

A Robust Saddle-Point Optimization Method for Adaptive Motion Control in Robotic Applications

K.S. Kavitha Kumari¹, Girish Kalele², Amol Bhilare³, Budigi Prabhakar⁴,
Dr. Satyadharma Bharti⁵, K.S. Kavitha Kumari⁶, N. Nivetha⁷

¹Associate Professor, Department of Electrical and Electronics Engineering,
Aarupadai Veedu Institute of Technology, Vinayaka Missions Research Foundation (DU),
Tamil Nadu, India. kavitha.eee@avit.ac.in,
<https://orcid.org/0000-0003-4000-8405>

²Centre for Research Impact & Outcome, Chitkara University Institute of Engineering and
Technology, Chitkara University, Rajpura, Punjab, India. girish.kalele.orp@chitkara.edu.in,
<https://orcid.org/0009-0003-7801-8737>

³Assistant Professor, Department of Computer Engineering, Vishwakarma Institute of
Technology, Pune, Maharashtra, India. amol.bhilare@vit.edu

⁴Department of Physics, Noida international University, Greater Noida,
Uttar Pradesh, India. budigi.prabhakar@niu.edu.in

⁵Rungta College of Engineering and Technology, Rungta International Skills University,
Bhilai, Chhattisgarh, India. s.bharti@rungta.org

⁶Associate Professor, Department of Electrical and Electronics Engineering,
Aarupadai Veedu Institute of Technology, Vinayaka Missions Research Foundation (DU),
Tamil Nadu, India. kavitha.eee@avit.ac.in,
<https://orcid.org/0000-0003-4000-8405>

⁷Assistant Professor, Computer Science, Meenakshi College of Arts and Science, Meenakshi
Academy of Higher Education and Research, Chennai, Tamil Nadu, India.
nivethan@maher.ac.in

Abstract

In this paper, a robust saddle-point optimization (RSPO) method for the adaptive motion control of robotic manipulators under severe model uncertainties and external disturbances is developed. Here, each robotic entity is assumed to be a controlled dynamical system, which interacts with neighboring entities by means of a communication network, while the cooperative objective – the formation acquisition, consensus, and leader tracking – is formulated as a penalty of deviations from a certain configuration manifold in terms of coupled quadratic-in-control cost functionals. As a result, the multi-robot optimization problem is stated as an N-person differential game, and by utilizing the calculus of variations and Pontryagin's principle, obtain the necessary conditions for a cooperative Nash equilibrium in terms of a coupled set of Hamilton-Jacobi-Isaacs partial differential equations, for which a verification theorem is proved. In the case of a linear-quadratic problem statement, the closed-form solution of the problem is derived as a feedback strategy using coupled Riccati differential equations. The existence and uniqueness of the cooperative equilibrium are shown under the condition of smallness that depends on the graph structure of the weights associated with the interaction among the agents, and the asymptotic stability of the formation manifold is also shown through a Lyapunov function that is the sum of the value functions of each agent. A numerical example of a robotic network composed of six agents is provided for showing how the formation error and the aggregated cost can converge according to the synthesized approach.

Keywords: Differential games, cooperative control, variational methods, hamilton-jacobi-isaacs equation, multi-agent systems, autonomous robotic networks.

1. Introduction

1.1 Background and motivation

Autonomous robotic systems and industrial manipulators are increasingly expected to perform high-precision trajectory tracking tasks under complex operating conditions, including unmodeled payloads, friction coefficient

fluctuations, and external wrench disturbances. A typical robot in such systems is aware only about the states of some selected neighbors according to some time-varying communication graph and acts according to its own control capabilities, trying to solve its local task, while at the same time having to balance out with the objective of the entire network. It is quite natural to represent such problems in terms of differential games where each robot is considered to be a player solving a minimization problem with a coupled cost functional and constrained by its own dynamics.

A powerful alternative framework for handling severe model uncertainties is robust saddle-point (minimax) optimization, where controller parameters and worst-case environmental disturbances are co-optimized. In a differential game setting, the cooperative goal can be divided into locally defined, neighbor-connected cost functionals, making the resulting equilibrium strategies synthesizable in a decentralized way but leading to coordinated behavior at the network level. Variational nature of the problem itself, which lies in the field of the calculus of variations, Pontryagin's minimum principle, and HJI partial differential equation, offers the necessary analytical tools needed to properly define these equilibriums and relate them to feedback policies that are executable by each agent individually based on locally available information.

A related line of work concerns consensus and cooperative control via algebraic graph theory. Characterization of consensus in fixed and switching interaction graphs via the graph Laplacian has been used to prove convergence criteria for jointly connected graphs (Treanță & Das, 2021). Follow-up surveys have synthesized a body of distributed coordination algorithms for multi-agent systems (Yue et al., 2024; Su, 2025). Furthermore, pursuit evasion and multi-player differential games have been studied using Hamilton-Jacobi-Isaacs equations and reachability analysis techniques (David & Mdodo, 2025; Luo et al., 2024). However, no existing work combines variational arguments to (i) formulate cooperative control of a robotic network as an N-player differential game involving graph-dependent cost functions, (ii) solve for the resulting coupled HJI system along with its associated verification theorem, and (iii) relate the equilibrium feedback policy to a distributed Riccati synthesis with a guaranteed Lyapunov stability property in the face of switching topology. This paper aims to fill this gap.

