



---

**Innovations Deserving  
Exploratory Analysis Programs**

***Rail Safety IDEA Program***

---

# **AUTONOMOUS DETECTION OF COMPRESSED AIR LEAKS ON TRAINS**

Final Report for  
Rail Safety IDEA Project 48

Prepared by:  
Christopher Stoos  
Southwest Research Institute

***September 2022***

---

**NATIONAL** Sciences  
**ACADEMIES** Engineering  
Medicine

 **TRANSPORTATION RESEARCH BOARD**

## **Innovations Deserving Exploratory Analysis (IDEA) Programs Managed by the Transportation Research Board**

This IDEA project was funded by the Rail Safety IDEA Program.

The TRB currently manages the following three IDEA programs:

- The NCHRP IDEA Program, which focuses on advances in the design, construction, and maintenance of highway systems, is funded by American Association of State Highway and Transportation Officials (AASHTO) as part of the National Cooperative Highway Research Program (NCHRP).
- The Rail Safety IDEA Program currently focuses on innovative approaches for improving railroad safety or performance. The program is currently funded by the Federal Railroad Administration (FRA). The program was previously jointly funded by the Federal Motor Carrier Safety Administration (FMCSA) and the FRA.
- The Transit IDEA Program, which supports development and testing of innovative concepts and methods for advancing transit practice, is funded by the Federal Transit Administration (FTA) as part of the Transit Cooperative Research Program (TCRP).

Management of the three IDEA programs is coordinated to promote the development and testing of innovative concepts, methods, and technologies.

For information on the IDEA programs, check the IDEA website ([www.trb.org/idea](http://www.trb.org/idea)). For questions, contact the IDEA programs office by telephone at (202) 334-3310.

IDEA Programs  
Transportation Research Board  
500 Fifth Street, NW  
Washington, DC 20001

The project that is the subject of this contractor-authored report was a part of the Innovations Deserving Exploratory Analysis (IDEA) Programs, which are managed by the Transportation Research Board (TRB) with the approval of the National Academies of Sciences, Engineering, and Medicine. The members of the oversight committee that monitored the project and reviewed the report were chosen for their special competencies and with regard for appropriate balance. The views expressed in this report are those of the contractor who conducted the investigation documented in this report and do not necessarily reflect those of the Transportation Research Board; the National Academies of Sciences, Engineering, and Medicine; or the sponsors of the IDEA Programs.

The Transportation Research Board; the National Academies of Sciences, Engineering, and Medicine; and the organizations that sponsor the IDEA Programs do not endorse products or manufacturers. Trade or manufacturers' names appear herein solely because they are considered essential to the object of the investigation.

# **AUTONOMOUS DETECTION OF COMPRESSED AIR LEAKS ON TRAINS**

## **IDEA Program Final Report**

### **Rail Safety IDEA - 48**

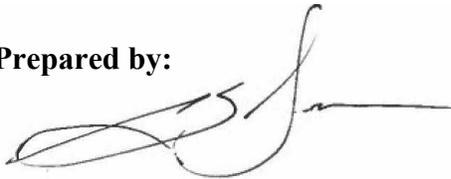
**SwRI® Project No. 03.26792**

**Prepared for:  
Rail Safety IDEA Program  
Transportation Research Board  
The National Academies of Sciences, Engineering, and Medicine**

*Christopher Stoos, Lead Engineer  
Southwest Research Institute  
6220 Culebra Road  
San Antonio, TX 78238*

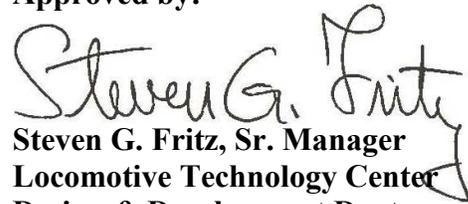
*September 2022*

**Prepared by:**



**Christopher Stoos, Lead Engineer  
Locomotive Technology Center  
Design & Development Dept.**

**Approved by:**



**Steven G. Fritz, Sr. Manager  
Locomotive Technology Center  
Design & Development Dept.**

#### **POWERTRAIN ENGINEERING DIVISION**

This report shall not be reproduced, except in full, without the written approval of Southwest Research Institute®.  
Results and discussion given in this report relate only to the test items described in this report.

## **ACKNOWLEDGEMENTS**

Many people have contributed to the success of this project. The SwRI project team would like to thank the following for their support along the way:

Inam Jawed from the National Academy of Sciences for guidance and support throughout the project.

Brad Kerchof (Advanced Rail Management) and Tom Bartlett (ARMS) as project advisors.

Mike Iden (Tier 5 Locomotive), Sean Cronin (Metra), Michael Cleveland (Progress Rail), and Steven Albright (Canadian National Railway) for participating in the project as our Expert Advisory Panel.

Rick Scholte and Paul van Dooren from Sorama, who provided insight, expertise and guidance throughout the project.

Martin Budweg, Simon Powell, Dylan Erwin, Craig Haase and Rob Hornsblow from Fluke who provided the SV600 sensors used in the project, as well as technical information and guidance.

## **PROJECT TEAM**

### **National Academy of Sciences**

Gwen Chisholm-Smith, TCRP Manager  
Inam Jawed, Program Manager

### **Project Advisors**

Brad Kerchof, Advanced Rail Management  
Tom Bartlett, ARMS

### **Expert Advisory Panel**

Michael Iden, Tier 5 Locomotive LLC  
Michael Cleveland, Progress Rail  
Steven Albright, Canadian National  
Sean Cronin, Metra

### **Southwest Research Institute**

Christopher Stoos, Lead Engineer  
Shane Siebenaler, Director  
Thomas Reinhart, Institute Engineer  
Matthew Hoffmeyer, Sr. Research Engineer  
Heath Spidle, Sr. Research Engineer  
Jake Janssen, Research Engineer  
Steven Fritz, Senior Manager

### **Fluke Process Instruments**

Simon Powell, Business Unit Leader  
Martin Budweg, Manager  
Craig Haase, Account Sales Manager  
Dylan Erwin, Technical Sales Manager

### **Sorama**

Rick Scholte, Founder and CEO  
Paul van Dooren, Product Owner Listener Platform



**RAIL SAFETY IDEA PROGRAM  
COMMITTEE**

**CHAIR**

MARTITA MULLEN  
*Canadian National Railway*

**MEMBERS**

TOM BARTLETT  
*Automated Railroad Maintenance Systems*  
MELVIN CLARK  
*LTK Engineering Services*  
MICHAEL FRANKE  
*Retired Amtrak*  
BRAD KERCHOF  
*Norfolk Southern Railway*  
STEPHEN M. POPKIN  
*Volpe National Transportation Systems  
Center*

**APTA LIAISON**

NARAYANA SUNDARAM  
*American Public Transportation Association*

**FRA LIAISON**

TAREK OMAR  
*Federal Railroad Administration*

**TRB LIAISON**

SCOTT BABCOCK  
*Transportation Research Board*

**IDEA PROGRAMS STAFF**

GWEN CHISHOLM-SMITH, *Manager, Transit  
Cooperative Research Program*  
INAM JAWED, *Senior Program Officer*  
VELVET BASEMERA-FITZPATRICK, *Senior Program  
Officer*  
DEMISHA WILLIAMS, *Senior Program Assistant*

**EXPERT REVIEW PANEL  
SAFETY IDEA PROJECT**

**48**  
BRAD KERCHOF, *Norfolk Southern Railway (Retd.)*  
TOM BARTLETT, *Automated Railroad Maintenance  
Systems*  
MICHAEL IDEN, *Tier 5 Locomotive, LLC*  
MICHAEL CLEVELAND, *Progress Rail*  
STEVEN ALBRIGHT, *Canadian National Railway*  
SEAN CRONIN, *Metra*



## TABLE OF CONTENTS

Acknowledgements .....	ii
Project Team.....	iii
Executive Summary.....	vi
1.0 Project Background.....	1
2.0 Concept and Innovation .....	1
3.0 Investigation.....	2
3.1 Hardware Selection and Acquisition.....	2
3.2 Initial Testing .....	3
3.3 AI Integration and Final Testing .....	8
4.0 Plans for Implementation .....	14
5.0 Conclusions .....	15
6.0 Lessons Learned.....	15
7.0 Lead Investigator Profiles .....	17
8.0 References.....	18

## APPENDICES

Appendix A: Research Results .....	No. of Pages 2
------------------------------------	-------------------

