

Research Experience for Undergraduates (REU): Smart Cities – Advancing Mobility

August 2026

A Research Report from the Pacific Southwest
Region University Transportation Center

Brendan Tran Morris, University of Nevada, Las Vegas



TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No. PSR-24-44	2. Government Accession No. N/A	3. Recipient's Catalog No. N/A	
4. Title and Subtitle Research Experience for Undergraduates (REU): Smart Cities – Advancing Mobility		5. Report Date 08/17/2026	
		6. Performing Organization Code N/A	
7. Author(s) Brendan Tran Morris, ORCID ID # 0000-0002-8592-8806		8. Performing Organization Report No. N/A	
9. Performing Organization Name and Address METTRANS Transportation Center University of Southern California University Park Campus, RGL 216 Los Angeles, CA 90089-0626		10. Work Unit No. N/A	
		11. Contract or Grant No. USC sub# SCON-00005219 USDOT # 69A3552348309	
12. Sponsoring Agency Name and Address U.S. Department of Transportation Office of the Assistant Secretary for Research and Technology 1200 New Jersey Avenue, SE, Washington, DC 20590		13. Type of Report and Period Covered Final report (09/29/2025-05/31/2026)	
		14. Sponsoring Agency Code USDOT OST-R	
15. Supplementary Notes https://smartcities.sites.unlv.edu/ https://doi.org/10.25554/0x3f-cc97			
16. Abstract This project accepted three Las Vegas students into the Spring 2026 Research Experience for Undergraduates (REU): Smart Cities – Advancing Mobility program as an extension of an existing NSF REU program that has been in operation since 2021. The students worked closely with faculty mentors to gain experience in research to help develop the next generation of transportation graduate students and professionals. The student projects included i) Development of an OSHA construction hazard narrative dataset for LLM analysis; ii) Development of a rail anomalies dataset and benchmarking semantic segmentation algorithms for early detection of defects; and iii) X-ray-based part tracking for body kinematics. The students completed their research alongside classes and presented their work at the UNLV Office of Undergraduate Research Symposium. These projects have been instrumental in furthering mentor research agendas and are expected to result in useful products such as public datasets and publications.			
17. Key Words Undergraduate research, transportation safety, machine learning, semantic segmentation, rail anomaly detection, construction safety, computer vision, smart cities.		18. Distribution Statement No restrictions.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 18	22. Price N/A

Form DOT F 1700.7 (8-72)

Reproduction of completed page authorized

Contents

Acknowledgements	5
Abstract	6
Executive Summary	7
Introduction	9
Project 1: Modeling Latent Human Factors in Construction Reports (Alexa Alisangco, Jay Park)	9
OSHA Accident Narrative Database Development	9
Outcomes	10
Project 2: Semantic Segmentation Rail Anomalies Detection (Saarah Majid, Brendan Morris)	11
Rail Anomalies Dataset (RAD)	11
Evaluation Protocol	12
Experimental Results	13
Outcomes	15
Project 3: Automated Rigid Body Tracking for Xray-Based 3D Kinematics (Fabrizio Valdivia, Sarah Harris, David Lee)	15
Experimental Implementation & Results	16
Future Work	16
Conclusion	17
References	17
Data Management Plan	18

About the Pacific Southwest Region University Transportation Center

The Pacific Southwest Region University Transportation Center (UTC) is the Region 9 University Transportation Center funded under the US Department of Transportation’s University Transportation Centers Program. Established in 2016, the Pacific Southwest Region UTC (PSR) is led by the University of Southern California and includes seven partners: Long Beach State University; University of California, Davis; University of California, Irvine; University of California, Los Angeles; University of Hawaii; Northern Arizona University; Pima Community College.

The Pacific Southwest Region UTC conducts an integrated, multidisciplinary program of research, education and technology transfer aimed at *improving the mobility of people and goods throughout the region*. Our program is organized around four themes: 1) technology to address transportation problems and improve mobility; 2) improving mobility for vulnerable populations; 3) Improving resilience and protecting the environment; and 4) managing mobility in high growth areas.

