B*: Efficient and Optimal Base Placement for Fixed-Base Manipulators

Website Paper

Code Video

Zihang Zhao[1,2,3,6,*], Leiyao Cui[2,5,*], Sirui Xie[1,2,3,*], Saiyao Zhang[2,5], Zhi Han[5], Lecheng Ruan[4], and Yixin Zhu[1,2,3,†]

[1] Institute for Artificial Intelligence, Peking University [2] School of Psychological and Cognitive Sciences,Peking University
[3] Beijing Key Laboratory of Behavior and Mental Health, Peking University [4] College of Engineering, Peking University
[5] University of Chinese Academyof Sciences, Beijing [6] LeapZenith AI Research, Shanghai [*] Equal contributors [†] Corresponding author

Motivation

Base placement is task-dependent

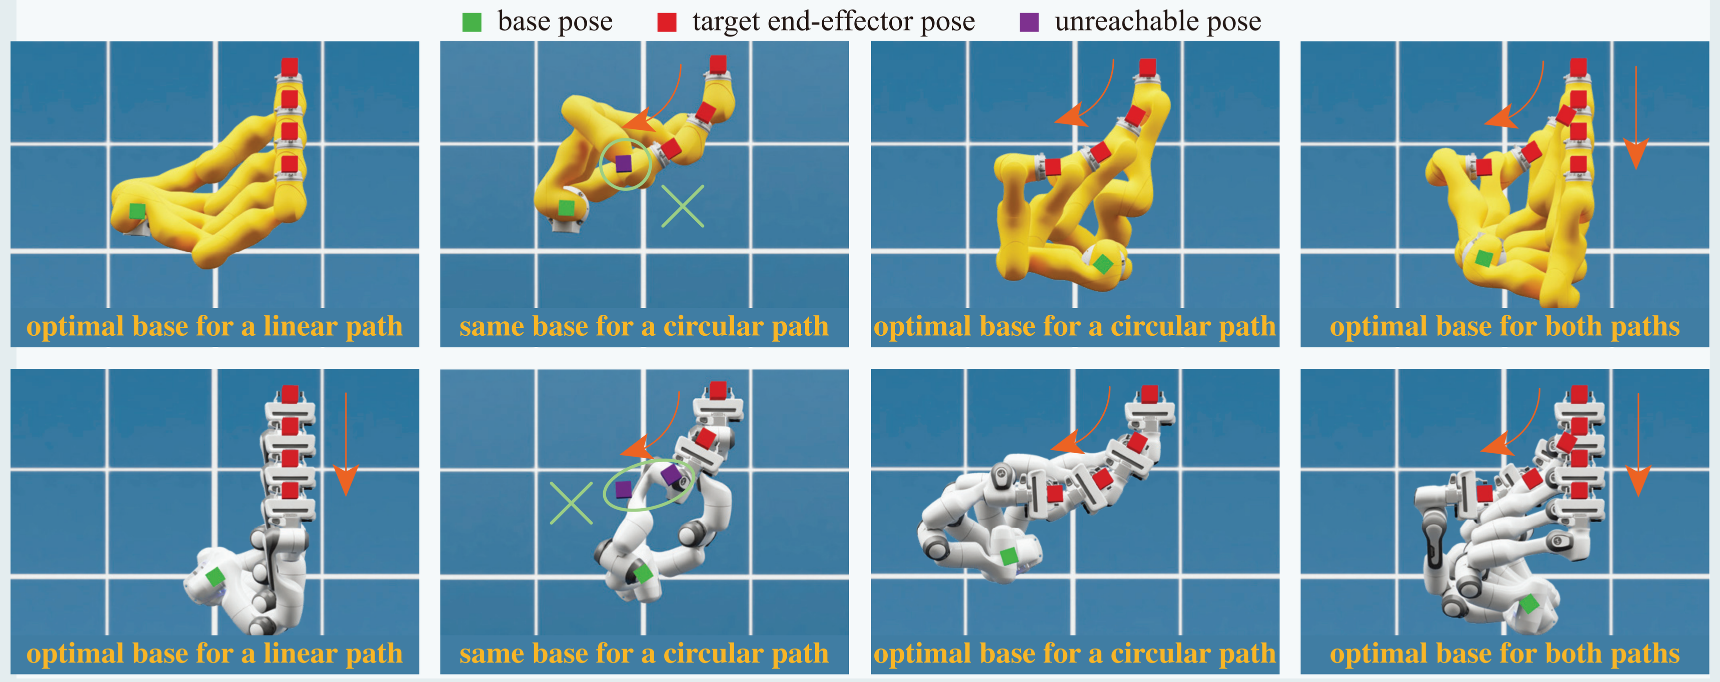

- **Input:** An ordered sequence of desired end-effector poses $\boldsymbol{x}_{1:t} = [\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_t] \in \mathbb{R}^{t \times 6}$
- **Output:** An optimal base placement $\boldsymbol{q}^b = [x^b, y^b, \theta^b]^T \in \mathrm{SE}(2)$ and feasible joint configurations $\boldsymbol{q}^m_{1:t} = [\boldsymbol{q}_1, \boldsymbol{q}_2, \cdots, \boldsymbol{q}_t] \in \mathbb{R}^{t \times n}$

