

Novelty Assessment Report

Paper: Differentiable Model Predictive Control on the GPU

PDF URL: <https://openreview.net/pdf?id=bFYfV6c9zu>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-27

Abstract

Differentiable model predictive control (MPC) offers a powerful framework for combining learning and control. However, its adoption has been limited by the inherently sequential nature of traditional optimization algorithms, which are challenging to parallelize on modern computing hardware like GPUs. In this work, we tackle this bottleneck by introducing a GPU-accelerated differentiable optimization tool for MPC. This solver leverages sequential quadratic programming and a custom preconditioned conjugate gradient (PCG) routine with tridiagonal preconditioning to exploit the problem's structure and enable efficient parallelization. We demonstrate substantial speedups over CPU- and GPU-based baselines, significantly improving upon state-of-the-art training times on benchmark reinforcement learning and imitation learning tasks. Finally, we showcase the method on the challenging task of reinforcement learning for driving at the limits of handling, where it enables robust drifting of a Toyota Supra through water puddles.

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **GPU-Accelerated Differentiable Model Predictive Control**

A total of **36 papers** were analyzed and organized into a taxonomy with **28 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Gradient-Based Trajectory Optimization Methods**
- **Sampling-Based and Hybrid Methods**
- **Learning-Integrated MPC Frameworks**
- **Domain-Specific MPC Applications**
- **Specialized Optimization and Solver Techniques**
- **Differentiable Simulation and Surrogate Modeling**
- **Hardware Design and VLSI Optimization**
- **Conceptual and Survey Contributions**

Complete Taxonomy Tree

- GPU-Accelerated Differentiable Model Predictive Control Survey Taxonomy
- Gradient-Based Trajectory Optimization Methods
 - Iterative Linear System Solvers
 - Preconditioned Conjugate Gradient Approaches ★ (2 papers)
 - [0] Differentiable Model Predictive Control on the GPU (Anon et al., 2026) [View paper](#)
 - [2] Mpcgpu: Real-time nonlinear model predictive control through preconditioned conjugate gradient on the gpu (Emre Adabag, 2024) [View paper](#)
 - Condensed-Space Interior-Point Methods (1 papers)
 - [26] Exploiting GPU/SIMD Architectures for Solving Linear-Quadratic MPC Problems* (David Cole, 2022) [View paper](#)
 - Parallel Shooting and Direct Methods
 - Primal-Dual and KKT System Solvers (1 papers)
 - [7] Primal-Dual iLQR for GPU-Accelerated Learning and Control in Legged Robots (Amatucci, 2025) [View paper](#)
 - Sequential Quadratic Programming Methods (1 papers)
 - [19] Multilevel parallel GPU implementation of SQP solvers for Nonlinear MPC (P. C. N. Verheijen, 2024) [View paper](#)
 - Batched and Scalable Trajectory Optimization (2 papers)
 - [1] GATO: GPU-Accelerated and Batched Trajectory Optimization for Scalable Edge Model Predictive Control (Plancher, 2025) [View paper](#)
 - [4] GPU Acceleration for Real-Time, Whole-Body, Nonlinear Model Predictive Control (Plancher, 2022) [View paper](#)
 - Input Parameterization and Dimensionality Reduction (1 papers)
 - [24] Parameterized and GPU-Parallelized Real-Time Model Predictive Control for High Degree of Freedom Robots (Killpack, 2020) [View paper](#)
- Sampling-Based and Hybrid Methods
 - Path Integral and Stochastic Control (2 papers)
 - [10] On the Benefits of GPU Sample-Based Stochastic Predictive Controllers for Legged Locomotion (Giulio Turrisi, 2024) [View paper](#)
 - [13] Real-time whole-body control of legged robots with model-predictive path integral control (Juan Alvarez-Padilla, 2025) [View paper](#)
 - Hybrid Gradient-Sampling Approaches (2 papers)

