



US DOT National
University Transportation Center for Safety

Carnegie Mellon University



Autonomous Driving in Hazardous Conditions: Pushing the Limits of Vehicle Control

Safe Iterative Learning Control and Failure-Aware Invariant Set
Learning for Vehicles on Rough Terrain and in Emergencies

Project #612

Rahul Mangharam, Principal Investigator

<https://orcid.org/0000-0002-3388-8283>

Antonio Loquercio, Co-Principal Investigator

ORCID: <https://orcid.org/0000-0002-8410-3933>

University of Pennsylvania, School of Engineering and Applied Science
200 South 33rd Street, Moore Building, Philadelphia, PA 19104

Final Report – October 2026

Federal Grant #69A3552344811

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, under grant number 69A3552344811 from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

Technical Report Documentation Page

12. Sponsoring Agency Name and Address U.S. Department of Transportation, Office of the Assistant Secretary for Research and Technology, University Transportation Centers Program (UTC grant, Bipartisan Infrastructure Law) 1200 New Jersey Avenue SE, Washington, DC 20590 Safety21 National University Transportation Center, Carnegie Mellon University (UTC grant recipient) 5000 Forbes Avenue, Pittsburgh, PA 15213		13. Type of Report and Period Covered Final Report (July 1, 2025 – June 30, 2026)
		14. Sponsoring Agency Code USDOT
15. Supplementary Notes Conducted in cooperation with the U.S. Department of Transportation and the Safety21 University Transportation Center. Federal share: \$110,000 in UTC funds. Non-federal match: \$110,000 University of Pennsylvania cost share (1:1). Project record: https://ppms.cit.cmu.edu/projects/detail/612 . Project websites: https://xlab.upenn.edu/SIT-LMPC/ and https://xlab.upenn.edu/learning-from-failure/ . Publication DOIs and identifiers: https://doi.org/10.1109/LRA.2025.3634881 , https://doi.org/10.48550/arXiv.2604.06776 . Full list of research products in Section 8.		
16. Abstract Project 612 develops control methods that let automated vehicles drive fast and stay safe on rough terrain and in hazardous conditions such as the flooded, debris-covered, and washed-out roads left by Hurricanes Helene and Milton. Models of such conditions are imperfect, and failures are costly, so the project builds controllers that learn from experience without giving up safety. Two methods are reported. SIT-LMPC (IEEE Robotics and Automation Letters, 2026) learns a safe set and a value function from previous attempts at a task and uses a GPU-parallel, information-theoretic model predictive controller with an adaptive penalty to keep the vehicle inside its constraints. On a 1/5-scale off-road car it raised average speed by 75 percent over the initial run and stayed inside the track boundaries, while the baseline did not converge without crashing. FAIL (IEEE Conference on Decision and Control, 2026) learns state-control invariant sets directly from failing trajectories, without knowing the dynamics, and provably converges to the maximal invariant set. In a double-integrator test, six failing trajectories recovered all 14 constraints of the exact set. Remaining work toward the proposal, including a hazardous-driving dataset, software release, and pilot sites with the Delaware Valley Regional Planning Commission, continues in the next phase.		
17. Key Words Autonomous vehicles; hazardous conditions; off-road driving; learning model predictive control; safe learning; invariant sets; learning from failure; model predictive path integral control; Autoware		18. Distribution Statement No restrictions. This document is available through the Safety21 project record.
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 18
		22. Price

Contents

1	Project Overview	1
1.1	Project Funding and Matching Funds	1
1.2	Summary of Results	2
2	Why Hazardous Conditions Need Learning with Safety Guarantees	2
3	SIT-LMPC: Safe Learning Control for Iterative Tasks	3
3.1	Method	3
3.2	Simulation Results	3
3.3	Real-World Off-Road Results	4
4	FAIL: Learning Invariant Sets from Failures	5
4.1	Method	5
4.2	Results	6
5	Status Against the Project Proposal	7
6	Partners, Deployment, and Workforce	8
6.1	Partners and Path to Deployment	8
6.2	Impact on the Transportation System	8
6.3	Workforce Development	8
6.4	Dissemination	8
7	Next Phase and Recommendations	8
8	Research Products, Data, and Dissemination	9
8.1	Publications	9
8.2	Software, Data, and Websites	9
8.3	Datasets and Metadata	10
8.4	Research Personnel and ORCID Information	10
A	Abbreviations	11

