

TAPNAV: Humanoid Navigation through Tactile Active Perception

Anonymous Author(s)



Fig. 1: TAPNAV enables vision-denied humanoid navigation through tactile active perception. The humanoid actively makes contact with the surrounding environment at informative locations, using contact observations to reduce localization uncertainty and achieve reliable navigation.

Abstract—Navigation in vision-denied environments is challenging for humanoid robots because proprioceptive odometry drifts and localization uncertainty accumulates rapidly. We present TAPNAV, a tactile active-perception framework that enables humanoid navigation toward a goal by actively probing surrounding structures without relying on vision. TAPNAV maintains a pose belief from odometry, IMU, and tactile contact observations, and couples uncertainty-aware global route planning with information-gain-driven local probing. The global planner searches for routes that keep predicted localization uncertainty bounded by exploiting opportunities for tactile correction, while the local planner selects probe actions that maximize expected information gain. A whole-body controller coordinates the humanoid’s locomotion and end-effector contact to execute the planned navigation and probe motions. We evaluate TAPNAV in simulation and on a Unitree G1 across different floor plans and obstacle geometries. TAPNAV achieves lower state estimation error and a higher task completion rate than baselines. These results demonstrate that actively planning physical interactions with the environment can provide localization cues for reliable humanoid navigation without vision. More results are available at TAPNAV project page.

I. INTRODUCTION

Humanoid robots hold great promise for bringing intelligent physical assistance to human-centered environments [1]. To succeed in indoor environments, humanoid robots require not only robust locomotion but also reliable localization and navigation under safety-critical and perceptually degraded conditions. This capability is particularly important in hazardous or inaccessible scenarios such as fire emergency

response, industrial gas-leak inspection, and engine-room operation with degraded lighting.

Recent progress in humanoid perception and navigation has largely relied on RGB-D cameras and LiDAR for environment sensing and state estimation [2], [3]. These sensing modalities, however, can become unreliable in visually degraded environments, compromising both localization and navigation capabilities. In contrast, humans can compensate for limited vision through haptic interactions, which offer an alternative source of geometric information when visual observations are unreliable.

Motivated by how people navigate with haptic feedback, this study equips a humanoid robot with tactile sensors for contact-based navigation. Traditionally, robotic tactile sensing has been widely researched in fine-grained tasks such as identifying grasped objects, estimating in-hand pose, or reconstructing local shape [4], [5]. Extending tactile sensing to humanoid navigation requires selecting informative contacts while maintaining whole-body balance. Sparse tactile observations must then be integrated with motion estimates to support localization throughout navigation.

We introduce TAPNAV: Humanoid Navigation through Tactile Active Perception, as shown in Fig. 1. TAPNAV is an integrated framework that guides a humanoid to a goal using tactile sensing without relying on cameras or LiDAR. To the best of our knowledge, TAPNAV is the first system to use actively planned tactile probing, rather than reactive or incidental contact, to navigate a humanoid through an office-

scale indoor environment without vision.

We evaluate TAPNAV in simulation and on a Unitree G1 humanoid, covering indoor layouts with long hallways, corners, pillars, and free-standing islands. Our contributions are summarized as follows:

- A belief-space tactile navigation planner that balances route progress with localization uncertainty by exploiting opportunities for tactile correction.
- An active tactile probing framework that selects contact actions based on their expected information gain, accounting for both contact and contact-free outcomes to actively reduce robot pose uncertainty.
- An integrated vision-denied humanoid navigation system that couples belief-space planning and active tactile perception with reinforcement learning-based whole-body locomotion control, validated in both simulation and real-world hardware experiments in office-scale environments.

II. RELATED WORK

A. Tactile Localization and Active Sensing

Tactile sensing has been widely applied to object-centric state estimation. Learned tactile representations and probabilistic estimation account for ambiguous contact observations, while additional contacts and sequential measurements constrain the set of plausible poses [6], [7]. Joint pose and shape estimation has also been studied through visuo-tactile neural reconstruction and tactile-only geometric modeling, allowing object geometry to be estimated alongside its motion [8]. Active sensing couples these estimation problems with the selection of subsequent interactions. Learned exploration policies use tactile feedback to select motions that improve surface coverage or object recognition [9], [10]. Information-based approaches evaluate candidate interactions through expected information gain or contact-aware information, selecting informative motions for estimating object pose, geometry, or physical parameters [8], [11]. On the other hand, contact sensing has also been used for robot localization, where geometric and terrain information from gait-induced contacts is incorporated without active action selection [12].

B. Belief-Space Planning and Active Localization

Belief-space planning explicitly reasons about state uncertainty when selecting navigation actions, accounting for how motion and future observations affect the robot’s belief [13]. Active localization further couples navigation with sensing decisions, selecting informative viewpoints or deciding when to acquire observations to improve pose estimation [14], [15]. These approaches, however, primarily assume that localization information can be obtained through remote perceptual observations. In contrast, TAPNAV enables active localization in vision-denied settings by planning motions that leverage contact observations to maintain an accurate robot-state belief.