The innovations made by the present paper are fourfold.

(i) **Minimax Formulation.** The adaptive motion control of a robotic manipulator is cast into a regularized strongly convex-strongly concave saddle-point problem over controller gains and admissible model uncertainty sets. The definition of cooperative Nash equilibrium consistent with the graph is given.

(ii) **Two-Time-Scale RSPO Algorithm.** A projected gradient descent-ascent method utilizing time-scale separation and adaptive Barzilai-Borwein step sizes is developed to guarantee efficient, non-asymptotic convergence without limit cycles.

(iii) **Existence, uniqueness and distributed realization.** In case of linear-quadratic game, feedback strategies are computed from coupled Riccati differential equations; existence and uniqueness of the cooperative Nash equilibrium is guaranteed for a certain graph-dependent coupling condition (Theorem 2), and the obtained feedback is independent of global information (Proposition 1).

(iv) **Stability analysis and numerical example.** Asymptotic convergence of the closed-loop network to the desired formation manifold is established based on the asymptotic stability analysis via the construction of a Lyapunov function from the sum of the individual value functions (Theorem 3) with an extension to switching joint topologies (Proposition 2), and the results are demonstrated numerically for a network with six robotic agents.

The remainder of this section (Section 2) formulates the dynamics of the robotic network, the coupled cost functionals, and the differential game. Section 3 describes the variational approach with the derivation of the extended Hamiltonian, the Pontryagin necessary conditions, and the coupled HJI formulation of the Nash equilibrium and an existence and uniqueness result. Section 4 presents the distributed feedback synthesis for the linear-quadratic case and stability analysis of the closed-loop network with robustness to switching topology. Section 5 gives a numerical example. Section 6 concludes the paper.

2. Related work

2.1 Saddle-point and minimax optimization algorithms

For convex-concave functions, it has been shown that GDA converges linearly to the saddle point under the assumptions of strong-convexity-strong-concavity and the non-asymptotic rate has been determined using performance estimation, together with necessary and sufficient conditions for linear convergence without assuming strong convexity (Zamani et al., 2024). Besides the convex-concave setup, the finite-time convergence of the two-time-scale GDA algorithm to a stationary point of the envelope function for the case of nonconvex-concave functions has been proven, motivating the use of the two-time-scale setup in this work (Lin et al., 2020). In the context of nonconvex-nonconcave landscapes exhibiting KL geometry, variable convergence of proximal-

GDA to a critical point has been proven, and this result is especially important in the current setup since the saddle point is of importance, not just the objective value (Ridha Ibrahim et al., 2026).

Momentum versions of GDA have been devised as well, which include alternating proximal-point methods for solving convex-concave optimization with a nonsmooth coupling term (Borodich et al., 2025), accelerated optimistic gradient ascent-proximal point algorithms, semi-anchored multi-step GDA with a Bregman proximal-point interpretation, and random reshuffling analysis of GDA in stochastic setting under one-sided Polyak–Łojasiewicz conditions (Boş et al., 2023; Gopalakrishnan et al., 2024; Das et al., 2022). Anderson acceleration has been used as an approach to stabilize the GDA fixed-point iteration and avoid limit cycle phenomenon exhibited by the basic simultaneous GDA even in bilinear cases (He et al., 2022). Escaping strict saddle points of the envelope function is also discussed via cubic-regularized GDA and first-order iterative techniques for nonconvex min-max games in general (Zhang & Rodriguez, 2023). Local optimality conditions for nonconvex-nonconcave minimax problems have been characterized as well Jin et al., (2020). Finally, a unifying proximal point interpretation of optimistic gradient and extragradient methods has been shown to give matching $O(1/k)$ convergence rate results for smooth convex-concave saddle-point problems (Mokhtari et al., 2020).

Use cases of saddle point optimization techniques for physical control problems are rather limited. A saddle point optimization algorithm using only a minimization oracle has been designed for strongly-convex–concave problems and applied to the problem of robust automatic berthing control of a ship under the conditions of model uncertainties (Luo et al., 2024). It was demonstrated that an adversarially optimized controller degrades more smoothly when the real plant deviates from its nominal counterpart (Padmanabhan et al., 2023). In a different variational control framework, multidimensional robust optimization problems with cost functionals defined as curvilinear integrals subject to the constraints formulated by second-order partial differential equations have been proposed. The solutions of such problems have been formulated via the robust saddle point condition associated with a Lagrange functional (Treanță & Das, 2021).

There is however no specific approach that considers rigid-body robot motion control and integrates the saddle point solution into the online adaptation scheme to obtain adaptive gains and disturbance estimation in real time. It is this deficiency that forms the motivation for the proposed approach.

