Generating Unsafe Road Activity from LLMs and Road Imagery to Learn Better Safety

Robert Tamburo[1]

Srinivasa Narasimhan[2]

Final Report – August 31, 2026

[1] https://orcid.org/0000-0002-5636-9443

[2] https://orcid.org/0000-0003-0389-1921

Technical Report Documentation Page

<table>
<tr>
<td>1. Report No.

Enter the report number assigned by the sponsoring agency.</td>
<td>2. Government Accession No.</td>
<td>3. Recipient’s Catalog No.</td>
</tr>
<tr>
<td colspan="2" rowspan="2">4. Title and Subtitle

Generating Unsafe Road Activity from LDMs and Road Imagery to Learn Better Safety</td>
<td>5. Report Date

August 31, 2026</td>
</tr>
<tr>
<td>6. Performing Organization Code

Enter any/all unique numbers assigned to the performing organization, if applicable.</td>
</tr>
<tr>
<td colspan="2">7. Author(s)

Robert Tamburo, https://orcid.org/0000-0002-5636-9443

Srinivasa Narasimhan, https://orcid.org/0000-0003-0389-1921</td>
<td>8. Performing Organization Report No.

Enter any/all unique alphanumeric report numbers assigned by the performing organization, if applicable.</td>
</tr>
<tr>
<td colspan="2" rowspan="2">9. Performing Organization Name and Address

Carnegie Mellon University

5000 Forbes Ave

Pittsburgh, PA 15213</td>
<td>10. Work Unit No.</td>
</tr>
<tr>
<td>11. Contract or Grant No.

Federal Grant # 69A3552344811 / 69A3552348316</td>
</tr>
<tr>
<td colspan="2" rowspan="2">12. Sponsoring Agency Name and Address

Safety21 University Transportation Center (UTC)

5000 Forbes Ave

Pittsburgh, PA 15213</td>
<td>13. Type of Report and Period Covered

Final Report (July 1, 2025 – July 31, 2026)</td>
</tr>
<tr>
<td>14. Sponsoring Agency Code

USDOT</td>
</tr>
</table>

15. Supplementary Notes

Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration.

Enter information not included elsewhere, such as translation of (or by), report supersedes, old edition number, alternate title (e.g. project name), hypertext links to documents or related information in the form of URLs, PURLs (preferred over URLs - https://archive.org/services/purl/help), DOIs (https://www.doi.org/), insertion of QR codes, copyright or disclaimer statements, etc. Edit boilerplate FHWA statement above if needed.

16. Abstract

Enter a brief factual summary of the most significant information, including the purpose, methods, results, and conclusions of the work. When appropriate, the abstract should include advice on how the results of the research can be used. For guidance, please see ANSI/NISO Z39.14-1997 (R2015) Guidelines for Abstracts (https://www.niso.org/publications/ansiniso-z3914-1997-r2015-guidelines-abstracts).

17. Key Words	18. Distribution Statement
Autonomous land vehicles, autonomous vehicle guidance, digital simulation, vulnerable road users, navigation, planning, long-tailed data, training data, unsafe roadway activity	No restrictions. This document is available through the National Technical Information Service, Springfield, VA 22161. Enter any other agency mandated distribution statements. Remove NTIS statement if it does not apply.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	9	Refers to the price of the report. Leave blank unless applicable.

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

Project Investigators and Contributors

This is a summary report combining multiple sub-projects conducted at the University Transportation Center at Carnegie Mellon University in relation to the overall project titled "Generating Unsafe Road Activity from LLMs and Road Imagery to Learn Better Safety". The participants who contributed to the different sub-projects are listed below.

- Srinivasa Narasimhan (Carnegie Mellon University, ORCID: 0000-0003-0389-1921)
- Robert Tamburo (Carnegie Mellon University, ORCID: 0000-0002-5636-9443)
- Anurag Ghosh (Carnegie Mellon University, ORCID: 0000-0002-1617-0851)
- Shen Zheng (Carnegie Mellon University)
- Gaurav Parmar (Carnegie Mellon University, ORCID: 0000-0001-8640-2243)
- Khiem Vuong (Carnegie Mellon University, ORCID: 0009-0006-2474-2270)
- Angela Chen (Carnegie Mellon University, ORCID: 0009-0000-7345-604X)
- Manmohan Chandraker (NEC Labs America and UC San Diego)
- Francesco Pittaluga (NEC Labs America)
- Juan Alvarez-Padilla (Massachusetts Institute of Technology)

Description of the Problem

The FHWA zero death vision is a strategic goal to eliminate all traffic fatalities and severe injuries. One of the pillars of this strategy is to advance life-saving technology in vehicles. Many such in-vehicle technological advances have been developed and deployed in the past decade for example backup cameras, lane departure warning systems, adaptive cruise control, emergency brake systems, driver attention monitoring, etc., and more recently, autonomous driving. Standard machine learning models often fail in detecting and planning around unsafe activity because they are primarily trained on typical driving data, which is readily available. Atypical driving data is limited because unsafe activity is uncommon, highly diverse, and occurs randomly. Using this type of atypical data can be used to train improved models and better simulate existing in-vehicle technologies and allow for the development and deployment of new vision-based methods and technologies.