- [14] Feedback-MPPI: Fast Sampling-Based MPC via Rollout Differentiation--Adios low-level controllers (Tommaso Belvedere, 2025) [View paper](#)
- [17] GPU-Enabled Parallel Trajectory Optimization Framework for Safe Motion Planning of Autonomous Vehicles (Yeongseok Lee, 2024) [View paper](#)
- Constrained Sampling Methods (1 papers)
- [12] Constrained sampling-based MPC using path integral for collision-free robot manipulation (Xingfang Wang, 2025) [View paper](#)
- Sample-Based Newton-Like Methods (1 papers)
- [21] Improving the Accuracy of Sample-based Model Predictive Control via Sample-based Newton-like method with Approximated Hessian and Gradient by Quadratic $\hat{\alpha}_i$ (S Nakatani, 2021) [View paper](#)
- Learning-Integrated MPC Frameworks
 - Differentiable MPC for End-to-End Learning
 - Actor-Critic and Reinforcement Learning Integration (2 papers)
 - [11] Actor-critic model predictive control: Differentiable optimization meets reinforcement learning (Romero, 2024) [View paper](#)
 - [33] Dynamic Grasping Based on Reinforcement Learning and Differentiable MPC (Jianzhi Lyu, 2025) [View paper](#)
 - Differentiable Optimization Layers (1 papers)
 - [15] Differentiable optimization-based modeling for machine learning (Amos, 2019) [View paper](#)
 - Kinematics-Aware and Task-Specific Learning (2 papers)
 - [8] iKap: Kinematics-aware planning with imperative learning (Qihang Li, 2025) [View paper](#)
 - [25] Differentiable GPU-Parallelized Task and Motion Planning (William Shen, 2024) [View paper](#)
 - Residual and Hybrid Learning-Control Architectures (1 papers)
 - [22] Residual MPC: Blending Reinforcement Learning with GPU-Parallelized Model Predictive Control (Lee Ho Jae, 2025) [View paper](#)
 - Adaptive and Iterative Learning MPC (2 papers)
 - [29] SIT-LMPC: Safe Information-Theoretic Learning Model Predictive Control for Iterative Tasks (Zirui Zang, 2026) [View paper](#)
 - [30] Adaptive DiffTune MPC for Marine Robots: An OceanSim-Based Study (ZH Ismail, 2025) [View paper](#)
 - Neural Network-Based Model Learning (1 papers)
 - [34] Learning Robot Control: From Reinforcement Learning to Differentiable Simulation (Song, 2024) [View paper](#)
- Domain-Specific MPC Applications
 - Legged Robot Locomotion (1 papers)
 - [3] Physics-guided neural network and GPU-accelerated nonlinear model predictive control for quadcopter (Seong Hyeon Hong, 2022) [View paper](#)
 - Autonomous Vehicle Control (1 papers)
 - [9] Model predictive control for aggressive driving over uneven terrain (Tyler Han, 2023) [View paper](#)
 - Robot Manipulation and Grasping (1 papers)
 - [18] Neural Joint Space Implicit Signed Distance Functions for Reactive Robot Manipulator Control (Mikhail Koptev, 2023) [View paper](#)
 - Soft and Continuum Robots (2 papers)
 - [16] Adaptive Model-Predictive Control of a Soft Continuum Robot Using a Physics-Informed Neural Network Based on Cosserat Rod Theory (Habich, 2025) [View paper](#)
 - [35] DIFFERENTIABLE PARALLEL SIMULATION FOR PLANNING AS DESIGN FOR SOFT GROWING ROBOTS (Gao, 2025) [View paper](#)
 - Process Control and Industrial Systems (2 papers)
 - [27] GPU-based model predictive control of nonlinear parabolic partial differential equations system and its application in continuous casting (Yuan Wang, 2019) [View paper](#)
 - [32] Differentiable Predictive Control of Permanent Magnet Synchronous Motors (Ali Abdelwanis, 2025) [View paper](#)
- Specialized Optimization and Solver Techniques
 - Neural Network-Reformulated Solvers (1 papers)
 - [23] ReLU-QP: A GPU-Accelerated Quadratic Programming Solver for Model-Predictive Control (Arun L. Bishop, 2023) [View paper](#)
 - Dual Gradient-Projection Methods (1 papers)
 - [31] GPU-Accelerated Dual Gradient-Projection Algorithm for Embedded Linear Model Predictive Control (Luke Nukulaj, 2025) [View paper](#)
- Differentiable Simulation and Surrogate Modeling
 - Computational Fluid Dynamics and FSI (1 papers)
 - [5] Diff-FlowFSI: A GPU-Optimized Differentiable CFD Platform for High-Fidelity Turbulence and FSI Simulations (Fan, 2025) [View paper](#)
 - Neural Operator-Based Surrogate Models (1 papers)
 - [28] Channel-assisted Fourier neural operator for high-fidelity far-wake prediction of a wind turbine under yawed conditions (Eunchan Lee, 2025) [View paper](#)
 - Cellular Automata and Discrete Event Simulation (1 papers)
 - [20] PyTorchFire: A GPU-accelerated wildfire simulator with Differentiable Cellular Automata (Xia Zeyu, 2025) [View paper](#)
- Hardware Design and VLSI Optimization (1 papers)
 - [6] LEGO-Size: LLM-Enhanced GPU-Optimized Signoff-Accurate Differentiable VLSI Gate Sizing in Advanced Nodes (Yi-Chen Lu, 2025) [View paper](#)
- Conceptual and Survey Contributions (1 papers)
 - [36] Prior Work: Hardware Acceleration for Reliable Realtime Optimal Control (Plancher, n.d.) [View paper](#)

