

Embodying Multi-Hand Manipulation Policies by Searching the Assignment and Null Spaces

Yorai Shaoul, Jiaoyang Li, Maxim Likhachev

The Robotics Institute, Carnegie Mellon University

Abstract

Learned manipulation policies increasingly predict motions for abstract “hands” and are attractive in practice because they rely on easily collected demonstrations and transfer across robot platforms. Executing these trajectories on multi-arm robots, however, is not trivial. Multi-hand policy outputs must be assigned to physical arms, each arm must realize a configuration-space motion that tracks its prescribed end-effector trajectory, and all arms must respect kinematic limits and avoid collisions. In the absence of algorithms that directly address this problem, practitioners typically extend single-arm inverse-kinematics (IK) pipelines in an ad hoc way, with no guarantees of feasibility or safety. In this work, we close this execution gap with a search-based framework that is theoretically complete for grounding policy-generated multi-hand trajectories onto physical multi-arm systems. Building on Conflict-Based Search, our method explicitly searches over both the discrete assignment of trajectories to arms and the continuous Jacobian null spaces of redundant manipulators, using redundancy to avoid inter-arm collisions while tracking the prescribed motions. This unified treatment of assignment and null-space motion yields a practically efficient planner that safely realizes coordinated manipulation-policy outputs on multi-arm robots. See omcbsa.github.io for more.

Introduction

Modern manipulation policies increasingly output end-effector motions rather than joint commands, making demonstrations easy to collect and enabling transfer across robot platforms (Black et al. 2024; Chi et al. 2024, 2023). This abstraction works well for single-arm systems, where differential IK can reliably track the policy’s motions. In multi-arm settings, however, it introduces an execution gap: policies do not specify which arm should execute each motion, nor how several manipulators should follow prescribed paths without colliding. As a result, standard single-arm strategies fail even when multi-arm realizations exist.

We close this gap by casting multi-arm policy embodiment as a search over assignments and null-space motions, ensuring that each arm tracks a prescribed trajectory while avoiding inter-arm collisions. Succinctly, we contribute (1) an efficient planner, with two effective variants, for safely realizing multi-hand trajectories on multi-arm systems; (2)

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Policy outputs often specify only end-effector trajectories, leaving a gap to determine arm-motion assignments and to reason over redundancy for collision-free embodiment. **Left:** visualizing the redundancy manifold $\mathcal{M}$ with multiple configurations q , some safe and some in collision, all with the same end-effector pose $x \in SE(3)$, i.e., $FK(q) = x$. **Right:** three arms collaboratively flip or push.

a theoretical analysis establishing resolution-completeness over the search space induced by motion primitives; and (3) a physical and simulated experimental study.

Background

Let us begin by formally defining the problem at hand, reviewing relevant fundamentals, and surveying related work.

On-Manifold Anonymous Multi-Robot-Arm Motion Planning (OM-AMRAMP)

We consider N manipulators $\{R^i\}_{i=1}^N$, each with d joints (also termed DOF: degrees of freedom) and configuration space $\mathcal{Q}_R^i \subset \mathbb{R}^d$. A configuration $q^i \in \mathcal{Q}_R^i$ is a choice of joint angles for R^i , which determines the pose of its end-effector (i.e., “hand”) via the forward-kinematics map $FK^i : \mathcal{Q}_R^i \rightarrow SE(3)$ and the volume it occupies in the workspace, denoted $R^i(q^i) \subset \mathbb{R}^3$. We denote the (static) obstacle region with $\mathcal{O} \subset \mathbb{R}^3$.

