Belief-Guided Interactive Contact Exploration with Proprioceptive Sensing on Low-Cost Robot Arms

Author Names Omitted for Anonymous Review.

Abstract—**Planning and control in real-world manipulation often rely on perception that is incomplete, noisy, or unavailable. This issue is especially pronounced for low-cost robot arms, which may lack tactile skins, external force-torque sensors, reliable torque sensing, or robust workspace vision. In such settings, contact is often treated as a collision failure. We instead study contact as an imperfect but useful source of information. This extended abstract frames low-cost manipulation without exteroceptive sensing as a two-stage problem: weak proprioceptive interaction signals must first be converted into an actionable contact belief, and a belief-conditioned policy must then use this state to decide whether to probe, reach, retreat, or gently push. The belief does not aim to reconstruct the scene precisely; it estimates task-relevant regions such as free, blocked, uncertain, and potentially pushable areas. Our current prototype instantiates this framework with simple contact-event abstractions and a hand-designed policy. Preliminary simulation and real-robot prototypes suggest that weak contact cues can support belief updates for constrained reach-and-interaction tasks, while motivating future work on optimized signal-to-belief updates, belief-conditioned policy learning, and systematic ablations against no-belief, random-probing, and binary-contact baselines.**

I. Introduction

Robot manipulation is moving beyond structured industrial environments into homes and cluttered workspaces [2, 1]. Low-cost robot arms can reduce hardware barriers and make manipulation systems easier to reproduce and deploy [14]. However, low cost also weakens the usual sensing and control assumptions: these platforms often have lower precision, lower stiffness, noisier actuation, and no tactile skin, wrist force-torque sensor, reliable torque sensing, or robust workspace vision [14, 10, 5]. Unstructured scenes further add occlusion, clutter, and unpredictable contact.

Humans face a related problem when vision is incomplete. When reaching into a dark cabinet, contact is not necessarily a failure: light touch, resistance, compliance, motion outcomes, and proprioception help infer open directions, obstacles, and whether an object may be movable [9]. Although human arms provide much richer sensing and control, the underlying principle is useful for robotics: uncertain contact can be information for action.

We ask whether a low-cost robot arm (USD 4.4k) can use only proprioceptive signals to select actions without exteroceptive sensing. These signals are not sufficient for precise contact localization, object identification, or full scene reconstruction. We therefore treat low-cost manipulation as a belief-construction and belief-conditioned control problem: the robot should infer only the state variables needed for the next safe action, rather than the full scene.

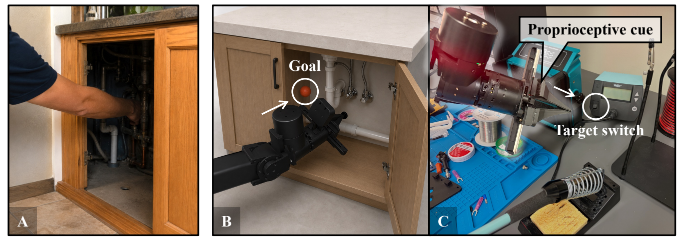

Fig. 1. Motivation and target settings. (A) Human exploratory touch in a dark cabinet. (B) Simulated cabinet reach to a goal. (C) Hardware reach to an occluded switch on a cluttered workbench.

Prior work has shown that contact and active information gathering can be useful rather than purely adversarial in manipulation. Whole-arm tactile systems have reached through clutter [6]; contact-driven, POMDP-based, and contact-only methods have used contact to estimate occupancy, reduce pose uncertainty, or retrieve objects [13, 12, 15]; active-perception grasping handles heavy clutter and difficult optical properties with camera motion [3, 8]; and sensorless or proprioception-only methods recover contact events or memories without tactile skins [10, 11, 5]. Our setting is lower-observability: sparse, noisy proprioceptive events must produce only the belief variables needed for action selection.

