Dual Control Reference Generation for Single-Shot Robotic Task Execution under Parametric Uncertainty

Victor Vantilborgh[1], Hrishikesh Sathyanarayan[2], Guillaume Crevecoeur[1], Ian Abraham[3] and Tom Lefebvre[1]
[1] Ghent University, [2] Yale University, [3] University of Sydney

***Abstract*— We consider the problem of *single-shot* robotic task execution under parametric uncertainty: a partially unknown robot must complete a task accurately on its first attempt, without a prior identification phase. We propose a generic framework for reference trajectory generation that deliberately excites only task-relevant parameters during execution, enabling concurrent online identification and control. We propose to solve a robust trajectory optimization problem, integrating adaptive closed-loop dynamics with an ad-hoc adaptation law, and use the reference as an artificial input. We show that our formulation naturally reasons about Fisher information. We validated the solution on a 7-DoF Franka FR3 pick-and-place task. Our results suggest that control-aware trajectories yield significantly improved task accuracy and identification compared to conventional approaches.**

I. Introduction

Many robotic applications require accurate task execution from the very first attempt. A manipulator cannot afford trial runs for every new payload; a drone with uncertain mass must complete its trajectory without a calibration flight; a legged robot on unfamiliar terrain cannot pause to estimate contact parameters before committing to a gait. In all cases, the robot must execute a task using a model-based controller whose model is wrong, and must fix that model on the fly.

Contrastingly, state-of-the-art approaches *do* separate this problem into two phases [1]–[3]. First, excite the system, often via optimal experiment design (OED) [4], [5], and identify a model. Second, execute the task using a model based control strategy (MPC, MBRL, etc.), recycling the model from the previous step. Such a dedicated exploration phase is impractical when environments change frequently or execution time windows are short. Moreover, conventional OED objectives are *task-agnostic*, meaning they maximize information about all parameters indiscriminately, wasting effort on directions with negligible task relevance.

In this work we propose an alternative strategy grounded in dual control [6], [7]. Rather than solving the intractable general dual control problem, our main idea is to design reference trajectories that excite *task-relevant* parameters whilst, *simultaneously*, executing the task. For the reference generation strategy we adopt a robust trajectory optimization approach, adopting the adaptive closed-loop dynamics and using the reference as an artificial input. We validate our approach on a 7-DoF Franka FR3 pick-and-place task under payload uncertainty in simulation.

II. Problem Formulation

Consider a parametric system $\dot{x} = f(x, u, \theta)$ with state $x \in \mathbb{R}^{n_x}$, control $u \in \mathbb{R}^{n_u}$, and unknown parameters $\theta \in \mathbb{R}^{n_\theta}$ with distribution $\theta \sim \mathcal{P}$ over the admissible set Θ. We consider an optimization problem of the form $\min_u \mathbb{E}_\theta[J(u)]$ where J equals a final state cost, m, at time time T and plus running cost, l, accumulated over time. The exact solution to this problem is given by an optimal feedback policy. This is a dual optimal control problem due to the expectation over θ [6]. The optimal policy simultaneously pursues task performance and information gain about uncertain parameters but only insofar as increasing parameter certainty would improve task performance [7].

To reduce the computational complexity without fully decoupling identification and control, we propose to prestructure the feedback policy to include an adaptation mechanism, i.e. $u = \pi(t, x, d, \hat{\theta})$ and $\dot{\hat{\theta}} = \rho(t, x, d, \hat{\theta})$ (e.g. [8]). Here $d \in \mathbb{R}^{n_d}$ represents a deterministic reference state trajectory (the design variable), π a control policy, and ρ an adaptation law. The framework is agnostic to the specific choice of π and ρ; the only requirement is that the closed-loop system admits a Lyapunov stability guarantee under parametric uncertainty,

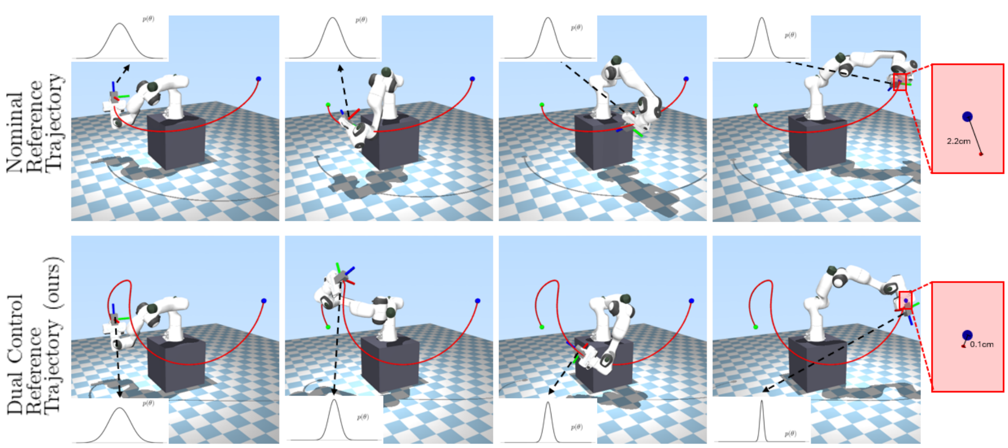

Fig. 1. Nominal reference trajectory (top) vs. dual-control trajectory from our framework (bottom). The robot transports a payload with unknown parameters θ from start (green) to target (blue). The nominal trajectory insufficiently excites the dynamics, leading to poor identification and large target deviation. The proposed trajectory actively excites task-relevant parameters during execution, improving estimates and enabling precise tracking.