The input to OM-AMRAMP is a set of N end-effector trajectories of H steps, $X^j = \{x_1^j, \dots, x_H^j\}$ for $j = 1, \dots, N$, with $x_t^j \in SE(3)$, which must be realized by the arms. We seek a *feasible* solution—a bijection $\sigma : \{1, \dots, N\} \rightarrow \{1, \dots, N\}$ from robots to trajectories and configuration-space trajectories $\tau^i = \{\tau_1^i, \dots, \tau_H^i\}$ with $\tau_t^i \in \mathcal{Q}_R^i$, such that for all robots R^i , timesteps t , and all $k \neq i$, each arm exactly tracks a trajectory without colliding. I.e.,

$$FK^i(\tau_t^i) = x_t^{\sigma(i)}, \quad R^i(\tau_t^i) \cap R^k(\tau_t^k) = \emptyset, \quad R^i(\tau_t^i) \cap \mathcal{O} = \emptyset.$$

We leave reaching τ_1^i to standard multi-arm motion planners (Mishani et al. 2025; Shaoul et al. 2024a; Shome et al. 2020; Huang et al. 2025; Kuffner and LaValle 2000).

Fundamentals and Related Work

The prevailing approach to embodying single-hand manipulation trajectories on a robotic arm is via inverse kinematics (IK) tracking (Yao and Gupta 2005). This process typically begins by computing a first satisfying configuration τ_1 where $FK(\tau_1) = x_1$, and then proceeding to find the remaining τ_t such that $FK(\tau_t) = x_t$, often attempting to minimize $|\tau_t - \tau_{t-1}|$. Solving the IK problem, where we seek a configuration q whose end-effector pose matches a desired $x \in SE(3)$, is often done numerically via the geometric Jacobian $J(q)$, which relates joint velocities to end-effector velocities through $\dot{x} = J(q)\dot{q}$. Given an initial guess q_{seed} with pose error $\Delta x = x - FK(q_{\text{seed}})$, a local IK update is obtained by $\Delta q = J^+(q_{\text{seed}}) \Delta x$, where J^+ is the pseudoinverse. Iterating $q_{\text{seed}} \leftarrow q_{\text{seed}} + \Delta q$ drives the end-effector toward the target. This process, which we denote $IK(x, q_{\text{seed}})$, is used for tracking by attempting random seeds q_{rand} until finding $\tau_1 \leftarrow IK(x_1, q_{\text{rand}})$ and then $\tau_t \leftarrow IK(x_t, \tau_{t-1})$.

For common manipulators with more than six joints, $FK(q) = x_t$ admits infinitely many solutions (Fig. 1), forming the low-dimensional manifold $\mathcal{M}_t = \{q \in \mathcal{Q}_R \mid FK(q) = x_t\}$. Because numerical IK returns only a *single* solution determined by its seed, IK-tracking commits early to one trajectory solution and cannot explore configurations on $\mathcal{M}_t$ that might avoid collisions later in execution.

Advanced single-robot pose-constrained planners explore $\mathcal{M}_t$ more broadly, e.g., by sampling multiple IK solutions per pose and searching over their connections (De Maeyer, Moyaers, and Demeester 2017), by projection-based sampling on the manifold (Stilman 2007; Berenson et al. 2009; Oriolo and Mongillo 2005), or by systematically searching redundant joint values (Cohen, Chitta, and Likhachev 2014). These methods improve upon IK-tracking, but they still do not address assignments and coordination jointly.

For redundant manipulators, a general way to explore $\mathcal{M}_t$ arises from the null space of the Jacobian. At a configuration q , the null space $\text{Null}(J(q))$ consists of all joint-space directions that leave the end-effector pose unchanged. I.e., with $N(q)$, a matrix whose columns span $\text{Null}(J(q))$, any update of the form $\Delta q = N(q) \alpha$, where α is a small coefficient vector, moves the robot locally without altering its end-effector pose. These null-space directions provide systematic “near neighbors” on the constraint manifold and allow local exploration of redundant configurations. In this work, we adopt this null-space-based exploration to generate feasible on-manifold motions while simultaneously reasoning about trajectory–arm assignments and inter-arm collisions.