We make three preliminary contributions: (i) a formulation of low-cost manipulation without exteroceptive sensing as a belief-construction problem from weak proprioceptive interaction signals; (ii) an actionable contact belief interface over free, blocked, uncertain, and pushable regions; and (iii) a prototype framework and planned ablation protocol that separate two future algorithmic targets: signal-to-belief map construction and belief-conditioned policy learning.

II. From Weak Contact Cues to Actionable Belief

We study contact-rich manipulation without exteroceptive sensing in cluttered workspaces, where a robot must move an end-effector or simple tool toward a target while hidden structure or clutter may block the path. Without vision or tactile skin, observations are weak proprioceptive events: a commanded motion stalls, tracking error grows, motor current changes, or motion succeeds without contact. These cues are useful but under-specified; they do not identify object shape, contact force, or an exact contact point.

The useful state is therefore smaller than a reconstructed scene. We maintain a region-level contact belief whose labels are chosen by the action they support: free regions enable reaching, blocked regions call for avoidance or retreat, un-

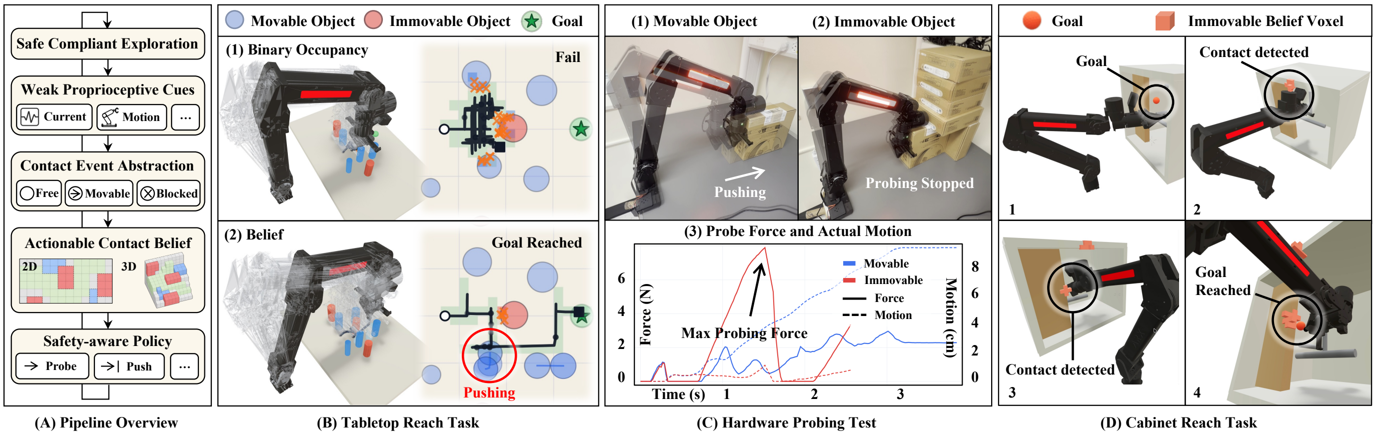

Fig. 2. Pipeline and preliminary experiments. (A) Safe compliant exploration converts weak proprioceptive cues into contact events, updates an actionable contact belief, and passes it to a safety-aware policy. (B) In tabletop reaching, binary occupancy treats contact as blockage and fails, while the belief distinguishes pushable from blocked regions and reaches the goal. (C) Hardware force and motion traces separate movable from immovable contacts under a bounded probing force. (D) In cabinet reaching, detected contacts update immovable belief voxels for retreat, replanning, and goal reaching.

certain regions call for probing, and pushable regions may justify a gentle push or sweep. Unlike planning with known object arrangements or detailed pushing dynamics [4, 7], the planner must first build this task-relevant state from weak contact evidence.

Viewed as a framework, the belief map is not a representation of the world; it represents what the robot should do next under low-cost sensing constraints. This creates two algorithmic subproblems: signal-to-belief construction, which fuses tracking error, motor-current changes, stall events, and motion outcomes into label probabilities; and belief-conditioned control, which selects actions that trade off task progress, information gathering, and contact safety.

