Jacobian-Guided Active Learning for Gaussian Process-Based Inverse Kinematics

Shibao Yang[1] and Pengcheng Liu[1] and Nick Pears[1]

***Abstract*— Frequent replanning in dynamically changing environments often pushes robot manipulators towards singular configurations and joint limits, causing traditional inverse kinematics (IK) solvers to fail and hindering adaptability. We address this with an enhanced Gaussian Process IK (GP-IK) framework that uses a Jacobian-guided acquisition strategy for robust planning. This method adapts its exploration-exploitation balance in real-time based on local sensitivity and mechanical constraints, ensuring the planner can find reliable solutions even near manipulator limits. By enabling robust performance in challenging configurations, our approach allows for a tighter integration of perception and planning, fostering more adaptive and resilient robots, as demonstrated on a 7-DOF Franka robot.**

I. INTRODUCTION

For robots to operate effectively in dynamically changing environments, they must frequently replan their motions, a process that often drives manipulators towards singular configurations and joint limits [1], [2], [3]. While data-driven methods like Gaussian Process Inverse Kinematics (GP-IK) are promising for such adaptive planning [4], [5], [6], they falter precisely in these critical regions where the robot's Jacobian matrix becomes ill-conditioned [7], [2]. In these states, the IK solution becomes highly sensitive and numerically unstable, leading to unpredictable behaviour. Compounding this issue, conventional GP-IK exploration strategies are computationally inefficient, wasting resources by indiscriminately amassing data from many unfruitful attempts to find a solution. This highlights a critical need for a more intelligent exploration mechanism that can efficiently navigate these challenging configurations to ensure robust and reliable robot motion in response to real-world changes.

The integration of Jacobian-based sensitivity analysis with active learning principles presents a novel opportunity to address these fundamental challenges. We identify three key areas where this integration could substantially improve system performance: 1) Lack of local mechanical constraint awareness leads to inefficient exploration in challenging configurations. 2) Crucial balance between exploration and exploitation needs to account for both task-space objectives and configuration-space limitations, particularly when operating near joint limits or singular positions. 3) Adaptation of sampling strategies must consider both probabilistic uncertainty and mechanical feasibility to ensure reliable operation across diverse conditions.

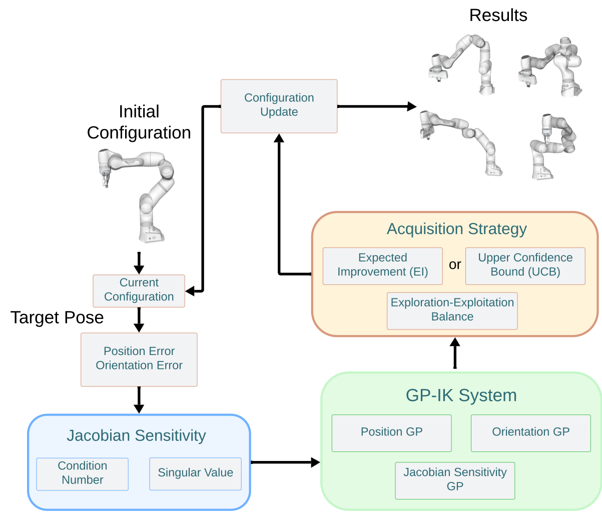

Fig. 1: The Jacobian-Guided GP-IK Framework that uses GP models to learn optimal movements, guided by an acquisition strategy that balances exploration and exploitation.

To address these challenges, we propose a novel Jacobian-Guided GP-IK framework that fundamentally advances the capability of active learning strategies in robotic manipulation. Our approach introduces a sophisticated Jacobian-guided acquisition mechanism that dynamically adapts the exploration-exploitation balance based on local sensitivity measures and mechanical constraints. This innovation is integrated with *Expected Improvement* (EI) and *Upper Confidence Bound* (UCB) acquisition functions [8], enabling the system to intelligently navigate complex configuration spaces while maintaining robustness near singular configurations and joint limits. By incorporating Jacobian-based sensitivity information into the acquisition process, our method maintains high solution quality while significantly reducing computational overhead in challenging scenarios.

We make the following contributions:

- A novel Jacobian-Guided GP-IK framework is proposed, integrating mechanical sensitivity analysis with an active learning strategy.
- A comprehensive analysis is conducted on how the Jacobian-guided exploration strategies affects the GP-IK solutions, revealing distinct advantages for different operational requirements.

