TRANSPORTATION POOLED FUND PROGRAM
QUARTERLY PROGRESS REPORT

Lead Agency (FHWA or State DOT): Kansas DOT

INSTRUCTIONS:

Lead Agency contacts should complete a quarterly progress report for each calendar quarter during which the projects are active. Please provide a project schedule status of the research activities tied to each task that is defined in the proposal; a percentage completion of each task; a concise discussion (2 or 3 sentences) of the current status, including accomplishments and problems encountered, if any. List all tasks, even if no work was done during this period.

Transportation Pooled Fund Program Project # *(i.e, SPR-2(XXX), SPR-3(XXX) or TPF-5(XXX)* **TPF-5(535)**	**Transportation Pooled Fund Program - Report Period:** ☒Quarter 1 (January 1 – March 31) ☐Quarter 2 (April 1 – June 30) ☐Quarter 3 (July 1 – September 30) ☐Quarter 4 (October 1 – December 31)

TPF Study Number and Title: **TPF-5(535): Human-centered Steel Bridge Inspection enabled by Augmented Reality and Artificial Intelligence**		
Lead Agency Contact: **David Behzadpour**	**Lead Agency Phone Number:** **785-291-3847**	**Lead Agency E-Mail** David.Behzadpour@ks.gov
Lead Agency Project ID: Click or tap here to enter text.	**Other Project ID (i.e., contract #):** Click or tap here to enter text.	**Project Start Date:** 10/1/2024
Original Project Start Date: **10/1/2024**	**Original Project End Date:** **9/30/2027**	**If Extension has been requested, updated project End Date:** Click or tap to enter a date.

Project schedule status:

☒On schedule	☐On revised schedule	☐Ahead of schedule	☐ Behind schedule

Overall Project Statistics:

Total Project Budget	Total Funds Expended This Quarter	Percentage of Work Completed to Date
$600,000	$95	10%

Project Description:

The main objective of this proposed research is to provide state DOTs practical tools for supporting human-centered steel bridge inspection with real-time defect (e.g., fatigue cracks and corrosion) detection, documentation, tracking, and decision making. The proposed research will not only bridge the gaps identified in the IDEA project, but also expand the existing capability by developing AI algorithms for crack and corrosion detection. In addition to AR headsets, the project will also develop AR-based inspection capability using tablet devices. The tablet device can be used to perform AR-based inspection directly in a similar way to the AR headset. It can also leverage Unmanned Aerial Vehicles (UAV) for remote image and video acquisition during inspections, enabling bridge inspections from a distance in a human-centered manner.

Progress this Quarter
(includes meetings, work plan status, contract status, significant progress, etc.):

1. **Task 1: CV and AI algorithms for crack and corrosion inspection**

In the previous report, we evaluated the performance of YOLOv11 and DeepLabV3+ models for detecting steel fatigue cracks using the IPC-SHM 2020 dataset. Since YOLOv11 demonstrated superior performance compared to DeepLabV3+, we proceeded with the YOLO model for simultaneous detection of both fatigue cracks and corrosion on newly captured image data from DOT. Inaccurate labeling can mislead the model during training, leading to poor generalization and performance. However, producing high-quality labeled datasets is labor-intensive and time-consuming, especially for pixel-level tasks like crack and corrosion segmentation.

To address the challenges, this project adopts a flexible hybrid framework incorporating elements from transfer learning, domain adaptation, unsupervised and semi-supervised learning. We begin by leveraging pre-trained models trained on publicly available datasets and fine-tune them using field-collected images of steel bridges. This transfer learning step enables the model to benefit from generalized feature representations while adapting to the target domain features.

Following initial fine-tuning on the source domain, the model is deployed to generate predictions on the unlabeled target domain. High-confidence predictions are retained using a confidence threshold, and uncertain samples are filtered out. These pseudo-labeled target samples are iteratively incorporated into the training pipeline, forming the basis for self-training and potential integration with active learning strategies.