## LIST OF FIGURES

Figure ES-1. Air Leak at 3.0, 4.5, 6.0, and 7.5 Meter Distances .....	vi
Figure ES-2. Example Composite Image with Leak Locations Highlighted.....	vii
FIGURE 1. Fluke SV600 Fixed Acoustic Imager .....	2
FIGURE 2. Joule-Thompson Effect Cooling at Various JT Coefficients.....	3
FIGURE 3. Air Leak at 3.0, 4.5, 6.0, and 7.5 Meter Distances .....	4
FIGURE 4. Compressed Air Leaks at 51dB Minimum Detection Threshold .....	4
FIGURE 5. Compressed Air Leaks at 40dB Minimum Detection Threshold .....	5
FIGURE 6. Compressed Air Leaks at 45dB Minimum Detection Threshold .....	5
FIGURE 7. Compressed Air Leak on Piping Using Thermal and Visual Spectrum Cameras .....	6
FIGURE 8. Compressed Air Leak at Drain Valve Using Thermal and Visual Spectrum Cameras .....	7
FIGURE 9. Train Air Leak Detection Using 3 Cameras .....	7
FIGURE 10. Annotated frame in CVAT .....	8
FIGURE 11. Pixelwise Semantic Segmentation Mask of Air Leak .....	9
FIGURE 12. YOLOv5 Detection of an Air Leak .....	9
FIGURE 13. Example Composite Image of Passing Locomotive with Detected Air Leaks .....	10
FIGURE 14. Example Air Leak Notification Alert .....	10
FIGURE 15. SSD Training and Validation Metrics .....	11

FIGURE 16. SSD Confusion Matrix .....	12
FIGURE 17. SSD Validation Examples .....	13
FIGURE 18. Spatial and Temporal Tracking of an Air Leak .....	13
FIGURE 19. Air Leaks in Air Brake Compartment Below Cab (left) and Rear Sander Mag Valve (right).....	16
FIGURE A1. Detection of an Air Leak .....	2
FIGURE A2. Example Composite Image of Passing Locomotive with Detected Air Leaks .....	2

### **LIST OF TABLES**

TABLE 1. Fluke SV600 Fixed Acoustic Imager Specs.....	2
TABLE 2. Camera Specs.....	3
TABLE 3. Model Validation Results .....	12
TABLE 4. Air Leak Tracking Metrics.....	14

## EXECUTIVE SUMMARY

The purpose of this Rail Safety IDEA project was to develop a proof-of-concept system to autonomously detect compressed air leaks on moving trains. Since the introduction of the air brake system, the rail industry has considered minor air leaks to be the cost of doing business due to the time and expense required to find and fix the leaks, but minimizing compressed air leaks is one of the key areas in which railroads can make significant gains towards reducing fuel consumption and exhaust emissions. Past practices with regards to air leaks have not focused on fuel efficiency and have generally focused on operability by allowing some level of leaks on certain compressed air systems on the train.

While brake tests can sometimes tell you if you have a leak, they cannot tell you where within the train the leak is. Finding these leaks requires an employee to manually search, often unnecessarily going on, under, or between rolling stock to listen or feel for leaks. It is an inefficient, labor intensive, and time-consuming process, which is why the Federal Railroad Administration (FRA) and railroads have defined acceptable leak rates. Automated detection of locomotive and train air leaks could reduce the time and labor necessary to find air leaks, reduce locomotive fuel consumption and exhaust emissions, and at the same time improve employee safety.

During stage 1 of the system development, Southwest Research Institute (SwRI) sourced the hardware necessary for the system. This included a visual spectrum camera and a thermal imaging camera owned by SwRI as well as two Fluke SV600 fixed acoustic sensors. The Fluke sensors are a commercially available sensor that contains a 64-microphone array which, using technology developed by Fluke and their technology partner Sorama, provide an audio mask overlay on a visual spectrum image to show where sound sources are coming from. The sensor allows the user to define upper and lower limits for decibel as well as frequency.

During initial testing, the optimal settings to detect air leaks while mitigating other sounds were determined. A versatile stand to support the SV600 sensors and cameras was built and various positions of the sensors in relation to the track were investigated (Figure ES-1). It was determined that air leaks were best captured with a minimum decibel threshold of 45dB and a frequency range of 30kHz-45kHz.



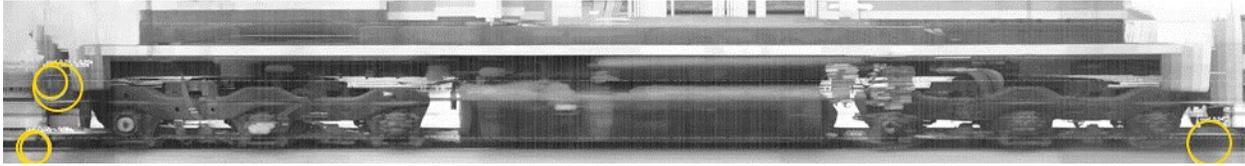
**Figure ES-1. Air Leak at 3.0, 4.5, 6.0, and 7.5 Meter Distances**

Stage 2 of the development focused on implementation of the machine vision system. Various models were implemented to see which was best suited to the task. The data was initially annotated using Computer Vision Annotation Tool (CVAT), manually labeling known leaks and gladhands in the initial round, while reserving 20% of the initial data as a validation set. Two types of convolutional neural network (CNN) were then trained and evaluated for the detection of air leaks. Both mobilenetV3 and YoloV5 were assessed, with the YoloV5 model producing better overall results.

The model was still triggering too many false positives, so SwRI added tracking logic which looked for leak persistence through time as it progresses through the field of view for the sensor. This spatial and temporal tracking allowed for the elimination of a large portion of the false positives that were seen before by making sure that the signal tracked through the entire field of view.

Once detection (and false detection) had been accomplished at an acceptable level, a basic notification system was developed. When a leak is detected, a composite image of the piece of equipment is created with the location of the detected leaks marked. Figure ES-2 shows an example of this composite image. In addition to the images, a

notification is generated and sent to mechanical department personnel with the number of leaks, the leak locations, and the system confidence in each leak detected.



**Figure ES-2. Example Composite Image with Leak Locations Highlighted**

This Rail Safety IDEA project has resulted in the successful completion of a proof-of-concept system that can autonomously detect compressed air leaks on moving trains and notify mechanical personnel of the exact leak location. At this point, the system accurately detects and flags air leaks 84.6% of the time, with false positives occurring only 0.03% of the time. With further field development and data collection, it is expected that the system accuracy on air leak detection and the false positive rate would both improve significantly.

## **1.0 PROJECT BACKGROUND**

Freight trains use compressed air for various important functions, including air brakes, radiator shutters, valve actuation, horns, and more. The compressed air for these systems is generated using compressors on each locomotive in the train. Air leaks are a significant problem, but due to the difficulty in detecting them, 49 CFR 232 [1] defines a 60 standard cubic feet per minute (SCFM) upper limit on brake pipe leakage for trains utilizing head end power only, and a 90 SCFM upper limit on brake pipe leakage for trains using Distributed Power (DP).

Southwest Research Institute (SwRI) has performed testing which showed that air leaks have a significant negative impact on the fuel efficiency and total exhaust emissions. Increases in Net Traction Specific Fuel Consumption (NTSFC) penalties due to air leaks have been shown to be 2% over the linehaul duty cycle and over 4% on the switcher duty cycle with a 30 SCFM brake pipe leak, and over 5% using the linehaul duty cycle and an astonishing 14% NTSFC penalty over the switcher duty cycle with a 60 SCFM leak [2].

NTSFC can be looked at as a measure of overall vehicle efficiency for locomotives. It is a measure of how much fuel is used by mass per unit of tractive effort, usually expressed in units of pounds per traction horsepower-hour (lb/thp-hr). By reducing NTSFC, you increase the amount of work a locomotive can do pulling freight or passenger cars per gallon of fuel.

While air brake tests can sometimes tell you if you have a leak, it cannot tell you where within the train the leak is. Finding these leaks requires an employee to manually search, often going on, under, or between rolling stock to listen or feel for leaks. It is an inefficient and time-consuming process, which is why railroads have acceptable leak rates. Automated detection of locomotive and train air leaks could reduce the time and labor necessary to find air leaks, reduce locomotive fuel consumption and exhaust emissions, and at the same time improve employee safety. Given the additional factor of safety that reduced air leakage would have on the air brake systems on trains, this could also improve operational rail safety for the nation's transportation sector.

The purpose of this Rail Safety IDEA project was to develop and test a proof-of-concept autonomous air leak detection system which can notify mechanical personnel of the exact location of air leaks.