Narrative

Core task: GPU-accelerated differentiable model predictive control. The field organizes around several complementary directions. Gradient-Based Trajectory Optimization Methods focus on iterative solvers and direct differentiation through dynamics, often leveraging automatic differentiation frameworks to compute sensitivities efficiently. Sampling-Based and Hybrid Methods explore stochastic rollouts and path-integral formulations that naturally parallelize on GPUs, trading off gradient precision for robustness in high-dimensional or uncertain settings. Learning-Integrated MPC Frameworks blend neural network components with classical control loops, using differentiable MPC layers to enable end-to-end training. Domain-Specific MPC Applications tailor these techniques to robotics, autonomous driving, and soft-body systems, while Specialized Optimization and Solver Techniques develop custom algorithms for constrained quadratic programs and nonlinear solvers. Differentiable Simulation and Surrogate Modeling emphasizes physics engines

and reduced-order models that support backpropagation, and Hardware Design contributions address VLSI and embedded implementations.

A particularly active line of work targets real-time control for legged robots and manipulators, where GPU parallelism enables rapid trajectory recomputation at high frequencies—examples include Whole Body MPC GPU[4] and GPU Stochastic Predictive Legged[10]. Another contrasting theme is the integration of learned components, as seen in Physics Neural MPC Quadcopter[3] and Adaptive DiffTune MPC[30], which adapt model parameters online. The original paper, Differentiable MPC GPU[0], sits squarely within the gradient-based optimization branch, specifically employing preconditioned conjugate gradient solvers to handle large-scale linear systems arising from sequential quadratic programming. This places it close to Mpcgpu[2], which similarly exploits GPU-accelerated iterative methods, but Differentiable MPC GPU[0] emphasizes end-to-end differentiability to support learning pipelines. Compared to sampling-heavy approaches like Feedback-MPPI[14], it trades stochastic exploration for deterministic gradient descent, reflecting a broader tension between computational efficiency and algorithmic flexibility across the taxonomy.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Mpcgpu: Real-time nonlinear model predictive control through preconditioned conjugate gradient on the gpu

Authors: Emre Adabag, Miloni Atal, William Gerard, Brian Plancher | **Year/Venue:** 2024 | **URL:** [View paper](#)

Abstract

Nonlinear Model Predictive Control (NMPC) is a state-of-the-art approach for locomotion and manipulation which leverages trajectory optimization at each control step. While the performance of this approach is computationally bounded, implementations of direct trajectory optimization that use iterative methods to solve the underlying moderately large and sparse linear systems, are a natural fit for parallel hardware acceleration. In this work, we introduce MPCGPU, a GPU-accelerated, real-time NMPC s...