This work was supported in part by the Engineering and Physical Sciences Research Council (EPSRC) under Grant EP/Y000773/1. (Corresponding author: Pengcheng Liu)

[1]Shibao Yang, Pengcheng Liu and Nick Pears are with Department of Computer Science, University of York, YO10 5GH, United Kingdom. E-mails: {shibao.yang, pengcheng.liu, nick.pears}@york.ac.uk

- An enhanced GP update mechanism has been developed, leveraging Jacobian information to ensure robust performance under various operational conditions.

II. METHOD

GP models have become a principled way to tackle IK under redundancy, uncertainty, and nonlinearity: multivariate Gaussian IK captures distributions over feasible joint configurations for smooth motion, and Gaussian-distributed damping within Damped Least Squares improves singularity robustness [4]. GP regression has also been used to learn forward/inverse kinematics, particularly in soft robots showcasing strong modelling of complex mappings [9]. At the same time, Jacobian-based optimisation remains the workhorse for IK and singularity analysis, yet effectively blending probabilistic GP predictions with joint constraints and task priorities is still challenging [10]. To improve data efficiency, active-learning GP frameworks have been explored for reward learning and kinematic calibration, leveraging Bayesian uncertainty to guide informative sampling [8]. Building on these ideas, our Jacobian-guided GP-IK incorporates GP-driven adaptive weighting and a sensitivity metric to balance exploration and exploitation while improving robustness and accuracy near joint limits and singularities [11].

A. Active Learning

Active learning with GPs iteratively fits a surrogate to observed data and selects new query points via an acquisition function that trades off exploitation (promising mean predictions) and exploration (high uncertainty), reducing samples needed to model the objective or locate optima [12]. In our GP-IK setting, we employ two complementary choices, EI to target configurations likely to improve accuracy and UCB to probe uncertain regions, supporting efficient search under redundancy and near difficult areas [13], [8]. However, these standard acquisitions rely solely on GP means and variances and are agnostic to kinematic structure, joint limits, and singularities; notably, they provide no mechanism to respond when the Jacobian becomes ill-conditioned near singular configurations.

B. Jacobian-Guided Acquisition

Jacobian-Guided Acquisition (Algorithm 1) presents a novel approach to IK that enhances solution quality and convergence reliability by incorporating configuration sensitivity analysis with machine learning techniques. The configuration quality assessment employs a weighted combination of fundamental manipulability metrics. The primary formulation integrates the condition number of the Jacobian matrix [14], manipulability measure [15], and kinematic isotropy [16]. The condition number is the ratio of the largest to smallest singular value of the Jacobian matrix which indicates numerical stability of the IK solution. The manipulability measure represents the volume of the manipulability ellipsoid and is calculated as the square root of the product of all singular values, which quantifies the robot's ability to generate velocities in arbitrary directions.

Algorithm 1 Jacobian-Guided Acquisition

Input: *target_pose*, *initial_config*
Output: *final_config*, *success*
Initialize:
  $q \leftarrow$ *initial_config*
  $GP \leftarrow$ *InitializeGaussianProcess*()
Main Loop:
while not converged **do**
  // 1. Compute current state and error
  $current\_pose \leftarrow$ *ForwardKinematics*(q)
  $error \leftarrow$ *ComputeError*(*target_pose*, $current\_pose$)
  if $error <$ *tolerance* **then**
    return (q, true)
  end if
  // 2. Sensitivity Analysis
  $J \leftarrow$ *ComputeJacobian*(q)
  sensitivity $\leftarrow$ *AnalyzeJacobianSensitivity*(J)
  UpdateGPModel(q, $sensitivity$)
  // 3. Compute and apply step
  $weights \leftarrow$ *GetWeights*(q, $sensitivity$)
  $\lambda \leftarrow sensitivity \times$ *damping_factor*
  $dq \leftarrow$ *DampedLeastSquares*(J, $weights$, λ, $error$)
  // 4. Update configuration
  $q \leftarrow q + dq$
  $q \leftarrow$ *ClampToJointLimits*(q)
end while
return (q, *CheckSuccess*(q, *target_pose*))

We employ GPR [17] to model the mapping between configuration space and sensitivity metrics. The covariance function uses a Matérn kernel with $\nu = 2.5$, chosen for its smoothness properties. This specific level of smoothness is particularly appropriate for modelling physical systems like robotic manipulators. The resulting GP prior assumes functions that are smooth enough to capture the underlying physics while still allowing for the non-linear behaviours approaching singularities.