The current controller follows the loop in Fig. 2A. Safe, compliant exploration generates weak cues, contact-event abstractions update the belief, and a hand-designed safety-aware policy selects from probe, reach, retreat, and gentle push. This makes the prototype an evaluation scaffold: the belief should be judged by whether it changes action choices under the same weak observations, not by how completely it reconstructs the scene.

III. Prototype Suite and Planned Evaluation

Figure 2 summarizes the current prototype suite and the planned evaluation path. The tabletop reach task in Fig. 2B tests whether a contact-derived belief changes decisions beyond binary occupancy. A free/occupied map can mark blockage, but it has no state for whether contact should be avoided or used; in mixed movable and immovable clutter, this causes the robot to stop short of the goal. The actionable belief records contact outcomes as free, blocked, uncertain, or pushable regions, allowing the policy to probe around hard contacts and apply a bounded push when movable clutter blocks an otherwise useful path.

The hardware force-estimation experiment in Fig. 2C checks whether the low-cost arm provides enough proprioceptive separation to support those labels. We run bounded probes against movable and immovable objects using force estimates from onboard signals, without tactile skin or an external force-torque sensor. Movable contacts show continued motion under lower resistance, while immovable contacts quickly reach the maximum probing force and stop. These force and motion traces calibrate the contact labels used in hardware reach settings such as the cluttered switch task in Fig. 1C.

The cabinet reach task in Fig. 2D stresses the same interface under constrained 3D contact. Contacts may occur on the gripper, wrist, or arm body, and the system only needs a coarse action-level update: detected contacts deposit immovable belief voxels near the inferred region, after which the robot retreats and replans toward the goal. This tests whether contact memory reduces repeated collisions with the same surface while still allowing progress.

These prototypes provide preliminary evidence and define the planned evaluation. We will compare stop-on-contact, random probing, a free/blocked map, the full actionable belief, and an oracle map across sparse clutter, mixed movable/immovable clutter, and constrained cabinet reaching. The ablations separate belief construction from belief use: contact-label errors and hardware failure modes evaluate the signal-to-belief module, while success, probes, repeated blocked contacts, and completion time evaluate control.

IV. Discussion

The current prototypes do not constitute a complete manipulation system; they test whether weak proprioceptive events can induce coarse action labels that affect behavior without scene reconstruction. Future work will replace hand-designed contact updates and policies with uncertainty-aware signal-to-belief models and belief-conditioned controllers.

The planned ablations will identify both performance gains and failure modes. False blocked labels waste motion, false pushable labels make contact too aggressive, and persistent uncertainty should trigger probing rather than repeated collisions.

References

[1] Jeannette Bohg, Karol Hausman, Bharath Sankaran, Oliver Brock, Danica Kragic, Stefan Schaal, and Gaurav S. Sukhatme. Interactive Perception: Leveraging Action in Perception and Perception in Action. *IEEE Transactions on Robotics*, 33(6):1273–1291, 2017. doi: 10.1109/TRO.2017.2721939. URL https://doi.org/10.1109/TRO.2017.2721939.

[2] Sachin Chitta, E. Gil Jones, Matei Ciocarlie, and Kaijen Hsiao. Mobile Manipulation in Unstructured Environments: Perception, Planning, and Execution. *IEEE Robotics & Automation Magazine*, 19(2):58–71, 2012. doi: 10.1109/MRA.2012.2191995. URL https://doi.org/10.1109/MRA.2012.2191995.

[3] Yixiang Dai, Siang Chen, Kaiqin Yang, Dingchang Hu, Pengwei Xie, Guosheng Li, Yuan Shen, and Guijin Wang. Active-Perceptive Language-Oriented Grasp Policy for Heavily Cluttered Scenes. *IEEE Robotics and Automation Letters*, 10(11):11094–11101, 2025. doi: 10.1109/LRA.2025.3604750. URL https://doi.org/10.1109/LRA.2025.3604750.