List of Figures

1	SIT-LMPC architecture. The learning loop updates the safe set and value function after each run, and the control loop solves many MPPI problems in parallel. Figure from [1].	3
2	Convergence of lap time over iterations for three simulated experiments, with the fastest lap trajectories colored by velocity. Figure from [1].	4
3	Real-vehicle experiment: (a) experimental platform, (b) track layout, (c) lap time per iteration, and (d) boundary violations. Figure from [1].	5
4	The MSCI set computed from the known dynamics (left) and its state projection compared with the MCI set (right). Figure from [2].	6
5	The final learned polytope, which coincides with the MSCI set (left), and its state projection with a safe trajectory (right). Figure from [2].	7

List of Tables

1	Project 612 funding sources.	2
2	Main quantitative results.	2
3	Proposal outputs and their status at the end of the project period.	7
4	Publications supported by the project.	9
5	Websites and supporting material.	9
6	Project investigators, publication authors, and ORCID information.	10
7	Abbreviations used in this report.	11

1 Project Overview

Safety21 Project 612, Autonomous Driving in Hazardous Conditions: Pushing the Limits of Vehicle Control, addresses a gap between how automated vehicles are tested and the conditions in which people need them most. Automated driving systems are developed and validated on dry, well-marked roads. After a hurricane, the road may be flooded, covered in debris, or partly washed out, and a vehicle that serves emergency responders must still drive quickly and stay under control. Hurricanes Helene and Milton in 2024 showed how much road networks can be disrupted in a short time. In these conditions no model of the vehicle and the surface is accurate, and a failure can be costly.

Led by Principal Investigator Rahul Mangharam and Co-Principal Investigator Antonio Loquercio at the University of Pennsylvania, the project develops controllers that learn from experience while keeping safety guarantees. It reports two results.

1. **SIT-LMPC** is a safe, information-theoretic learning model predictive controller for tasks a vehicle repeats [1]. It learns a safe set and a value function from earlier attempts and runs in parallel on a graphics processing unit (GPU) at more than 100 Hz. It was demonstrated on a 1/5-scale off-road car on grass, and it was published in IEEE Robotics and Automation Letters (volume 11, issue 1, January 2026).
2. **FAIL** is a failure-aware iterative learning method that learns state-control invariant sets directly from failing trajectories, without knowing the vehicle dynamics [2]. It is accepted to the 2026 IEEE Conference on Decision and Control (CDC).

Section 2 explains why hazardous conditions need learning-based control with safety guarantees. Section 3 describes SIT-LMPC and its results, and Section 4 describes FAIL and its results. Section 5 compares the project outputs with the proposal and states which items are complete and which continue. Section 6 covers partners, deployment, and workforce, Section 7 describes the next phase, and Section 8 lists all research products.

Overall project outcome

Project 612 delivered a safe learning controller that raised average speed by 75 percent on a real off-road vehicle while keeping it inside the track boundaries, and a failure-aware method that learns exact invariant sets from failures alone. Both methods are published or accepted at leading IEEE venues. Together they give the field a foundation for vehicles that learn to drive hard terrain safely. Hazard-specific datasets, a public software release, patent filings, and pilot sites with regional planning partners continue in the next phase.

1.1 Project Funding and Matching Funds

Project 612 was funded through the Safety21 National University Transportation Center under U.S. Department of Transportation grant 69A3552344811 for the project period July 1, 2025 through June 30, 2026. Table 1 states the source of all project funds. The University of Pennsylvania provided the full non-federal match, equal to the federal award, as institutional cost share. No other matching sources were used. The federal share was fully expended by the end of the project period. Both publications also acknowledge support from National Science Foundation awards CISE-2431569 and 2514584, which are separate from the funds in Table 1.

Table 1: Project 612 funding sources.
kvoptions-patch[2019/11/29]

1.2 Summary of Results

Table 2 lists the main quantitative results reported by the project.

Table 2: Main quantitative results.