The Jacobian sensitivity analysis provides crucial information about the robot configuration, which examines singular values and their ratios to detect potential problematic configurations. The system calculates a sensitivity metric that considers multiple factors: manipulability (how easily the end-effector can move in any direction), isotropy (uniformity of movement capability), and proximity to joint limits.

We employ GPR to learn from previous attempts and predict sensitivity at new configurations. This creates a probabilistic model of the workspace, allowing the system to anticipate challenging regions before encountering them. The GP model is continuously updated as new configurations are explored, improving its predictions over time.

The acquisition function combines the learned sensitivity model with traditional error metrics to guide the optimisation process. This function balances exploration and exploitation, adapting its behavior based on the current state and predicted sensitivity. When the system encounters high-sensitivity regions, it automatically adjusts its step size and damping parameters to maintain stability.

III. EXPERIMENTAL EVALUATION

A. *Performance Analysis of Active Learning in GP-IK*

This study evaluates three inverse kinematics methods: a standard Gaussian Process-based approach (GP-IK) and two variants enhanced with active learning acquisition functions—Expected Improvement (EI) and Upper Confidence Bound (UCB). The experiments compare their performance in terms of success rate, convergence speed, and accuracy, first in a general case and subsequently in a more complex application involving a 7-DOF robotic arm.

TABLE I: Comparison of GP-IK Methods.

Metrics	**Normal GP-IK**	**EI**	**UCB**
Success Rate (%)	57.14	57.14	57.14
Iteration	289.29	200.00	69.29
Position Error (m)	0.1557	0.0935	0.1425
Orientation Error (rad)	0.3945	0.2769	0.3228

TABLE II: Overall Performance Metrics

Metrics	Normal GP-IK	EI	UCB
Success Rate (%)	93.33	100.00	93.30
Iteration	39.07	200.00	19.07
Position Error (mm)	3.04e−6	2.87e−6	2.27e−5
Orientation Error (rad)	1.56e−3	6.79e−6	1.40e−3

The initial results, shown in Table I, established that while all methods had an identical success rate (57.14%), their efficiency and accuracy varied significantly. The UCB-enhanced method converged fastest (69.29 iterations), whereas the EI-based approach yielded the highest accuracy with the lowest position (0.0935 m) and orientation (0.2769 rad) errors. Both active learning strategies clearly outperformed the standard GP-IK baseline.

To further assess these methods in a more constrained and complex scenario, a second experiment was conducted on a 7-DOF Franka Emika robot, focusing on challenges like redundancy resolution and joint limit handling. As detailed in Table II, EI demonstrated superior reliability with a 100% success rate, proving robust even in difficult configurations near joint limits where other methods faltered. UCB was again the most computationally efficient, requiring only 19.07 iterations on average. In contrast, EI consistently reached its 200-iteration cap, indicating a distinct trade-off between solution reliability and speed. All methods achieved high positional accuracy, but EI's exhaustive search resulted in substantially better orientation accuracy (6.79e-6 rad).

Across both experiments, a clear performance trade-off emerges: UCB provides the fastest convergence, making it ideal for time-sensitive applications, while EI ensures the highest accuracy and reliability, suiting it for high-precision tasks like assembly. These findings suggest that a hybrid approach—leveraging UCB's rapid initial convergence before switching to EI for fine-grained refinement—could offer

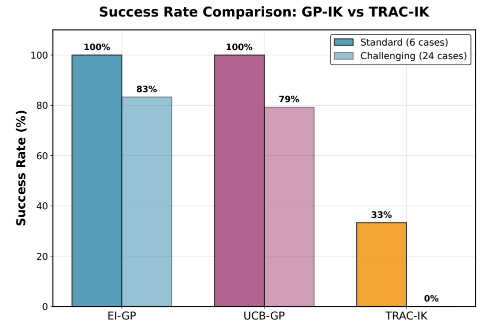

Fig. 2: The success rates within challenging cases.