On-Manifold CBS with Assignments

We now introduce our method, *On-Manifold Conflict-Based Search with Assignments* (Ω -CBSA, pronounced “Om” CBSA). At its core, Ω -CBSA relies on a single-robot planner, Ω -A*, which performs search directly on the pose-constraint manifolds induced by a desired end-effector trajectory. We first present Ω -A*_{single} for the single-trajectory

setting and then extend it to a multi-goal variant that allows a robot to implicitly solve the assignment problem by selecting which trajectory to embody during search. Our main multi-robot algorithm, Ω -CBSA, builds on this multi-goal Ω -A* to jointly reason about trajectory–arm assignments and inter-arm collisions. We also introduce two prioritized planners (PP), OM-PP-A*_{single} and OM-PP-A*, using the single-goal and multi-goal variants of Ω -A* in a prioritized scheme (Erdmann and Lozano-Perez 1987).

On-Manifold A*. Ω -A*_{single} is a single-robot, single-hand-trajectory planner that performs A* search (Nilsson 1980) directly on the pose-constraint manifold induced by an end-effector path $X = \{x_1, \dots, x_H\}$. The algorithm maintains an OPEN list (a priority queue) of states $s = (q, t)$, where q is a configuration with $FK(q) = x_t$. OPEN is initialized with multiple IK solutions for x_1 , obtained by running IK from diverse seeds to roughly sample the manifold $\mathcal{M}_1 = \{q \mid FK(q) = x_1\}$. Each state stores a cost-to-come $g(s)$, a heuristic $h(s) = H - t$, a collision count $c(s)$ with respect to any provided other-robot paths, and a priority value $f(s) = g(s) + w_h h(s) + w_c c(s)$ as in Veerapaneni, Kusnur, and Likhachev (2023). At each step, Ω -A*_{single} pops from OPEN the state with minimal $f(s)$; if this state is at $t = H$, the algorithm terminates and returns the corresponding on-manifold path. Otherwise, from a popped state (q, t) , it generates successors s' for x_{t+1} by first computing a nominal projection $q' = IK(x_{t+1}, q)$, and then exploring redundant realizations via nullspace motion: letting $N(q')$ denote a basis of $\text{Null}(J(q'))$, for each basis direction $n \in N(q')$ and step size $\pm\epsilon$, the algorithm proposes $q'' = q' + \epsilon n$ followed by a correction step $q'' \leftarrow IK(x_{t+1}, q'')$ to return to the manifold. Each valid q'' (collision-free with obstacles and within joint limits) becomes a successor state $(q'', t+1)$ with updated¹ $g(s') = g(s) + \|q'' - q\|$, h , c , and f , and is inserted into OPEN. We set $\epsilon = 0.2$, $w_h = 10$, $w_c = 0.1$.

On-Manifold Multi-Goal A*. To handle the setting where a robot can choose among multiple end-effector trajectories $\{X^1, \dots, X^M\}$, we design Ω -A*. Here, instead of initializing OPEN with IK solutions for the first pose of a single trajectory, we initialize it with IK solutions for the first pose of *every* candidate trajectory. Each state now has the form (q, t, j) , encoding a timed configuration tracking X^j , and the search proceeds exactly as in Ω -A*_{single}, with the only difference that successors $s' = (q', t+1, j)$ generated along X^j . This has assignment decisions handled implicitly, allowing the planner to determine both *which* trajectory to embody and *how* to track it without committing prematurely.

On-Manifold Prioritized Planning. A simple way to coordinate multiple manipulators is to use Ω -A* as a low-level planner in a prioritized scheme. Robots are ordered by a priority and planned sequentially: the highest-priority robot plans first, and each subsequent robot runs Ω -A* (or Ω -A*_{single}, if assignments are available) while treating previously planned trajectories as moving obstacles, disallow-

¹For initial states $g(s) = \|FK(q) - x_1\|$ for q the current robot configuration. $c(s') = c(s) + \text{COUNTCOLLISIONS}(q')$.

ing configurations q_t^i for which $R^i(q_t^i) \cap R^k(\tau_t^k) \neq \emptyset$ for some higher-priority R^k . We refer to the resulting planners as OM-PP-A* and OM-PP-A*_{single}. These on-manifold prioritized planners are efficient and often effective in practice, but—as in standard prioritized schemes—are incomplete.

