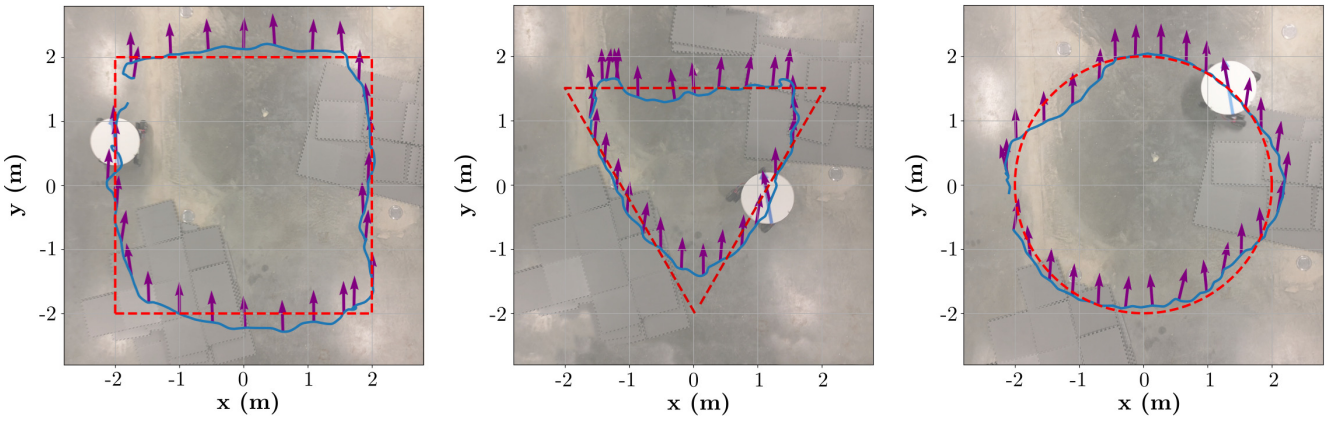

Fig. 7: Physical experiments of the square (a), triangular (b), and circular (c) paths with the reference path in red, the real path in blue, and purple arrows represent the turret's yaw at the corresponding position.

Fig. 8: COMPA during outdoor evaluation on a hill.

controlled experiments and showcase the robot's ability to navigate an unstructured environment.

The evaluation was conducted on a sloped, uneven surface with varying ground textures and multiple rocks and obstacles, in order to evaluate its holonomic control in a real-world setting, the effectiveness of the rocker suspension, and the active stabilization provided by the gimbal. During the evaluation, COMPA successfully navigated up and down a hill, showcasing its traction and stability on a slope as the rocker suspension system allowed the wheels to conform to the terrain. As shown in Fig. 8, the actively stabilized gimbal and turret correctly kept the payload level with the horizon and pointed in a constant heading.

While the rocker suspension effectively managed moderate terrain irregularities, a key limitation was identified concerning the rear caster wheels. Since the ability of the robot to traverse distinct steps and uneven rock formations is highly dependent on the caster angle at the moment of impact, when the caster wheel's rotational axis is not perpendicular to the obstacle, the wheel is unable to roll over the obstruction and gets stuck. One potential area for improvement is to include the caster angle as a parameter in the path planning algorithm. This would allow the robot to anticipate and adjust its trajectory or speed to ensure the casters are at a favorable angle when encountering obstacles, thus improving its overall performance and reliability on uneven terrain.

VI. Conclusion

In this paper, we introduce COMPA, a holonomic off-road DDROT platform with a rocker suspension and an actively stabilized gimbal designed to maintain payload stability on uneven terrain. This architecture successfully decouples the turret's attitude from the chassis, enabling the platform to execute arbitrary trajectories while the gimbal compensates for roll and pitch disturbances and the turret maintains its commanded heading. Our analysis using mobility ellipsoid reveals that COMPA's effective holonomicity is state dependent, contracting on slopes as actuation is diverted for attitude control. Extensive experiments in both simulation and the real world validate COMPA's holonomicity by demonstrating its ability to follow arbitrary trajectories, all the while maintaining a stable payload with minimal roll, pitch, and yaw tracking error.

References

[1] J. T. Costa and M. Yim, "Designing for uniform mobility using holonomicity," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 2448–2453.

[2] C. Min, S. Si, X. Wang, H. Xue, W. Jiang, Y. Liu, J. Wang, Q. Zhu, Q. Zhu, L. Luo *et al.*, "Autonomous driving in unstructured environments: How far have we come?" *arXiv preprint arXiv:2410.07701*, 2024.

[3] J. Seo, S. Sim, and I. Shim, "Learning off-road terrain traversability with self-supervisions only," *IEEE Robotics and Automation Letters*, vol. 8, no. 8, pp. 4617–4624, 2023.

[4] J. Frey, M. Patel, D. Atha, J. Nubert, D. Fan, A. Agha, C. Padgett, P. Spieler, M. Hutter, and S. Khattak, "Roadrunner-learning traversability estimation for autonomous off-road driving," *IEEE Transactions on Field Robotics*, 2024.