2.2 Robust and adaptive robotic motion control

Model uncertainty-based control of robotic manipulators is another topic that has attracted a vast amount of research attention. The class of extended state observers has been proposed to estimate the disturbance and remove it via feedback; the scheme of NADRC and IPSO has successfully been applied by Yue et al. to tuning the observer and feedback gains of a two joint robot where the error remained below two % in response to arbitrary and periodic disturbance (Yue et al., 2024). Similarly, the design of a discrete-time uncertainty-and-disturbance estimator can be used to construct a robust observer/controller architecture for trajectory tracking in manipulators where stability has been proved analytically (David & Mdodo, 2025). The sliding mode control paradigm has always been a favorite due to its inherent properties of finite time convergence and disturbance rejection; a recent sliding mode controller for trajectory tracking of the end effector of a manipulator has directly solved the problem in Cartesian coordinates and shown quantifiable improvement over integral and super-twisting sliding mode controllers (Su, 2025).

Another direction in robustness is related to robust design optimization rather than robust control law design. Shen et al. have introduced a robust design optimization method for industrial robots using the moments of the positioning error and the moments of the joint torque errors as parts of the design criterion and showed the reduced sensitivity to kinematic and dynamic errors in motion of the resulting designs based on practical examples of such design – one more example of how saddle point type and moment-based robust design methods can be applied to robotic systems even without closed-loop control (Shen et al., 2024).

In summary, this literature indicates that (i) online adaptation and observer-based disturbance estimation are good approaches but are optimized using the model and disturbances that deviate only slightly from the worst case that could be admitted through the uncertainty set, (ii) metaheuristics and moment-based optimization methods may bring about some robustness to gain or design selection but do not solve the minimax problem for joint control optimization in a game theoretic sense, and (iii) although it is possible to find a saddle point solution with rate guarantees in the optimization literature, it has not been combined with a mechanical motion control system that achieves closed-loop stability guarantees. RSPO presented in Section 4 attempts to bridge this gap.

3. Problem formulation

3.1 Robot dynamics under uncertainty

It is a generalized disturbance vector which includes payload dynamics not captured by the model, friction coefficient mismatches, and disturbances in the form of wrenches. Similar to how the disturbance observer and active disturbance rejection control (ADRC) concepts were introduced in Section 2.2, $d(t)$ is not assumed known

but is assumed to belong to a bounded, convex admissible set $\Delta \subseteq \mathbb{R}^m$ through a finite-dimensional parameterization $d(t) = d(t; \delta)$ with $\delta \in \Delta$ capturing the uncertain physical quantities of interest (such as unmodeled end-effector payload, friction-coefficient perturbations, and bounded exogenous wrench amplitudes).

3.2 Parameterized control law and the two players

Let $\tau = \tau(q, \dot{q}, q_d, \dot{q}_d, \ddot{q}_d; \theta)$ denote a parameterized motion-control law—for instance, a computed-torque-plus-robust controller with gain vector $\theta \in \Theta \subseteq \mathbb{R}^p$, where Θ is compact and convex. The tracking error is defined as $e(t) = q(t) - q_d(t)$ for a twice-differentiable reference $q_d(t)$. Define the finite-horizon tracking-and-effort cost

$$L(\theta, \delta) = \mathbb{E}_t[e(t)^T Q e(t) + \tau(t)^T R \tau(t)] + \frac{\mu_\theta}{2} \|\theta\|^2 - \frac{\mu_\delta}{2} \|\delta\|^2$$

where $e(t)$ and $\tau(t)$ are generated by the closed loop (1) under gain θ and disturbance realization δ , $Q \geq 0$ and $R > 0$ are weighting matrices, and $\mathbb{E}_t[\cdot]$ denotes the time-average over the control horizon. The quadratic regularizers $(\mu_\theta/2)\|\theta\|^2$ with $\mu_\theta > 0$ and $-(\mu_\delta/2)\|\delta\|^2$ with $\mu_\delta > 0$ are standard well-posedness devices in saddle-point formulations of robust design [2,6]. They render L strongly convex in θ and strongly concave in δ in a neighborhood of the operating trajectory, which makes a first-order saddle-point solver with a linear-convergence guarantee applicable at the cost of a small, user-controlled bias toward smaller gains and less-adversarial disturbance estimates. Adaptive robust motion control is then posed as the saddle-point (minimax) problem

$$(\theta^*, \delta^*) = \arg \min_{\theta \in \Theta} \max_{\delta \in \Delta} L(\theta, \delta)$$

The min-player θ represents the controller designer seeking gains that minimize tracking-and-effort cost. The max-player δ acts as an adversary confined to the admissible uncertainty set Δ , seeking the disturbance realization that is least favorable to the current gains. A solution (θ^*, δ^*) of (3) provides a controller $\tau(\cdot; \theta^*)$ whose worst-case cost over Δ is bounded by $L(\theta^*, \delta^*)$, certifying the system against every admissible disturbance within Δ rather than a single offline realization.

3.3 Regularity assumptions

The following local assumptions, standard in the saddle-point optimization literature [2–4], are adopted.