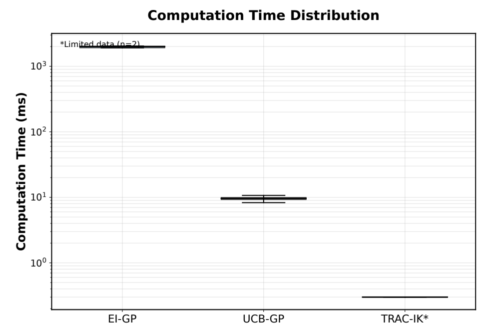

Fig. 3: Three computational regimes.

an optimal balance. Future work could explore adaptive strategies, potentially guided by reinforcement learning, to dynamically select the best acquisition function based on task requirements and convergence behaviour.

B. *Comparative Evaluation Against TRAC-IK*

We evaluated our probabilistic, GP-based IK solver against TRAC-IK [18], a state-of-the-art baseline that combines KDL's numerical IK with sequential quadratic programming using two complementary acquisition strategies: EI to emphasise exploitation and UCB to encourage exploration. Experiments on a 7-DOF manipulator used two datasets: six standard test cases (including home and ready) and a comprehensive set of 24 cases comprising 20 randomly generated configurations spanning the workspace plus four deliberately challenging scenarios (near-singularity, extended reach, elbow-up, and twisted). Success was defined as achieving a position error below 10^{-3} m and an orientation error below 10^{-2} rad. Particularly noteworthy is our method's performance on pathological test cases where traditional solvers struggle. For the near-singularity configuration at position $[0.554, 0.0, 0.521]$ m, both GP variants successfully converged while TRAC-IK failed after exhausting its iteration budget. Similarly, for the extended reach pose at $[-0.310, 0.0, 0.589]$ m, approaching the manipulator's workspace boundary, our probabilistic weighting mechanism adaptively adjusted the Jacobian regularisation to maintain numerical stability.

The results demonstrate that the learned GP models effectively capture the relationship between joint configurations and task-space errors, enabling robust solving even in regions

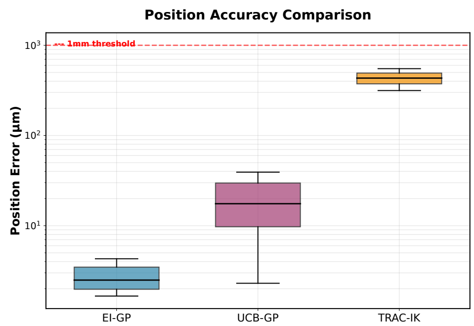

Fig. 4: The log-scale position errors.

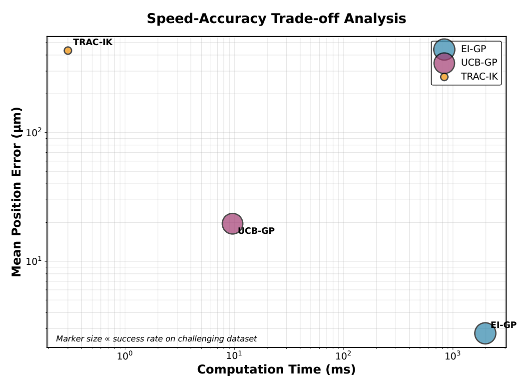

Fig. 5: The speed-accuracy trade-off.

where traditional geometric intuition fails. The adaptive weighting, informed by GP uncertainty estimates, provides a principled mechanism for handling ill-conditioned Jacobians without requiring manual tuning or special-case handling. The two acquisition functions demonstrated distinct behavioural characteristics aligned with their theoretical properties. The EI strategy exhibited conservative convergence behaviour with: Stricter convergence thresholds (0.99 for position, 0.98 for orientation); Slower learning rate ($\alpha = 0.5$) enabling fine-grained optimization; Multiple restart policy (5 attempts) ensuring global optimum discovery.

Conversely, the UCB strategy prioritised rapid convergence through: Relaxed thresholds (0.85 for position, 0.80 for orientation); Aggressive learning rate ($\alpha = 1.5$) with momentum term ($\beta = 0.3$); and Single-pass optimisation with early termination.