U.S. Department of Transportation (USDOT) 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, by a grant from the U.S. Department of Transportation’s University Transportation Centers Program. However, the U.S. Government assumes no liability for the contents or use thereof.

Disclosure

Brenden Tran Morris, Principal Investigator, conducted this research titled, “Research Experience for Undergraduates (REU): Smart Cities – Advancing Mobility” at Electrical and Computer Engineering, University of Nevada, Las Vegas. The research took place from September 29, 2025, to May 31, 2026, and was funded by a grant from the U.S. Department of Transportation in the amount of \$26,464. The research was conducted as part of the Pacific Southwest Region University Transportation Center research program.

Acknowledgements

This project supplemented the NSF Award # 2349616 “REU Sites: Smart Cities – Advancing Mobility” project.

Abstract

This project accepted three Las Vegas students into the Spring 2026 Research Experience for Undergraduates (REU): Smart Cities – Advancing Mobility program as an extension of an existing NSF REU program that has been in operation since 2021. The students worked closely with faculty mentors to gain experience in research to help develop the next generation of transportation graduate students and professionals. The student projects included i) development of an OSHA construction hazard narrative dataset for LLM analysis; ii) development of a rail anomalies dataset and benchmarking semantic segmentation algorithms for early detection of defects; and iii) X-ray-based part tracking for body kinematics. The students completed their research alongside classes and presented their work at the UNLV Office of Undergraduate Research Symposium. These projects have been instrumental in furthering mentor research agendas and are expected to result in useful products such as public datasets and publications.

Research Report: Research Experience for Undergraduates (REU): Smart Cities – Advancing Mobility

Executive Summary

With continued population growth and urbanization, transportation networks are stressed from aging infrastructure, higher demands, and limited resources. The next generation transportation engineers and operators cannot rely on traditional methods but look to data and machine learning intensive approaches to analyze and understand transportation needs. To address this growing need, the Transportation Research Center (TRC) at the University of Nevada, Las Vegas (UNLV) held a 15-week Research Experience for Undergraduates (REU) program in Smart Cities which paired undergraduates with faculty mentors to gain hands on experience in advanced mobility topics.

While this project was intended to supplement UNLV's existing National Science Foundation (NSF) REU Site (in operation since 2021) taking place during the summer of 2025, students conducted the research while taking classes during the Spring 2026 semester due to the PSR funding pause. The program was adapted from the 10-week program to the 15-week semester program. To accommodate this change, students were selected from a more limited pool of those who were living in the Las Vegas area. Instead of bringing students into the campus dorms, the students used their existing accommodations. The cost saving allowed the REU site to take on three total students instead of the planned two.

The three students came from varied backgrounds, but all worked on computationally challenging projects. Alexa Alisangco was a sophomore in Civil Engineering and mentored by Prof. Jay Park from Civil Engineering. Alexa's research explored the use of large language models (LLMs) to extract important accident factors from OSHA construction accident narratives to develop a new structured dataset from the unstructured narratives. Saarah Majid was a junior in Computer Science and mentored by Prof. Brendan Morris from Electrical Engineering. Saarah collected railway images online and manually segmented all objects of interest, including the target anomalies which could lead to track degradation. The new Rail Anomalies Dataset (RAD) is expected to support a new benchmark for visual rail inspection for the computer vision community. Finally, Fabrizio Valdivia was a senior in computer science mentored jointly by Prof. Sarah Harris from Electrical and Computer Engineering and David Lee in Life Sciences. Fabrizio's research examined the computer vision approach of scene flow to develop a marker less X-ray-based part tracking pipeline for studying body kinematics during movement. All three REU projects supported larger on-going research agendas that included other graduate students.

The REU research has led to the specific outcomes of poster presentations during the Spring 2026 Undergraduate Research Symposium:

- HNSE-P6-6: Alexa Alisangco, Unmesa Ray, and Jee Woong Park, "Modeling Latent Human Factors in Construction Reports through ReAct-Based Reasoning".

- HNSE-P6-1: Saarah Majid, Arian Izadi, Paul Stanik, and Brendan Morris, “Rail Anomalies Dataset for Semantic Segmentation Analysis.
- Poster for Fabrizio Valdivia will be presented during the Summer Undergraduate Symposium on July 31, 2026.