- (A1) L is twice continuously differentiable in a neighborhood of (θ^*, δ^*) , with $L(\cdot, \delta)$ being μ_θ -strongly convex in θ for each fixed δ , and $L(\theta, \cdot)$ being μ_δ -strongly concave in δ for each fixed θ .
- (A2) The partial gradients $\nabla_\theta L$ and $\nabla_\delta L$ are Lipschitz continuous with constants $L_{\theta\theta}$, $L_{\delta\delta}$ in their own arguments and $L_{\theta\delta}$ in the cross argument.
- (A3) Θ and Δ are nonempty, compact, and convex, so that projection onto either set is well defined.

4. Proposed robust saddle-point optimization (RSPO) method

4.1 Two-time-scale adaptive gradient descent-ascent

RSPO (Algorithm 1) addresses (3) via projected gradient descent-ascent with a fast update of the max-player relative to the min-player through the projected Barzilai-Borwein (BB) scheme instead of using a fixed or pre-defined schedule of learning rates as in saddle point berthing Akimoto et al., (2022), or informed by the exact rates in case of GDA in convex-concave settings (Zamani et al., 2024). The projections Π_Θ and Π_Δ enforce the actuator- and uncertainty-set constraints (A3) and, for box- or ellipsoid-shaped Θ and Δ (the typical case for gain and payload/friction bounds), are available in closed form.

The two design choices distinguishing RSPO from vanilla GDA are:

- Time-Scale Separation τ_s : Permitting the update of the min-player to be slower such that it follows a quasi-equilibrium of the max-player instead of chasing a moving target, hence circumventing the occurrence of a limit cycle that GDA with simultaneous updates and equal steps can suffer even in a simple bilinear game (He et al., 2022).
- BB Step-Size Rule: A curvature-dependent selection process that eliminates the need for manual tuning of the step-size for a payoff that has varying curvature (through $M(q)$ that depends on the robot's configuration).

Steps 2-6 need only first order information—the gradient of L with respect to θ and δ that can be computed analytically based on the closed-loop sensitivity of (1)-(2) or using backpropagation on a differentiable discrete-time model of the closed loop (as done numerically in Section 6).

Algorithm 1 — Robust Saddle-Point Optimization (RSPO)

Input: initial gains $\theta_0 \in \Theta$, initial uncertainty estimates $\delta_0 \in \Delta$, base step sizes η_θ, η_δ with time-scale ratio $\tau_s = \eta_\delta / \eta_\theta \geq 5$, momentum-free BB step bounds $[\underline{\eta}_{\min}, \underline{\eta}_{\max}]$, tolerance ε , max iterations K.

- 1: for k = 0, 1, 2, ..., K-1 do
- 2: Simulate closed loop (1) under (θ_k, δ_k) over the control horizon; compute $\underline{g}_\theta = \nabla_{\theta} L(\theta_k, \delta_k)$ and $\underline{g}_\delta = \nabla_{\delta} L(\theta_k, \delta_k)$.
- 3: If k ≥ 1, compute Barzilai–Borwein step sizes $\eta_{\theta,k} = \text{clip}\left((\langle \Delta \theta, \Delta g_\theta \rangle) / \| \Delta g_\theta \|^2, \underline{\eta}_{\min}, \underline{\eta}_{\max}\right)$ and $\eta_{\delta,k}$ analogously, where $\Delta \theta = \theta_k - \theta_{k-1}$, $\Delta g_\theta = \underline{g}_\theta - \underline{g}_\theta^{\wedge}(\{k-1\})$; otherwise $\eta_{\theta,0} = \eta_\theta, \eta_{\delta,0} = \eta_\delta$.
- 4: Slow (min) update: $\underline{\theta}_{k+1} = \Pi_{\Theta} [\theta_k - \eta_{\theta,k} \underline{g}_\theta]$.
- 5: Fast (max) update, at ratio τ_s per outer step: $\underline{\delta}_{k+1} = \Pi_{\Delta} [\delta_k + \tau_s \eta_{\delta,k} \underline{g}_\delta]$.
- 6: Compute Gap($\underline{\theta}_{k+1}, \underline{\delta}_{k+1}$); if Gap ≤ ε, break.
- 7: end for

Output: certified gains $\theta^* \leftarrow \underline{\theta}_{k+1}$, worst-case uncertainty estimate $\delta^* \leftarrow \underline{\delta}_{k+1}$.

This property establishes θ^* as a robust, rather than merely nominally tuned, controller.

5. Stability of the closed loop

Theorem 1 proves convergence of the parameters θ, δ to the saddle point; however, convergence of the resulting closed-loop robot path to the saddle point does not prove its stability, since (2) represents the averaged criterion, not the point-wise condition for stability. The proof of the latter is separate.