C. Contact-Rich Humanoid Loco-Manipulation

Recent advances in humanoid loco-manipulation have enabled whole-body coordination of locomotion and manipulation for tasks such as object transport and door opening [16], [17], dynamic catching [18], payload transport [19], [20], and force-regulated interaction [21]. These methods demonstrate increasingly capable coordination between lower-body locomotion and upper-body manipulation, including maintaining stability under interaction forces. Unlike conventional loco-manipulation, where contact typically serves the manipulation objective, TAPNAV uses contact as a sensing action: the planner dynamically selects when and where the robot should probe to reduce localization uncertainty.

III. SYSTEM OVERVIEW

We formulate the problem as a vision-denied humanoid navigation task on a known and static map M . The robot starts from a known initial state $\mathbf{x}_0$, its state at time t is represented by $\mathbf{x}_t = (p_t^x, p_t^y, \psi_t) \in \text{SE}(2)$. We localize the robot using odometry, an IMU-based heading estimate, and tactile contact with mapped surface. The goal is to reach a target region $\mathcal{X}_g$ while balancing route progress with active tactile probing to control localization uncertainty.

During navigation, the system coordinates a global route planner, a local tactile action planner, and a pose estimator, as shown in Fig. 2. At timestep t , given the map M and the current state belief b_t , the global planner computes a collision-free route to $\mathcal{X}_g$ that keeps the predicted pose uncertainty bounded. The route is represented by a sequence of subgoals, each specifying a desired displacement $\Delta \mathbf{x}_t$ from the preceding subgoal.

At each subgoal, the local planner uses the latest belief to select a tactile probe that maximizes the expected information gain (EIG) about the robot pose. A probe moves the end-effector j along a straight trajectory in a direction α relative to the torso heading until contact is detected or a travel limit is reached. Each probe action a is defined as $a = (j, \alpha) \in \mathcal{A}$ where $j \in \{L, R\}$ specifies the left or right end-effector, $\mathcal{A} = \{L, R\} \times \mathcal{A}_\alpha$, and $\mathcal{A}_\alpha$ is a uniformly discretized set of probing directions in $[-40^\circ, 110^\circ]$, subject to the arm’s kinematic feasibility. For both end-effectors, $\alpha = 0$ points straight forward and $\alpha > 0$ points outward.

The probe action a_t selected by the local planner produces a tactile observation $\mathbf{o}_t^{\text{tac}} = (\mathbf{s}_t, c_t)$, where $\mathbf{s}_t \in \mathbb{R}^2$ is the end-effector stopping position in the body frame, and $c_t \in \{0, 1\}$ is the contact outcome. $c_t = 1$ indicates contact and $c_t = 0$ indicates a contact-free event. The estimator maintains a pose belief $b_t(\mathbf{x}_t) = p(\mathbf{x}_t \mid \mathbf{u}_{1:t}, \mathbf{o}_{1:t}, M)$, where $\mathbf{u}_t = (\Delta p_t^x, \Delta p_t^y, \Delta \psi_t) \in \text{SE}(2)$ is the body-frame odometry increment $\mathbf{x}_t$ from step $t - 1$ to t , and $\mathbf{o}_t = (\mathbf{o}_t^{\text{tac}}, \tilde{\psi}_t)$ collects the observations, where $\tilde{\psi}_t$ denotes the IMU heading measurement. Finally, a low-level controller executes the displacement $\Delta \mathbf{x}_t$ and the probing action a_t . In summary, TAPNAV actively uses touch to reduce localization uncertainty, enabling reliable goal-directed navigation without vision.




Fig. 2: Overview of the TAPNAV framework. Odometry, IMU, and tactile observations together update the robot’s SE(2) pose belief. The global planner computes an uncertainty-bounded route with predicted tactile corrections, the local planner selects informative probes based on expected information gain, and the low-level controller executes locomotion and probing.

IV. METHOD

A. Tactile Sensing and Contact Processing




Fig. 3: Tactile end-effector design and contact-processing pipeline. The tactile image is aggregated into a total activation $T(t)$, which is calibrated to a normal force and converted to a binary contact outcome using a force threshold.

The Unitree G1 is equipped with custom tactile end-effectors to detect contact with the environment, as shown in Fig. 3. Each end-effector integrates a TouchTronix SensX 160 tactile sensor covered by a silicone pad made from Ecoflex 00-50 [22]. The SensX 160 consists of a 20×8 taxel array with an active sensing area of 20×50 mm and a force range of up to 100 N. The sensor is mounted on a curved, rigid housing to enable contact sensing across a range of approach angles.

The tactile image is first offset-corrected and filtered to suppress sensor noise. We define the total activation $T(t)$ as the sum of the filtered taxel activations at time t . The total activation is then converted to an estimated normal force $\hat{f}(t)$ using the inverse exponential response model introduced in WT-UMI [21]. Finally, the estimated normal force is thresholded to produce a binary contact outcome,

which is provided to the planner as a tactile observation.