## **2.0 CONCEPT AND INNOVATION**

Autonomous compressed air leak detection on trains is a difficult endeavor. Railroads are loud environments with noises coming from running engines, horns, bells, brakes, wheels squealing against the rail, and much more. This design uses the state-of-the-art audio beamforming technology available in the SV600 sensor and marries it with SwRI's expertise in leak detection, machine vision, and machine learning.

The machine vision and machine learning models will read the output of the Fluke sensor and determine if the detected audio signature matches what is seen in compressed air leaks from a frequency perspective. It will then determine if it is in a location where an air leak is possible to reduce or eliminate false positives from non-air leak related noises.

This novel approach to air leak detection should allow railroads to more easily address air leak repair, resulting in the following:

- Reduction of overall fuel consumption and harmful exhaust emissions
- Increase effectiveness of Auto-Engine Start Stop systems (AESS)
- Increase the lifespan of components such as air dryers, air compressors, and starters
- Improve employee safety by reducing the number of times employees must go on, under, or between rolling stock
- Improve operation efficiency by reducing air leak related train delays

### 3.0 INVESTIGATION

This section will discuss the process used to develop the proof-of-concept system.

#### 3.1 Hardware Selection and Acquisition

Through discussions between SwRI, Fluke, and Sorama it was determined that the Fluke SV600 Fixed Acoustic Imager (Figure 1) would be the ideal sensor option for this prototype development project. Table 1 shows the technical specifications of the SV600, which is a small form factor microphone array with an integrated camera that functions at the frequencies needed for compressed air leak detection (30-45kHz).

**TABLE 1. Fluke SV600 Fixed Acoustic Imager Specs**

Dimensions (LxWxD)	170 x 170 x 65 mm
Weight	0.85 kg
Frequency Range	0-55 kHz
Microphone Count	64
Camera Resolution	720p at 30fps
Operating Temperatures	-20°C to 50°C

Two of these sensors consigned and shipped to the Southwest Research Institute (SwRI) Locomotive Technology Center (LTC) and were received on 17 November 2021.



**FIGURE 1. Fluke SV600 Fixed Acoustic Imager**

The Fluke SV600 (Figure 1) is a commercially available fixed acoustic imager that is designed to enable users to detect, locate, and visualize sounds at specified frequency and decibel ranges. It is powered through power over ethernet (PoE) and allows for communication through a standard ethernet cable, providing a html based user interface. This allows significant customization options to narrow the detection range, eliminating non-air leak related sounds as much as possible which was needed for this project.

While the acoustic imager is the main sensor, two other imagers were eventually added to the system. A SwRI owned Basler 1920-40gc visual spectrum camera was added to obtain an alternate image without audio overlay provided by the Fluke SV600. The camera is equipped with a Kowa LM25HC-SW 25mm F/1.4 lens.

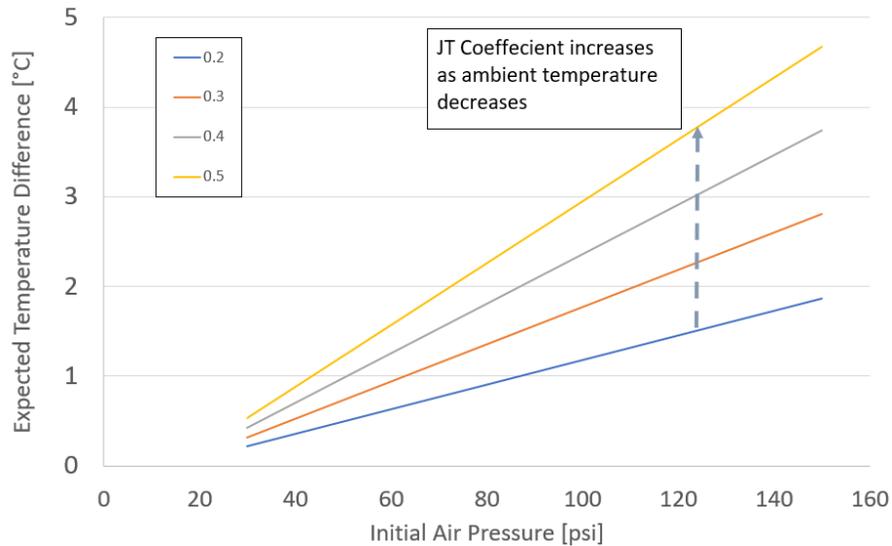
In addition to the Basler, a SwRI owned thermal imaging camera was also added to the system during the stage 1 testing. The thermal camera used was a Boson Compact longwave infrared (LWIR) Model 640 with a 24mm lens.

The specifications for the visual spectrum camera and the thermal camera used during the project are listed in Table 2.

**TABLE 2. Camera Specs**

	Visual Spectrum Camera	Thermal Camera
Manufacturer	Basler	Teledyne FLIR
Model	1920-40gc	Boson LWIR 640
Lens	Kowa LM25HC-SW, 25mm	FLIR 24mm, 18° HFOV
Camera Resolution	1920x1200	640x512
Frame Rate	42 fps	60/30 fps

The visual spectrum camera was used to allow for better part identification and leak location clarity, it also allowed for viewing behind the sound overlay map provided by the Fluke SV600 Sensor output. As mentioned above, the thermal imaging camera was being used to provide additional information in possible leak detection. Due to the Joule-Thomson effect, compressed air leaks create a cooling effect. This cooling effect varies based on initial pressure and ambient temperature, as shown in Figure 2. As the initial pressure of the air increases, the change in temperature becomes more severe [3]. Additionally, that temperature difference also becomes greater as the ambient temperature drops because it increases the Joule-Thompson coefficient.



**FIGURE 2. Joule-Thompson Effect Cooling at Various JT Coefficients**

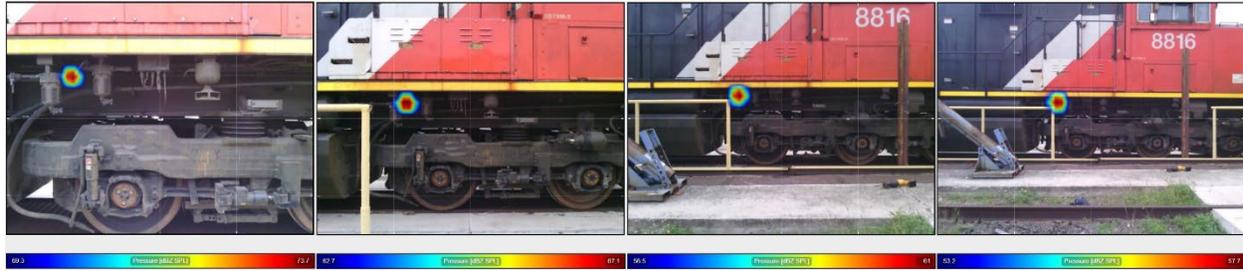
This indicated that there was some potential for thermal imaging to add to the accuracy of the system. The initial findings will be discussed in more detail in the following sections of this report.

SwRI staff constructed a versatile mounting stand that was able to orient the sensor and cameras in different orientations for testing. With such large, desired ranges of motion, it was determined that a fixed stand would not be appropriate for the initial design. The stand also needed to be heavy enough so it would not be affected by wind or vibrations from the train itself. Therefore, a portable stand was designed which could be moved and adjusted quickly and accurately to test various sensor positions in a rapid manner.

### 3.2 Initial Testing

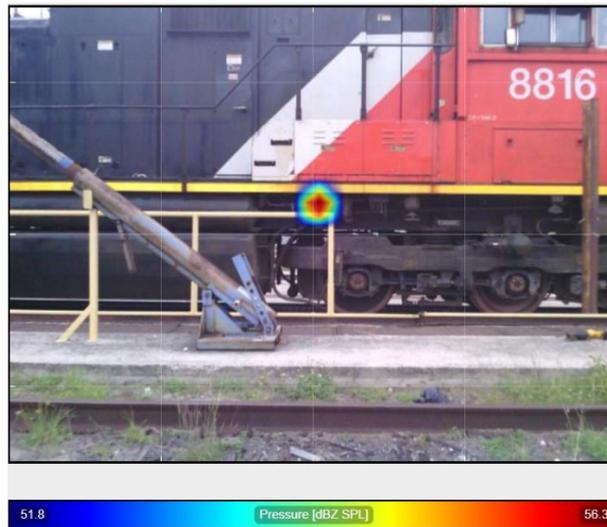
Once the Fluke SV600 sensors arrived at SwRI, one sensor was sent to the machine learning team to get familiar with sensor and its capabilities and the other was devoted to initial testing at the Locomotive Technology Center. During initial testing, air leaks were induced on a locomotive at multiple locations. As seen in Figure 3, the sensor was

capable of detecting a compressed air leak with the locomotive engine running in self load at Notch 3 at various distances. The sensor height was set at 1.2 meters above the rail and the sensor distance was varied.