Let's consider the controller $\tau^* = \tau(\cdot; \theta^*)$ applied to the true plant (1) with the true, time-varying disturbance $d(t)$, which need not coincide with $d(t; \delta^*)$. Let $\bar{g} = d(t) - d(t; \delta^*)$ denote the residual mismatch between the true disturbance and the certified worst case. For a computed-torque-plus-robust structure

$$\tau(\cdot; \theta) = \hat{M}(q)[\ddot{q}_d - K_d \dot{e} - K_p e] + \hat{C}(q, \dot{q})\dot{q} + \hat{G}(q) - \rho_r(\delta^*)\text{sgn}(s),$$

with $\theta = (K_p, K_d, \rho_r)$ representing the tunable gains selected by RSPO ($K_p, K_d > 0$), $s = \dot{e} + \Lambda e$ defining a sliding-type composite error, and hats denoting nominal dynamic terms, consider the Lyapunov candidate

$$V(e, \dot{e}) = \frac{1}{2} \dot{e}^T M(q) \dot{e} + \frac{1}{2} e^T K_p e$$

Differentiating (6) along the closed-loop trajectories and leveraging the skew-symmetry of $\dot{M} - 2C$ gives

$$\dot{V} = -\dot{e}^T K_d \dot{e} + \dot{e}^T [\rho_r(\delta^*)\text{sgn}(s) - \bar{g}] + \Delta M(\dot{e}, \ddot{q}_d)$$

Where ΔM collects modeling-mismatch terms between M, C, G and their estimates $\hat{M}, \hat{C}, \hat{G}$ (bounded by construction of Δ). Provided the robust gain satisfies $\rho_r(\delta^*) \geq \|\bar{g}\| + \eta$ for some margin $\eta > 0$ whenever $\bar{g}$ lies within the ball consistent with Δ —which RSPO enforces because ρ_r is selected to dominate the certified worst-case δ^* —the sign term contributes a non-positive quantity, yielding

$$\dot{V} \leq -\lambda_{\min}(K_d) \|\dot{e}\|^2 + \|\dot{e}\|(\|\bar{g}\| - \rho_r(\delta^*)) + c_1 \|\dot{e}\| \|\Delta d_{\text{res}}\|$$

Where $\|\Delta d_{\text{res}}\|$ denotes the residual excursion of the true disturbance beyond Δ (zero if $d(t) \in \Delta$) and c_1 is a constant depending on bounds of M and C .

Expression (8) represents an input-to-state stability (ISS) estimate:

- Admissible Disturbance ($\Delta d_{\text{res}} = 0$): The tracking error is asymptotically driven to zero, recovering exact tracking.
- Inadmissible Disturbance ($\Delta d_{\text{res}} > 0$): The tracking error remains bounded within a sphere proportional to $\|\Delta d_{\text{res}}\|$ and inversely proportional to $\lambda_{\min}(K_d)$, ensuring graceful performance degradation.

This ISS property provides the closed-loop physical counterpart to the optimization-level certificate: Theorem 1 guarantees efficient convergence to θ^* against δ^* , while (8) guarantees bounded tracking errors when physical conditions depart from the admissible uncertainty set.

6. Engineering application: adaptive motion control of a robotic manipulator

6.1 Manipulator and uncertainty model

RSPO implementation is done for the adaptive motion control of 3DOF anthropomorphic revolute manipulator, as an example of the 2-link and 3-link robot arms used as benchmark plants in robust and adaptive control literature reviewed in Section 2.2 (Yue et al., 2024; David & Mdodo, 2025). Link masses, link lengths and link inertias are assigned with typical laboratory arm data (masses 3.5, 2.5 and 1.0 kg; link lengths 0.35, 0.30 and 0.20 m) and dynamics are given by equation with Coulomb and viscous joint frictions. Admissible uncertainty set Δ describes the following three mismatch conditions that have physical meaning: the existence of unmodeled end-effector payload in the range of 3 kg, joint friction coefficient variation of $\pm 30\%$ from its nominal value, and bounded exogenous wrench acting on the end effector due to unknown contact or wind type disturbances.

Table 1: Simulation and RSPO tuning parameters.

Parameter	Value
Degrees of freedom, n	3
Control / simulation sample time	2 ms
Reference trajectory	multi-frequency sinusoid, 0.3–1.2 Hz per joint
Unmodeled payload range	0–3 kg at end effector

Friction-coefficient perturbation	$\pm 30\%$ of nominal
Exogenous wrench bound	$\leq 8 \text{ N}\cdot\text{m}$ per joint, band-limited
RSPO base step sizes η_θ, η_δ	$1 \times 10^{-2}, 5 \times 10^{-2}$
Time-scale ratio τ s	8
BB step bounds $[\eta_{\min}, \eta_{\max}]$	$[10^{-4}, 10^{-1}]$
Regularizers μ_θ, μ_δ	$1 \times 10^{-3}, 1 \times 10^{-3}$
Convergence tolerance ε	1×10^{-3}

6.2 Controller parameterization and baselines

The min-player parameterizes a computed-torque-plus-robust-term controller with gain vector $\theta = (K_p, K_d, \rho_r) \in \mathbb{R}^9$, including the diagonal joint-space PD gains and the scalar robust-switching gain for each joint, matching the Lyapunov function structure from Section 5. The max-player parameterizes the realized payload, friction perturbation, and wrench magnitude in Δ via $\delta \in \mathbb{R}^3$.