Additionally, there is one journal paper currently under preparation with a publication planned for the other two projects:

- Paul Stanik III, Saarah Majid, Arian Izadi, and Brendan T Morris, “Railroad Anomalies Dataset for Semantic Segmentation Benchmarking and Analysis”, in Trans. Intell. Vehicles, Aug. 2026, to submit.

The REU: Smart Cities – Advancing Mobility program during the successfully introduced three students to advanced data analysis and machine learning approaches during the Spring 2026 semester. The mentorship and hands-on training the students received will better prepare them for graduate studies or to be competitive in the workplace.

Introduction

The Research Experience for Undergraduates (REU) program was specifically designed to provide undergraduate students interested in research an opportunity to learn from close collaboration with faculty mentors. The students receive research training and mentorship specifically in transportation related fields. The goal is to prepare students for advanced graduate degrees and/or to be successful in industry positions.

Due to the change in project dates upon project reinstatement, the scope of the REU program was modified. The project was modified from a 10-week Summer 2025 program into a 15-week Spring 2026 program to account for the May end date. Three Las Vegas locals (rather than the originally planned two out-of-state students) were selected to participate:

- Alexa Alisangco, a Sophomore, Civil Engineering, University of Nevada, Las Vegas; Mentor: Associate Professor Jay Park, Civil Engineering; Project: Modeling Latent Human Factors in Construction Reports
- Saarah Majid, Junior, Computer Science, University of Nevada, Las Vegas; Mentor: Professor Brendan Morris, Electrical and Computer Engineering; Project: Semantic Segmentation for Rail Anomalies Detection
- Fabrizio Valdivia, Senior, Computer Science, University of Nevada, Las Vegas; Mentor: Professor Sarah Harris, Electrical and Computer Engineering; Project:

The REU students took part in mentored research with faculty during the Spring 2026 semester to support broader PSR-relevant research objectives.

Project 1: Modeling Latent Human Factors in Construction Reports (Alexa Alisangco, Jay Park)

Construction accident reports contain important information about work activities, hazards, equipment, and circumstances surrounding occupational injuries. Much of this information is recorded as unstructured narrative text, which makes systematic comparison and analysis difficult. This project focused on developing an Occupational Safety and Health Administration (OSHA) accident narrative database for future construction-safety research. The work included collecting OSHA accident narratives, identifying struck-by hazard cases, extracting relevant information, converting narrative text into a structured binary dataset, developing data visualizations, and organizing literature related to artificial intelligence in construction safety.

OSHA Accident Narrative Database Development

Approximately 500 OSHA accident narratives were collected and compiled as part of a larger dataset containing approximately 1,000 narratives. The data were systematically extracted from official OSHA sources. The narratives were reviewed to identify struck-by hazard cases using predefined keywords and domain-specific guidance. Because accident reports may contain lengthy descriptions, relevant information was separated from unrelated details.

The extracted information was converted into a structured tabular dataset containing more than 15 variables. Binary coding was used to indicate whether selected accident characteristics were present or absent in each narrative. This process transformed unstructured accident descriptions into a format suitable for future statistical and computational analysis.

This process conducted the following activities:

1. Collecting comprehensive accident narratives from OSHA databases.
2. Screening narratives for struck-by hazards.
3. Reviewing reports using keywords and construction-safety guidance.
4. Extracting accident-related information.
5. Excluding information unrelated to the accident analysis.
6. Organizing the extracted information into predefined variables.
7. Converting the variables into a binary tabular database.

Data Visualization

Interactive charts and visual dashboards were developed to compare accident categories with selected risk attributes. These visualizations supported preliminary exploration of the dataset and helped identify patterns or areas that may require further safety analysis. The visualizations were intended to improve understanding of the coded accident data rather than establish causal relationships.

Future Work

The initial goal of the project was to implement ReAct-based reasoning to identify factors from the collected and processed datasets. As the project progressed, the scope was refined to focus on dataset development, attribute extraction, and visualization-based mapping of the extracted attributes.