On-Manifold Conflict-Based Search with Assignments. To obtain a complete multi-robot planner for pose-constrained policy embodiment, we integrate our multi-goal Ω -A* with Conflict-Based Search (CBS) (Sharon et al. 2015). Ω -CBSA begins by planning an on-manifold trajectory for each robot²; these initial trajectories may contain two types of conflicts: *assignment conflicts*, in which two robots select the same candidate trajectory X^j , and *geometric conflicts*, in which two robots collide at time t .

As in CBS, Ω -CBSA organizes the search in a constraint tree (CT). Each CT node contains (i) a set of per-robot constraints, (ii) the trajectories $\{\tau^i\}_{i=1}^N$ computed under those constraints using Ω -A*, and (iii) the conflicts detected in these trajectories. At each iteration, Ω -CBSA selects the CT node with the fewest conflicts; if the node has no conflicts, its trajectories are returned as a solution. Otherwise, Ω -CBSA resolves one conflict by generating two child CT nodes, each created by copying the parent node and adding a different constraint. Each child then replans one robot, the robot to which the new constraint applies, using Ω -A*.

When an assignment conflict is encountered (e.g., robots R^i and R^k both select trajectory X^j) Ω -CBSA adds *assignment constraints* (inspired by Motes et al. (2020)). One child forbids R^i from selecting X^j , and the other forbids R^k . In the corresponding multi-goal Ω -A* replanning call, this is enforced by omitting all IK initializations for the forbidden trajectory when constructing the OPEN list.

If no assignment conflicts remain but a geometric conflict exists—say robots R^i and R^k collide at time t — Ω -CBSA selects a collision point $p \in R^i(\tau_t^i) \cap R^k(\tau_t^k)$ and creates two child CT nodes: one forbidding robot R^i from occupying p at time t , and one forbidding robot R^k from the same. In practice, during low-level validation, a state (q, t) violates such a constraint if the occupied volume of the robot at timestep t contains the witness point p . This is the same constraint model used in (Shaoul et al. 2024b).

By combining CBS’s high-level reasoning over assignment and geometric constraints with low-level on-manifold feasibility checks, Ω -CBSA provides a complete planner for pose-constrained multi-arm policy embodiment.

Theoretical Analysis

As is customary for search-based planning algorithms, Ω -CBSA adopts a finite-resolution discrete-time abstraction of the underlying continuous problem: low-level motion is generated with a finite set of motion primitives. This abstraction allows us to establish the *resolution-completeness* of Ω -CBSA for OM-AMRAMP. We do so by first analyzing Ω -A*, showing that it is resolution-complete, and then showing that its integration into a CBS framework with assignment and geometric constraints maintains this guarantee.

²In practice, the root-node construction can be warm-started with repeated IK-tracking, with all later planning using Ω -A*.

Lemma 1 (Resolution-Completeness of Ω -A*). *If a feasible pose-constrained trajectory τ^i following a fixed end-effector path $X^j \in \{X^1, \dots, X^N\}$ exists for a R^i under motion primitive discretization resolution, then Ω -A* will find it.*

Proof. Ω -A* searches a finite graph whose nodes are triples (q, t, j) with $t \in \{1, \dots, H\}$, q drawn from the finite set of configurations generated by the initialization set and motion primitives, and j a trajectory index. Since Ω -A* systematically explores this finite graph, every reachable node will eventually be expanded. Therefore, if a feasible trajectory exists in this induced graph and starts from one of the initialized states, Ω -A* will eventually reach and return it. $\square$

For a CBS-based planner to be complete, its single-robot planner must be complete, and the constraints it uses to resolve conflicts must be *mutually disjunctive*: for any pair of constraints c^1, c^2 generated from a single conflict between R^1 and R^2 , there does not exist a conflict-free joint solution in which both c^1 and c^2 are violated (Li et al. 2019).

Lemma 2 (Constraints in Ω -CBSA). *The assignment and geometric point constraints are mutually disjunctive.*

Proof. Consider first an assignment conflict where R^1 and R^2 both select X^j , and the two child CT nodes add constraints $c^1 = \{R^1 \text{ may not select } X^j\}$ and $c^2 = \{R^2 \text{ may not select } X^j\}$. Any joint solution that violates both c^1 and c^2 necessarily assigns X^j to both R^1 and R^2 . This contradicts the requirement that the assignment σ in OM-AMRAMP is a bijection between robots and trajectories, and therefore such a solution is invalid and cannot be conflict-free. Hence the assignment constraints are mutually disjunctive. Geometric point constraints are mutually disjunctive, as shown in Shaoul et al. (2024b). $\square$

Theorem 1 (Resolution-Completeness of Ω -CBSA). *Ω -CBSA is resolution-complete for OM-AMRAMP.*

Proof. By Lemma 2, all constraints added by Ω -CBSA are mutually disjunctive, and by Lemma 1, the low-level planner Ω -A* is resolution-complete. Since CBS with a complete low-level planner and mutually disjunctive constraints is complete, Ω -CBSA is resolution-complete. $\square$

This establishes that Ω -CBSA fully reasons over assignments, redundant-joint realizations, and inter-arm collisions, and will succeed whenever a valid embodiment exists.

Experimental Analysis

We simulated 450 benchmark problems to stress assignment reasoning, redundant-joint coordination, and collision avoidance. As illustrated in Fig. 2, each problem comprised up to six end-effector trajectories and robots arranged in diverse layouts with obstacles. The test suite included motions resembling learned manipulation-policy outputs and humanoid-inspired setups requiring tightly coordinated execution (Black et al. 2024; Shaoul et al. 2025). We also demonstrated Ω -CBSA on a physical 3-arm setup performing cloth rotation, box flip, and planar pushing (Fig. 1).




Figure 2: **Left images:** illustrations of our test cases visualizing six robots and six end-effector trajectories. The left image shows all robots moving from right to left in a constrained space, and the right image shows six robots performing a box-flipping motion. **From middle to the right:** overall success rates, success rates by robot count, and average overall runtime.

Baselines

To the best of our knowledge, no algorithm directly tackles the OM-AMRAMP problem. As such, we compared Ω -CBSA to three families of methods adapted from existing work: IK-tracking, pose-constrained roadmaps, and multi-robot implicit graph search.

IK-Tracking. This baseline reflects the standard practical workflow. Assuming a given assignment σ , each robot samples an IK solution for the first pose on their associated X^j and tracks the remainder of the trajectory using IK. Our implementation iterates over all σ , proceeding to the next assignment if one fails. This approach is efficient but myopic.

PP-Descartes. We adapt Descartes (De Maeyer, Moyaers, and Demeester 2017) to multi-robot settings. As with IK-Tracking, we enumerate all robot-trajectory assignments, and for each assignment we iterate over all priority orderings. Robots then plan in priority order: each constructs a pose-constrained roadmap by sampling IK solutions and connecting them with collision-free edges, requiring all samples and edges to also avoid the already planned higher-priority robots’ trajectories. If all robots successfully build and search their roadmaps, a joint solution is returned; otherwise, the method attempts the next ordering or assignment.

Composite A*. Inspired by search on implicit pose-constrained graphs (Cohen, Chitta, and Likhachev 2014), we implement a “composite” A* planner whose states are joint tuples $s = (q^1, \dots, q^N, t)$. For each robot R^i in state s , we define its single-robot successor set $\mathcal{S}^i$ exactly as in Ω -A*, i.e., using an IK step and null-space exploration to move from time t to $t + 1$ along its assigned trajectory. The successor set of s is then the Cartesian product of $\{\mathcal{S}^i\}_{i=1}^N$ at time $t + 1$, restricted to collision-free joint configurations. The OPEN list is initialized with IK seeds for the first pose of all end-effector trajectories, implicitly covering all assignments. While expressive, this search now reasons over a state space that grows exponentially with the number of robots.

Experimental Results

We focused our analysis on success rate (fraction of problems solved within 5 seconds) and planning time. Across all 450 benchmark problems, our proposed planners consistently outperformed all baselines in success rate. As shown in Fig. 2, Ω -CBSA achieves the highest overall success, followed closely by its prioritized variants OM-PP-A*

and OM-PP-A*_{single}. We attribute this to the combination of lightweight single-robot planning over redundant configurations and flexible reasoning over trajectory–arm assignments. When disaggregated by robot count (Fig. 2, second from right), the scalability gap becomes clearer: as the number of arms increases, the proposed methods remain consistently more successful than all baselines. While performing generally well, we observed that Descartes’ performance notably suffers due to the cost of roadmap construction. Composite A* suffers from rapid growth in branching factor as assignments are introduced, making search increasingly difficult. IK-Tracking, while simple and fast when it succeeds, is highly myopic and rarely found feasible solutions.

Our proposed planners remain fast: Ω -CBSA solved two-robot problems in 150 ms on average, and across all benchmarks achieved a mean runtime of roughly 640 ($\sigma = 990$ ms, min 16 ms, max 4.9 s, Fig. 2, rightmost plot). Our prioritized variants solved most problems in under one second. Among baselines, IK-Tracking was the fastest when it succeeded (about 160 ms on average), but its low success rate limits its practical utility. Overall, these results demonstrate that Ω -CBSA can solve more difficult planning problems than baselines, with the prioritized variants offering favorable speed–performance tradeoffs.

Implementation details. We used Pinocchio (Carpentier et al. 2019) for kinematics, Kinova Gen3 arms, and an Intel Core i9-12900H CPU (5.2 GHz). Within Ω -CBSA, root CT nodes were initialized using repeated IK-tracking (capped at 500 ms), producing initial trajectories with fewer collisions without affecting theoretical guarantees.

Conclusion

We introduced Ω -CBSA, a theoretically complete planner for embodying multi-hand manipulation policies on multi-arm robots. By combining on-manifold single-robot planning with a CBS-based coordination scheme, our method systematically reasons about redundancy, collisions, and assignments. Experiments with up to six arms show that Ω -CBSA and its prioritized variants achieve higher success rates than IK, sampling-based, and implicit graph-search baselines while maintaining real-time performance. We hope this work will lower the practical barriers associated with deploying multi-hand manipulation policies.