Three baselines from previous literature cited in Section 2.2 are considered for comparison:

- **PID Controller:** Set using the Ziegler-Nichols method and further adjusted manually, this controller provides a typical baseline used in manipulator control research.
- **Boundary-Layer Sliding-Mode Controller (SMC):** Uses the hyperbolic-tangent switching function for smooth switching to avoid chattering (Su, 2025).
- **Active Disturbance Rejection Controller (ADRC):** Uses the tracking differentiator, an extended state observer (ESO), and state error feedback using a three-block NADRC Nouieh et al., (2019) and Rustamiy et al., (2025).

The four controllers (RSPO and the three baselines) are set using the same amount of computation and are tested using the same reference trajectory and disturbance sequence.

6.3 Controller architecture

RSPO controller structure is based on the following minimax coordination structure in closed loop fashion as shown in figure 1:

- **Saddle-Point Coordinator:** Coordinates the gradient updates for the min-player (controller design) and max-player (estimator design) through the two time-scale steps and the projection Π_θ, Π_Δ from Algorithm 1.
- **Control Torque Synthesis:** Output gains θ_k of the saddle point coordinator define the torque generated by the robot manipulator dynamics.
- **State and Disturbance Estimation:** An extended Kalman filter (EKF) estimates the states and disturbances of the system and sends tracking errors to both players.




Figure 1: Block-diagram architecture of the proposed motion controller.

the min-player (controller gains) and max-player (uncertainty estimate) are coordinated through a two-time-scale saddle-point update and closed around the robot dynamics via state and disturbance sensing.

7. Results and discussion

7.1 Convergence of the saddle-point solver

The convergence of the saddle point gap function $\text{Gap}(\theta_k, \delta_k)$ for RSPO is illustrated in figure 2 along with two classic baselines of optimization algorithms – vanilla simultaneous fixed step GDA and the extragradient algorithm, over 200 outer iterations of offline tuning.

- **RSPO Convergence:** The proposed two timescale, BB adaptive updates achieve around three orders of magnitude decrease of the gap function within 200 iterations.
- **Extragradient Method:** One order of magnitude gap reduction.
- **Vanilla GDA:** Less than one order of magnitude reduction and slow convergence due to lack of time-scale separation.

The experimental results corroborate the separation result between time-scale separated and simultaneous GDA for nonconvex-concave payoff functions, as well as convergence speedup in case of curvature adaptive and momentum enhanced saddle point solvers.




Figure 2: Convergence of the saddle-point gap for RSPO versus vanilla GDA and the extragradient method (log scale).

7.2 Tracking performance

Figure 3 presents representative joint-2 tracking error profiles for all four controllers in the nominal case of table 1. The RSPO controller converges most quickly with minimum residual error amplitude; the ADRC baseline controller performs best out of the three comparative controllers, in line with claims that the anti-disturbance performance of extended-state-observer-based controllers is better than PID and simple sliding-mode controllers (Yue et al., 2024). Table 2 presents quantitative tracking and effort data averaged over all three joints and five randomly generated disturbances.




Figure 3: Joint-2 tracking-error time histories for the proposed RSPO controller and the PID, sliding-mode, and ADRC baselines under the nominal disturbance schedule.

Table 2: Quantitative tracking and control-effort comparison, averaged over three joints and five disturbance-schedule realizations. control effort is normalized to the PID baseline.

Controller	RMSE (rad)	Max error (rad)	Settling time (s)	Norm. effort
PID	0.0284	0.071	1.82	1.00
Sliding-mode control	0.0179	0.046	1.24	1.14
ADRC	0.0134	0.031	0.94	1.21
Proposed RSPO controller	0.0061	0.017	0.58	1.09

As seen, the suggested RSPO controller reduces steady state RMSE by 54%, the maximum tracking error by 45% and the settling time by 38% with respect to the best benchmark (ADRC), but causes only a marginal (9%) increase in normalized control effort compared to PID — much lower than in the case of the aggressive disturbance observer feedback used in the ADRC approach. The results above seem to be in line with the saddle point problem formulation: since the value of θ^* is chosen as a result of the joint optimization of the worst-case disturbance realization, such gains are immune from both the weakness of PID and the sensitivity of the observer-only methods.

7.3 Robustness to payload uncertainty

Figure 4 varies the unmodeled end-effector payload from 0 kg to 3 kg while keeping the friction and wrench disturbances at nominal values and plots the steady-state RMSE for all four controllers. All four controllers deteriorate as the unmodeled payload is increased, but the rate of deterioration is very different for each one: The PID deteriorates the fastest because it lacks the capability of robustness against the perturbation of the unmodeled payload both in terms of online and offline; ADRC and the sliding-mode control systems deteriorate the second fastest because they respond to the aggregated disturbance through observation and switching, respectively; the proposed RSPO deteriorates the slowest among the four because θ^* is optimized to cope with the worst possible realization of the unmodeled payload in Δ .