Approach

We took a multifaceted approach to the problem by developing methods for data generation, navigation, planning, and simulation. To accelerate development, we made use of a prior UTC project that produced the RoadWork dataset[3]. First, to address the lack of data of unsafe activity, we developed a method that could translate an image taken during safe driving conditions into an image during unsafe driving conditions, specifically bad weather and bad light. Next, a framework for a navigation planner was developed for real-time driving autonomy in unsafe roadway conditions. Finally, a pipeline was developed for generating data that can be used in simulators from uncalibrated, monocular dashcams.

Methodology and Findings

Changing the Weather and Time-of-Day in Images

We have developed an image-to-image translation method for creating unsafe weather and light conditions in driving images using latent diffusion models (LDMs). Our method preserves fine-detail loss with LDMs with a novel saliency-guided warp/unwarp framework. The framework applies a non-uniform backward resampling transformation to enlarge import image regions before encoding to better preserve structural details without increasing the overall image resolution. The warped image is then processed by an image-to-image translation model, the result of which is then mapped back via an inverse warp. The image can then be used in our outpainting-based synthetic data generation pipeline to add poor weather conditions or poor light conditions to the image (or video with frame-by-frame processing). Example generated images are shown in Figure 1. The method is independent of model, does not require architectural modification, and introduces negligible additional computational overhead.

---

[3] Ghosh, A., Zheng, S., Tamburo, R., Vuong, K., Alvarez-Padilla, J., Zhu, H., Cardei, M., Dunn, N., Mertz, C., Narasimhan, S.G.: Roadwork: A dataset and benchmark for learning to recognize, observe, analyze and drive through work zones. In: ICCV (2025).

Figure 1: Example of adding fog and golden sunlight to an image. Images have semantic consistency, good image quality, and lighting quality. Small details like traffic arrows and text remain visible.

Real-Time Autonomous Driving Planner

We developed a dual-approach framework for autonomous driving motion planning that separates navigation challenges into geometric maneuverability and semantic reasoning. The framework relies on a real-time rule-based autonomous driving action planner that extends standard planners by adding dynamic topological replanning at every time step and a language-based autonomous driving action planner that produces valid trajectory plans in a single forward pass or outputs textual reasoning before planning. The framework enables a hybrid approach that provides a practical pathway for integrating structured planning and foundation models in real-time autonomy. Examples of this framework are shown in Figure 2.

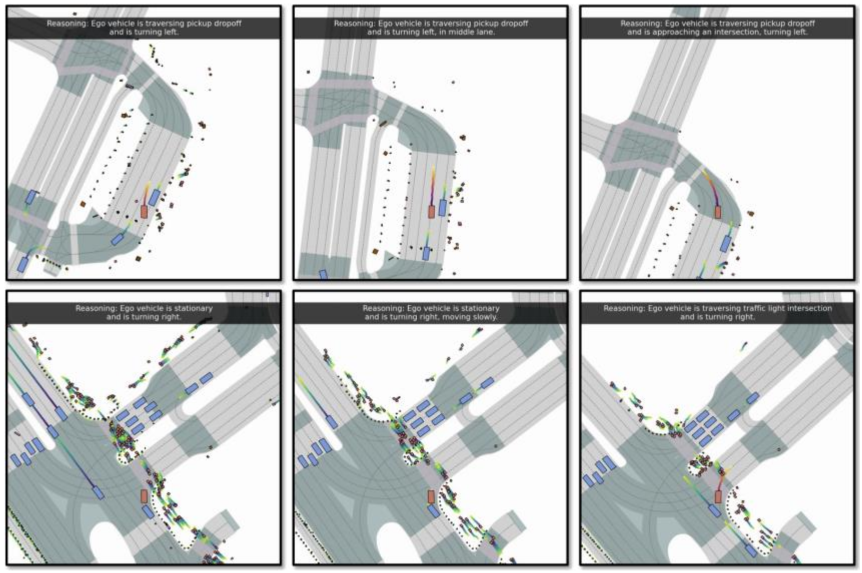

Figure 2: Two examples of our planning framework. Top row shows the ego vehicle (red) navigating around a vehicle in a pickup/dropoff zone. Bottom row shows the ego vehicle making a right turn in a crowded intersection.

Dashcam Videos to Driving Simulations Through Work Zones

Self-driving simulations generally rely on sparse data in a small number of cities or hand-curated synthetic scenarios. Dashcam videos provide a very broad range of locations and situations suitable for capturing unsafe roadway conditions like work zones. We developed a framework that takes uncalibrated, monocular dashcam videos and recovers accurate metric, geo-anchored 4D scenes that can be used in existing self-driving simulations. Shown in Figure 3 is an example frame showing reconstruction of the scene, detection of roadwork objects, and the simulated road environment.