a property satisfied by a broad class of adaptive controllers [8]–[10]. The reference d influences the closed-loop behavior both directly, by shaping the state trajectory through π, and indirectly, by determining the informativeness of the data available for parameter adaptation via ρ. We optimize d to minimize $\mathbb{E}_{\theta\sim\mathcal{P}}[J(u)]$, making the trajectory design informed about the controller and adaptation dynamics. This distinguishes our approach from prior work [11]–[14] that designs trajectories without knowledge of the tracking controller, guaranteeing neither tractability under mismatch nor informativeness for the specific adaptation scheme.

III. Methodology

A. Closed-Loop Uncertainty Propagation

Let $\xi = (x, \hat{\theta})$ denote the augmented state and $g(\xi, \theta)$ the closed-loop dynamics incorporating f, π, and ρ. Assuming $\theta \sim \mathcal{N}(\bar{\theta}, Q)$, a first-order Taylor expansion around the nominal trajectory $(\mu_\xi, \bar{\theta})$ yields the covariance dynamics

$$\dot{\Sigma}(t) = G(t)\Sigma(t) + \Sigma(t)G(t)^\top \tag{1}$$

where G contains Jacobians $A = \partial g/\partial \xi$ and $B = \partial g/\partial \theta$ evaluated along the nominal closed-loop trajectory. $\Sigma(t)$ captures how parameter uncertainty propagates into the state under the specific controller and adaptation law.

B. Control-Aware Trajectory Generation

We propose two methods:

1) Robust Optimization (RO): This method directly minimizes the expected task cost $\min_d \mathbb{E}_\theta[J(d)]$ subject to the closed-loop dynamics. A second-order expansion yields

$$\mathbb{E}_\theta[J] \approx \underbrace{J(\mu_x)}_{\text{nominal cost}} + \underbrace{\tfrac{1}{2}\operatorname{tr}(\Sigma_{xx}(T)M) + \tfrac{1}{2}\operatorname{tr}\Big(\int_0^T \Sigma_{\xi\xi} L \,\mathrm{d}t\Big)}_{\text{task-weighted uncertainty cost}} \tag{2}$$

where $M = \nabla_x^2 m$ and $L = \nabla_\xi^2 l$ are the curvatures of the terminal and running costs. This introduces a task-weighted uncertainty cost: parameter uncertainty is penalized strictly in directions that degrade task performance. Directions not affecting the cost receive zero weight.

2) Optimality Loss (OL): Alternatively, we build on the concept of *optimality loss* [11]. If the true θ were known, the optimal reference would be $d^\star(\theta) = \arg\min_d J(d)$. Using a prior estimate $\bar{\theta}$ instead incurs the optimality loss $\Delta(\bar{\theta}) = J(d^\star(\bar{\theta}), \theta) - J(d^\star(\theta), \theta) \geq 0$. A second-order expansion yields $\Delta \approx \frac{1}{2}\tilde{\theta}^\top D(\theta)\tilde{\theta}$, where D captures cost sensitivity to each parameter direction. Combining with the nominal task cost gives

$$\mathbb{E}_\theta[J_{\text{DUAL}}] = J(\mu_x) + \tfrac{1}{2}\operatorname{tr}\big(D(\bar{\theta})\,\Sigma_{\theta\theta}(T)\big) \tag{3}$$

This is a task-weighted A-optimality criterion: D directs exploration toward parameter dimensions whose uncertainty most degrades task performance.

C. Connection to Fisher Information and Stability

Unlike classical OED, which assumes an optimal estimator, both formulations propagate uncertainty through the *actual* closed-loop dynamics, including the Lyapunov-stable controller π and adaptation law ρ. The same Jacobians A and B that govern the covariance dynamics (1) also define the state sensitivity $\Psi = \partial\xi/\partial\theta$ via $\dot{\Psi} = A\Psi + B$, from which the Fisher information matrix is constructed [15]. The Cramér–Rao inequality, $\Sigma_{\theta\theta}(T) \succeq I^{-1}$, then establishes that both RO and OL reason about Fisher information not as an auxiliary objective, but as a direct consequence of closed-loop uncertainty propagation.

Furthermore, because the optimization simulates the full closed-loop system, trajectories leading to instability incur high expected cost and are naturally avoided. The generated references thus inherit the stability guarantees of the underlying adaptive controller, unlike task-agnostic OED which may violate them.