ReAct-based reasoning is therefore presented as a direction for future work that builds on the activities completed during this project. Future research will apply ReAct-based reasoning to the structured OSHA accident narrative database. ReAct combines language-model reasoning with actions such as retrieving relevant information, examining narrative evidence, and organizing intermediate findings.

This approach may support the identification of latent human factors that are not explicitly represented in standard accident classifications. The structured database and original narratives will provide the foundation for developing and evaluating the proposed reasoning framework.

Outcomes


This REU project resulted in a presentation at UNLV's Office of Undergraduate Research's Spring Undergraduate Research Symposium.

- HNSE-P6-6: Alexa Alisangco, Unmesa Ray, and Jee Woong Park, "Modeling Latent Human Factors in Construction Reports through ReAct-Based Reasoning".

The project produced structured OSHA accident narrative datasets that will be leveraged in Dr. Jay Park's ongoing research.

Project 2: Semantic Segmentation Rail Anomalies Detection (Saarah Majid, Brendan Morris)

Rail maintenance is a core necessity for upkeep of large rail systems and to prevent deterioration. Doing so in an orderly and uniform way, however, has proven difficult with the thousands of miles that railways span. Use of computer vision on locomotives during normal operation can allow for a more streamlined process that could eliminate the need for costly specialized equipment and disruptions of normal operations. This work evaluated popular semantic segmentation approaches for identification of railway anomalies (Figure 1), including vegetation overgrowth, mud-pumping, and standing water, which could be early indicators of rail track failure.



(a) Vegetation overgrowth



(b) Mud-pumping



(c) Standing water

Figure 1: Rail Track Anomalies. Anomalies indicate locations where water may be causing degradation to the rail track along ties, ballast, or the subgrade before damage to the rail itself.

Rail Anomalies Dataset (RAD)

Railroads can be broken down into four main components from top to bottom as rail, ties, ballast, and subgrade. In our approach for analyzing railroads, we focus on the ballast because the most concerning rail anomalies are located in this bottom layer. The ballast is the typical track bed used under most railroad tracks. It is made of crushed stones that transfer the pressure that the rail and ties experience to the subgrade beneath, and it fixes the ties to maintain the correct position of the track, promotes water drainage, and increases the elasticity of the track to allow the rail to return to its original position after a locomotive pass.

Some anomalies affect railroad track health and condition and require maintenance to find and make repairs. Anomalies can include explicit rail issues/defects or proxy measures from tie wear, surface ballast fouling, vegetation growth, poor drainage, and mud pumping. Our work focuses on three anomalies as illustrated in Figure 1: (a) vegetation growth, (b) standing water, and (c) mud pumping.

Poor drainage can cause standing water in the rail tracks, which can cause wooden ties to degrade through rot. Vegetation growth is indicative of weed killer not working properly and a possible sign of poor drainage. Mud pumping is a type of surface ballast fouling that occurs when wet beds appear in clay-like subgrade, causing a clay slurry to emerge from ballast compression when a locomotive passes through. This surface fouling inhibits drainage and prevents the ballast from correctly absorbing the pressure from the ties due to the space created [1]. Maintenance measurements to fix poor drainage and mud pumping can be costly as the ballast and ties must be repaired or replaced.

Although there are many datasets from the ego-vehicle perspective for self-driving car research (e.g., Cityscapes [2]) there is only a single publicly available semantic segmentation dataset for railroads in the RailSem19 dataset [3]. However, these datasets focus on typical scenarios and not the rail anomalies we are interested in. To address this shortcoming, a new Rail Anomalies Dataset (RAD) was curated to further research in the area.

Annotations were manually produced using Supervisely [4] and using the labels in Table 2. The 19 labels were designed to “match” RailSem19 in size and approximate complexity. The labels are similar to RailSem19, but we did modify the labels by combining their **truck** and **car** as a broader **automotive** class because there were not enough trucks in RAD. Since RailSem19 treated all the area between the rails as a single class, it was not suitable for fine analysis of anomalies between the rails or ties. Therefore, we replaced **rail-track** with **tie** and added **mud-pumping**. A total of 152 images were manually annotated and split with 106 for training, 30 for validation, and 16 for testing. Depending on the point of view and how busy the image is, an image typically took thirty minutes to four hours to fully annotate. Example images from RAD are presented in Figure 2 where the left gives the original image and the right the ground truth segmentation mask with vegetation overgrowth in dark green and mud pumping in yellow.




Figure 2: Rail Anomalies Dataset (RAD) Examples

Evaluation Protocol

The full evaluation protocol called for i) the curation of a new dataset for railway anomalies; ii) train a strong baseline segmentation model for transportation/city scenes; iii) transfer the city model to a rail specific context; and iv) fine-tuning the rail model to detect railway anomalies.

Benchmark Models

This project evaluated a number of popular semantic segmentation models using the open-source MMSeg library [5]. In particular, this work focused on convolutional networks BiSeNetV1, DANet, DeepLabV3+, and PSPNet. All approaches used ResNet for the backbone network.

Transfer Learning / Fine-Tuning

Since railway specific image datasets are limited, large publicly available automotive datasets for semantic segmentation were used to pretrain models before adapting to railway specific scenes. Automotive datasets have similar view as locomotives when traveling through city areas which include railway intersections and trams using the same roads as cars providing a reasonable initialization.

Each model was pretrained on the Cityscapes dataset [6] which contains 5000 fully annotated and 20,000 coarsely annotated images from the ego-perspective of an automobile. The dataset contains 35 class labels but only two of the were specific to the rail domain in “rail track” and “train”.

The RailSem19 dataset [3] was used to adapt the Cityscapes model to the rail context. RailSem19 was the first public semantic segmentation dataset in the rail domain. The annotations are a mix of manual and weakly-supervised automated annotations. The dataset consists of a total of 8,500 images from 38 countries in all four seasons and different weather conditions. Many of the class labels (or combinations) are consistent with Cityscapes, and resulted in 19 railway categories. Each model was trained for an additional 80,000 epochs with the modified class labels using the RailSem19 images.

Finally, the models were fine-tuned over 40,000 epochs on the smaller, but more focused RAD. This helped the model retain the railway context of RailSem19 but provided emphasis on rail anomalies. Through this adaption and fine-tune process, the models understood railway scenes and could identify areas of surface ballast fouling, vegetation overgrowth, and standing water.

Experimental Results

Segmentation results were evaluated using the mean intersection over union (mIoU) score which compares the overlap between model prediction and ground truth (higher better). Table 1 provides a comparison between the benchmark models for the CityScapes autonomous car dataset, the railway dataset RailSem19, and the rail anomalies (RAD). Overall, the performance was mixed between models with BiSeNetV1 having the highest mIoU for CityScapes, DeeplabV3+ for RailSem19, and DANet for RAD. These three approaches had similar performance and separated themselves from PSPNet.

Table 1: Benchmark mIoU Performance

Model	CityScapes	Railsem19	RAD
bisenetv1_r18	70.67	56.25	32.21
danet_r101	67.37	53.4	34.88
deeplabv3_plus_r50	68.41	57.36	33.03
pspnet_r50	59.72	48.09	28.95

A more detailed comparison for each class with the BiSeNetV1 network is provided in Table 2. While most of the classes were consistent during the transfer learning and fine-tuning processes, the bolded items were changed to develop the RAD model. With the largest dataset in CityScapes, the performance was pretty consistent between classes. After transferring to RailSem19, there was a fairly significant drop in performance (e.g., road going from 98.14 to 56.58) for some classes while introducing the rail specific objects (e.g., rail-track). Unfortunately, this same issue was observed when moving to RAD, where some specific objects with few examples were completely lost (e.g., traffic-light or traffic-sign). This may indicate overtraining occurred and, thus, more exploration is required.

Table 2: IoU Performance with BiSeNetV1. Bolded classes change meaning during the training process.