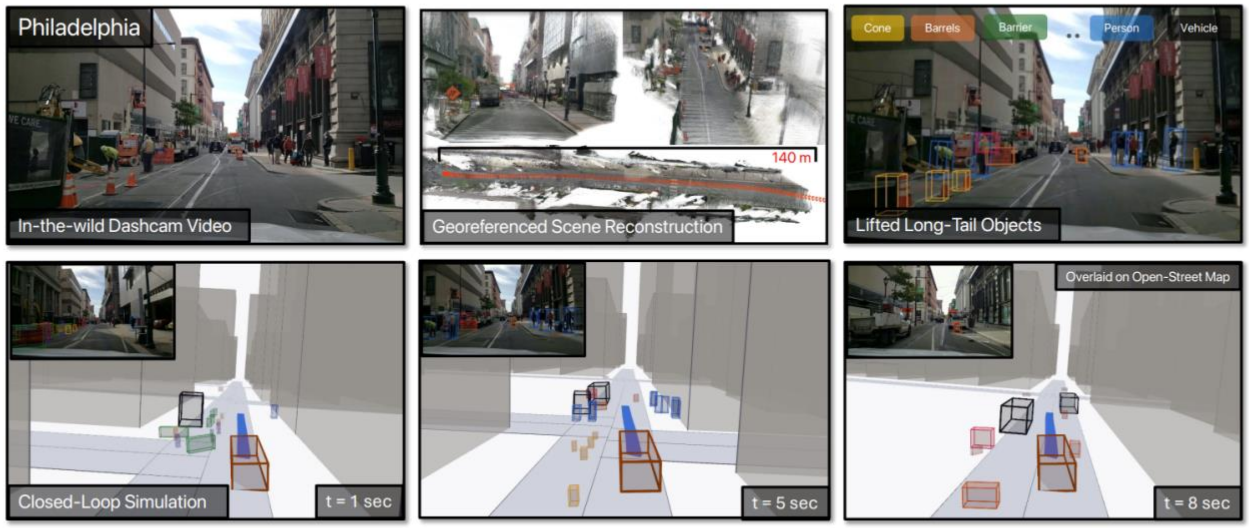

Figure 3: Example of self-driving simulation for unsafe activity. Top row (left to right) shows a frame from a dashcam video, the scene reconstruction, and detected road work objects. Bottom row shows the use of data extracted from the dashcam video in a simulation.

Conclusions

Our method for generating images with bad weather and bad lighting were shown to preserve details and create high-quality synthetic images compared to other methods. Our planning framework 1) improves geometric maneuvering capabilities while retaining the reliability and interpretability of structured planning, 2) introduces the first real-time language-action planner for closed-loop driving, and 3) capable of systems that are robust in routine driving and adaptable to edge cases in real-world driving scenarios. Our method for generating data for self-driving simulations can improve self-driving vehicle navigation by enabling learning unsafe activity from dashcam videos.

Recommendations

- In-vehicle safety technologies will greatly benefit from the development and deployment of methods that generate unsafe activity data and use the data, e.g., planning and navigation of self-driving vehicles through roadwork zones, crowded intersections, pedestrian crossings, etc.
- Methods that can be adapted to infrastructure technologies would be invaluable with the growing sensor instrumentation of roadway infrastructure.
- Public methods (code), datasets, and machine learning models would result in accelerated development in the community, more varied datasets and generalized models, efficiencies in data collection and testing, and reduced resource needs.

Publications

- Shen Zheng, Anurag Ghosh, Gaurav Parmar, and Srinivasa Narasimhan. “WarpI2I: Image Warping for Image-to-Image Translation,” European Conference on Computer Vision, 2026.
    - Project Page: https://shenzheng2000.github.io/WarpI2I.github.io/
    - Publication: https://arxiv.org/abs/2606.31018
    - Software: https://github.com/ShenZheng2000/WarpI2I
- Anurag Ghosh, Srinivasa Narasimhan, Manmohan Chandraker, and Francesco Pittaluga. “RAD-LAD: Rule and Language Grounded Autonomous Driving in Real-Time,”
    - Project Page: https://www.cs.cmu.edu/~rad-lad/
    - Paper: https://arxiv.org/abs/2603.28522
- Anurag Ghosh, Francesco Pittaluga, Khiem Vuong, Angela Chen, Juan Alvarez-Padilla, Manmohan Chandraker, Srinivasa Narasimhan. “Dash2Sim: Closed-Loop Driving Simulation from in-the-wild Dashcam Videos,”
    - Project Page: https://www.cs.cmu.edu/~dash2sim/
    - Paper: https://arxiv.org/abs/2606.07366