Relationship Analysis

Both papers belong to the Preconditioned Conjugate Gradient Approaches category, implementing PCG solvers with custom preconditioning strategies for GPU-accelerated MPC. They share the core approach of using tridiagonal (symmetric stair) preconditioning to exploit the sparse structure of trajectory optimization problems for efficient GPU parallelization. The key difference is that the original paper (DiffMPC) focuses on differentiable optimization with automatic differentiation through the solver for learning applications (RL/IL), while the candidate paper (MPCGPU) focuses on real-time non-differentiable NMPC for direct control applications without gradient computation through the solver.

Contributions Analysis

Overall novelty summary. The paper introduces a GPU-accelerated differentiable MPC solver combining sequential quadratic programming with a custom preconditioned conjugate gradient routine featuring tridiagonal preconditioning. It resides in the 'Preconditioned Conjugate Gradient Approaches' leaf, which contains only two papers total (including this one and one sibling). This is a notably sparse research direction within the broader taxonomy of 36 papers across 28 leaf nodes, suggesting the specific combination of GPU acceleration, differentiability, and PCG-based iterative solvers remains relatively underexplored compared to sampling-based or direct shooting methods.

The taxonomy reveals several neighboring directions: the sibling 'Condensed-Space Interior-Point Methods' leaf focuses on eliminating state variables rather than iterative PCG solvers, while the parallel 'Parallel Shooting and Direct Methods' branch employs primal-dual KKT solvers or multilevel SQP without emphasizing PCG preconditioning. Nearby 'Differentiable MPC for End-to-End Learning' nodes integrate MPC with actor-critic frameworks but do not necessarily prioritize iterative linear system solvers. The paper's position bridges gradient-based optimization infrastructure with learning-integrated applications, sitting at the intersection of solver design and end-to-end training pipelines.

Among 16 candidates examined across three contributions, none clearly refute the core claims. The GPU-accelerated differentiable optimization tool examined 7 candidates with 0 refutations; the tridiagonal PCG routine examined 1 candidate with 0 refutations; and the robust drifting application examined 8 candidates with 0 refutations. This limited search scope—focused on top-K semantic matches—suggests that within the examined literature, the specific combination of SQP, custom PCG preconditioning, and differentiability for learning tasks appears distinct. However, the small candidate pool means the analysis does not capture exhaustive prior work in iterative MPC solvers or GPU optimization.

Given the sparse taxonomy leaf and absence of refutations among 16 examined candidates, the work appears to occupy a relatively novel niche within GPU-accelerated MPC. The emphasis on tridiagonal preconditioning for differentiable SQP distinguishes it from both sampling-heavy methods and direct KKT approaches. Nonetheless, the limited search scope and small sibling set mean this assessment reflects positioning within a focused literature subset rather than comprehensive field coverage.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: GPU-accelerated differentiable optimization tool for MPC

Description: The authors introduce DiffMPC, a differentiable solver for model predictive control that runs efficiently on GPUs. It uses sequential quadratic programming combined with a custom preconditioned conjugate gradient routine that exploits the sparse-in-time structure of optimal control problems to enable parallelization over time steps.

This contribution was assessed against **7 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Multilevel parallel GPU implementation of SQP solvers for Nonlinear MPC

URL: [View paper](#)

Brief Assessment

Multilevel Parallel SQP[19] focuses on parallel GPU implementation of SQP solvers for NMPC without differentiability capabilities, whereas the original paper introduces a differentiable solver with custom PCG routines and implicit differentiation for learning applications.

2. Towards safe and tractable Gaussian process-based MPC: Efficient sampling within a sequential quadratic programming framework

URL: [View paper](#)

Brief Assessment

Safe Gaussian Process MPC[48] focuses on Gaussian process-based uncertainty propagation using sequential quadratic programming for robust MPC, not on GPU-accelerated differentiable optimization. The candidate uses SQP for sampling-based GP-MPC with different objectives (safety guarantees under epistemic uncertainty) rather than general differentiable MPC optimization on GPUs.