Non-convex workspace

Robot workspaces are riddled with discontinuities, singularities, and uneven dexterity, making base placement inherently non-convex.

Task & robot dependency

The optimal base varies with both robot kinematics and task — one task's best placement can make another infeasible.

Multi-constraint requirements

Valid placement must jointly satisfy path-wide reachability, joint continuity, collision avoidance, and task-specific goals across long horizons.

Current methods' limitations

Solution optimality

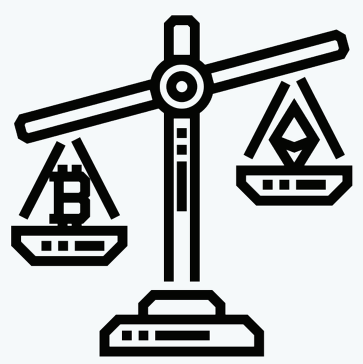

Computational efficiency

Method

Base placement formulation

$$
\begin{aligned}
\underset{\boldsymbol{q}^b, \boldsymbol{q}^m_{1:t}}{\text{minimize}} \quad & f(\boldsymbol{q}^b, \boldsymbol{q}^m_{1:t}) = \sum_{i=1}^{t-1} \left\| \boldsymbol{q}^m_{i+1} - \boldsymbol{q}^m_i \right\|_1 \\
\text{subject to} \quad & \psi(\boldsymbol{q}^b, \boldsymbol{q}^m_i) = \boldsymbol{x}_i, \qquad i = 1, 2, \ldots, t \\
& \boldsymbol{q}^{bm} \preccurlyeq \boldsymbol{q}^b \preccurlyeq \boldsymbol{q}^{bM}, \\
& \boldsymbol{q}^{mm} \preccurlyeq \boldsymbol{q}^m_i \preccurlyeq \boldsymbol{q}^{mM}, \qquad i = 1, 2, \ldots, t \\
& \mathrm{sd}(\boldsymbol{q}^b, \boldsymbol{q}^m_i) \geq 0, \qquad i = 1, 2, \ldots, t
\end{aligned}
$$

Base relaxation

$$\boldsymbol{q}^b \rightarrow \boldsymbol{q}^b_{1:t}$$

Relax problem by treating the fixed base as mobile

Constraint tightening

$$f'(\boldsymbol{q}^b_{1:t}, \boldsymbol{q}^m_{1:t}; \mu(j)) = f(\boldsymbol{q}^b_{1:t}, \boldsymbol{q}^m_{1:t}) + \mu(j) \sum_{i=1}^{t} \|\boldsymbol{q}^b_i - \bar{\boldsymbol{q}}^b\|_1$$