B. Belief-Space Tactile Navigation Planning

1) *Pose belief:* At each time step t , we maintain a pose belief b_t over the robot’s current pose $\mathbf{x}_t$ using a particle filter, which is then acted on by the global and local planners. During locomotion, odometry propagates the pose hypotheses, while tactile and IMU observations correct the accumulated drift. The belief is represented by N_p weighted particles $\{\mathbf{x}_t^{(1)}, \dots, \mathbf{x}_t^{(i)}, \dots, \mathbf{x}_t^{(N_p)}\}$,

$$b_t(\mathbf{x}_t) = p(\mathbf{x}_t | \mathbf{u}_{1:t}, \mathbf{o}_{1:t}, M) \approx \sum_{i=1}^{N_p} w_t^{(i)} \delta(\mathbf{x}_t - \mathbf{x}_t^{(i)}),$$

where $\mathbf{x}_t^{(i)} = (p_t^{x(i)}, p_t^{y(i)}, \psi_t^{(i)}) \in \text{SE}(2)$ is the i -th pose hypothesis, and $w_t^{(i)}$ is its normalized weight with $\sum_{i=1}^{N_p} w_t^{(i)} = 1$. The particles are initialized around the known initial pose. For computational efficiency, we approximate the belief as a Gaussian $b_t(\mathbf{x}_t) \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ whose mean $\boldsymbol{\mu}_t = (\mu_t^x, \mu_t^y, \mu_t^\psi) \in \mathbb{R}^3$ and covariance $\boldsymbol{\Sigma}_t \in \mathbb{R}^{3 \times 3}$ are computed from the weighted particles.

Given the body-frame increment $\mathbf{u}_t$ from odometry, each particle i is propagated by composing $\mathbf{x}_{t-1}^{(i)}$ with $\mathbf{u}_t$ perturbed by Gaussian noise:

$$\mathbf{x}_t^{(i)} = \mathbf{x}_{t-1}^{(i)} \oplus (\mathbf{u}_t + \boldsymbol{\varepsilon}_t^{(i)}), \quad \boldsymbol{\varepsilon}_t^{(i)} \sim \mathcal{N}(\mathbf{0}, Q(\mathbf{u}_t)), \quad (1)$$

where $\oplus$ denotes SE(2) pose composition and $Q(\mathbf{u}_t)$ models odometry noise. After propagation, each particle is reweighted using tactile and heading likelihoods based on




Fig. 4: Tactile likelihood. (a) The residual $\delta_t^{(i)}$ under hypothesis i and the nearest point on the map surface $\mathbf{c}_t^a$ defined in Sec. IV-C; (b) Under hypothesis 1, the end-effector would penetrate the surface ($\delta_t^{(1)} < 0$); (c) Under hypothesis 2, it would not reach the surface ($\delta_t^{(2)} > 0$).

the tactile probe and measured IMU heading as

$$\begin{aligned} w_t^{(i)} &\propto w_{t-1}^{(i)} p(\mathbf{o}_t | \mathbf{x}_t^{(i)}, M) \\ &= w_{t-1}^{(i)} p(\mathbf{o}_t^{\text{tac}} | \mathbf{x}_t^{(i)}, M) p(\tilde{\psi}_t | \psi_t^{(i)}). \end{aligned} \quad (2)$$

The tactile likelihood $p(\mathbf{o}_t^{\text{tac}} | \mathbf{x}_t^{(i)}, M)$, described in Sec. IV-B.2, evaluates the consistency between the probe observation and the known map. The heading likelihood $p(\tilde{\psi}_t | \psi_t^{(i)})$ is a wrapped Gaussian likelihood for $\tilde{\psi}_t$ with mean $\psi_t^{(i)}$ and variance $\sigma_{\psi,t}^2 = \sigma_{\psi,t-1}^2 + [Q(\mathbf{u}_t)]_{\psi\psi}$, which accumulates noise over time. The particle set is resampled when the effective sample size $N_{\text{eff}} = 1 / \sum_i (w_t^{(i)})^2$ falls below $0.5N_p$, with small Gaussian roughening [23].

2) *Tactile likelihood*: A tactile probe provides geometric information about the robot pose by relating the observed end-effector stopping position to the known map.

For each particle $\mathbf{x}_t^{(i)}$, we transform the observed body-frame stopping position $\mathbf{s}_t$ into the map frame. We denote the transformed stopping point's signed normal distance to the nearest map surface by $\delta_t^{(i)}$. $\delta_t^{(i)} = 0$ indicates that the stopping point lies on the surface, $\delta_t^{(i)} < 0$ indicates predicted penetration beyond the surface, and $\delta_t^{(i)} > 0$ indicates that the probe stops before reaching the surface, as illustrated in Fig. 4.

For a contact observation ($c_t = 1$), a consistent pose hypothesis should place the stopping point on the surface, and therefore hypotheses are penalized according to $|\delta_t^{(i)}|$. For a contact-free observation ($c_t = 0$), a hypothesis is inconsistent only when $\delta_t^{(i)} < 0$, since it predicts that the end-effector should have encountered the surface before reaching its observed stopping position. We therefore define the tactile likelihood as

$$p(\mathbf{o}_t^{\text{tac}} | \mathbf{x}_t^{(i)}, M) \propto \begin{cases} \exp\left(-\frac{(\delta_t^{(i)})^2}{2\sigma_c^2}\right), & c_t = 1, \\ \exp\left(-\frac{[-\delta_t^{(i)}]_+^2}{2\sigma_c^2}\right), & c_t = 0, \end{cases} \quad (3)$$

where $[\cdot]_+ = \max(\cdot, 0)$ and σ_c is the measurement noise for contact. Thus, contact observations favor pose hypotheses

that place the stopping point on a map surface, whereas contact-free observations downweight hypotheses that predict a surface should have been encountered during the probe. The resulting likelihood is used in (2) to update and normalize the particle weights.

3) *Uncertainty-bounded global planner*: In vision-denied navigation, the geometrically shortest route may not be the most reliable route, as localization uncertainty can grow rapidly when the robot traverses regions without opportunities for tactile correction. The global planner searches for a short route subject to predicted localization uncertainty bounds, favoring contact with map surfaces when tactile probing correction is needed.

We impose position and heading uncertainty bounds $(\bar{\sigma}_p, \bar{\sigma}_\psi)$ along a route and limit route displacements to $|\Delta \mathbf{x}_t| \leq 0.4$ m. Let $\Delta \mathbf{x}_t^{xy}$ denote the position entries of a displacement $\Delta \mathbf{x}_t$, $\Delta \mathbf{x}_t^\psi$ the yaw entry, Σ_t^{xy} the position block of the covariance Σ_t , and $\Sigma_t^{\psi\psi}$ the yaw variance. The bounds $(\bar{\sigma}_p, \bar{\sigma}_\psi)$ are imposed at every displacement $\Delta \mathbf{x}_t$:

$$\sqrt{(\Delta \mathbf{x}_t^{xy})^\top \Sigma_t^{xy} \Delta \mathbf{x}_t^{xy}} \leq \bar{\sigma}_p, \quad \sqrt{(\Delta \mathbf{x}_t^\psi)^2 \Sigma_t^{\psi\psi}} \leq \bar{\sigma}_\psi. \quad (4)$$

To evaluate (4) during route search, the planner predicts how the pose belief evolves along each candidate route. Propagating the full particle set for every candidate is computationally intractable. For each candidate displacement $\Delta \mathbf{x}_t$, our planner first propagates (μ_t, Σ_t) through (1), accounting for the uncertainty accumulated during locomotion. We also chose the robot heading to face the nearest reachable map surface to allow tactile probing after a displacement.

The planner then predicts the effect of tactile probing using a simplified observation model. For global route evaluation, when a surface is within the probing range of a predicted mean pose, it assumes contact with a local plane of normal $\mathbf{n}_t$ and applies a Kalman update that reduces the covariance Σ_t along $\mathbf{n}_t$ according to the contact likelihood in (3). Otherwise, the propagated covariance is retained without a tactile update. This approximation captures how proximity to environmental structure can constrain localization uncertainty without simulating actual probing during route search.

A constrained A* search applies this belief propagation, predicts the tactile update at each candidate displacement, and prunes candidates that violate (4). It returns the shortest remaining route as a sequence of displacement subgoals. The local planner in Sec. IV-C then selects a probe action at each subgoal. During execution, actual observations update the belief, and the route is replanned when the estimated pose deviates from it beyond a threshold.

C. Expected-Information-Gain-Driven Action Selection

At each subgoal along the global route, the robot probes the environment with its tactile end-effectors to bound the localization uncertainty. We determine the probe action based on the probabilities and information gains of the possible outcomes of each candidate action to maximize the expected reduction in uncertainty. Under the belief $b_t \sim \mathcal{N}(\mu_t, \Sigma_t)$, each candidate action $a \in \mathcal{A}$ can result in one of two possible outcomes: either a contact or a contact-free observation.




Fig. 5: Information gain of one probe action. (a) The margin uncertainty, modeled as $\tilde{m}_t^a \sim \mathcal{N}(\mu_{m_t^a}, \sigma_{m_t^a}^2)$. (b) Pose belief along $\mathbf{n}_t^a$ before the probe, after a contact, and after a miss (prior truncated at $(\mathbf{n}_t^a)^\top \mu_t^{xy} - \mu_{m_t^a}$).

1) *Contact geometry of an action candidate*: For each candidate action a , we first compute the hypothetical end-effector position at the travel limit of its straight-line probing trajectory under the current mean pose $\mu_t = (\mu_t^x, \mu_t^y, \mu_t^\psi)$. We find the map-surface point closest to this predicted end-effector position and express it in the body frame as $\mathbf{c}_t^a$, as illustrated in Fig. 4(a). The surface around $\mathbf{c}_t^a$ is locally approximated as a plane with outward normal $\mathbf{n}_t^a \in \mathbb{R}^2$ in the map frame M .

To quantify how the predicted probe-to-surface geometry affects pose uncertainty, we linearize the signed normal distance to this local plane around μ_t . Its Jacobian with respect to the robot pose is $H_t^a = [(\mathbf{n}_t^a)^\top, (\mathbf{n}_t^a)^\top R(\mu_t^\psi + \frac{\pi}{2}) \mathbf{c}_t^a] \in \mathbb{R}^{1 \times 3}$, where $R(\mu_t^\psi)$ denotes the rotation from the body frame to the map frame at the mean heading μ_t^ψ . Thus, H_t^a maps errors in robot pose to changes in the signed normal distance.

2) *Contact probability*: At the current mean pose μ_t , we model the margin $\tilde{m}_t^a$, defined as the signed normal distance from the end-effector position at its travel limit to the local surface at $\mathbf{c}_t^a$, as a Gaussian random variable. Under this definition, $\tilde{m}_t^a < 0$ indicates that the surface lies within the probe's reach. $\tilde{m}_t^a$ is subject to uncertainty arising from both the pose uncertainty Σ_t and the control noise, modeled as $\tilde{m}_t^a \sim \mathcal{N}(\mu_{m_t^a}, \sigma_{m_t^a}^2)$ with

$$\sigma_{m_t^a}^2 = H_t^a \Sigma_t H_t^{a\top} + \sigma_{\text{control}}^2, \quad (5)$$

where $H_t^a \Sigma_t H_t^{a\top}$ is the margin variance of the end-effector's normal coordinate induced by the pose uncertainty Σ_t and $\sigma_{\text{control}}^2$ accounts for controller noise. We can model the probability of contact as $p_t(a) = p(\tilde{m}_t^a < 0) = \Phi(-\mu_{m_t^a}/\sigma_{m_t^a})$, where Φ is the standard normal cumulative distribution function, as shown in Fig. 5(a).

3) *Expected information gain*: Next, we quantify the reduction in pose uncertainty for each possible probe outcome. For a contact observation, the signed normal distance to the surface provides a Gaussian measurement with Jacobian H_t^a and contact measurement noise variance σ_c^2 . The resulting in-

formation gain is $\text{IG}_t^c(a) = \frac{1}{2} \log(1 + H_t^a \Sigma_t H_t^{a\top} / \sigma_c^2)$ [24].

A contact-free observation establishes that the probe did not encounter the surface, corresponding to the inequality $\tilde{m}_t^a > 0$. This truncates the margin distribution in (5) at zero. Its information gain is $\text{IG}_t^f(a) = -\frac{1}{2} \log[1 - f_t^a(1 - \varrho_t^a)]$, where $f_t^a = H_t^a \Sigma_t H_t^{a\top} / \sigma_{m_t^a}^2$ is the fraction of margin variance attributable to pose uncertainty and ϱ_t^a is the variance ratio of the truncated normal distribution relative to its prior. Fig. 5(b) shows the resulting belief reduction under contact and contact-free outcomes.

We then score each candidate action by its expected information gain (EIG) over these two outcomes:

$$\begin{aligned} a_t &= \arg \max_{a \in \mathcal{A}} \text{EIG}_t(a), \\ &= \arg \max_{a \in \mathcal{A}} [p_t(a) \text{IG}_t^c(a) + (1 - p_t(a)) \text{IG}_t^f(a)], \end{aligned} \quad (6)$$

where $\text{IG}_t^c(a)$ and $\text{IG}_t^f(a)$ denote the information gains of contact and contact-free outcomes, respectively. After a_t is executed, the resulting tactile observation o_t^{tac} is fed back into the particle filter to update the belief. Action selection and execution are repeated on the updated belief until the maximum EIG falls below a threshold, at which point the robot proceeds to the next subgoal.

D. Low-level Controller

To achieve navigation with frequent stop-and-go maneuvers while maintaining contact with the surroundings, we learn a whole-body control policy that tracks end-effector pose and force as well as footstep placement [16]. As shown in Fig. 7, similar to previous works [20], we adopt a decoupled architecture in the controller. At time t , the upper-body goal is denoted as $\mathbf{g}_t^u = [\{^B \hat{\mathbf{T}}_t^{\text{ee}j}, \mathbf{f}_t^{\text{ee}j}, c_t^{\text{ee}j}\}_{j \in \{L,R\}}]$, which is specified per end-effector j and consists of the target end-effector pose $^B \hat{\mathbf{T}}_t^{\text{ee}j} \in \text{SE}(3)$ in the robot base frame $\{B\}$, the target normal contact force $\mathbf{f}_t^{\text{ee}j} \in \mathbb{R}_{\geq 0}$, and a binary contact indicator $c_t^{\text{ee}j} \in \{0, 1\}$ for the force command. The lower-body goal $\mathbf{g}_t^l = [^S \mathbf{p}_t^f, T_t^f, c_t^f]$ contains the target footstep $^S \mathbf{p}_t^f \in \text{SE}(2)$ expressed in the current stance-foot frame $\{S\}$, the footstep countdown $T_t^f \in \mathbb{R}$, and a binary foot-swing indicator $c_t^f \in \{0, 1\}$ of the current swing foot. The shared proprioception data $\mathbf{s}_t^p = [\mathbf{q}_{t-H:t}^r, \dot{\mathbf{q}}_{t-H:t}^r, \boldsymbol{\omega}_{t-H:t}^r, \mathbf{g}_{t-H:t}^r, \mathbf{a}_{t-H:t-1}^\pi]$ contains a history of H -timestep joint positions $\mathbf{q}_t^r \in \mathbb{R}^n$, joint velocities $\dot{\mathbf{q}}_t^r \in \mathbb{R}^n$, robot angular velocity $\boldsymbol{\omega}_t^r \in \mathbb{R}^3$, robot projected gravity $\mathbf{g}_t^r \in \mathbb{R}^3$, and whole-body actions $\mathbf{a}_t^r \in \mathbb{R}^n$. The upper-body policy π^u takes the upper-body goal $\mathbf{g}_t^u$ and outputs upper-body actions $\mathbf{a}_t^u \in \mathbb{R}^{n_u}$. The lower-body policy π^l takes the locomotion goal $\mathbf{g}_t^l$ and outputs lower-body actions $\mathbf{a}_t^l \in \mathbb{R}^{n_l}$. Formally, we learn an upper-body policy $\pi^u : \mathbf{s}_t^p \times \mathbf{g}_t^u \rightarrow \mathbf{a}_t^u$ and a lower-body policy $\pi^l : \mathbf{s}_t^p \times \mathbf{g}_t^l \rightarrow \mathbf{a}_t^l$.

V. EXPERIMENTS

We evaluate TAPNAV in simulation and on a Unitree G1 humanoid using the parameters in Table I. The experiments address three questions:




Fig. 6: TAPNAV navigation performance in simulation. (a)–(d) correspond to four evaluated maps, each with its own start (purple) and goal (green star). For each map, (i) shows the map layout with the estimated trajectory (blue), the ground-truth trajectory measured from the simulator (black), and contact points from probing (red), where each grid cell is $1\text{ m} \times 1\text{ m}$; (ii), (iii) show position and yaw uncertainty (blue) and error (black) over time, respectively, with vertical red lines marking contact events. The red and yellow boxes highlight the largest decreases in both estimation error and uncertainty, which occur when the robot contacts non-planar wall surfaces.




Fig. 7: The low-level controller is composed of an upper- and a lower-body policy. The upper-body policy tracks the target end-effector pose and force, while the lower-body policy tracks footstep commands. The two policies are jointly trained with shared whole-body proprioception.

Q1: Does TAPNAV provide reliable and safely bounded state estimation throughout navigation across different maps?

Q2: How accurately does the low-level controller adjust the robot’s position throughout navigation?

Q3: Can TAPNAV perform reliably in real-world navigation tasks?

A. Simulation Results

All simulation experiments are conducted in MuJoCo. To address **Q1**, TAPNAV is evaluated on four maps shown in Fig. 6. For each map, TAPNAV navigates from a known starting point to the goal through active tactile probing. Across all environments, TAPNAV maintains estimated trajectories closely following the ground-truth states measured from the simulator, indicating accurate localization during navigation. In Fig. 6(a)–(d)(i), the robot regularly makes contact with the environment. The contact points along the navigation paths are associated with uncertainty and error reduction. The largest uncertainty reductions in Fig. 6(a)–(d)(ii), (iii) are closely related to the contact geometry; contacts with

Parameter	Setting
Controller Frequency	200 / 50 Hz
Replanning Deviation Margin	$> 0.1\text{ m}$
Global Planner Grid Size	0.1 m
Minimal Wall Clearance	0.2 m
$ \Delta \mathbf{x}_t _{\max}$	0.40 m
$(\bar{\sigma}_p, \bar{\sigma}_\psi)$	$(0.18\text{ m}, 20^\circ)$
N_p	600
Initial Pose Standard Deviation	$0.03\text{ m/axis}, 1.5^\circ\text{ yaw}$
Number of Probes per Subgoal	Max. 4
EIG Threshold	0.15 nats
σ_c	0.018 m
σ_{control}	0.045 m
Trials per Map and Method	20 (seeds 0 – 19)
Simulation Time Limit	1800 s

TABLE I: Simulation and navigation parameters.

different surface normals around corners and pillars provide complementary position constraints, significantly reducing position uncertainty. On the other hand, flat-wall contacts only constrain position along the surface normal but can reduce heading uncertainty. Together, these results show that TAPNAV exploits active and informative tactile contacts to reduce localization uncertainty.

Table II reports an ablation study comparing the navigation performance of TAPNAV and the baselines across the four maps in Fig. 6. The ablation study baselines include (i) Odometry-only, (ii) Sweep-touch, and (iii) Random-touch. Odometry-only disables tactile probing and relies entirely on the odometry and IMU updates for navigation, while Random-touch and Sweep-touch replace the EIG-driven action selection of Sec. IV-C with uniformly sampled probing actions and a grid sweep over the probing directions, respectively. We report two metrics: the **pose estimation error**, given as the position and yaw errors between the torso pose estimated from the robot’s belief and the ground truth, averaged over successful trials, and the **success rate**, defined as the fraction of trials where the robot reaches the goal region. The Odometry-only baseline achieves the lowest

Method	Map a			Map b			Map c			Map d		
	PosErr (cm) ↓	YawErr (deg) ↓	SR (%) ↑	PosErr (cm) ↓	YawErr (deg) ↓	SR (%) ↑	PosErr (cm) ↓	YawErr (deg) ↓	SR (%) ↑	PosErr (cm) ↓	YawErr (deg) ↓	SR (%) ↑
TAPNAV	6.5±2.5	1.4±0.4	95	9.4±3.6	2.8±1.8	100	5.9±1.9	2.0±1.3	100	16.8±5.0	2.3±1.6	85
Random-touch	9.6±3.8	2.1±1.4	90	11.4±3.9	2.9±1.4	70	8.7±3.6	2.9±1.7	75	16.6±6.0	2.7±1.8	75
Sweep-touch	11.7±4.5	2.0±1.1	50	11.2±3.9	2.5±1.5	50	10.9±4.2	2.4±1.6	30	13.7±3.5	1.32±0.1	25
Odometry-only	23.4±11.2	3.5±2.9	35	36.5	4.0	5	13.2	0.6	5	28.5±12.2	4.0±4.6	45

TABLE II: MuJoCo simulation results. The navigation performance of TAPNAV and the baselines is evaluated across four maps. Errors are reported as mean $\pm$ std; entries without a standard deviation correspond to settings in which only a single trial succeeded.




Fig. 8: Yaw-noise sensitivity. (a) position error; (b) yaw error; (c) success rate. Bands and bars show mean $\pm$ standard error in (a)–(b) and 95% confidence intervals in (c).

average success rate of 22.5%, highlighting the importance of active tactile sensing for reliable navigation. Compared with the two active tactile sensing baselines, TAPNAV achieves a higher average success rate of 95%, compared with 77.5% for Random-touch and 38.8% for Sweep-touch, while achieving lower position error and comparable yaw error. These results support the effectiveness of TAPNAV in maintaining localization accuracy and completing navigation tasks. To test whether these gains persist under increased heading uncertainty, we add yaw noise and compute errors over all completed trials including failures. Fig. 8 shows that TAPNAV achieves the lowest mean position and yaw errors for all IMU noise. It also retains the highest success rate, maintaining its navigation advantage even under additional heading noise.

To answer **Q2**, we compare the performance of the TAPNAV low-level footstep controller with a velocity-tracking controller. The two controllers track a sequence of global planar torso targets under open-loop execution. Consecutive targets are separated by at most 0.5 m, and the commanded velocity is limited to 0.2 m/s. As shown in Fig. 9, heading errors accumulate across successive targets, causing the velocity-tracking controller to drift from the desired route. In contrast, the footstep controller reliably follows the target sequence and achieves minimal drift.

B. Real-World Experiments

To answer **Q3**, we evaluate TAPNAV on a Unitree G1 equipped with tactile end-effectors, as shown in Fig. 3. We conduct experiments in an office hallway and a kitchen environment to evaluate the integrated estimation, planning, and control pipeline under real-world sensing and control uncertainties. Fig. 10 illustrates successful navigation trials in both environments. During execution, the robot moves around the counter, makes successive turns, and reaches the goal. Along the route, it actively probes different surfaces




Fig. 9: Simulated open-loop goal tracking with velocity- and footstep-tracking policies (30 trials per policy). (a) Executed trajectories, highlighting the trial with median final position error for each policy; (b) Actual and commanded heading over time.

to reduce localization uncertainty, demonstrating effective tactile sensing across changes in contact geometry.

VI. CONCLUSION

In this study, we presented TAPNAV, a vision-denied humanoid navigation framework integrating tactile-based active probing, belief-space localization, uncertainty-aware global planning, and whole-body control. Simulation and hardware experiments demonstrate goal-directed navigation using tactile feedback, with improved navigation success and lower pose error over baselines from the ablation study in simulation. A current limitation is that TAPNAV assumes a known static floor plan and approximates the particle belief as a unimodal Gaussian during planning, which we plan to relax by extending the framework to multimodal beliefs and partially known maps.

VII. ACKNOWLEDGMENT

ChatGPT was used in Fig. 1 to generate the uncertainty ellipses, in Fig. 6 for anti-aliasing using our original plot, and in Fig. 10 for generating series number and individual facial mosaics, which we then manually added to the plots to cover the human faces.

REFERENCES

- [1] Z. Gu, J. Li, W. Shen, W. Yu, Z. Xie, S. McCrory, X. Cheng, A. Shamsah, R. Griffin, C. K. Liu *et al.*, “Humanoid locomotion and manipulation: Current progress and challenges in control, planning, and learning,” *IEEE/ASME Transactions on Mechatronics*, vol. 31, no. 2, pp. 2300–2330, 2026.
- [2] W. Zhu, A. Raju, A. Shamsah, A. Wu, S. Hutchinson, and Y. Zhao, “Emobipednav: Emotion-aware social navigation for bipedal robots with deep reinforcement learning,” *IEEE/ASME Transactions on Mechatronics*, 2026.
- [3] K. Muenprasitvej, J. Jiang, A. Shamsah, S. Coogan, and Y. Zhao, “Bipedal safe navigation over uncertain rough terrain: Unifying terrain mapping and locomotion stability,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2024, pp. 11 264–11 271.




Fig. 10: TAPNAV deployed on a Unitree G1 humanoid equipped with our tactile end-effectors, demonstrating real-world navigation capabilities in (a) a kitchen environment and (b) an office hallway. For each environment, (i) shows snapshots of the real-world deployment; (ii) shows the robot's path and active tactile probes in the robot's egocentric view.

- [4] Y. Han, K. Yu, R. Batra, N. Boyd, C. Mehta, T. Zhao, Y. She, S. Hutchinson, and Y. Zhao, "Learning generalizable vision-tactile robotic grasping strategy for deformable objects via transformer," *IEEE/ASME Transactions on Mechatronics*, vol. 30, no. 1, pp. 554–566, 2024.
- [5] M. Van der Merwe, K. Ota, D. Berenson, N. Fazeli, and D. K. Jha, "Simultaneous extrinsic contact and in-hand pose estimation via distributed tactile sensing," *IEEE Robotics and Automation Letters*, vol. 11, no. 3, pp. 2394–2401, 2026.
- [6] S. Suresh, Z. Si, S. Anderson, M. Kaess, and M. Mukadam, "Midas-touch: Monte-carlo inference over distributions across sliding touch," in *Proceedings of The 6th Conference on Robot Learning*, vol. 205. PMLR, 2023, pp. 319–331.
- [7] M. Bauza, A. Bronars, and A. Rodriguez, "Tac2pose: Tactile object pose estimation from the first touch," *The International Journal of Robotics Research*, vol. 42, no. 13, p. 1185–1209, 2023.
- [8] E. K. Gordon, B. Baraki, H. Bui, and M. Posa, "Active tactile exploration for rigid body pose and shape estimation," in *IEEE International Conference on Robotics and Automation*, 2026.
- [9] A.-H. Shahidzadeh, S. J. Yoo, P. Mantripragada, C. D. Singh, C. Fermüller, and Y. Aloimonos, "Actexplore: Active tactile exploration on unknown objects," in *IEEE International Conference on Robotics and Automation*, 2024, pp. 3411–3418.
- [10] J. Xu, H. Lin, S. Song, and M. Ciocarlie, "Tandem3d: Active tactile exploration for 3d object recognition," in *IEEE International Conference on Robotics and Automation*, 2023, p. 10401–10407.
- [11] H. Sathyanarayan and I. Abraham, "Behavior Synthesis via Contact-Aware Fisher Information Maximization," in *Proceedings of Robotics: Science and Systems*, 2025.
- [12] R. Buchanan, J. Bednarek, M. Camurri, M. R. Nowicki, K. Walas, and M. Fallon, "Navigating by touch: haptic monte carlo localization via geometric sensing and terrain classification," *Autonomous Robots*, vol. 45, no. 6, pp. 843–857, 2021.
- [13] G. Puthumanaim, P. Padrao, J. Fuentes, L. Bobadilla, and M. Ornik, "Guided agents: Enhancing navigation policies through task-specific uncertainty abstraction in localization-limited environments," *IEEE Robotics and Automation Letters*, vol. 11, no. 3, pp. 2999–3006, 2026.
- [14] L. Di Giammarino, B. Sun, G. Grisetti, M. Pollefeys, H. Blum, and D. Barath, "Learning where to look: Self-supervised viewpoint selection for active localization using geometrical information," in *Computer Vision – ECCV*. Springer Nature Switzerland, 2025, pp. 188–205.
- [15] C. L. Shek, K. Torshizi, T. Williams, and P. Tokekar, "When to localize? a risk-constrained reinforcement learning approach," in *American Control Conference*, 2025, pp. 194–199.
- [16] Z. Gu, Y. Chen, Z. Chai, A. Cueva, T. Nguyen, Y. Wu, H. Xue, M. Kim, I. Legene, F. Liu, K. Kim, A. Barula, Y. Chen, and Y. Zhao, "Refine-dp: Diffusion policy fine-tuning for humanoid loco-manipulation via reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 11, no. 10, pp. 11 617–11 624, 2026.
- [17] F. Liu, Z. Gu, Y. Cai, Z. Zhou, H. Jung, J. Jang, S. Zhao, S. Ha, Y. Chen, D. Xu, and Y. Zhao, "Opt2skill: Imitating dynamically-feasible whole-body trajectories for versatile humanoid loco-manipulation," *IEEE Robotics and Automation Letters*, 2025.
- [18] A. Shukla, K. Fay, H. Liu, D. A. Djapri, A. Buynitsky, and N. Hansen, "Catch it all!: Generalizable dynamic catching with a dexterous hand," in *4th RSS Workshop on Dexterous Manipulation: Scalable Learning for Human-Level Skills*, 2026.
- [19] A. Huang, Z. Wu, S. Atar, Y. Zhi, and M. Yip, "SteadyTray: Learning object balancing tasks in humanoid tray transport via residual reinforcement learning," *arXiv preprint arXiv:2603.10306*, 2026.
- [20] J. Jang, Z. Wang, Z. Zhou, F. Wu, and Y. Zhao, "SEEC: Stable end-effector control with model-enhanced residual learning for humanoid loco-manipulation," *arXiv preprint arXiv:2509.21231*, 2025.
- [21] J. Jang, Z. Gu, A. Cueva, Z. Chai, J. Sheng, T. Nguyen, H. Galundia, Y. Wu, H. Xue, I. Legene, O. Mediratta, D. Doan, A. Collins, S. Sadegh, K. Kim, R. Dhalbiso, Z. Chen, and Y. Zhao, "Wt-umi: Tactile-based whole-body manipulation via force-supervised contact-aware planning," in *Conference on Robot Learning*, 2026.
- [22] TouchTronix Robotics Inc., *SensX thin-film tactile sensors*, <https://www.touchtronix.io/>.
- [23] N. J. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proceedings F (Radar and Signal Processing)*, vol. 140, no. 2, pp. 107–113, 1993.
- [24] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Wiley-Interscience, 2006.