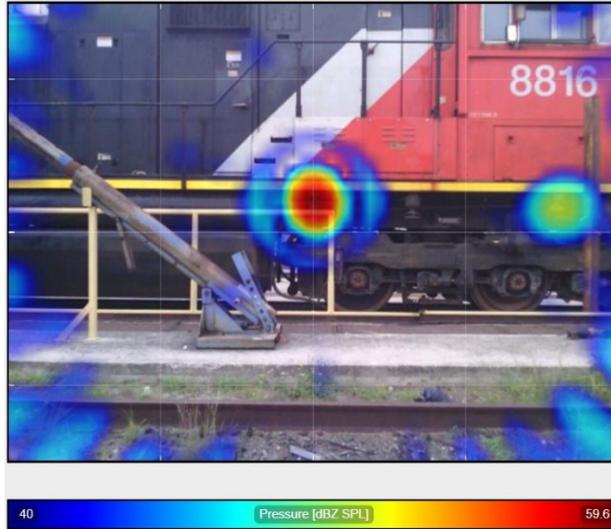
**FIGURE 3. Air Leak at 3.0, 4.5, 6.0, and 7.5 Meter Distances**

In Figure 3, one thing that isn't noted is that there is an additional air leak that is not being detected using the auto-range for the decibel settings on the sensor. The decibel minimum was tested at various levels to determine where the bottom cutoff was to detect smaller compressed air leaks. In Figure 4 the unit was auto-ranged with a minimum of 51 decibels. The second air leak was not detected at all with this minimum setting.



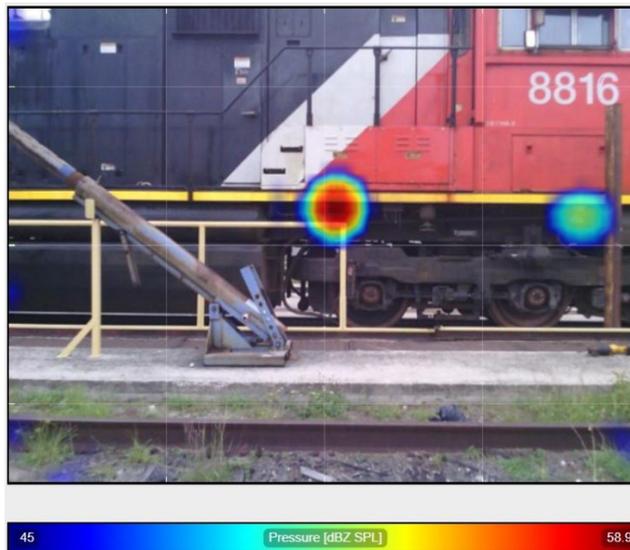
**FIGURE 4. Compressed Air Leaks at 51dB Minimum Detection Threshold**

In Figure 5, the minimum range was set to 40 decibels, which is slightly louder than a whisper. The second leak becomes obvious in this range but note that there were significant other detections displayed on the sound map with this minimum which can be seen at the edge of the field of view below.



**FIGURE 5. Compressed Air Leaks at 40dB Minimum Detection Threshold**

The minimum detection threshold was then set to 45 decibels (Figure 6). This minimized the additional detections and gave a good, clean detection signal for both air leaks in the image. While various other settings were tested throughout the course of the development, a 45dB minimum detection limit was consistently shown to be the optimal setting.



**FIGURE 6. Compressed Air Leaks at 45dB Minimum Detection Threshold**

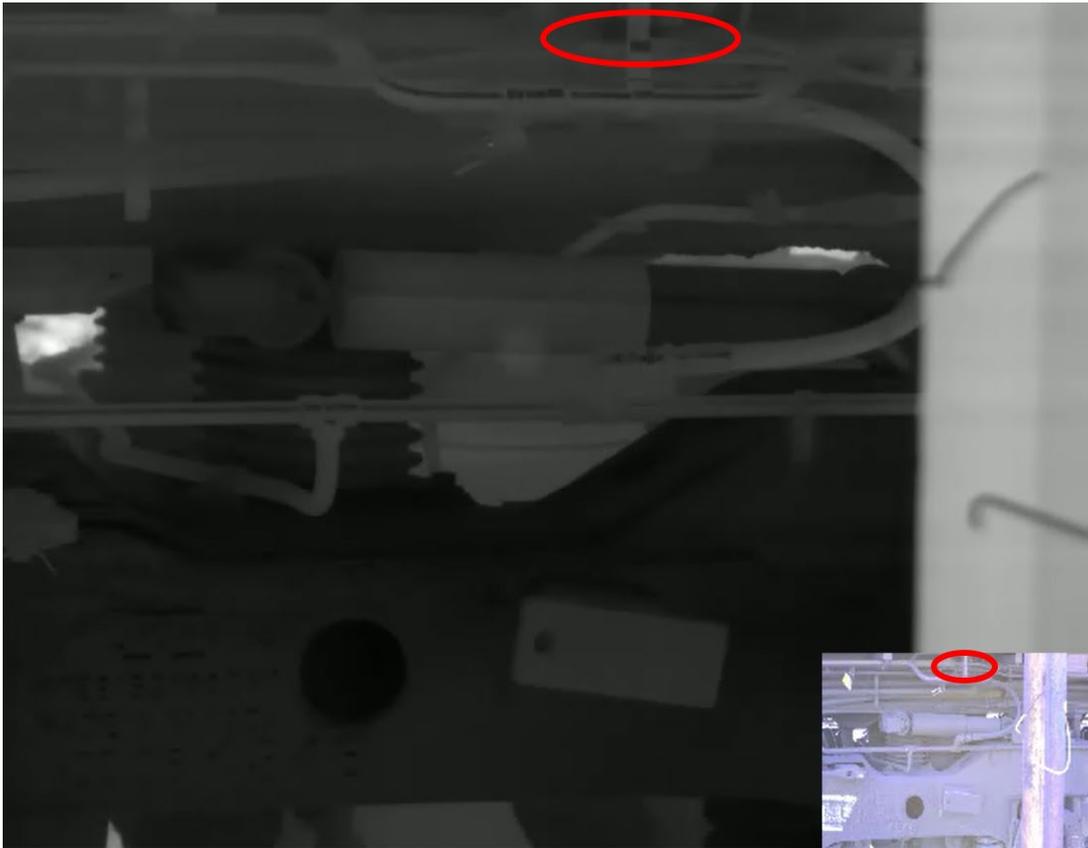
At this point in the initial testing a meeting was held between SwRI, Fluke, and Sorama on January 18, 2022. At the meeting some technical details of the SV600 sensors were discussed. Based on input from Sorama, the decision was made to add a separate visual spectrum camera to the system. Heath Spidle noted that SwRI has a visual spectrum camera that we could use at no cost, along with a thermal imaging camera and that the camera combination had been used on previous SwRI projects. This meant that the code existed already to process and align the thermal and visual spectrum cameras, allowing for the addition of the thermal camera on a trial basis.

Further testing was then completed using the SV600 sensor in concert with the visual spectrum and thermal cameras. The sensor distance was set to 3.0 meters, as this was determined to be the most likely distance based on earlier testing

and discussions with rail industry representatives including feedback from Project Advisor, Brad Kerchof, who recommended focusing on lower track speeds at this early stage.

Additional testing focused on generating data for Stage 2 of this project, including initial time alignment of the visual spectrum camera with the SV600 output. This data was generated with various passes with known air leak locations. During this data generation, the thermal imaging camera was also looked at to determine what, if any, cooling was occurring at these leak locations.

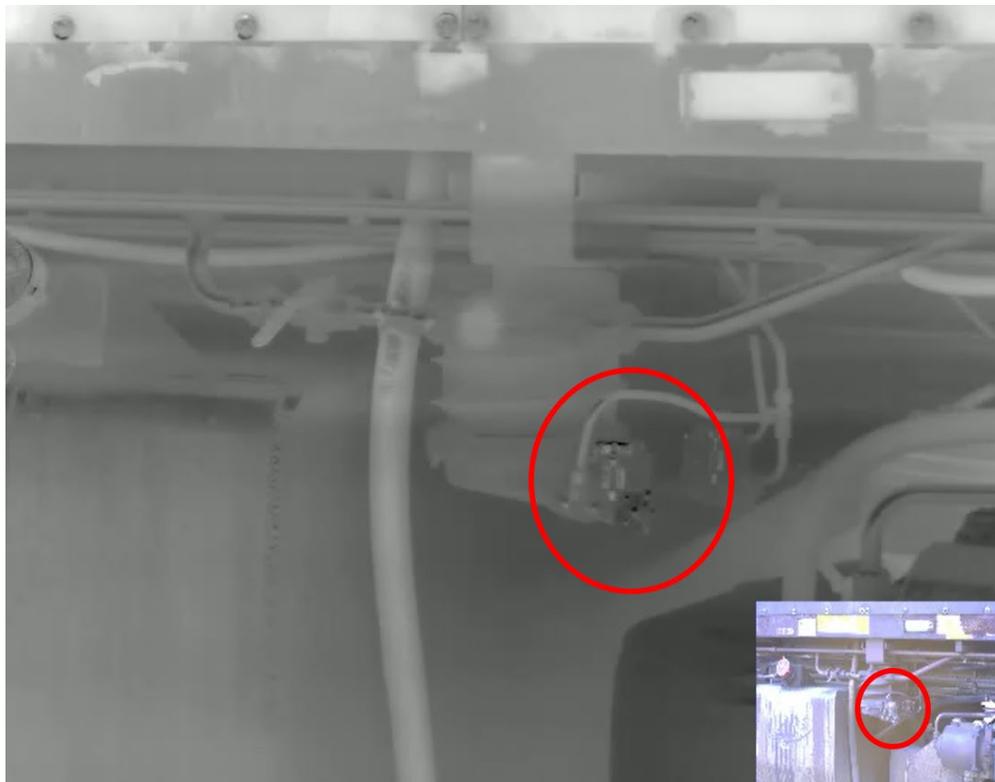
Images produced by both the thermal and visual spectrum cameras, showing air leaks in piping and a drain valve, are shown in Figures 7 and 8. The larger photo in each figure is the thermal image; the inset photo is the visual spectrum image. In each photo the leak location, indicated by a circle, shows to be slightly cooler than the surrounding area. Figure 7 shows a significant cooling of the clamp on which the leak was blowing; behind the clamp, note a dark area which is the leaking pipe fitting.



**FIGURE 7. Compressed Air Leak on Piping Using Thermal and Visual Spectrum Cameras**

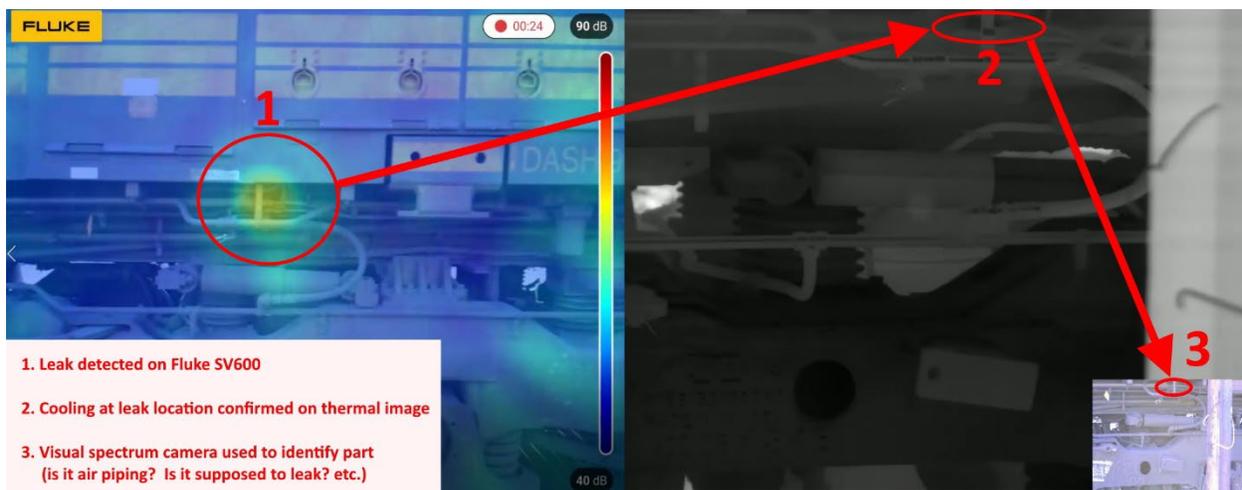
Figure 8 shows similar dark areas where the leak was, but in this case the leak itself was not blowing on to any surrounding part as the drain valve was simply blowing down to atmosphere. You can see some cooling occurring on the top of the valve body as well as at the handle. Obviously, these images themselves are not enough to detect the leaks, as there are temperature variations throughout the images. It is notable though that the cooling is detectable via the thermal imaging, which means it may assist in positive identification of compressed air leaks.

The results were positive enough that use of the thermal imaging camera was continued throughout the project in order to get the data necessary to determine if the benefit provided by the addition of the thermal imaging camera justifies the additional cost that it would add to the system.



**FIGURE 8. Compressed Air Leak at Drain Valve Using Thermal and Visual Spectrum Cameras**

Using the three imagers in combination gives a good understanding of how the sensors can work together. In Figure 9, you can see that the leak was detected on the Fluke SV600. Cooling was also detected on the thermal imaging camera at the same time and location while the visual camera helped identify the suspect component. At this point, the system was ready to proceed with the machine vision development and integration.



**FIGURE 9. Train Air Leak Detection Using 3 Cameras**

### 3.3 AI Integration and Final Testing

The most important component for any machine learning and deep learning model to be successful is having a large and robust dataset with representative data for the task at hand. In this case we needed data from the SV600 of leaking compressed air components along with data from moving locomotives. Gathering data from operational equipment is very important in building a system that can not only detect leaks but reject other sources of noise, e.g. track ringing, compressor noise, engine noise, bells, horns, etc. To accomplish this task, testing was completed at SwRI's Locomotive Technology Center (LTC) in central San Antonio, TX. Data was collected using both a single locomotive and locomotive consist, over multiple testing sessions. Air leaks were introduced in known locations to record and catalog ground truth data. The induced air leaks were sealed and opened as well as increased/decreased in intensity as the locomotive(s) were moved back and forth across the field of view of the sensors. In this instance more data was able to be collected on the intra-train connection components, such as knuckles and gladhands, which will serve as the basis for determining where a piece of rolling stock ends and the next one begins.

#### 3.3.1 Data Annotation

Prior to training and evaluating a machine learning model, images collected from the SV600 were manually annotated using CVAT, an open-source annotation tool. Regions including and around known leaks were labeled with closed polygons and the gladhands were labeled with bounding boxes (Figure 10) in order for the system to understand where pieces of equipment meet. To separate unique runs of data collection, the start and end times of each data collection run was recorded for future use in the training pipeline. A held-out validation set was created by randomly selecting 20% of these runs, leaving the rest of the annotations as example for training the model. Images collected from the Basler and Boson sensors were then registered to the SV600 image for each run to spatially align all images to the labels.

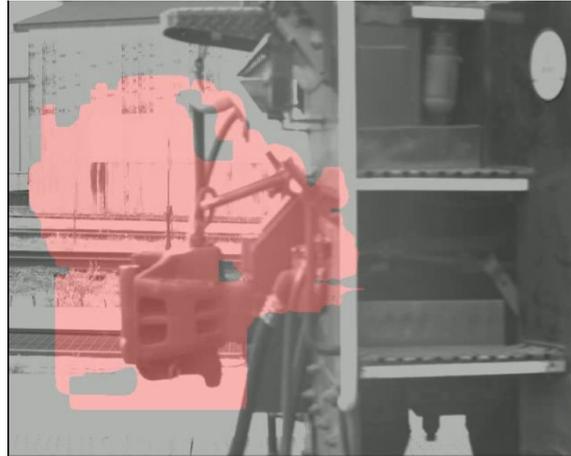


FIGURE 10. Annotated frame in CVAT

#### 3.3.2 Machine Learning Model

Convolutional neural networks (CNNs) have showed state-of-the-art performance in a variety of image-based tasks in the past decade due to their ability to learn hierarchical patterns in image data with significantly fewer parameters than fully connected networks. For this project, two types of CNN architectures were trained and evaluated for the detection of air leaks. The first was a mobilenetV3 based segmentation model, which classifies each pixel in an image as either belonging to an air leak or a part of the background as shown in Figure 11. Segmentation models have the benefit of producing sharp predictions with the precise location of leaks at the cost of requiring advanced post-processing algorithms to determine which groups of pixels belong to an individual instances of a leak. The second model implemented was YOLOv5, an object detection architecture that holistically locates the centers of the leaks and bounds them using a standard geometry (boxes or circles). This method allows for the identification of individual instances to be a part of the deep learning process which reduces the granularity of the predictions.

Development of the leak detection algorithm began with training the segmentation model on the images from the first data collection. While this model was able to detect larger leaks with confidence, it frequently reported smaller leaks as background. These false positives were determined to be inherent to the segmentation approach as larger leaks were given significantly more importance than smaller ones. This issue was further compounded by the ambiguity of what constitutes a leak on a pixel level. The visible and longwave images were also introduced as inputs to the model during this iteration of model development. While telling features of a leak were present in this data, poor overlap in the field of view of the cameras did not provide sufficient information for model to increase performance.



**FIGURE 11. Pixelwise Semantic Segmentation Mask of Air Leak**

Following the second round of data collection, the YOLOv5 model was implemented and trained. A python script was written to convert the air leak labels into the bounding box format required by the model (Figure 12). To reduce the errors from ambiguous labels seen previously, the identification of leaks was weighed relatively higher than bounding task in the training process. Results from this model showed a large improvement in the detection of smaller leaks and a higher confidence when detecting larger leaks. More false positives were found than with the previous model, though they lacked consistency between frames and could be filtered out with post-processing methods. The YOLOv5 architecture also allowed new features to be detected, such as gladhands, with only a marginal increase to the model size. Additional annotations were created for the past data collections to include the presence of gladhands and the model was retrained using both classes of labels.



**FIGURE 12. YOLOv5 Detection of an Air Leak**

An important feature of leaks not captured by the machine learning models is their persistence through time. While potential false positives could be filtered out by the models based on their shape and relative location to the locomotive, post-processing methods can be used to increase the confidence in detected leaks by looking at frames sequentially in time and tracking them while within the frames. An object tracking algorithm was used to identify and track detections across multiple frames. This technique allows the system to build an improved confidence score and measure the amount of time a leak is seen in the frame (Figure 13). Leaks with a probability over 50% were tracked by this algorithm and had to be present for 0.5 seconds to be a reported leak. While these thresholds were used for the reported metrics, they could be easily adjusted to fit different use cases and sensitivities.

An alerting system was developed to notify users of a detected leak and provide visualizations of their locations. A composite image is generated for each piece of rolling equipment, with the post-processed leaks being overlaid on the image. This image is then reported to the appropriate mechanical personnel for further investigation and repair as time permits. Figure 14 shows an example of the automated notification system. Note that this notification system is just an example and can be adapted as necessary to meet the needs of each railroad.



**FIGURE 13. Example Composite Image of Passing Locomotive with Detected Air Leaks**

### AirLeak Detected NOTIFICATION

L

Loco-Air-Leak@noreply.org

To Spidle, Heath A.; Janssen, Jake A.; Stoos, Christopher R.

Air-Leak\_22\_06\_03\_113257.jpg

921 KB

2022/06/03 113257 - 5 Leaks Detected

Leak-0: Seen for 0.55s - 0.54 Confidence

Leak-1: Seen for 0.55s - 0.70 Confidence

Leak-2: Seen for 0.55s - 0.37 Confidence

Leak-3: Seen for 0.55s - 0.49 Confidence

Leak-4: Seen for 0.55s - 0.68 Confidence

**FIGURE 14. Example Air Leak Notification Alert**

### 3.3.3 Results

Figure 15 shows validation metrics for the object detection deep learning network. The Y-axis is the normalized score between 0 and 1; the X-axis is the epoch. As the model is trained and subsequently evaluated after each epoch, the

precision, recall, and mean average precision improve and change. As the model learns, it starts to reach a stabilized state where the validation metrics stop changing which means the model has fit itself to the data as much as it can with the currently available input data, hyperparameters and architectures used. These results are summarized in Table 3. The metrics that are used in reporting are defined below:

True Positive ( $t_p$ ): Algorithm identified a leak when there was a leak present (1)

True Negative ( $t_n$ ): Algorithm did not trigger when there was no leak present (2)

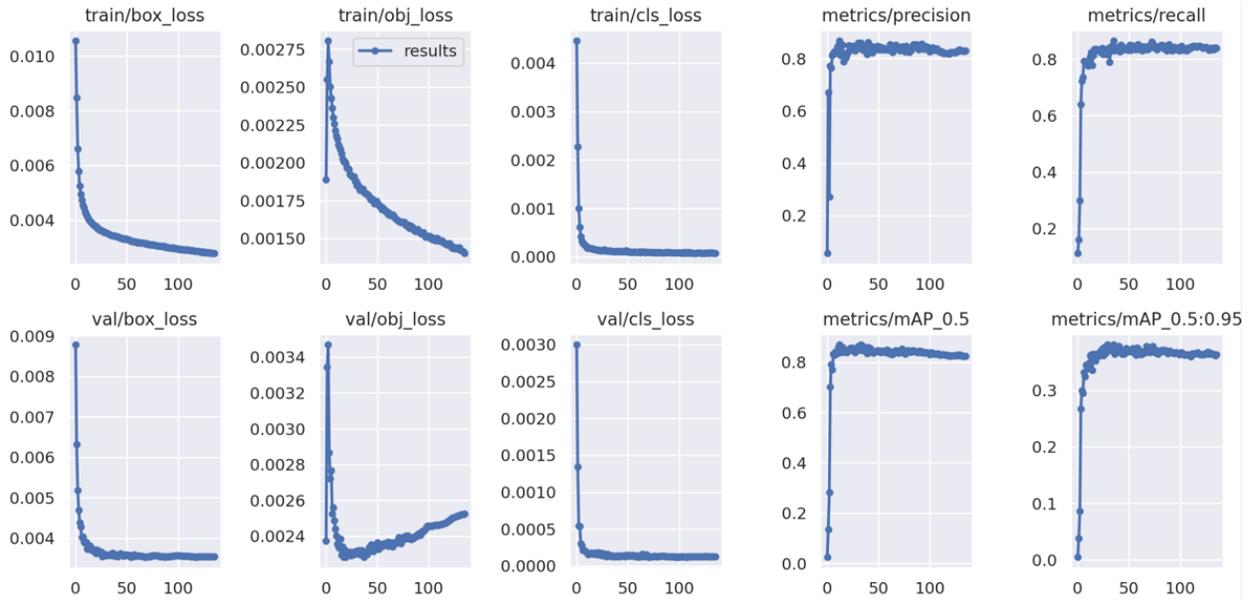
False Positive ( $f_p$ ): Algorithm Identified a leak when there was no leak present (3)

False Negative ( $f_n$ ): Algorithm did not trigger when there was a leak present (4)

Precision =  $\frac{t_p}{t_p + f_p}$ : measures the positive predictive value, i.e. the percentage of your predictions are correct. Maximizing precision will minimize the false-positive errors (5)

Recall =  $\frac{t_p}{t_p + f_n}$ : measures how well you find all the positives. Maximizing recall will minimize the false-negative errors (6)

Mean Average Precision (mAP) =  $\int_0^1 p(r)dr$ : Area under the precision recall curve. (7)



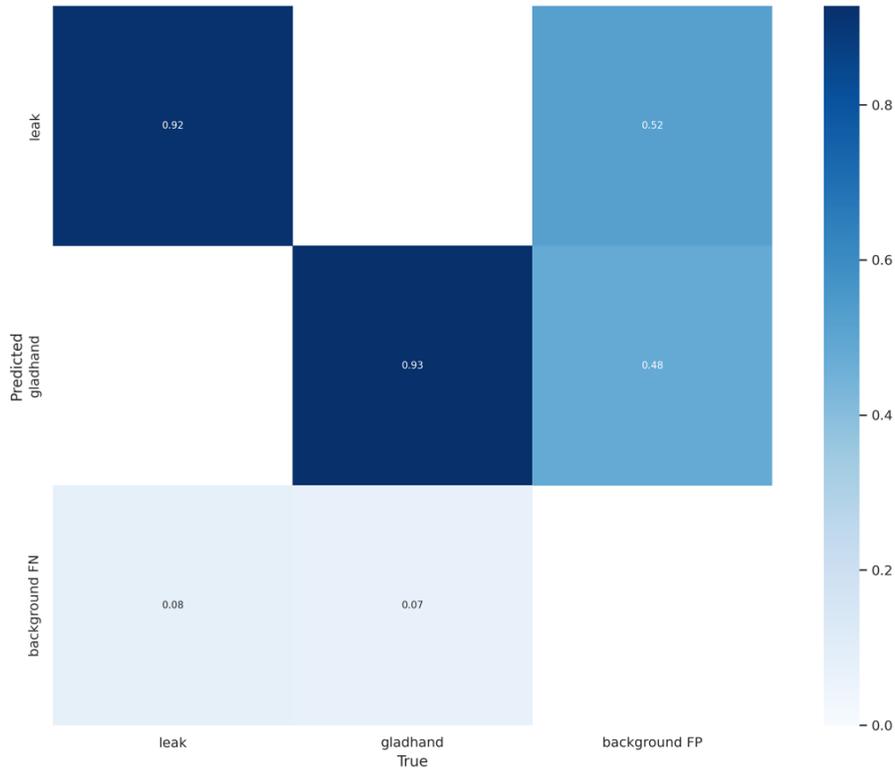
**FIGURE 15. SSD Training and Validation Metrics**

The selected model has a precision of 0.57 and recall of 0.63 with a mAP of 0.54. mAP\_0.5 refers to the model's ability to detect and classify objects with at least a 50% overlap in the object detection region vs. the ground truth label. mAP\_0.5:0.95 is the average of the mAP scores between 50% and 95% overlap with steps of 5%. Table 3 shows results from the validation process.