Class Comparison			IOU		
Cityscapes	Railsem19	RAD	Cityscapes	RailSem19	RAD
road	road	road	98.14	56.58	30.81
sidewalk	sidewalk	sidewalk	84.68	55.67	4.33
building	construction	construction	92.41	71.61	27.58
wall	tram-track	tram-track	50.17	63.64	0
fence	fence	fence	56.48	44.54	8.97
pole	pole	pole	65.45	54.45	34.7
traffic light	traffic-light	traffic-light	72.33	35.38	0
traffic sign	traffic-sign	traffic-sign	80.32	27.72	0
vegetation	vegetation	vegetation	92.45	82.5	79.27
terrain	terrain	terrain	63.63	61.62	21.03
sky	sky	sky	95.19	94.17	92.65
person	human	person	81.84	44.55	0
rider	rail-track	tie	60.44	80.69	61.11
car	car	automotive	95.12	65.28	0
truck	truck	trackbed	78.74	8.87	72.7
bus	trackbed	on-rails	85.59	67.41	2.1
train	on-rails	rail-raised	78.42	63.73	81.9
motorcycle	rail-raised	rail-embedded	66.74	52.38	0
bicycle	rail-embedded	mud-pumping	77.71	37.91	94.86
mIoU			70.67	56.25	32.21

Table 2 gives detailed IoU results for each individual class at each stage of the training process for the BiSeNetV1 network.

Despite the poor IoU scores, the segmentation results are generally acceptable as evidenced in the Figure below.



Figure 3: Example BiSeNetV1 segmentation results. Left: raw image. Middle: Cityscapes model. Right: Railsem19 model.

Outcomes

This REU project resulted in a presentation at UNLV’s Office of Undergraduate Research’s Spring Undergraduate Research Symposium:

- HNSE-P6-1: Saarah Majid, Arian Izadi, Paul Stanik, and Brendan Morris, “Rail Anomalies Dataset for Semantic Segmentation Analysis.

A paper is expected for submission to the IEEE Transactions on Intelligent Vehicles at the end of the 2026 summer

- Paul Stanik III, Saarah Majid, Arian Izadi, and Brendan T Morris, “Railroad Anomalies Dataset for Semantic Segmentation Benchmarking and Analysis”, in Trans. Intell. Vehicles, Aug. 2026, to submit.

Project 3: Automated Rigid Body Tracking for Xray-Based 3D Kinematics (Fabrizio Valdivia, Sarah Harris, David Lee)

Automatically tracking rigid bodies is useful for detecting detailed kinematics, for example, interacting limbs of humans or other animals. Such detection can be useful in security situations, such as identifying weapons or other threats, and also for more fully understanding locomotion dynamics and dysfunction. This project implements a tracking algorithm for detecting rigid bodies such as bones and predicting their interactions.

While this algorithm can be applied to any rigid body or group of rigid bodies, our current goal is to understand how ground reaction forces (GRFs) and locomotion patterns affect hip mechanics in trotting dogs, with the end goal of understanding the stages of dysfunction, including the existence of osteoarthritis or, in this case, the related canine disease called canine hip dysplasia (CHD) in dogs. However, the hip is difficult to visualize because it is deep-seated and the femoral head is deeply articulated with the pelvis, as a ball-and-socket. To quantify locomotor dynamics, we use a method that

combines marker-based X-ray motion analysis with machine learning, which we are training based on 3D range of motion of canine cadavers with implanted radiopaque markers.

The standard approach for high precision motion tracking of deep structures is to use X-ray motion analysis. However, this method requires surgically implanted bone markers which is an invasive, labor-intensive process that is impractical for canine subjects. To overcome this, we combine marker-based X-ray motion analysis of cadavers with machine learning, thus enabling the tracking of rigid bodies without implanted markers. We used a cadaver study to obtain a training dataset and completed marker-based X-ray motion analysis to determine precise translations and rotations of skeletal rigid bodies (i.e., the pelvis and femora). The machine learning component here is MS-RAFT-3D, a scene flow estimation method that calculates how a rigid 3D object rotates and translates in space across multiple scales of motion [7, 8]. This method allows us to quantify rigid body dynamics and visualize deep structures without invasive implants.