3. Physics-guided neural network and GPU-accelerated nonlinear model predictive control for quadcopter

URL: [View paper](#)

Brief Assessment

Physics Neural MPC Quadcopter[3] focuses on physics-guided neural networks for quadcopter control, not on developing GPU-accelerated differentiable optimization solvers for MPC using sequential quadratic programming and preconditioned conjugate gradient methods.

4. ReLU-QP: A GPU-Accelerated Quadratic Programming Solver for Model-Predictive Control

URL: [View paper](#)

Brief Assessment

ReLU-QP[23] focuses on GPU-accelerated quadratic programming using ADMM reformulated as a neural network, not on differentiable optimization with sequential quadratic programming and custom PCG routines for computing gradients through MPC problems.

5. Parallel Shooting Sequential Quadratic Programming for Nonlinear MPC Problems

URL: [View paper](#)

Brief Assessment

Parallel Shooting SQP[46] focuses on parallel shooting methods for nonlinear MPC using SQP but does not implement differentiability or gradient computation through the optimization problem. The candidate addresses parallelization for forward solving only, not the backward pass for computing sensitivities that is central to the original paper's contribution.

6. Structure-exploiting sequential quadratic programming for model-predictive control

URL: [View paper](#)

Brief Assessment

Structure-Exploiting SQP[47] focuses on structure-exploiting sequential quadratic programming methods for MPC but does not provide sufficient detail in the available context about GPU acceleration or differentiability features to challenge the novelty of DiffMPC's specific approach combining GPU-accelerated differentiable optimization with custom preconditioned conjugate gradient routines.

7. DEEP FLEXQP: ACCELERATED NONLINEAR PRO

URL: [View paper](#)

Brief Assessment

DEEP FLEXQP[49] focuses on accelerating quadratic programming (QP) solvers through deep unfolding for sequential quadratic programming, not on GPU-accelerated differentiable MPC with sparse-in-time structure exploitation via custom PCG routines as in the original paper.

Contribution 2: Preconditioned conjugate gradient routine with tridiagonal preconditioning

Description: The authors adapt and implement a PCG routine with tridiagonal preconditioning that solves the KKT linear systems by exploiting the block-tridiagonal structure of the Schur complement. This design enables parallelization over time steps and supports warm-starting, making it suitable for GPU execution.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Implementation of a distributed parallel in time scheme using PETSc for a parabolic optimal control problem

URL: [View paper](#)

Brief Assessment

Parallel Time PETSc[37] focuses on a parareal preconditioner for parabolic optimal control problems, not the tridiagonal preconditioning scheme for MPC KKT systems described in the original paper.

Contribution 3: Application to robust drifting via domain randomization and RL

Description: The authors demonstrate DiffMPC on a reinforcement learning task for autonomous drifting under model mismatch. They use domain randomization over nonlinear dynamics to learn MPC cost and vehicle parameters, achieving robust drifting through water puddles on a Toyota Supra.

This contribution was assessed against **8 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Drift Cornering Control and Real-Vehicle Deployment for Electric Vehicles

URL: [View paper](#)

Brief Assessment

Drift Cornering Deployment[43] focuses on trajectory optimization with TD3-based RL for real-vehicle drift deployment, not on domain randomization over nonlinear dynamics for robust MPC learning as in the original paper.

2. Reference-free formula drift with reinforcement learning: From driving data to tire energy-inspired, real-world policies

URL: [View paper](#)

Brief Assessment

Reference-Free Formula Drift[38] focuses on tire energy-inspired policies for formula drift racing without reference trajectories, whereas the original paper uses domain randomization with MPC for robust drifting through water puddles. The candidate's emphasis is on reference-free drift control rather than domain randomization techniques for handling model mismatch.

3. Towards safe reinforcement learning in the real world

URL: [View paper](#)

Brief Assessment

Safe Reinforcement Learning[41] focuses on safe RL methods for general autonomous vehicle control in slippery environments, not specifically on domain randomization for robust drifting with MPC cost learning as in the original paper.

4. MPC-Inspired Reinforcement Learning for Verifiable Model-Free Control

URL: [View paper](#)

Brief Assessment

MPC-Inspired Reinforcement Learning[45] focuses on learning QP controller parameters for vehicle drift maneuvering, not on using domain randomization with differentiable MPC for autonomous drifting through water puddles as in the original paper.

5. Learning to Drift with Individual Wheel Drive: Maneuvering Autonomous Vehicle at the Handling Limits

URL: [View paper](#)

Brief Assessment

Learning to Drift[42] focuses on individual wheel drive (IWD) vehicle control for drifting maneuvers using RL with domain randomization. The original paper applies DiffMPC to learn MPC parameters for drifting a Toyota Supra through water puddles, while Learning to Drift[42] develops an RL framework for a 1/10 scale RC car platform with independent wheel speed control. These represent different technical approaches and experimental platforms for autonomous drifting.

6. Enhance Generality by Model-based Reinforcement Learning and Domain Randomization

URL: [View paper](#)

Brief Assessment

Model-Based Domain Randomization[39] focuses on path tracking and general autonomous driving tasks using domain randomization with model-based RL, not specifically on the challenging task of autonomous drifting at the limits of handling through water puddles as demonstrated in the original paper.

7. High-Speed Cornering Control and Real-Vehicle Deployment for Autonomous Electric Vehicles

URL: [View paper](#)

Brief Assessment

High-Speed Cornering Control[40] focuses on RL-based drift control for consumer-grade electric vehicles using TD3 and MPC fusion, not on domain randomization over nonlinear dynamics or learning MPC cost parameters through model mismatch scenarios like water puddles.

8. Adversarial Reinforcement Learning for Circular Autonomous Drifting Under Drivetrain Uncertainty

URL: [View paper](#)

Brief Assessment

Adversarial Drifting[44] focuses on adversarial training for circular drifting under drivetrain uncertainty, not on domain randomization over nonlinear dynamics with MPC cost learning as in the original paper.

Appendix: Text Similarity Detection

Textual similarity detection checked 17 papers and found 1 similarity segment(s) across 1 paper(s).

The following **1 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. Mpcgpu: Real-time nonlinear model predictive control through preconditioned conjugate gradient on the gpu

Detected in: Core Task (sibling)

△ **Note:** This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] Differentiable Model Predictive Control on the GPU [View paper](#)
- [1] GATO: GPU-Accelerated and Batched Trajectory Optimization for Scalable Edge Model Predictive Control [View paper](#)
- [2] Mpcgpu: Real-time nonlinear model predictive control through preconditioned conjugate gradient on the gpu [View paper](#)
- [3] Physics-guided neural network and GPU-accelerated nonlinear model predictive control for quadcopter [View paper](#)
- [4] GPU Acceleration for Real-Time, Whole-Body, Nonlinear Model Predictive Control [View paper](#)
- [5] Diff-FlowFSI: A GPU-Optimized Differentiable CFD Platform for High-Fidelity Turbulence and FSI Simulations [View paper](#)
- [6] LEGO-Size: LLM-Enhanced GPU-Optimized Signoff-Accurate Differentiable VLSI Gate Sizing in Advanced Nodes [View paper](#)
- [7] Primal-Dual iLQR for GPU-Accelerated Learning and Control in Legged Robots [View paper](#)
- [8] iKap: Kinematics-aware planning with imperative learning [View paper](#)
- [9] Model predictive control for aggressive driving over uneven terrain [View paper](#)
- [10] On the Benefits of GPU Sample-Based Stochastic Predictive Controllers for Legged Locomotion [View paper](#)
- [11] Actor-critic model predictive control: Differentiable optimization meets reinforcement learning [View paper](#)
- [12] Constrained sampling-based MPC using path integral for collision-free robot manipulation [View paper](#)
- [13] Real-time whole-body control of legged robots with model-predictive path integral control [View paper](#)
- [14] Feedback-MPPI: Fast Sampling-Based MPC via Rollout Differentiation--Adios low-level controllers [View paper](#)
- [15] Differentiable optimization-based modeling for machine learning [View paper](#)
- [16] Adaptive Model-Predictive Control of a Soft Continuum Robot Using a Physics-Informed Neural Network Based on Cosserat Rod Theory [View paper](#)
- [17] GPU-Enabled Parallel Trajectory Optimization Framework for Safe Motion Planning of Autonomous Vehicles [View paper](#)
- [18] Neural Joint Space Implicit Signed Distance Functions for Reactive Robot Manipulator Control [View paper](#)
- [19] Multilevel parallel GPU implementation of SQP solvers for Nonlinear MPC [View paper](#)
- [20] PyTorchFire: A GPU-accelerated wildfire simulator with Differentiable Cellular Automata [View paper](#)
- [21] Improving the Accuracy of Sample-based Model Predictive Control via Sample-based Newton-like method with Approximated Hessian and Gradient by Quadratic $\hat{\alpha}$ [View paper](#)
- [22] Residual MPC: Blending Reinforcement Learning with GPU-Parallelized Model Predictive Control [View paper](#)
- [23] ReLU-QP: A GPU-Accelerated Quadratic Programming Solver for Model-Predictive Control [View paper](#)
- [24] Parameterized and GPU-Parallelized Real-Time Model Predictive Control for High Degree of Freedom Robots [View paper](#)
- [25] Differentiable GPU-Parallelized Task and Motion Planning [View paper](#)
- [26] Exploiting GPU/SIMD Architectures for Solving Linear-Quadratic MPC Problems* [View paper](#)
- [27] GPU-based model predictive control of nonlinear parabolic partial differential equations system and its application in continuous casting [View paper](#)