Result	Evidence	Section
Real-time control on embedded hardware	SIT-LMPC runs at more than 100 Hz on an NVIDIA Jetson Orin AGX on a 1/5-scale off-road car.	3.3
Faster driving on rough terrain	Average speed rose by 75 percent, up to 3.5 m/s, from an initial lap of 52.13 s on a grass track.	3.3
Safety while learning	SIT-LMPC stayed within the track while the baseline ABC-LMPC did not converge without crashing. Final lap time was 31.43 percent faster than the baseline.	3.3
Safe in simulation over many iterations	No safety violations over all 150 iterations in the simulated racing benchmark.	3.2
Learning invariant sets from failures	Six failing trajectories were enough to learn all 14 constraints of the maximal state-control invariant set, with zero Hausdorff distance and 0.0353 s of learning time.	4.2

2 Why Hazardous Conditions Need Learning with Safety Guarantees

Vehicle controllers rely on models of the vehicle and of the surface it drives on. On pavement these models can be identified accurately. On wet grass, mud, gravel, flood water, or debris, tire forces change from one patch to the next and cannot be known in advance. A controller built on a nominal model is then either too cautious to be useful or too aggressive to be safe.

Learning from experience is a natural remedy, because the vehicle can improve with every attempt. The difficulty is that a learning controller must explore, and exploration near the limits of grip can end in a crash. The project therefore studies two questions that arise in any hazardous setting.

- **How can a vehicle go faster with every attempt without leaving the safe region?** SIT-LMPC answers this for tasks that repeat, such as driving the same stretch of damaged road several times, and it handles noise, model error, and imperfect localization.
- **How can a vehicle learn what is unsafe from the failures it cannot avoid?** FAIL answers this by turning each observed failure into a constraint that prevents the same failure from happening again.

Both questions connect to the priorities of the U.S. Department of Transportation for automated driving systems (ADS), advanced driver-assistance systems (ADAS), artificial intelligence and machine learning, and robotics. They also support a longer-term goal in the proposal, which is to let vehicles operate in emergency and disaster response when roads are damaged.

3 SIT-LMPC: Safe Learning Control for Iterative Tasks

3.1 Method

Many vehicle tasks are iterative. The vehicle drives the same route many times, and each run can use data from the earlier runs. Learning model predictive control (LMPC) [3] builds a safe set from past trajectories and a terminal cost that estimates the cost to finish the task. Each new run must end inside the safe set, so the vehicle is never asked to do anything that earlier runs showed to be unsafe. The ABC-LMPC method [4] extends this to stochastic nonlinear systems with a sampling-based controller. In practice it encodes state constraints as a large constant cost, which makes it overly conservative, and its cross-entropy sampling is prone to infeasible solutions and mode collapse.

SIT-LMPC replaces the controller with an information-theoretic model predictive controller, known as model predictive path integral (MPPI) control [5], which samples many trajectories and requires no gradients. The project makes four contributions, shown in Figure 1.

1. A constrained MPPI formulation for general stochastic nonlinear systems that makes no assumption about the form of the dynamics or the state constraints.
2. An adaptive penalty method that samples penalty parameters online. It balances safety against performance, because a fixed high penalty is too conservative and a fixed low penalty makes the system unsafe.
3. A value function learned with normalizing flows from the trajectories of earlier iterations, which models uncertainty more richly than Gaussian priors or Bayesian neural networks.
4. A fully parallel implementation on a GPU, which allows 100 Hz or faster real-time control on an NVIDIA Jetson Orin AGX.




Figure 1: SIT-LMPC architecture. The learning loop updates the safe set and value function after each run, and the control loop solves many MPPI problems in parallel. Figure from [1].

3.2 Simulation Results

The method was tested in three experiments of increasing difficulty, each repeated five times per controller. The first is a deterministic point-mass navigation task with a circular obstacle. All methods improved over iterations, and SIT-LMPC converged fastest. Because the obstacle makes the problem nonconvex, the gradient-based LMPC converged to a local minimum, while the sampling-based methods escaped it.

The second experiment is a simulated race car with a dynamic single-track model, noise on observations and controls, and a lap-time objective with track-boundary constraints. Figure 2 shows that ABC-LMPC often crashed partway through the iterations, because choosing the single least-cost sample leaves the vehicle close to a constraint boundary where noise can push it out. SIT-LMPC converged to a lower lap time and stayed safe through all 150 iterations.




Figure 2: Convergence of lap time over iterations for three simulated experiments, with the fastest lap trajectories colored by velocity. Figure from [1].

An ablation study isolated the key components. Controllers that use cross-entropy sampling were always infeasible under noise. Normalizing flows gave better performance than Bayesian neural networks for both samplers. The adaptive penalty improved learning in every combination, and it was most effective as part of SIT-LMPC.