References

- Berenson, D.; Srinivasa, S. S.; Ferguson, D.; and Kuffner, J. J. 2009. Manipulation planning on constraint manifolds. In *2009 IEEE international conference on robotics and automation*, 625–632. IEEE.
- Black, K.; Brown, N.; Driess, D.; Esmail, A.; Equi, M.; Finn, C.; Fusai, N.; Groom, L.; Hausman, K.; Ichter, B.; et al. 2024. π : A Vision-Language-Action Flow Model for General Robot Control. *CoRR*.
- Carpentier, J.; Saurel, G.; Buondonno, G.; Mirabel, J.; Lami-reaux, F.; Stasse, O.; and Mansard, N. 2019. The Pinocchio C++ library – A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives. In *IEEE International Symposium on System Integrations (SII)*.
- Chi, C.; Xu, Z.; Feng, S.; Cousineau, E.; Du, Y.; Burchfiel, B.; Tedrake, R.; and Song, S. 2023. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*.
- Chi, C.; Xu, Z.; Pan, C.; Cousineau, E.; Burchfiel, B.; Feng, S.; Tedrake, R.; and Song, S. 2024. Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots. *arXiv preprint arXiv:2402.10329*.
- Cohen, B.; Chitta, S.; and Likhachev, M. 2014. Single-and dual-arm motion planning with heuristic search. *The International Journal of Robotics Research*, 33(2): 305–320.
- De Maeyer, J.; Moyaers, B.; and Demeester, E. 2017. Cartesian path planning for arc welding robots: Evaluation of the descartes algorithm. In *2017 22nd IEEE International conference on emerging technologies and factory automation (ETFA)*, 1–8. IEEE.
- Erdmann, M.; and Lozano-Perez, T. 1987. On multiple moving objects. *Algorithmica*, 2: 477–521.
- Huang, P.; Liu, R.; Liu, C.; and Li, J. 2025. Apex-mr: Multi-robot asynchronous planning and execution for cooperative assembly. *Robotics: Science and Systems (RSS)*.
- Kuffner, J.; and LaValle, S. 2000. RRT-connect: An efficient approach to single-query path planning. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, volume 2, 995–1001. IEEE. ISBN 978-0-7803-5886-7.
- Li, J.; Surynek, P.; Felner, A.; Ma, H.; Kumar, T. S.; and Koenig, S. 2019. Multi-agent path finding for large agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7627–7634.
- Mishani, I.; Shaoul, Y.; Natarajan, R.; Li, J.; and Likhachev, M. 2025. SRMP: Search-Based Robot Motion Planning Library. *arXiv preprint arXiv:2509.25352*.
- Motes, J.; Sandström, R.; Lee, H.; Thomas, S.; and Amato, N. M. 2020. Multi-robot task and motion planning with sub-task dependencies. *IEEE Robotics and Automation Letters*.
- Nilsson, N. 1980. Principles of Artificial Intelligence Tioga Publishing. Palo Alto, CA.
- Oriolo, G.; and Mongillo, C. 2005. Motion planning for mobile manipulators along given end-effector paths. In *Proceedings of the 2005 IEEE international conference on robotics and automation*.
- Shaoul, Y.; Chen, Z.; Mohamed, N. G.; Pecora, F.; Likhachev, M.; and Li, J. 2025. Collaborative Multi-Robot Non-Prehensile Manipulation via Flow-Matching Co-Generation. In *AAAI-26 Workshop on Multi-Agent Path Finding*.
- Shaoul, Y.; Mishani, I.; Likhachev, M.; and Li, J. 2024a. Accelerating Search-Based Planning for Multi-Robot Manipulation by Leveraging Online-Generated Experiences. In *34th International Conference on Automated Planning and Scheduling*.
- Shaoul, Y.; Veerapaneni, R.; Likhachev, M.; and Li, J. 2024b. Unconstraining multi-robot manipulation: Enabling arbitrary constraints in ecbs with bounded sub-optimality. In *Proceedings of the International Symposium on Combinatorial Search*.
- Sharon, G.; Stern, R.; Felner, A.; and Sturtevant, N. R. 2015. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219: 40–66.
- Shome, R.; Solovey, K.; Dobson, A.; Halperin, D.; and Bekris, K. E. 2020. drt*: Scalable and informed asymptotically-optimal multi-robot motion planning. *Autonomous Robots*, 44(3-4): 443–467.
- Stilman, M. 2007. Task constrained motion planning in robot joint space. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- Veerapaneni, R.; Kusnur, T.; and Likhachev, M. 2023. Effective Integration of Weighted Cost-to-Go and Conflict Heuristic within Suboptimal CBS. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(10): 11691–11698.
- Yao, Z.; and Gupta, K. 2005. Path planning with general end-effector constraints: Using task space to guide configuration space search. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.