[4] Mehmet R. Dogar and Siddhartha S. Srinivasa. A Planning Framework for Non-Prehensile Manipulation under Clutter and Uncertainty. *Autonomous Robots*, 33(3):217–236, 2012. doi: 10.1007/s10514-012-9306-z. URL https://doi.org/10.1007/s10514-012-9306-z.

[5] Wanjia Fu, Hongyu Li, Ivy X. He, Stefanie Tellex, and Srinath Sridhar. UniTac: Whole-Robot Touch Sensing Without Tactile Sensors. *arXiv preprint arXiv:2507.07980*, 2025. doi: 10.48550/arXiv.2507.07980. URL https://arxiv.org/abs/2507.07980.

[6] Advait Jain, Marc D. Killpack, Aaron Edsinger, and Charles C. Kemp. Manipulation in Clutter with Whole-Arm Tactile Sensing. *The International Journal of Robotics Research*, 32(4):458–482, 2013. doi: 10.1177/0278364912471865. URL https://doi.org/10.1177/0278364912471865.

[7] Julius Jankowski, Lara Brudermüller, Nick Hawes, and Sylvain Calinon. Robust Pushing: Exploiting Quasi-Static Belief Dynamics and Contact-Informed Optimization. *The International Journal of Robotics Research*, 44(12):1959–1980, 2025. doi: 10.1177/02783649251318046. URL https://doi.org/10.1177/02783649251318046.

[8] Ginga Kennis and Shogo Arai. HEAPGrasp: Hand-Eye Active Perception to Grasp Objects with Diverse Optical Properties. *IEEE Robotics and Automation Letters*, 11(3):3206–3213, 2026. doi: 10.1109/LRA.2026.3653331. URL https://doi.org/10.1109/LRA.2026.3653331.

[9] Susan J. Lederman and Roberta L. Klatzky. Hand Movements: A Window into Haptic Object Recognition. *Cognitive Psychology*, 19(3):342–368, 1987. doi: 10.1016/0010-0285(87)90008-9. URL https://doi.org/10.1016/0010-0285(87)90008-9.

[10] Jacky Liang and Oliver Kroemer. Contact Localization for Robot Arms in Motion without Torque Sensing. In *IEEE International Conference on Robotics and Automation*, 2021. doi: 10.48550/arXiv.2011.03142. URL https://arxiv.org/abs/2011.03142.

[11] Federico Pizarro Bejarano, Bryson Jones, Daniel Pastor Moreno, Joseph Bowkett, Paul G. Backes, and Angela P. Schoellig. ProDapt: Proprioceptive Adaptation using Long-Term Memory Diffusion. In *IEEE International Conference on Robotics and Automation*, pages 15588–15594, 2025. URL https://arxiv.org/abs/2503.00193.

[12] Muhammad Suhail Saleem, Rishi Veerapaneni, and Maxim Likhachev. A POMDP-Based Hierarchical Planning Framework for Manipulation under Pose Uncertainty. *arXiv preprint arXiv:2409.18775*, 2024. doi: 10.48550/arXiv.2409.18775. URL https://arxiv.org/abs/2409.18775.

[13] Muhammad Suhail Saleem, Lai Yuan, and Maxim Likhachev. A Contact-Driven Framework for Manipulating in the Blind. *arXiv preprint arXiv:2510.20177*, 2025. doi: 10.48550/arXiv.2510.20177. URL https://arxiv.org/abs/2510.20177.

[14] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, 2023. doi: 10.15607/RSS.2023.XIX.016. URL https://doi.org/10.15607/RSS.2023.XIX.016.

[15] Sheng Zhong, Nima Fazeli, and Dmitry Berenson. Soft Tracking Using Contacts for Cluttered Objects to Perform Blind Object Retrieval. *IEEE Robotics and Automation Letters*, 2022. doi: 10.48550/arXiv.2201.10434. URL https://arxiv.org/abs/2201.10434.