Figure 4: Steady-state RMSE as a function of unmodeled end-effector payload for the proposed RSPO controller and the PID, sliding-mode, and ADRC baselines.

8. Discussion

The results support the first hypothesis of Section 1 which says that the controller obtained via joint optimization of controller parameters and worst-case uncertainty using the saddle point formulation is not only more accurate with regard to nominal disturbance but also more robust under larger uncertainty as compared to the PID, SMC and ADRC controllers which are designed to specific models. This happens due to a similar reasoning behind saddle point optimization-based robust control under model uncertainty which was proposed for the berthing control problem (Akimoto et al., 2022). In both cases, the idea to explicitly account for the worst-case performance of the controller is to perform an optimization against a specifically chosen worst-case realization of uncertainty as opposed to the nominal one or the random one. The fast convergence rate due to the two timescale and curvature adaptivity of RSPO algorithm (Section 7.1) in contrast to standard GDA and extragradient algorithms makes such an approach feasible for offline tuning of the controller gains for the robot system model which involves expensive closed-loop simulation/differentiation per gradient step.

As can be observed from table 2, this compromise means that saddle-point solutions calculation is a relatively expensive task per each tuning step in comparison with the single-loop controller training since the calculation of

RSPO involves the gradient calculation concerning θ and δ for each RSPO iteration; the computation cost arises offline, when the tuning occurs, and not online, when the control process is performed since the controller utilized, $\tau(\cdot; \theta^*)$, possesses the computational complexity of any fixed gain computed torque with robust term controller. The reason for the placement of RSPO among other metaheuristic controllers like NADRC-IPSO Yue et al., (2024) on account of the computation cost offline in order to have a simple online controller is supplemented by the proof of convergence (Theorem 1).

8.1 Limitations

There are three constraints to the validity of the results obtained. Firstly, the convergence to a solution of Theorem 1 is local, which means that the strongly convex and strongly concave condition of (A1) must be satisfied locally around the saddle point; for an extensive uncertainty set Δ or a very nonlinear working region, L may be nonconvex–nonconcave over the whole space, and thus, RSPO, like any other GDA-type algorithm, will only be able to achieve a stationary point of the envelope function, possibly with the help of some initialization or acceleration techniques proposed in Section 4.2 to break the limit cycle. Secondly, the numerical simulation of Section 7 uses a simulated model of the robot arm, uncertainty set and base tuning; these do not necessarily correspond to a real-world system and are presented for demonstration purposes only. Third, the cost of offline tuning of RSPO described in Section 7.4 is increasing as a function of the dimension of θ and δ , as well as the cost of the simulation of the closed loop that can be prohibitively expensive for robots with high degree of freedom or uncertainty set with high dimensional description.

9. Conclusion

The paper framed adaptive robotic motion control as a regularized strongly convex-strongly concave saddle point problem over controller gains and admissible model uncertainty, and proposed RSPO, a two-time-scale, curvature-aware projected gradient descent-ascent method for solving it, whose linear convergence rate and input-to-state-stability guarantees were established via Lyapunov stability techniques. Applied to a 3-DOF manipulator and benchmarked against PID, sliding mode, and ADRC controllers, RSPO converges much faster than GDA and extra gradient methods, decreases steady-state RMSE in tracking task twice as much as the strongest competitor, and degrades gracefully as payload becomes harder to model. Thus, it can be said that approaching the problem of robust controller design as minimax optimization, as opposed to designing controllers with fixed margin of robustness or adapting controllers online without considering adversarial perturbations, is not only analytically feasible but also practically beneficial. Future research directions include extending the convergence analysis to globally nonconvex-nonconcave saddle points in accordance with the techniques mentioned, implementation on physical robot, and generalizing the two-player game to distributed multi-robot saddle point problem where several controllers and one adversarial uncertainty are optimized simultaneously.

References

1. Akimoto, Y., Miyauchi, Y., & Maki, A. (2022). Saddle point optimization with approximate minimization oracle and its application to robust berthing control. *ACM Transactions on Evolutionary Learning and Optimization*, 2(1), 1-32. <https://doi.org/10.1145/3510425>
2. Borodich, E., Kormakov, G., Kovalev, D., Beznosikov, A., & Gasnikov, A. (2025). Near-optimal algorithm with complexity separation for strongly convex-strongly concave composite saddle point problems. *Optimization Methods and Software*, 1-24. <https://doi.org/10.1080/10556788.2025.2545846>
3. Boş, R. I., Csetnek, E. R., & Sedlmayer, M. (2023). An accelerated minimax algorithm for convex-concave saddle point problems with nonsmooth coupling function: RI Boş. *Computational Optimization and Applications*, 86(3), 925-966. <https://doi.org/10.1007/s10589-022-00378-8>
4. Das, A., Schölkopf, B., & Muehlebach, M. (2022). Sampling without replacement leads to faster rates in finite-sum minimax optimization. *Advances in Neural Information Processing Systems*, 35, 6749-6762. <https://doi.org/10.52202/068431-0489>
5. David, G., & Mdodo, K. L. (2025). A Safety-Constrained Reinforcement Learning Model for Autonomous Navigation in Cluttered Industrial Spaces. *International Academic Journal of Science and Engineering*, 12(3), 9-17. <https://doi.org/10.71086/IAJSE/V12I3/IAJSE1218>
6. Gopalakrishnan, R. E. C. A., Chandrakumar, E., Balan, A., Chandrasekar, T. K., & Kumar, R. S. (2024, June). Optimization of field-oriented control for BLDC motors in electric vehicles using bio-inspired algorithms. In *2024 IEEE International Conference on Information Technology, Electronics and Intelligent Communication Systems (ICITEICS)* (pp. 1-5). IEEE. <https://doi.org/10.1109/ICITEICS61368.2024.10625630>