3.3 Real-World Off-Road Results

To test the method on a real vehicle, SIT-LMPC was deployed on a 1/5-scale autonomous off-road race car on a grass track. Localization came from a Fixposition Vision-RTK 2 GNSS system. The initial safe set came from an MPPI controller tracking the centerline at 2 m/s, which gave an initial lap time of 52.13 s. The vehicle model was identified empirically and is not fully accurate, and localization and control are noisy, as they would be after a disaster.

Figure 3 shows the results. SIT-LMPC improved lap time over the iterations and raised the average velocity by 75 percent, up to 3.5 m/s, which was the acceleration limit of the car. It stayed within the track boundaries. The baseline ABC-LMPC did not converge without crashing, and the final SIT-LMPC lap time was 31.43 percent faster than the baseline. The result shows that the method works on a real stochastic system with model mismatch, imperfect localization, and the control noise of an off-road surface.





Figure 3: Real-vehicle experiment: (a) experimental platform, (b) track layout, (c) lap time per iteration, and (d) boundary violations. Figure from [1].

4 FAIL: Learning Invariant Sets from Failures

4.1 Method

A safe controller needs to know from which states it can keep the vehicle inside its limits forever. The set of such states is the maximal control invariant (MCI) set. Classical algorithms compute it from an accurate dynamics model, which is exactly what is missing in hazardous conditions. Existing data-driven methods find the safe states but not the inputs that keep the system safe at each state.

FAIL introduces state-control invariance, which extends invariance from the state space to the joint space of states and controls. The maximal state-control invariant (MSCI) set holds the MCI set and, for every state in it, the inputs that keep the system invariant. The paper proves that the state projection of the MSCI set is the MCI set, and that the sections of the MSCI set at each state are the admissible invariance-preserving inputs [2].

For deterministic linear time-invariant systems with polytopic constraints, the FAIL algorithm starts with the known constraint polytope and learns one new halfspace from each one-step failing state and input pair. Each failure reveals a violated halfspace of the predecessor of the current constraint set, and a regression on the failing trajectory recovers that halfspace without knowing the system matrices. The learned set is intersected with the current one, which excludes the failure from repeating. The paper proves that the learned set shrinks monotonically and converges to the MSCI set in finitely many iterations, given trajectories with sufficient

excitation.

4.2 Results

The method was tested on a double integrator with position, velocity, and input limits. The initial polytope has six constraints. Computed from the known dynamics, the MSCI set has 14 halfspaces, of which eight come from the predecessor recursion, and the MCI set has eight. Figure 4 shows that the state projection of the MSCI set equals the MCI set, with a Hausdorff distance of zero.




Figure 4: The MSCI set computed from the known dynamics (left) and its state projection compared with the MCI set (right). Figure from [2].

With the dynamics unknown, FAIL learned all eight predecessor halfspaces from six failing trajectories, generated by two constant-input controllers and random admissible controllers. The learned polytope equals the MSCI set at termination, and each learned halfspace matches a halfspace of the exact set to within a numerical tolerance of 10^{-3} . Learning took 0.0353 s. At termination, none of 1200 random admissible trajectories left the learned set. Figure 5 shows the final learned polytope, its projection, and a safe trajectory. Scaling the projection step to higher-dimensional systems and designing the controllers that generate informative failures are the topics of a journal extension.




Figure 5: The final learned polytope, which coincides with the MSCI set (left), and its state projection with a safe trajectory (right). Figure from [2].

5 Status Against the Project Proposal

Table 3 compares the outputs named in the proposal with the work completed in the project period. The table states plainly which items are complete and which continue.

Table 3: Proposal outputs and their status at the end of the project period.

Proposal output	Status	Evidence or plan
Research impact and publications	Complete	Two publications: SIT-LMPC in IEEE RA-L (January 2026) and FAIL accepted to IEEE CDC 2026.
Controllers that learn from imperfect models in rough terrain	Complete for the iterative-task case	SIT-LMPC on a 1/5-scale off-road car at 100 Hz or more on embedded hardware (Section 3).
Learning safety limits from failure	Complete for linear systems	FAIL, with convergence proof and numerical validation (Section 4).
Training of a doctoral student	Supported	Penn researchers carried out the research reported in both papers (Section 6.3).
Open-source hazardous-driving dataset	Continues in the next phase	Data collection on hazardous terrain is planned (Section 7).
Released software	Continues in the next phase	Project websites are public. A public code release is planned.
Patent filings	Continues in the next phase	No filing was made in the project period.
Pilot developments at Philadelphia-area off-road test sites	Continues in the next phase	Planning with the Delaware Valley Regional Planning Commission (DVRPC).
Multiple-hypothesis model-based reinforcement learning	Continues in the next phase	SIT-LMPC and FAIL are the safe-learning foundation for this work.