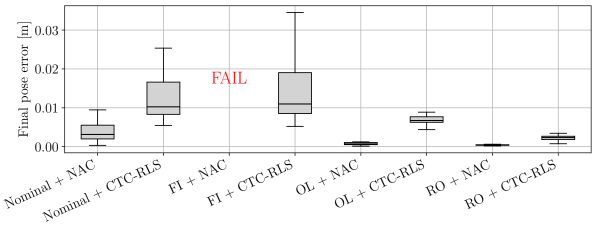

Fig. 2. Final pose error (20 sampled payloads). RO and OL achieve lower and more consistent errors across both controllers. The task-agnostic FIM trajectory is unstable for NAC.

IV. Experimental Results

We validate on a pick-and-place task with a 7-DoF Franka FR3 in MuJoCo. The robot moves a 2 kg cube (10 cm side, CoM offset $[0.04, 0.04, 0.10]$ m) from $p_0 = [-0.2, -0.5, 0.1]$ m to $p_T = [0.5, 0.5, 0.4]$ m, assessed over 20 sampled payloads (50% covariance). To demonstrate generality w.r.t. the control law, we evaluate on two distinct Lyapunov-stable adaptive controllers: a passivity-based controller with natural adaptation law [9] (NAC), performing updates on a Riemannian manifold for physical consistency, and a computed torque controller with recursive least squares [10] (CTC-RLS). Baselines: *nominal* (task-only) and *FIM-maximizing* (T-optimality, task-agnostic) trajectories.

Fig. 2 shows that RO and OL consistently reduce pose error mean and variability across both controllers. The nominal trajectory provides insufficient excitation for online identification, resulting in persistent mismatch. The FIM-based trajectory yields the best global parameter estimates, but its aggressive, task-agnostic excitation compromises the primary objective and is even unstable for the NAC. This instability is a direct consequence of designing excitation without accounting for the controller dynamics.

In contrast, RO and OL focus excitation on parameters that actually matter for this task: mass and center-of-mass offsets, which dominate the pick-and-place dynamics. They initially inject targeted oscillations for rapid identification, then converge conservatively as uncertainty decreases—yielding both improved identification and precise execution. Since the optimization propagates the actual closed-loop dynamics, it inherently respects the controller's Lyapunov stability, explaining why the proposed trajectories avoid the instability of the FIM baseline. RO provides a slight edge over OL; NAC outperforms CTC-RLS due to physically consistent Riemannian updates. A simplified variant without ρ also substantially reduced error for fixed-parameter controllers via passive trajectory robustness.

References

[1] B. Zhang, Z. Zhou, and R. Vasudevan. Provably-safe, online system identification. *arXiv preprint:2504.21486*, 2025.

[2] N. Sobanbabu, G. He, T. He, Y. Yang, and G. Shi. Sampling-based system identification with active exploration for legged robot sim2real learning. *arXiv preprint:2505.14266*, 2025.

[3] H. Sathyanarayan and I. Abraham. Behavior synthesis via contact-aware fisher information maximization. *arXiv preprint:2505.12214*, 2025.

[4] R. A. Fisher. Theory of statistical estimation. *Proceedings of the Cambridge Philosophical Society*, 22(5):700–725, 1925.

[5] V.V. Federov. Optimal experimental design. *Wires Computational Statistics*, 2(5):581–589, 2010.

[6] A. A. Feldbaum. Dual control theory. *Avtomatika i Telemekhanika*, 21(9):1240–1249, 1960.

[7] A. Mesbah. Stochastic model predictive control with active uncertainty learning: A survey on dual control. *Annual Reviews in Control*, 45:107–117, 2018.

[8] J. J. Slotine and W. Li. *Applied nonlinear control*, volume 199. Prentice hall Englewood Cliffs, 1991.

[9] T. Lee, J. Kwon, and F. C. Park. A natural adaptive control law for robot manipulators. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–9, 2018.

[10] Mark W Spong, Seth Hutchinson, and M Vidyasagar. Robot modeling and control. *John Wiley &amp*, 2020.

[11] B. Houska, D. Telen, F. Logist, M. Diehl, and J. Van Impe. An economic objective for the optimal experiment design of nonlinear dynamic processes. *Automatica*, 51:98–103, 2015.

[12] X. Feng and Y. Jiang. A scheme for simultaneous optimal tracking control and experiment design. *IEEE Access*, 8:25364–25371, 2020.

[13] R. Jackson, D. Varagnolo, and S. Knorn. A dual control approach to solve exploration versus exploitation trade-offs in the design of personalized physical exercise sessions. *IEEE Control Systems Letters*, 7:2383–2388, 2023.

[14] A. Srour, A. Franchi, P. R. Giordano, and M. Cognetti. Experimental validation of sensitivity-aware trajectory planning for a redundant robotic manipulator under payload uncertainty. *IEEE Robotics and Automation Letters*, 10(2):1561–1568, 2025.

[15] A. D. Wilson, J. A. Schultz, and T. D. Murphey. Trajectory synthesis for fisher information maximization. *IEEE Transactions on Robotics*, 30(6):1358–1370, 2014.