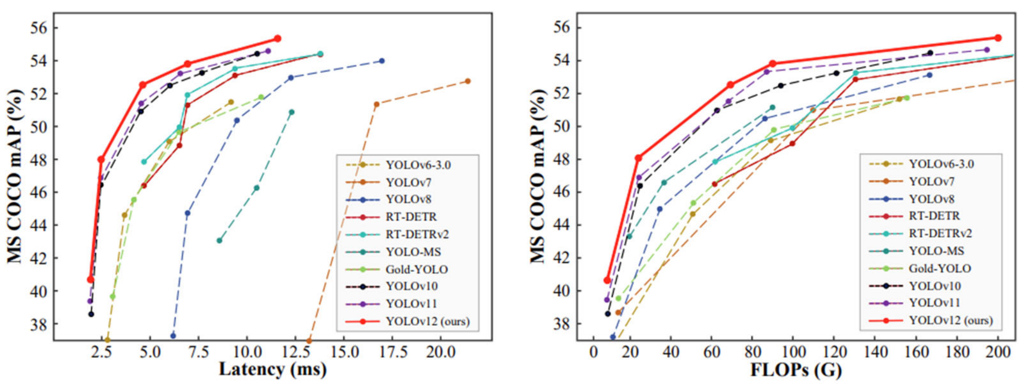

Fig. 1 YoloV12 in comparison with other popular models (Tian et al., 2025).

For the damage segmentation task, we adopt YOLOv12, the latest advancement in the YOLO object detection family. YOLOv12 introduces a novel attention-centric architecture optimized for real-time performance, integrating area attention, R-ELAN (Residual Efficient Layer Aggregation Networks), and improved attention modules (Tian et al., 2025). These enhancements allow YOLOv12 to effectively balance computational efficiency and detection accuracy. According to benchmark evaluations (Figure 1), YOLOv12 consistently outperforms prior models across multiple latency and computational cost settings, achieving higher mean Average Precision (mAP) on MS COCO with reduced inference time and FLOPs.

The final output of this methodology is a robust, trained deep learning model capable of accurately segmenting cracks and corrosion. This model will be integrated into an Augmented Reality (AR)-based monitoring system, enabling real-time visualization and assessment of structural damage. Using AR devices, the inferred damage can be spatially anchored onto the 3D surface of bridge structures, allowing inspectors to interactively explore and assess critical areas in real-world settings.

2. **Task 2: AR-based software for human-centered bridge inspection**
Continued development is progressing with the AR infrastructure inspection app. We have made the decision to shift the platform of this app off of the Microsoft HoloLens due to Microsoft's planned end of support of the headset. As a result, the project will now transition to the Magic Leap 2 headset which provides all the essential functionality required for our inspection tasks. Fortunately, the overall tech stack for the app will remain largely unchanged, as Magic Leap 2 supports most of the features of Microsoft's MRTK3. The Magic Leap 2 utilizes the Android OS build platform, which allows us to use a wider variety of software packages that were not available to the HoloLens 2. As a result of these changes, we have revised our inference pipeline and will now utilize OpenCV for Unity, as it allows quick, native processing of Yolo models on the headset. This coupled with the collaboration between graduate students has led to successful improvements in the inference process, enabling accurate native inferencing to be achieved. To accommodate the new hardware, we will be using a different spatial persistence approach using Magic Leap's spatial anchors to ensure inspection data can be referenced across sessions.

Anticipated work next quarter:

1. **Task 1: CV and AI algorithms for crack and corrosion inspection**
Integrating additional publicly available data and implementing an active learning framework to generate high-confidence pseudo labels significantly contributes to improving the model's performance through iterative self-supervised refinement.

2. **Task 2: AR-based software for human-centered bridge inspection**
During the next reporting period, we aim to complete the transition of our application to Magic Leap 2. A key focus will be to integrate the Magic Leap's spatial anchors to achieve spatial persistence across sessions and inspections. Development will also begin on a database system that will store inspection results for long-term reuse. We hope to begin initial testing with inspectors to gather practical feedback and inform further development. We also hope to begin early development of a drone-based inspection application.