6 Partners, Deployment, and Workforce

6.1 Partners and Path to Deployment

The deployment plan names two partners. The Delaware Valley Regional Planning Commission (DVRPC) is the metropolitan planning organization for the Philadelphia region and is the planned partner for pilot work at off-road test sites in the area, such as the Philly Pump Track. The Autoware Foundation maintains the open-source software stack for automated driving. Because SIT-LMPC is a controller that takes a vehicle model and a cost, and not a vehicle-specific policy, it is a candidate control module for platforms that use Autoware. Integration is planned work. The 1/5-scale experiments are the first step in a path from small-scale vehicles to full-size vehicles on hazardous terrain.

6.2 Impact on the Transportation System

The results support transportation resilience in three ways. First, a vehicle that improves itself while staying safe can be commissioned on a new surface without a precise model, which lowers the cost and time of deployment. Second, the real-time performance on embedded hardware shows that the approach can run on vehicle-grade computers. Third, learning exact safety limits from failures offers a route to certify constraints for systems whose dynamics are only partly known, which matters for any vehicle that must operate after a disaster.

6.3 Workforce Development

The research was carried out by Penn researchers Ahmad Amine, Zirui Zang, and Nick-Marios Kokolakis, together with collaborators Truong X. Nghiem of the University of Central Florida and Ugo Rosolia. Zang and Amine contributed equally to SIT-LMPC. The work trained early-career researchers in safe learning, optimization-based control, GPU programming, and hardware experiments on a real vehicle, which are skills in demand for automated driving, robotics, and safety-critical autonomy.

6.4 Dissemination

Each publication has a public project website. The SIT-LMPC site is <https://xlab.upenn.edu/SIT-LMPC/> and the FAIL site is <https://xlab.upenn.edu/learning-from-failure/>. Both papers are available as preprints on arXiv [1, 2].

7 Next Phase and Recommendations

The results form a base for the work that remains in the proposal. The following steps are planned.

- **Hazard-specific data.** Collect and publish an open dataset of vehicle behavior on hazardous and off-road surfaces, with metadata that describe the vehicle, surface, sensors, and conditions.
- **Hybrid simulation.** Combine first-principles models with data-driven models so that simulations capture floods, debris, and washed-out roads.

- **Multiple-hypothesis model-based reinforcement learning.** Plan with several candidate models and weigh them by model uncertainty, using SIT-LMPC and FAIL as the safety layer.
- **Scaling and extension.** Extend FAIL to higher-dimensional and nonlinear systems, and apply SIT-LMPC to a wider range of platforms and to systems with unknown dynamics, as stated in the conclusions of the papers.
- **Pilot sites and software.** Run pilots with DVRPC at Philadelphia-area off-road test sites, release the software, and pursue patent protection where appropriate.

Barriers. Validation so far uses a 1/5-scale vehicle on grass and a numerical linear system. Full-size vehicles, hazardous surfaces such as water and debris, and real pilot sites remain to be tested, and FAIL needs further work for high-dimensional systems.

8 Research Products, Data, and Dissemination

The project produced two publications. Table 4 lists them with their identifiers.

8.1 Publications

Table 4: Publications supported by the project.

Short name	Citation and links	Venue and status
SIT-LMPC	Z. Zang, A. Amine, N.-M. T. Kokolakis, T. X. Nghiem, U. Rosolia, and R. Mangharam, “SIT-LMPC: Safe Information-Theoretic Learning Model Predictive Control for Iterative Tasks,” <i>IEEE Robotics and Automation Letters</i> , vol. 11, no. 1, pp. 986–993, January 2026. DOI: https://doi.org/10.1109/LRA.2025.3634881 arXiv: https://arxiv.org/abs/2602.16187 Website: https://xlab.upenn.edu/SIT-LMPC/	Journal article, published
FAIL	A. Amine, N.-M. T. Kokolakis, U. Rosolia, T. X. Nghiem, and R. Mangharam, “Failure-Aware Iterative Learning of State-Control Invariant Sets,” accepted to the <i>IEEE Conference on Decision and Control (CDC)</i> , 2026. arXiv: https://arxiv.org/abs/2604.06776 Website: https://xlab.upenn.edu/learning-from-failure/	Conference paper, accepted

8.2 Software, Data, and Websites

Table 5: Websites and supporting material.