IV. Conclusions

In conclusion, we introduce a novel Jacobian-Guided GP-IK framework that provides a robust and computationally efficient foundation for robots operating in dynamically changing environments. By integrating mechanical sensitivity into the active learning process, our method reliably solves for motions near singular configurations and joint limits, a critical requirement for the frequent replanning needed in adaptive robotics. By offering a foundational motion planning layer that enables tighter integration between perception and planning, our work makes a direct contribution to creating more adaptive and responsive robots; when the planner can be trusted in challenging configurations, higher-level systems can more confidently react to real-time sensory input. Looking forward, the enhanced stability provided by our framework is essential for deploying robots in complex, interactive scenarios, such as human-robot collaboration and dynamic obstacle avoidance, paving the way for the next generation of truly adaptive and resilient robotic systems.

References

[1] H. Wang, Z. Zhou, X. Zhong, and Q. Chen, "Singular configuration analysis and singularity avoidance with application in an intelligent robotic manipulator," *Sensors*, vol. 22, no. 3, p. 1239, 2022.

[2] S. Yang, P. Liu, and N. Pears, "Enhancing learning from demonstration with dls-ik and promps," in *2024 29th International Conference on Automation and Computing (ICAC)*. IEEE, 2024, pp. 1–6.

[3] S. Liu and P. Liu, "Benchmarking and optimization of robot motion planning with motion planning pipeline," *The International Journal of Advanced Manufacturing Technology*, vol. 118, no. 3, pp. 949–961, 2022.

[4] X. Wang, J. Cao, X. Liu, L. Chen, and H. Hu, "An enhanced step-size gaussian damped least squares method based on machine learning for inverse kinematics of redundant robots," *IEEE Access*, vol. 8, pp. 68 057–68 067, 2020.

[5] R. A. Shyam, P. Lightbody, G. Das, P. Liu, S. Gomez-Gonzalez, and G. Neumann, "Improving local trajectory optimisation using probabilistic movement primitives," in *2019 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2019, pp. 2666–2671.

[6] S. Yang, P. Liu, and N. Pears, "Benchmarking of robot arm motion planning in cluttered environments," in *2023 28th International Conference on Automation and Computing (ICAC)*. IEEE, 2023, pp. 1–6.

[7] Y. Wang, N. Dehio, A. Tanguy, and A. Kheddar, "Impact-aware task-space quadratic-programming control," *The International Journal of Robotics Research*, vol. 42, no. 14, pp. 1265–1282, 2023.

[8] E. Daş and J. W. Burdick, "An active learning based robot kinematic calibration framework using gaussian processes," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 11 495–11 501.

[9] G. Fang, X. Wang, K. Wang, K.-H. Lee, J. D. Ho, H.-C. Fu, D. K. C. Fu, and K.-W. Kwok, "Vision-based online learning kinematic control for soft robots using local gaussian process regression," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1194–1201, 2019.

[10] K. Wen and C. Gosselin, "Kinematically redundant hybrid robots with simple singularity conditions and analytical inverse kinematic solutions," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3828–3835, 2019.

[11] S. Tang, K. Fujimoto, and I. Maruta, "Actively learning gaussian process dynamical systems through global and local explorations," *IEEE Access*, vol. 10, pp. 24 215–24 231, 2022.

[12] G. P. Kontoudis and M. Otte, "Adaptive exploration-exploitation active learning of gaussian processes," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 9448–9455.

[13] X. Yue, Y. Wen, J. H. Hunt, and J. Shi, "Active learning for gaussian process considering uncertainties with application to shape control of composite fuselage," *IEEE Transactions on Automation Science and Engineering*, vol. 18, no. 1, pp. 36–46, 2020.

[14] J.-P. Merlet, "Jacobian, manipulability, condition number, and accuracy of parallel robots," 2006.

[15] S. Patel and T. Sobh, "Manipulator performance measures-a comprehensive literature survey," *Journal of Intelligent & Robotic Systems*, vol. 77, pp. 547–570, 2015.

[16] J. Angeles and C. S. López-Cajún, "Kinematic isotropy and the conditioning index of serial robotic manipulators," *The international journal of Robotics Research*, vol. 11, no. 6, pp. 560–571, 1992.

[17] G. Giacomuzzo, R. Carli, D. Romeres, and A. Dalla Libera, "A black-box physics-informed estimator based on gaussian process regression for robot inverse dynamics identification," *IEEE Transactions on Robotics*, 2024.

[18] P. Beeson and B. Ames, "Trac-ik: An open-source library for improved solving of generic inverse kinematics," in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015, pp. 928–935.