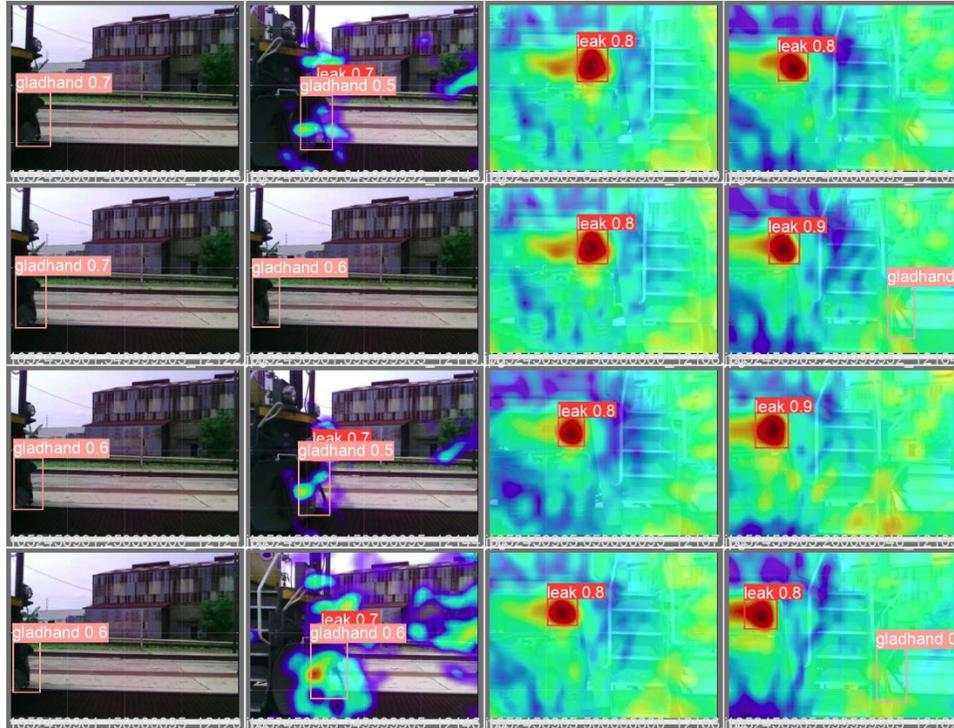
**TABLE 3. Model Validation Results**

Class	Images	Labels	P	R	mAP@.5	mAP@.5:.95
all	14,037	2,901	0.652	0.797	0.745	0.337
leak	14,037	852	0.468	0.813	0.687	0.354
gladhand	14,037	2,049	0.835	0.780	0.803	0.319

Figure 16 shows the confusion matrix over the selected classifier. This matrix is a heatmap of how the classifier is predicting based on the labels. The Y-axis is what the classifier determined was in the image. The X-axis is the ground truth. A perfect classifier would have a diagonal line from top left to bottom right. In the case of this classifier, we can see that it does very well at identifying gladhands with little false negatives. With Air leaks it predicts correctly 65% of the time, there is some confusion when the air leaks are on top of gladhands and there are cases where some air leaks were missed entirely. This matrix also shows that all the false positives that the model reported were for air leaks. The performance and metrics of the model could be improved over time with more training, testing and validation data.



**FIGURE 16. SSD Confusion Matrix**



**FIGURE 17. SSD Validation Examples**

The metrics discussed above are pointwise metrics, meaning they are calculated and aggregated over every image in the validation dataset and do not consider the temporal and spatial relationships of an air leak as it progresses through the field of view of the camera (Figure 17). Using the object tracking techniques described above, we were able to reduce the false positives caused by track noise, bells, and main reservoir blowdown events. Additionally, the number of false negatives was reduced by allowing for infrequent missed detections while tracking a leak. Figure 18 depicts the tracking of an air leak as the locomotive passes the ID value of the leak is tracked.



**FIGURE 18. Spatial and Temporal Tracking of an Air Leak**

The metrics below were calculated using the same validation set previously used for the pointwise metrics. The individual frames were first processed by the object detection model to create leak predictions for each image. These predictions and images were then used as inputs to the tracking algorithm to determine the presence of persistent leaks, the occurrence of a new locomotive, and the direction of the locomotive. Prior to running the tracking algorithm, 215 false positive leaks were predicted over the 1732 images, leading to a false positive per image rate of 12.4%. After

applying the tracking algorithm, the false positive rate was reduced to 0.03%. The false positives were found to be caused by persistent noise along the tracks which could potentially be filtered out by additional post-processing or higher confidence thresholds used within the tracking algorithm. The F1 score was additionally used to give a holistic evaluation of the full system's accuracy.

$$F_1 \text{ Score} = \frac{t_p}{t_p + \frac{1}{2}(f_p + f_n)} : \text{measures accuracy of predictions, i.e. the harmonic mean of precision and recall} \quad (8)$$

**TABLE 4. Air Leak Tracking Metrics**

Precision	0.786
Recall	0.846
F1 Score	0.815

As shown in Table 4, the accuracy of the system in identifying leaks in the reserved validation runs was 84.6%. After the addition of the various models and tracking system discussed above, the proof-of-concept detection system was able to detect air leaks with very limited false positives and autonomously notify staff of the leak locations. With further development of the system in the field, autonomous wayside detection of rail air leaks can be refined and should provide railroads with significant opportunity at reducing fuel consumption, emissions, and improving employee and operational safety. Those future development of the system is discussed in Section 4 below.

#### 4.0 PLANS FOR IMPLEMENTATION

The work completed in this project has demonstrated that autonomous detection of air leaks on moving trains is possible. Commercial implementation of this system will require further field development, as the air systems on trains are incredibly complicated with near infinite variations in train makeup. A high-level plan for final development and implementation of this design is outlined below.

1. Design and fabricate a field ready version of the system
  - a. Railroads are not laboratories – Needs to be *everything* proof
  - b. Final system will likely include two SV600 sensors per side of track at various heights and angles
  - c. One high-definition visual spectrum camera per sensor (if necessary)
  - d. AEI tag reader implemented for equipment identification
  - e. Work with Fluke and Sorama to develop software to ease implementation of sensors into the system
2. Deploy prototype system on active rail line
  - a. Preference would be a captive fleet
  - b. Work with maintenance personnel on the ground to verify findings and use that information to train system
3. Data Collection
  - a. Machine Learning systems require significant input data
  - b. Length of time required is dependent on traffic. More trains = less time.
4. Prototype refinement
  - a. The field trial will allow for improvements on both the system hardware and software designs
  - b. System refinement will occur concurrently with the data collection
5. Revenue service implementation
  - a. Once the technology is sufficiently matured through the field testing and system refinement, it will be available to railroads to implement

## 5.0 CONCLUSIONS

The proof-of-concept development of a wayside autonomous air leak detection system has showed great promise. The system was successfully able to detect air leaks at various locations with varying air leak rates with limited false positives and should greatly reduce the burden on mechanical personnel in finding air leaks on equipment. With further development the system could be improved by integrating AEI tag readers and implementing a machine learning system that can identify not just air leaks, but the individual components that are leaking further reducing the burden on mechanical personnel. These refinements will require significant data collection. This comprehensive data collection is proposed in the next phase of this development effort.

As railroads strive to reduce their greenhouse gas emissions to meet their SBTi targets, improving overall vehicle efficiency through NTSFC reduction and AESS improvement is crucial. This system, if implemented correctly, could make significant improvements to both. While the system will significantly reduce the time spent locating air leaks, it cannot fix the leaks autonomously. As with many projects related to locomotive fuel consumption, it will initially require significant time and labor from mechanical department personnel to properly address air leaks. How that is accomplished will be the decision of the individual railroad, but support of mechanical department personnel will be required.