Item	Purpose	Access
SIT-LMPC website	Paper, videos, and project description.	https://xlab.upenn.edu/SIT-LMPC/
FAIL website	Paper and project description.	https://xlab.upenn.edu/learning-from-failure/

8.3 Datasets and Metadata

The project reports simulation and hardware experiments, and the results are in the two papers. The experiment records include the vehicle model parameters, track layout, controller settings, noise levels, number of iterations and rollouts (five per controller), and the metrics reported in the papers. These records are maintained by the authors under the project Data Management Plan (DMP). A public hazardous-driving dataset is a planned output of the next phase (Section 7).

8.4 Research Personnel and ORCID Information

Table 6 gives ORCID identifiers for the investigators and for the authors of the project publications.

Table 6: Project investigators, publication authors, and ORCID information.

Name	Project or publication role	ORCID
Rahul Mangharam	Principal investigator; author of both publications	https://orcid.org/0000-0002-3388-8283
Antonio Loquercio	Co-principal investigator	https://orcid.org/0000-0002-8410-3933
Ahmad Amine	Author of both publications	https://orcid.org/0009-0008-6080-8411
Zirui Zang	Author (SIT-LMPC)	
Nick-Marios T. Kokolakis	Author of both publications	https://orcid.org/0000-0003-4316-610X
Truong X. Nghiem	Co-author of both publications, University of Central Florida	https://orcid.org/0000-0003-4841-3800
Ugo Rosolia	Co-author of both publications	https://orcid.org/0000-0002-1682-0551

References

- [1] Z. Zang, A. Amine, N.-M. T. Kokolakis, T. X. Nghiem, U. Rosolia, and R. Mangharam, “SIT-LMPC: Safe information-theoretic learning model predictive control for iterative tasks,” *IEEE Robotics and Automation Letters*, vol. 11, no. 1, pp. 986–993, 2026. doi: 10.1109/LRA.2025.3634881. [Online]. Available: <https://arxiv.org/abs/2602.16187>
- [2] A. Amine, N.-M. T. Kokolakis, U. Rosolia, T. X. Nghiem, and R. Mangharam, “Failure-aware iterative learning of state-control invariant sets,” accepted to *IEEE Conference on Decision and Control*, 2026. arXiv:2604.06776. [Online]. Available: <https://arxiv.org/abs/2604.06776>
- [3] U. Rosolia and F. Borrelli, “Learning model predictive control for iterative tasks. A data-driven control framework,” *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 1883–1896, 2018.
- [4] B. Thananjeyan, A. Balakrishna, U. Rosolia, J. E. Gonzalez, A. Ames, and K. Goldberg, “ABC-LMPC: Safe sample-based learning MPC for stochastic nonlinear dynamical systems with adjustable boundary conditions,” in *Algorithmic Foundations of Robotics XIV*. Springer, 2021, pp. 1–17.
- [5] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, “Information-theoretic model predictive control: Theory and applications to autonomous driving,” *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.

- [6] Safety21 University Transportation Center, “Project 612: Autonomous driving in hazardous conditions: Pushing the limits of vehicle control.” Accessed: Oct. 6, 2026. [Online]. Available: <https://ppms.cit.cmu.edu/projects/detail/612>

A Abbreviations

Table 7: Abbreviations used in this report.

Abbreviation	Definition
ABC-LMPC	Adjustable boundary condition learning model predictive control
ADAS	Advanced driver-assistance system
ADS	Automated driving system
CDC	IEEE Conference on Decision and Control
CEM	Cross-entropy method
DMP	Data Management Plan
DVRPC	Delaware Valley Regional Planning Commission
FAIL	Failure-aware iterative learning
GNSS	Global navigation satellite system
GPU	Graphics processing unit
LMPC	Learning model predictive control
MCI	Maximal control invariant
MPC	Model predictive control
MPPI	Model predictive path integral
MSCI	Maximal state-control invariant
PPMS	Safety21 project and performance management system
RA-L	IEEE Robotics and Automation Letters
SIT-LMPC	Safe information-theoretic learning model predictive control
USDOT	U.S. Department of Transportation
UTC	University Transportation Center