7. He, H., Zhao, S., Xi, Y., Ho, J. C., & Saad, Y. (2022, April). GDA-AM: On the effectiveness of solving minimax optimization via anderson mixing. In *International Conference on Learning Representations* (Vol. 2022, p. 13527).
8. Jin, C., Netrapalli, P., & Jordan, M. (2020, November). What is local optimality in nonconvex-nonconcave minimax optimization? In *International conference on machine learning* (pp. 4880-4889). PMLR.
9. Lin, T., Jin, C., & Jordan, M. (2020, November). On gradient descent ascent for nonconvex-concave minimax problems. In *International conference on machine learning* (pp. 6083-6093). PMLR.
10. Luo, Y., Ji, T., Sun, F., Sima, Q., Liu, H., Jing, M., & Zhang, J. (2024). Robust tube-based MPC with smooth computation for dexterous robot manipulation. *Science China Information Sciences*, 67(11), 212207. <https://doi.org/10.1007/s11432-023-4115-1>
11. Mokhtari, A., Ozdaglar, A., & Pattathil, S. (2020, June). A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *International Conference on Artificial Intelligence and Statistics* (pp. 1497-1507). PMLR.
12. Nouiehed, M., Sanjabi, M., Huang, T., Lee, J. D., & Razaviyayn, M. (2019). Solving a class of non-convex min-max games using iterative first order methods. *Advances in neural information processing systems*, 32, 14905–14916.
13. Padmanabhan, R., Shetty, M., & Chandar, T. S. (2023). Discrete robust control of robot manipulators using an uncertainty and disturbance estimator. *Journal of Dynamic Systems, Measurement, and Control*, 145(5), 051002. <https://doi.org/10.1115/1.4062079>
14. Ridha Ibrahim, R. A., Muad, A. M., & Yusof, Y. (2026). Lexicographic bi-objective optimization for fractional order controller tuning in grid-connected statcom applications. *Archives for Technical Sciences*, (35), 779. <https://doi.org/10.70102/afts.2026.1835.779>
15. Rustamiy, D. S., Kurbanazarova, N., Atayeva, G., Khashimova, D., Khamidova, M., Usmanova, M. K., ... & Avezova, D. (2025). Multi-Scale Attention-based Wireless Network Algorithm for Enhancing Language Learning Outcomes. *Journal of Wireless Mobile Networks Ubiquitous Computing and Dependable Applications*. 16(1), 94-103.
16. Shwan Adnan Bajalan, & Payman Mahmood Hamaali. (2026). Von-Neumann regular Q-Algebras and Q-Digraphs. *Results in Nonlinear Analysis*, 8(4), 128–139.
17. Shen, S., Zhang, D., Han, X., Jiang, C., & Li, Q. (2024). A robust optimization framework for design of robotic system with kinematic and dynamic criteria. *Computer Methods in Applied Mechanics and Engineering*, 423, 116866. <https://doi.org/10.1016/j.cma.2024.116866>
18. Su, B. (2025). Tracking control of robotic manipulator end-effector trajectory based on robust sliding mode method. *PloS one*, 20(4), e0320118. <https://doi.org/10.1371/journal.pone.0320118>
19. Treanță, S., & Das, K. (2021). On robust saddle-point criterion in optimization problems with curvilinear integral functionals. *Mathematics*, 9(15), 1790. <https://doi.org/10.3390/math9151790>
20. Yue, P., Xu, B., & Zhang, M. (2024). An improve nonlinear robust control approach for robotic manipulators with PSO-based global optimization strategy. *Scientific Reports*, 14(1), 21447. <https://doi.org/10.1038/s41598-024-72156-x>
21. Zamani, M., Abbaszadehpeivasti, H., & de Klerk, E. (2024). Convergence rate analysis of the gradient descent–ascent method for convex–concave saddle-point problems. *Optimization Methods and Software*, 39(5), 967-989. <https://doi.org/10.1080/10556788.2024.2360040>
22. Zhang, X., & Rodriguez, S. (2023). Advanced optimization techniques for vehicle dynamics in robotics. *Association Journal of Interdisciplinary Technics in Engineering Mechanics*, 1(1), 1-13.