- [28] Channel-assisted Fourier neural operator for high-fidelity far-wake prediction of a wind turbine under yawed conditions [View paper](#)
- [29] SIT-LMPC: Safe Information-Theoretic Learning Model Predictive Control for Iterative Tasks [View paper](#)
- [30] Adaptive DiffTune MPC for Marine Robots: An OceanSim-Based Study [View paper](#)
- [31] GPU-Accelerated Dual Gradient-Projection Algorithm for Embedded Linear Model Predictive Control [View paper](#)
- [32] Differentiable Predictive Control of Permanent Magnet Synchronous Motors [View paper](#)
- [33] Dynamic Grasping Based on Reinforcement Learning and Differentiable MPC [View paper](#)
- [34] Learning Robot Control: From Reinforcement Learning to Differentiable Simulation [View paper](#)
- [35] DIFFERENTIABLE PARALLEL SIMULATION FOR PLANNING AS DESIGN FOR SOFT GROWING ROBOTS [View paper](#)
- [36] Prior Work: Hardware Acceleration for Reliable Realtime Optimal Control [View paper](#)
- [37] Implementation of a distributed parallel in time scheme using PETSc for a parabolic optimal control problem [View paper](#)
- [38] Reference-free formula drift with reinforcement learning: From driving data to tire energy-inspired, real-world policies [View paper](#)
- [39] Enhance Generality by Model-based Reinforcement Learning and Domain Randomization [View paper](#)
- [40] High-Speed Cornering Control and Real-Vehicle Deployment for Autonomous Electric Vehicles [View paper](#)
- [41] Towards safe reinforcement learning in the real world [View paper](#)
- [42] Learning to Drift with Individual Wheel Drive: Maneuvering Autonomous Vehicle at the Handling Limits [View paper](#)
- [43] Drift Cornering Control and Real-Vehicle Deployment for Electric Vehicles [View paper](#)
- [44] Adversarial Reinforcement Learning for Circular Autonomous Drifting Under Drivetrain Uncertainty [View paper](#)
- [45] MPC-Inspired Reinforcement Learning for Verifiable Model-Free Control [View paper](#)
- [46] Parallel Shooting Sequential Quadratic Programming for Nonlinear MPC Problems [View paper](#)
- [47] Structure-exploiting sequential quadratic programming for model-predictive control [View paper](#)
- [48] Towards safe and tractable Gaussian process-based MPC: Efficient sampling within a sequential quadratic programming framework [View paper](#)
- [49] DEEP FLEXQP: ACCELERATED NONLINEAR PRO [View paper](#)