Significant Results:

1. **Task 1: CV and AI algorithms for crack and corrosion inspection**

To construct the pre-labeled dataset required for training the YOLOv12 model, we integrate two domain-specific datasets: the IPC-SHM 2020 dataset for fatigue crack detection, and a corrosion condition state dataset for bridge surface degradation analysis.

The IPC-SHM 2020 dataset comprises 200 high-resolution images (4928 × 3264 pixels) depicting steel fatigue cracks on welded joints. Of these, 120 images include pixel-wise binary annotations, where crack regions are labeled as 1 and the background (non-crack) as 0. The original binary mask annotations are preprocessed and converted into YOLO-compatible .txt format for polygonal segmentation.

The corrosion condition state dataset contains 440 images at resolutions of 1600×1200 and 1002×750 pixels. This dataset provides pixel-level annotations across multiple corrosion severity levels—Fair, Poor, and Severe. To standardize the input size and facilitate training, the images are partitioned into 512×512-pixel patches, and the corresponding-colored segmentation masks are converted into YOLO polygon format with appropriate class labels.

These two datasets form the foundation for hybrid approach in our training pipeline. Representative samples from both datasets are shown in the figures to illustrate the range of damage patterns captured across different classes and severity levels.

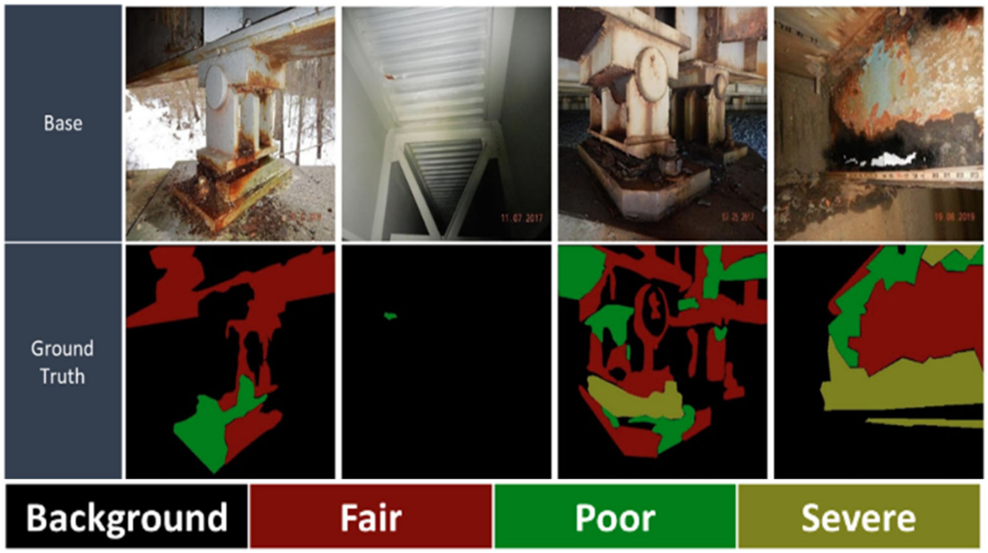

Fig. 2 Corrosion Dataset

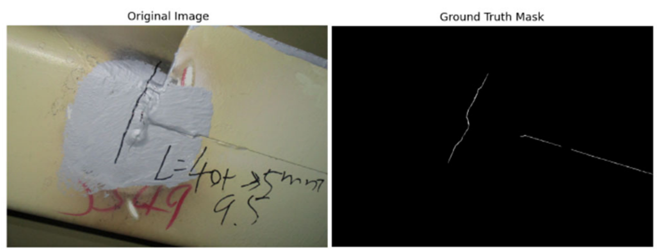

Fig. 3 Fatigue Crack Dataset

By combining both datasets, we generate a total of 1,496 image patches of size 512×512 pixels, encompassing two classes: crack and corrosion. To enhance model generalization and increase data diversity, we apply a range of data augmentation techniques, including random rotation, shear transformation, scaling, HSV color space jittering, and horizontal flipping. These augmentations simulate various real-world imaging conditions and structural variations.