[5] C. Pan, A. Datar, A. Pokhrel, M. Choulas, M. Nazeri, and X. Xiao, "Traverse the Non-Traversable: Estimating Traversability for Wheeled Mobility on Vertically Challenging Terrain," Sep. 2024.

[6] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.

[7] H. Karnan, E. Yang, G. Warnell, J. Biswas, and P. Stone, "Wait, that feels familiar: Learning to extrapolate human preferences for preference-aligned path planning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 13 008–13 014.

[8] X. Cai, J. Queeney, T. Xu, A. Datar, C. Pan, M. Miller, A. Flather, P. R. Osteen, N. Roy, X. Xiao *et al.*, "Pietra: Physics-informed evidential learning for traversing out-of-distribution terrain," *IEEE Robotics and Automation Letters*, 2025.

[9] X. Xiao, J. Biswas, and P. Stone, "Learning inverse kinodynamics for accurate high-speed off-road navigation on unstructured terrain," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 6054–6060, 2021.

[10] H. Karnan, K. S. Sikand, P. Atreya, S. Rabiee, X. Xiao, G. Warnell, P. Stone, and J. Biswas, "VI-IKD: High-speed accurate off-road navigation using learned visual-inertial inverse kinodynamics," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 3294–3301.

[11] A. Pokhrel, A. Datar, M. Nazeri, and X. Xiao, "CAHSOR: Competence-aware high-speed off-road ground navigation in SE (3)," *IEEE Robotics and Automation Letters*, vol. 9, no. 11, pp. 9653–9660, 2024.

[12] A. Datar, C. Pan, M. Nazeri, and X. Xiao, "Toward Wheeled Mobility on Vertically Challenging Terrain: Platforms, Datasets, and Algorithms," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 16 322–16 329.

[13] A. Datar, C. Pan, and X. Xiao, "Learning to model and plan for wheeled mobility on vertically challenging terrain," *IEEE Robotics and Automation Letters*, vol. 10, no. 2, pp. 1505–1512, 2025.

[14] A. Datar, C. Pan, M. Nazeri, A. Pokhrel, and X. Xiao, "Terrain-Attentive Learning for Efficient 6-DoF Kinodynamic Modeling on Vertically Challenging Terrain," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Abu Dhabi, United Arab Emirates: IEEE, Oct. 2024, pp. 5438–5443.

[15] M. Nazeri, A. Datar, A. Pokhrel, C. Pan, G. Warnell, and X. Xiao, "Verticoder: Self-supervised kinodynamic representation learning on vertically challenging terrain," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 6536–6543.

[16] P. Roth, J. Frey, C. Cadena, and M. Hutter, "Learned Perceptive Forward Dynamics Model for Safe and Platform-aware Robotic Navigation," in *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2025.

[17] T. Xu, C. Pan, M. B. Rao, A. Datar, A. Pokhrel, Y. Lu, and X. Xiao, "Verti-bench: A general and scalable off-road mobility benchmark for vertically challenging terrain," in *Robotics: Science and Systems (RSS) 2025*, 2025.

[18] S. Jiang, Z. Li, S. Lin, W. Shi, Z. Zhu, H. Che, S. Yin, C. Zhang, and Z. Jia, "Design, control and experiments of an agile omnidirectional mobile robot with active suspension," in *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, 2022, pp. 913–918.

[19] S. Yin, W. Wang, L. Hu, Z. Zhu, and Z. Jia, "Attitude control of an all-wheel-drive rover with integrated active suspension system," in *2023 9th International Conference on Mechatronics and Robotics Engineering (ICMRE)*. IEEE, 2023, pp. 58–64.

[20] S. Kim and P. M. Wensing, "Design of dynamic legged robots," *Foundations and Trends® in Robotics*, vol. 5, no. 2, pp. 117–190, 2017.

[21] L. Tagliavini, G. Colucci, A. Botta, P. Cavallone, L. Baglieri, and G. Quaglia, "Wheeled mobile robots: state of the art overview and kinematic comparison among three omnidirectional locomotion strategies," *Journal of intelligent & robotic systems*, vol. 106, no. 3, p. 57, 2022.

[22] G. Mourioux, C. Novales, G. Poisson, and P. Vieyres, "Omni-directional robot with spherical orthogonal wheels: concepts and analyses," in *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.* IEEE, 2006, pp. 3374–3379.

[23] K. Tadakuma, R. Tadakuma, and J. Berengeres, "Development of holonomic omnidirectional vehicle with "omni-ball": spherical wheels," in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007, pp. 33–39.

[24] G. Ishigami, K. Iagnemma, J. Overholt, and G. Hudas, "Design, development, and mobility evaluation of an omnidirectional mobile robot for rough terrain," *Journal of Field Robotics*, vol. 32, no. 6, pp. 880–896, 2015.

[25] T. Yoshikawa, "Manipulability of robotic mechanisms," *The international journal of Robotics Research*, vol. 4, no. 2, pp. 3–9, 1985.

[26] L. E. Kavraki, M. N. Kolountzakis, and J.-C. Latombe, "Analysis of probabilistic roadmaps for path planning," *IEEE Transactions on Robotics and automation*, vol. 14, no. 1, pp. 166–171, 1998.