Penalize base movement with an iteration-dependent coefficient

Sequential linearization

$$\phi_c(x) = \phi(x_0) + \dot{\phi}(x_0)(x - x_0)$$

Solve sequential LP subproblems inside adaptive trust regions

Visualization of B* optimization process

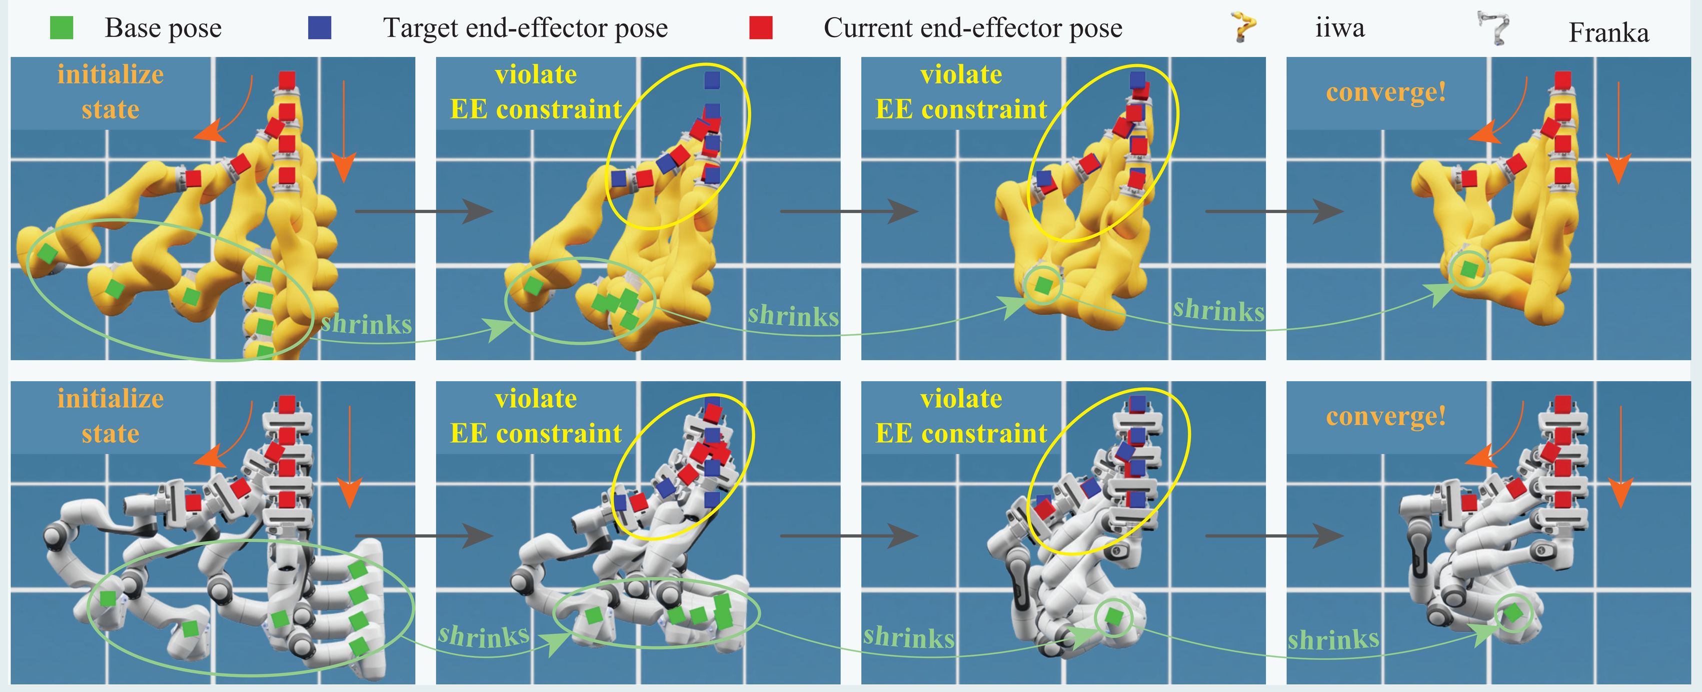

Validation

100%

valid solutions on all 2,400 generated tasks

10⁵x

lower base ATE than baselines

2,400

random trajectories, six complexity levels

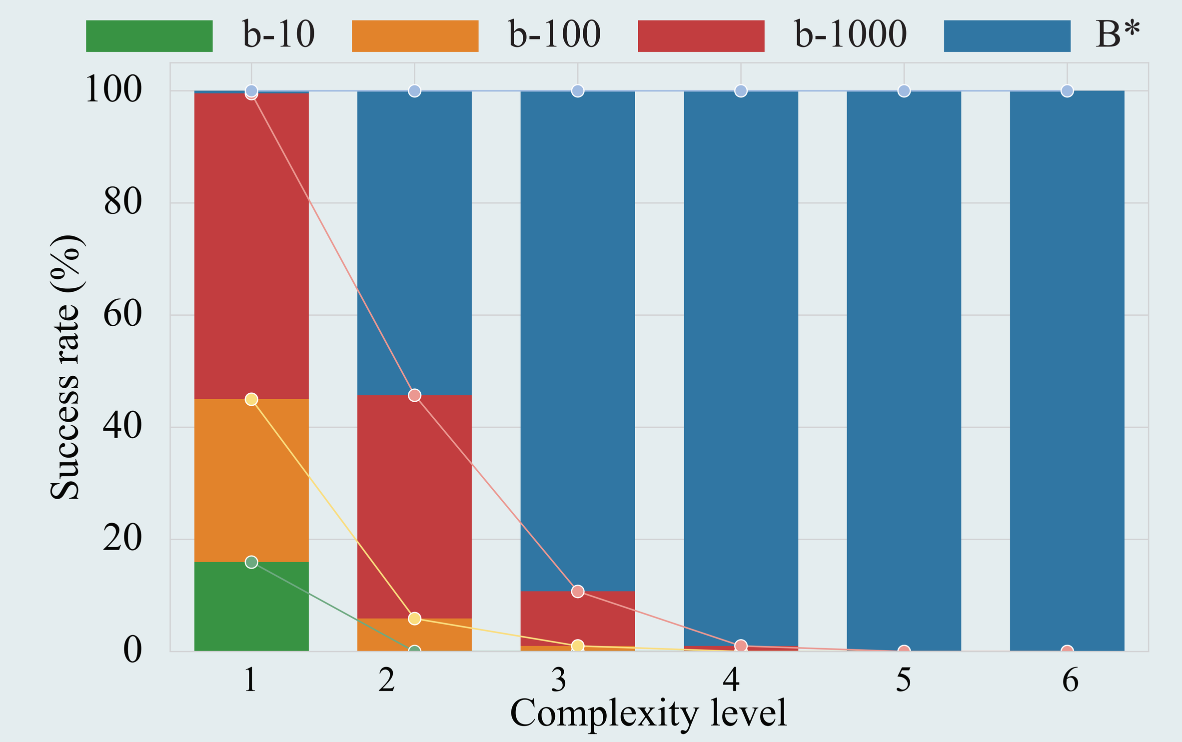

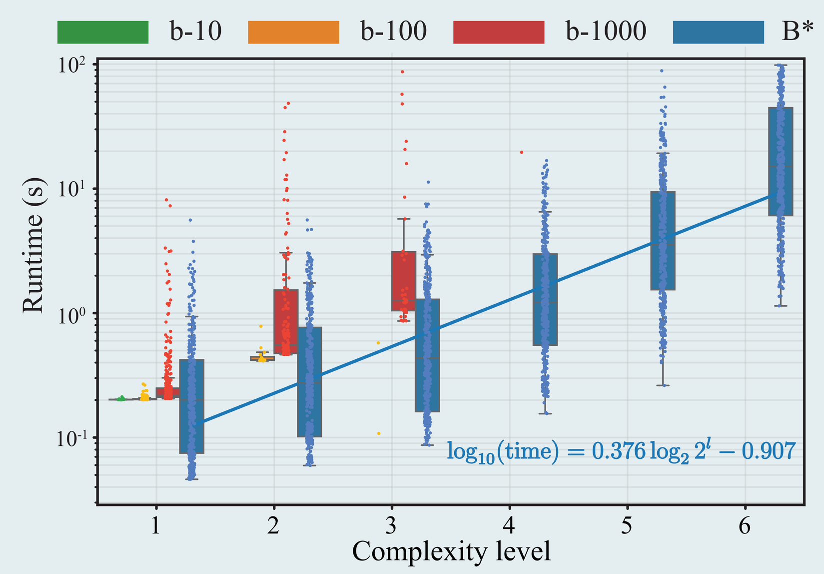

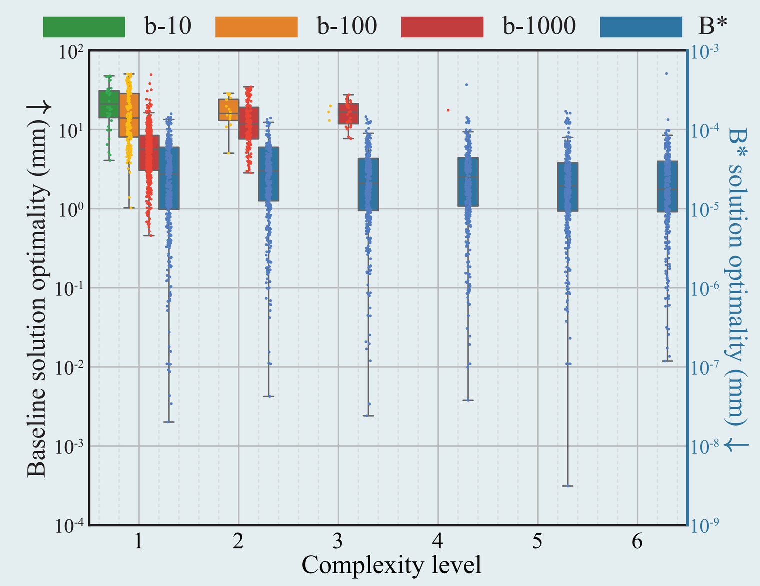

B* outperforms baselines, achieving a 100% success rate, higher precision, and superior runtime efficiency.

Real-world validation