Following dataset preparation, the YOLOv12 model is fine-tuned using the augmented dataset and trained for 400 epochs. The segmentation performance is quantitatively assessed using the mean Average Precision at an IoU threshold of 0.5 (mAP@0.5). As shown in the corresponding results table, the overall mAP@0.5 across all damage types is 45%. Notably, the segmentation performance on corrosion damage is lower compared to cracks, which can be attributed to the higher proportion of background pixels and the diffuse, less well-defined boundaries typical of corrosion features. This leads to greater ambiguity in localization, thus reducing detection accuracy for this damage type.

Table. 1 Performance of tuned Yolov12 on public datasets

mAP@0.5	
Crack	0.53
Corrosion	0.37
All	0.45

We utilize the pretrained YOLOv12 model to generate pseudo-labels on a new dataset collected by the Texas Department of Transportation. This dataset consists of 223 high-resolution images with dimensions of 5712×4284 pixels, capturing real-world bridge surfaces under various environmental conditions. As shown in the figure below, YOLOv12 demonstrates relatively good performance in segmenting corroded regions, indicating effective generalization from the initial training dataset. However, due to the limited presence of fatigue cracks in this dataset, the model shows no significant detection of cracks. Using active learning procedure, we will reduce the positive and negative errors gradually.

Fig. 4 Pseudo labels generated on field images

In the next step, we will apply confidence-based filtering to retain only those predictions with high-confidence scores, reducing the likelihood of false positives. These filtered predictions are then incorporated into the training pipeline as pseudo-labeled data to further fine-tune the model. To enhance the robustness and generalizability of the YOLOv12 model, we also plan to supplement the training data set with additional samples from other publicly available datasets, particularly those rich in fatigue crack instances.

2. Task 2: AR-based software for human-centered bridge inspection

Migration from the HoloLens 2 to the Magic Leap 2 platform has begun and much progress has been made. The Magic leap provides extensive and well-documented support for porting applications from the HoloLens platform. We have been steadily progressing through the feature list and using these guides throughout the process. As shown in Figure 5, a new Unity project was created with the build conditions and software packages that are required for the Magic Leap 2 Hardware.

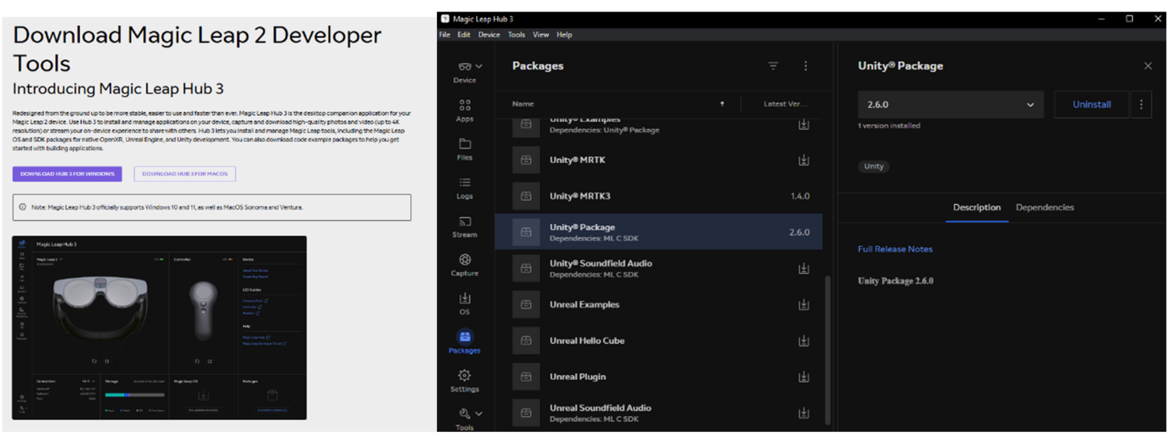

Fig. 5 Development tools for Magic Leap Unity Development

Similar to Microsoft's MRTK3 feature tool, used for installing MRTK plugins into a Unity project to allow for HoloLens development, the Magic Leap platform has the Magic Leap hub. The Magic Leap Hub allows developers to download and install software packages that allow Unity to target the Magic Leap and use its features during development. We have installed the Magic Leap Hub and all the relevant packages for the Magic Leap development to create the current Unity project.