It is expected that, as the existing air leaks are addressed and fixed over time, the time necessary to maintain the fleet will gradually reduce. In time, properly addressing air leaks may even reduce overall mechanical labor burden as components such as air compressors, brake valves, air dryers, and starters should require less maintenance and need to be changed out less often.

As covered in detail above, even with limited data from a machine learning perspective, the system was able to positively identify air leaks with 84.6% accuracy and has a false positive rate of 0.03%. With these results we believe the proof-of-concept to be a success, and believe that the system has the ability to greatly help railroads in their quest to improve employee safety, reduce emissions, and reduce fuel consumption.

## 6.0 LESSONS LEARNED

This proof-of-concept system was trained with a relatively small amount of data and the data contained only locomotives on a ~1/4 mile stretch of track. To improve the equipment range, performance and reliability of the model, more data is needed of longer trains with varied rail cars and equipment at various operation speeds. This cannot be achieved in a lab or shop setting and will require field testing on an active rail line.

Additionally, the inputs to the model were constrained by available access to the SV600 camera feeds. A notable improvement could be achieved with access to the SV600's raw camera stream and Audio mask. The current iteration relied on recording the screen from the sensors web dashboard and then cropping it and feeding it into the model, which prohibits any real-time capabilities. This means that all data had to first be saved, then post-processed. This likely wouldn't affect overall system capabilities as post processing can occur immediately after a train passed. In future development SwRI could work with Fluke and Sorama to implement those abilities on the SV600. Additionally, separating the audio mask and video streams could eliminate the need for the visual spectrum camera as a separate sensor.

The SV600 sensors were accurate enough that it was determined that the thermal imaging camera was not needed for leak verification, even though it did tend to show some cooling effect at leak locations. It may still be necessary in a final product for identification in darkness, as the visual spectrum camera and the fluke image will not show clear images at night. Which camera is used has little impact on the machine vision system design.

Initial testing was limited in track speed of <10mph. Additional testing at a suitable location will be required to determine the upper limits of track speed at which accurate detection could occur. Detection is limited by sensor frame rate, distance from the track, and train speed. As speed increases or distance from the sensor to the train decreases, the number of frames that a suspected leak will be in the detection zone of the sensor decreases. If it is possible to increase the frame rate of the SV600 (30fps) detection could occur at significantly higher track speeds.

There are additional complications with detecting leaks without direct line of sight and how to address those leaks. As you will note in Figure 19 below, the Fluke sensors are capable of detecting air leaks behind closed carbody doors. Leaks in areas like compressor compartments, air brake compartments, etc. are all very common on locomotives. In these situations, neither the SV600 nor the visual spectrum camera will be able to identify the exact location of these leaks, which means component identification is impossible. It is therefore difficult to determine if the leak is a purposeful leak such as the J1 Relay releasing brake cylinder air or an actual leak at a mag valve or something similar. These leaks may be detected on a single pass, but may require multiple passes of the same piece of equipment to determine if the leak is an undesirable leak or a necessary air release.

Similar situations will arise on rail cars with brake pipe piping that is not exposed. On rail cars we can assume that any leak is unintended (outside of brake release situations), so there is less concern with false positive in these situations. Work will need to go into determining how to approach this issue on locomotives though, as there are more complications in the compressed air system. While both leaks shown in Figure 19 were indeed unintended leaks, without further inspection this may be difficult to determine on a single pass by the sensors.



**FIGURE 19. Air Leaks in Air Brake Compartment Below Cab (left) and Rear Sander Mag Valve (right)**

## 7.0 LEAD INVESTIGATOR PROFILES

**Chris Stoos** – Lead Engineer, Southwest Research Institute  
M.E. Engine Systems, University of Wisconsin – Madison (2017)  
B.S. Mechanical Engineering, University of Texas at San Antonio (2011)

Mr. Stoos has a combined 21 years of rail experience. He has been with SwRI for 12 years specializing in locomotive engine performance, including fuel consumption, emissions, and dual fuel operation. In his time at SwRI, he has performed emissions and fuel consumption testing on well over 150 locomotives and has led various other projects including hydrogen and natural gas fueled engines, emissions development, and EPA certification related testing.

Before his time at SwRI, Mr. Stoos spent 3 years as a Diesel Machinist for Union Pacific Railroad and 8 years as a Locomotive Crewmember (88U) in the U.S. Army, working as a conductor and licensed locomotive Engineer.

**Heath Spidle** - Senior Research Engineer, Southwest Research Institute  
B.S. Electrical Engineering – University of Texas at San Antonio

Mr. Spidle has been with SwRI for 5 years and is currently leading the Intelligent Energy Inspections Program Area, focused on applying machine learning to optical and signals-based inspection technologies.

**Jake Janssen** - Research Engineer, Southwest Research Institute  
B.S. Computational Engineering – University of Texas at Austin

Mr. Janssen has been with SwRI for 3 years and is currently involved in research targeting efficient deep-learning based computer vision models.

## 8.0 REFERENCES

1. 49 C.F.R. 232 <https://www.ecfr.gov/current/title-49/subtitle-B/chapter-II/part-232>
2. Hedrick, J. and Fritz, S. “*AAR RP-589 Locomotive Compressor Load*” Locomotive Maintenance Officer Association Annual Meeting, 2020
3. Demirel, Y. “Nonequilibrium Thermodynamics, Second Edition” 2007 ISBN 978-0-444-53079-0

**APPENDIX A:**

**Research Results**

Program Steering Committee: NCHRP IDEA Program Committee  
Month and Year: September 2022  
Title: Autonomous Detection of Compressed Air Leaks on Trains  
Project Number: Rail Safety IDEA-48  
Start Date: October 1, 2021  
Completion Date: September 14, 2022  
Product Category: Wayside Detection  
Principal Investigator: Christopher Stoos, Lead Engineer - Southwest Research Institute  
E-Mail: cstoos@swri.org  
Phone: 210-522-2647

**TITLE:**

Autonomous Detection of Compressed Air Leaks on Trains

**SUBHEAD:**

Developed a proof-of-concept wayside detection system to locate air leaks on moving trains and notify mechanical personnel of the location of detected leaks.

**WHAT WAS THE NEED?**

Compressed air leaks are difficult to locate on a train as they contain miles of compressed air piping. Employees need to go on, under, or between rolling stock to listen and feel for leaks which is difficult and time consuming. This reality has caused regulations to be in place that allow certain levels of “allowable” leaks. Railroads needed a good method to identify and locate air leaks making repair easier and less time consuming.

**WHAT WAS OUR GOAL?**

The goal of this project was to develop and prove a proof-of-concept system to detect compressed air leaks autonomously to reduce the inspection burden on employees with the goal of increasing employee safety, reducing fuel consumption, and reducing overall emissions.

**WHAT DID WE DO?**

Using commercially available hardware from Fluke Process Instruments, SwRI developed and implemented a machine learning model to identify air leak signatures from the sensor output while greatly reducing false positives. Additionally, an automatic notification system was implemented

**WHAT WAS THE OUTCOME?**

The outcome of this project was a successful proof-of-concept system. The final system achieved a 84.6% success rate at identifying air leaks with only a 0.03% false positive rate. With further development we believe those numbers could be greatly improved.

**WHAT IS THE BENEFIT?**

This system benefits railroad employees through a reduction in the number of times employees must break the plane of the train by going on, under, or between rolling stock. Knowing where the leaks are could also reduce the distance employees must walk in order to get to the leak location. Employee safety is also benefited in that non-essential repairs can be noted and made the next time the equipment is in the shop for inspection, allowing for repair in a more controlled setting.

Railroad service will be benefitted through a reduction in train delays and unplanned stops due to air leaks.

Finally, the public overall will benefit as air leaks cause a significant increase in overall railroad fuel consumption and exhaust emissions. A reduction in air leaks will directly reduce harmful locomotive-based emissions including PM2.5, NO<sub>x</sub>, and GHG emissions.

### LEARN MORE

For more information and to see the full report, please visit the Transportation Research Board website at <https://www.trb.org/IDEAProgram/IDEASafety.aspx>.

For questions regarding this project please email the Principal Investigator:

Christopher Stoos  
[cstoos@swri.org](mailto:cstoos@swri.org)

### IMAGES



**FIGURE A1. Detection of an Air Leak**



**FIGURE A2. Example Composite Image of Passing Locomotive with Detected Air Leaks**