Experimental Implementation & Results

We have successfully collected an X-ray dataset in two planes with radiopaque markers implanted in the rigid bodies of canine cadavers, which are implanted in the pelvis, right femur, and left femur. We have collected movement through a range of motion that is tracked by X-ray motion analysis (XMA Lab software). These results act as the ground truth based on the precise tracking of radiopaque markers, which results are used to compare the results from our MS-RAFT-3D adaptation. Figure 1*** below shows example 2D X-ray images from the dataset that we generated.



Figure 4: 2D X-ray image of dog skeleton with radiopaque marker sets in rigid bodies (pelvis, left femur, right femur)

In addition to collecting the X-ray 2D video images, we have implemented MS-RAFT-3D in Python and have run example data sets on it. Adapting the MS-RAFT-3D algorithm to X-ray images is an advancement and offers a challenge because the rigid bodies appear translucent in the X-ray images. So, contours from both sides of the rigid body are visible, unlike rigid body images captured from a typical camera using visible light.

Future Work

Our future steps include continuing to adapt MS-RAFT-3D to be able to properly read in the tiff files from the canine datasets and capture the 3D translation and rotation of rigid objects based on skeletal

contours, followed by validation of these results movements known precisely from marker-based X-ray motion analysis. The approach developed here may be generalized to tracking other rigid bodies imaged with X-rays, such as in security applications.

Conclusion

Continuing the tradition started with the original NSF REU: Smart Cities – Intelligent, Safe, and Secure Mobility program in the summer of 2021, this PSR UTC project provided support for three local Las Vegas students to take part the Smart Cities REU in the Spring of 2026. The students worked closely with UNLV faculty mentors for 15 weeks to receive hands on experience and research training to support the development of future transportation leaders. The student presented their work with a poster at the UNLV Office of Undergraduate Research Symposium. Additionally, the students are contributing to future publications.

References

- [1] A. Hudson, G. Watson, L. Le Pen and W. Powrie, "Remediation of Mud Pumping on a Ballasted Railway Track," *Procedia Engineering*, vol. 143, pp. 1043-1050, 2016.
- [2] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth and B. Schiele, "The Cityscapes Dataset for Semantic Urban Scene Understanding," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [3] O. Zendel, M. Murschitz, M. Zeilinger, D. Steininger, S. Abbasi and C. Beleznai, "RailSem19: A Dataset for Semantic Rail Scene Understanding," in *Conference on Computer Vision and Pattern Recognition (CVPR) Workshop, 2019*, 2019.
- [4] Supervisely, "Supervisely," [Online]. Available: <https://supervise.ly>. [Accessed 2026].
- [5] M. Contributors, "MMSegmentation: OpenMMLab Semantic Segmentation Toolbox and Benchmark," 2020. [Online]. Available: <https://github.com/open-mmlab/mmdetection>.
- [6] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth and B. Schiele, "The Cityscapes Dataset for Semantic Urban Scene Understanding," 2016, *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [7] Z. Teed and J. Deng, "RAFT-3D: Scene Flow Using Rigid-Motion Embeddings," in *CVPR*, 2021.
- [8] J. Schmid, A. Jahedi, N. B. Senn and A. Bruhn, "MS-RAFT-3D: A Multi-Scale Architecture for Recurrent Image-Based Scene Flow," in *IEEE International Conference on Image Processing*, 2025.

Data Management Plan

Products of Research

This project is expected to produce the following datasets: (1) a structured OSHA accident narrative database for Project 1; (2) the Rail Anomalies Dataset (RAD) with semantic segmentation annotations for Project 2; and (3) X-ray video imagery with radiopaque marker tracking data for Project 3. Dataset generation is currently in progress and will be completed following the REU program conclusion.

Data Format and Content

Data will be provided in standard formats including CSV (tabular OSHA database), image files with JSON/XML annotation files (RAD semantic segmentation), and TIFF video files (X-ray imaging data).

Data Access and Sharing

Upon completion, datasets will be made publicly available through a recognized repository (e.g., Dryad or Dataverse). Access links will be provided upon deposit.

Reuse and Redistribution

Data will be made available for unrestricted public use for research and educational purposes, subject to applicable UNLV and sponsor policies. [Hyperlink for the dataset to be deposited]