So far as part of our migration, key components of the AR application have been successfully ported and adapted. This includes the user interface, camera management scripting, and projector components used for visualizing ML inferences. The spatial mapping functionality has also been migrated, enabling real-time surface mapping for inference outputs. In order to port prior progress, changes had to be made to scripts and to project settings as they were imported into the new project so that they reference the Magic Leap 2's feature set and AR management software package rather than Microsoft's.

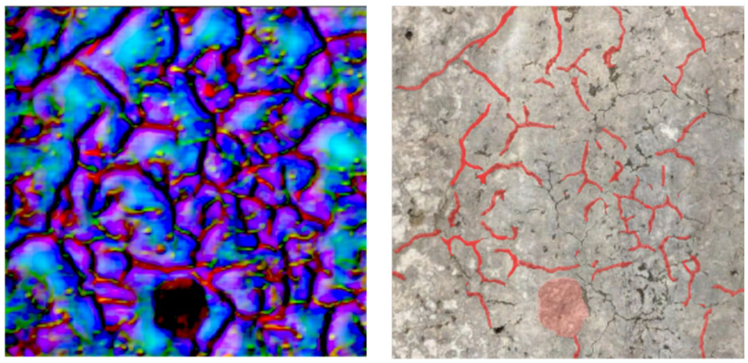

Fig. 6 Sentis inference outputs (previous) vs. OpenCV for Unity inference outputs (current)

A significant amount of development time has been spent troubleshooting incorrect inference results on HoloLens 2. These issues were largely due to compatibility problems between our YOLO models and the Unity ML packages available on that platform. The use of Unity's Sentis inference engine would consistently produce flawed outputs with incorrect color channels. With the shift to Magic Leap 2, an Android OS based system, we are now able to utilize OpenCV for Unity, which offers built-in support for some YOLO processing. This change has resulted in accurate inference outputs that clearly outline detected cracks with correct color channels and allows for the use of bounding boxes. The figure above compares the old inference outputs to the new ones processed with OpenCV for Unity.

To replace MRTK3's World Locking Tools used for spatial persistence on the HoloLens, the Magic Leap 2 platform offers its own system of spatial anchors. These anchors will enable us to maintain hologram positions from inspection results across sessions/inspections, allowing inspectors to revisit previous assessments and identify damage progress. We are currently in the process of identifying the relevant APIs and beginning the integration of spatial anchors into our application.

Although much of the porting work has been completed, we have only recently received our Magic Leap 2 headset. We are beginning our hands on testing with the device. We are currently focused on ensuring that the migrated functionality behaves identically to the original implantation, and that the new inference pipeline seamlessly integrates with the actual Magic Leap 2 hardware.

Circumstance affecting project or budget. (Please describe any challenges encountered or anticipated that might affect the completion of the project within the time, scope and fiscal constraints set forth in the agreement, along with recommended solutions to those problems).

Microsoft has officially ended support for the HoloLens 2, ceasing both the sale of new units and software support by 2027. To ensure our AR software package remains accessible and practical for field inspectors, we have decided to transition development to the Magic Leap 2 headset.

This shift has introduced a slight delay in the AR software timeline as we work to migrate the application to the new platform. Unlike the HoloLens, which operates on the Universal Windows Platform (UWP), the Magic Leap 2 runs on Android OS. This change brings both opportunities and challenges; while some applications previously inaccessible on UWP are now available on Android, adjustments are required to accommodate the new environment.

Additionally, student recruitment for this project has been affected by ongoing visa challenges for international students. Despite this, the project team is actively working to onboard graduate research assistants by summer 2025.

Potential Implementation: