U.S. Department of Transportation

Federal Railroad Administration

Extended Field Trials of LRAIL for Automated Track Change Detection

Office of Research, Development and Technology
Washington, DC 20590

DOT/FRA/ORD-20/14

Final Report
April 2020

REPORT DOCUMENTATION PAGE			*Form Approved OMB No. 0704-0188*
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE April 2020		3. REPORT TYPE AND DATES COVERED Technical Report – September 13, 2018 – December 31, 2019
4. Extended Field Trials of LRAIL for Automated Track Change Detection			5. FUNDING NUMBERS 693JJ618C000023
6. AUTHOR(S) Richard Fox-Ivey, Thanh Nguyen, John Laurent			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Pavemetrics Systems, Inc. 150 boul. Rene-Levesque, Suite 1820 Quebec, Canada G1R 5B1			8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Department of Transportation Federal Railroad Administration Office of Railroad Policy and Development Office of Research, Development and Technology Washington, DC 20590			10. SPONSORING/MONITORING AGENCY REPORT NUMBER DOT/FRA/ORD-20/14
11. SUPPLEMENTARY NOTES COR: Cameron Stuart			
12a. DISTRIBUTION/AVAILABILITY STATEMENT This document is available to the public through the FRA eLibrary.			12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 words) LRAIL is a 3-dimensional laser inspection system that combines high-speed laser triangulation sensors and artificial intelligence to automatically inspect the condition of, and detect safety-related changes in, railway infrastructure. For this project, LRAIL technology was installed on a specialized hi-rail inspection trailer and data were captured on Amtrak lines over a 4-month period in three separate deployments across an 8-mile test loop near the Delaware-Pennsylvania border. Repeat runs in the same direction and reverse runs were captured during each deployment to model repeatability, mean, and standard deviation of measurements. The noise floor for change detection was also modeled through analysis of repeat runs and the detection of known/deliberate changes. LRAIL change measurements were determined to be highly repeatable, with an overall average of 99.28 percent agreement between runs.			
14. SUBJECT TERMS Laser triangulation, algorithms, change detection, safety, inspection			15. NUMBER OF PAGES 71
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18
298-102

METRIC/ENGLISH CONVERSION FACTORS

ENGLISH TO METRIC

LENGTH (APPROXIMATE)

1 inch (in)	=	2.5 centimeters (cm)
1 foot (ft)	=	30 centimeters (cm)
1 yard (yd)	=	0.9 meter (m)
1 mile (mi)	=	1.6 kilometers (km)

AREA (APPROXIMATE)

1 square inch (sq in, in^2)	=	6.5 square centimeters (cm^2)
1 square foot (sq ft, ft^2)	=	0.09 square meter (m^2)
1 square yard (sq yd, yd^2)	=	0.8 square meter (m^2)
1 square mile (sq mi, mi^2)	=	2.6 square kilometers (km^2)
1 acre = 0.4 hectare (he)	=	4,000 square meters (m^2)

MASS - WEIGHT (APPROXIMATE)

1 ounce (oz)	=	28 grams (gm)
1 pound (lb)	=	0.45 kilogram (kg)
1 short ton = 2,000 pounds (lb)	=	0.9 tonne (t)

VOLUME (APPROXIMATE)

1 teaspoon (tsp)	=	5 milliliters (ml)
1 tablespoon (tbsp)	=	15 milliliters (ml)
1 fluid ounce (fl oz)	=	30 milliliters (ml)
1 cup (c)	=	0.24 liter (l)
1 pint (pt)	=	0.47 liter (l)
1 quart (qt)	=	0.96 liter (l)
1 gallon (gal)	=	3.8 liters (l)
1 cubic foot (cu ft, ft^3)	=	0.03 cubic meter (m^3)
1 cubic yard (cu yd, yd^3)	=	0.76 cubic meter (m^3)

TEMPERATURE (EXACT)

[(x-32)(5/9)] °F = y °C

METRIC TO ENGLISH

LENGTH (APPROXIMATE)

1 millimeter (mm)	=	0.04 inch (in)
1 centimeter (cm)	=	0.4 inch (in)
1 meter (m)	=	3.3 feet (ft)
1 meter (m)	=	1.1 yards (yd)
1 kilometer (km)	=	0.6 mile (mi)

AREA (APPROXIMATE)

1 square centimeter (cm^2)	=	0.16 square inch (sq in, in^2)
1 square meter (m^2)	=	1.2 square yards (sq yd, yd^2)
1 square kilometer (km^2)	=	0.4 square mile (sq mi, mi^2)
10,000 square meters (m^2)	=	1 hectare (ha) = 2.5 acres

MASS - WEIGHT (APPROXIMATE)

1 gram (gm)	=	0.036 ounce (oz)
1 kilogram (kg)	=	2.2 pounds (lb)
1 tonne (t)	=	1,000 kilograms (kg)
	=	1.1 short tons

VOLUME (APPROXIMATE)

1 milliliter (ml)	=	0.03 fluid ounce (fl oz)
1 liter (l)	=	2.1 pints (pt)
1 liter (l)	=	1.06 quarts (qt)
1 liter (l)	=	0.26 gallon (gal)
1 cubic meter (m^3)	=	36 cubic feet (cu ft, ft^3)
1 cubic meter (m^3)	=	1.3 cubic yards (cu yd, yd^3)

TEMPERATURE (EXACT)

[(9/5) y + 32] °C = x °F

QUICK INCH - CENTIMETER LENGTH CONVERSION

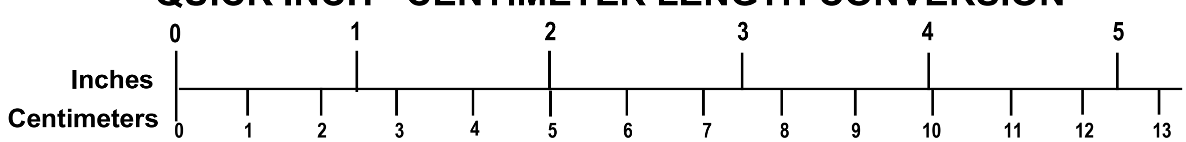

QUICK FAHRENHEIT - CELSIUS TEMPERATURE CONVERSIO

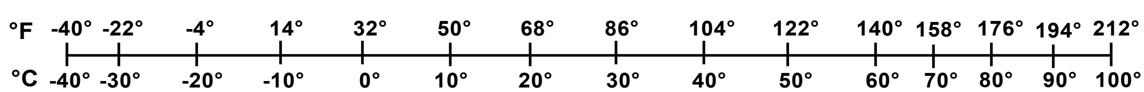

For more exact and or other conversion factors, see NIST Miscellaneous Publication 286, Units of Weights and Measures. Price $2.50 SD Catalog No. C13 10286

Updated 6/17/98

Acknowledgements for Technical Contributions

Mike Trosino, Amtrak

Carl Walker, Amtrak

Emily Greve, Amtrak

John Kotyla, Amtrak

Steven Chrismer, Amtrak

Contents

Illustrations

Tables

Executive Summary

This report details the deployment of Pavemetrics' Laser Rail Inspection System, "LRAIL," for the purposes of automated change detection. The project was conducted between September 2018 and December 2019 at filed locations on Amtrak property and at Pavemetrics' offices in Quebec, Canada.

The project involved a combination of field sensor data acquisition, deliberate manual changes in the field, office algorithm development, algorithm testing and validation, and system performance reporting.

Data were captured on Amtrak lines over a 4-month period. An 8-mile test loop, from MP 25 on the south leg of Wye-Landlith to MP 17 at the Hook Interlocking (near the Delaware-Pennsylvania border), was used for the test loop. This stretch was selected because it is active, receives regular traffic, and both natural deterioration and restorative maintenance were expected over the test period.

Data were collected over three separate field deployments in May, June, and September 2019. Both repeat runs in the same direction and reverse runs were captured during each deployment to model repeatability, mean, and standard deviation of measurements.

The noise floor for change detection was also modeled through analysis of repeat runs and the detection of known/deliberate changes.

The extended field trial proved successful. Repeatability, mean, and standard deviation of change measurements were determined and noise floors for each measured parameter were established.

LRAIL change measurements were determined to be highly repeatable, ranging between 95.72 percent (for skew angle) to 100 percent (for fasteners, joint bar bolting, tie rating, ballast level in gauge, and ballast level on left field side), with an overall average of 99.28 percent.

Most importantly, the ability to detect organic changes above the noise floor was validated with numerous examples, including ballast level, ballast fouling, tie condition, joint bar bolting, and joint gap being detected.

1. Introduction

This project continued research into automated track change technology for the purpose of enhancing the safety of railroad track. This work is directly applicable to the Federal Railroad Administration (FRA) Track Research Division's strategic priority of developing track inspection technologies that detect defects before they become failures in service.

1.1 Background

A Phase 1 project related to this research was completed under FRA contract DTFR5317C00005 and involved an initial proof of concept (POC). During the POC, Pavemetrics successfully demonstrated its ability to detect changes in fasteners, anchors, spikes, ties, joints, and ballast as well as record rail stamping information. The technical report from Phase 1, "Laser Triangulation for Track Change and Defect Detection."

As Phase 1 was a POC, the quantity of rail tested and the duration of testing performed was fairly limited. Additionally, it focused on the detection of artificial changes only as opposed to organic ones (i.e., those due to normal track deterioration or to restorative maintenance).This approach was suitable for proving the potential of the technology but was too limited to explore how it might perform in the field and be of use to a typical track inspector. The focus of this second project was a more realistic assessment of the potential value of deploying change detection technology under typical field inspection conditions.

1.2 Objectives

The objective of this project was to evaluate the potential for the Laser Rail Inspection System (LRAIL) to intelligently detect relevant changes in infrastructure, and/or unsafe conditions, and to notify stakeholders for immediate action.

The project demonstrated the potential for change detection technology to address a number of use cases in the industry:

- "Go/No-Go" safety assessment of the track based on the detection of significant condition changes between runs
- Detection of gradually changing track conditions over time to enhance mid- and long-term maintenance planning
- Detection of changing maintenance conditions to help validate that contracted maintenance work was performed as planned
- The augmentation of other railway datasets (GRMS, track geometry, rail flaw, etc.) with change detection data to support better decision making

Major milestones for the project included:

- Project kick-off meeting
- Sensor installation and onsite data collection
- Algorithm and software development
- Data analysis

- Project reporting

Expected project outcomes were:

- Streamlined operation of the technology to better support future field deployment
- Statistics regarding the number and magnitude of organic changes in track condition over time
- Modeling of the “noise floor” of the technology (a practical limit for the magnitude of change that can be detected)
- Validation of the technology’s ability to detect organic changes above the noise floor

1.3 Overall Approach

Researchers scanned rail conditions in the field over a period of 4 months, enhanced change detection algorithms, streamlined processing software, and finally reported detected changes and overall system performance.

As the project lead, Pavemetrics was responsible for supplying sensor hardware and developing algorithms and processing software for the project.

As the railroad testing partner, Amtrak was responsible for selecting a test site and providing guidance regarding railroad inspection practices.

1.4 Scope

This report documents field trials and data analysis efforts directed at the development of an automated track change detection system. The scope of this report includes a description of the measurement system and data processing methods as well as a detailed analysis of the data collected during the three field trials

1.5 Organization of the Report

This report is organized into five sections. Section 2 documents the field data collection efforts. Section 3 provides a detailed review of the processed data. Section 5 contains the conclusions, and Section 6 provides recommendations for further development and testing.

2. Field Data Collection

Pavemetrics' LRAIL is based on the principle of 3-dimensional laser triangulation. LRAIL integrates high-speed industrial lasers, cameras, precision optics, and inertial measurement units to capture a 13-foot-wide, high-resolution image and a geometrically accurate 3-dimensional scan of the entire track bed. Measurement resolution is approximately 0.2 mm in the z-axis and 1 mm in the x- and y- axes.

These sensors are designed to be mounted on a hi-rail vehicle or a revenue service car operating at speeds between 0 and 100 mph (0–160 km/h). As the vehicle moves, the sensor captures both intensity and range data (i.e., high-resolution images and 3-dimensional models) of the rails, ties, ballast, and other features. As such, this technology can provide the combined benefits of both imaging systems (e.g., line-scan cameras) as well as laser profiling systems in a single unit.

Computer algorithms are used to automatically analyze both high-resolution images and 3-dimensional profiles to make measurements as well as detect assets, defects, and changes between runs. For example, the system can detect and report changes related to tie condition, fasteners inventory, ballast volumes, and geometry parameters, such as gauge, crosslevel, and alignment in a single pass.

2.1 Sensor Installation

For this project, a specialized hi-rail trailer was constructed and fitted with the LRAIL sensor technology, battery-bank power supply, data storage computers, GPS receivers, and related hardware (Figure 1 and Figure 2).

Figure 1 – LRAIL Trailer

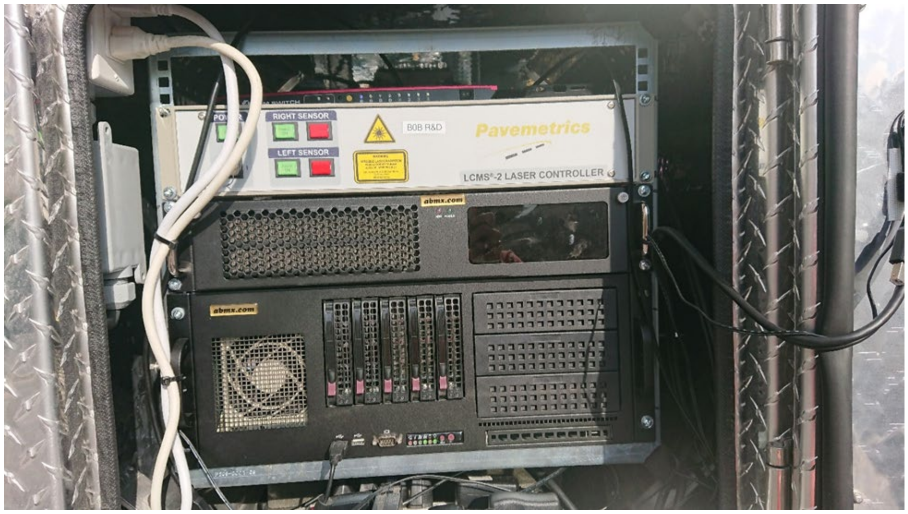

Figure 2 – Onboard Computing and Power Supply

This self-contained and self-powered configuration allowed researchers to quickly deploy the LRAIL onto track for testing without time-consuming installation on railcars or maintenance vehicles.

An optical encoder (Figure 3) was used to trigger image capture and images were sent to a frame grabber to be digitized and then processed by the CPU. Image compression was performed on-the-fly using lossless algorithms to minimize data storage without compromising the usefulness of the data.

Figure 3 – Optical Encoder (attached to rail wheel)

2.2 Data Capture

To streamline data collection in the field, a real-time software interface was developed. On-screen readouts include mile post and speed, real-time graphs for geometry, and display of track images (Figure 4).

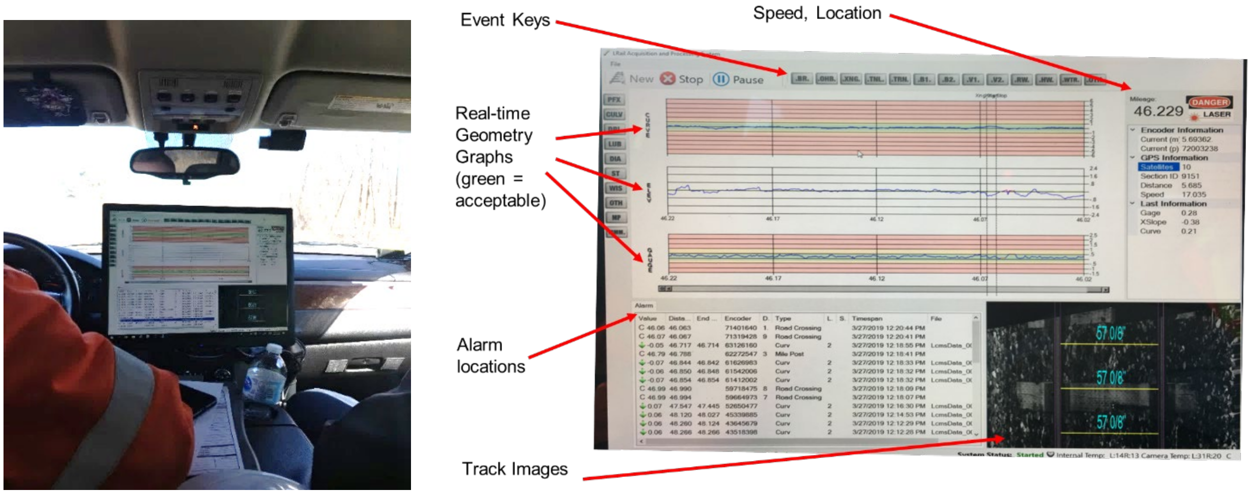

Figure 4 – Real-time User Interface

The LRAIL simultaneously captured both intensity and range images of railway assets. Intensity images (Figure 5) are produced by mapping the intensity of the reflected laser light, and range images (Figure 6) are produced by mapping the elevation of each measurement point.

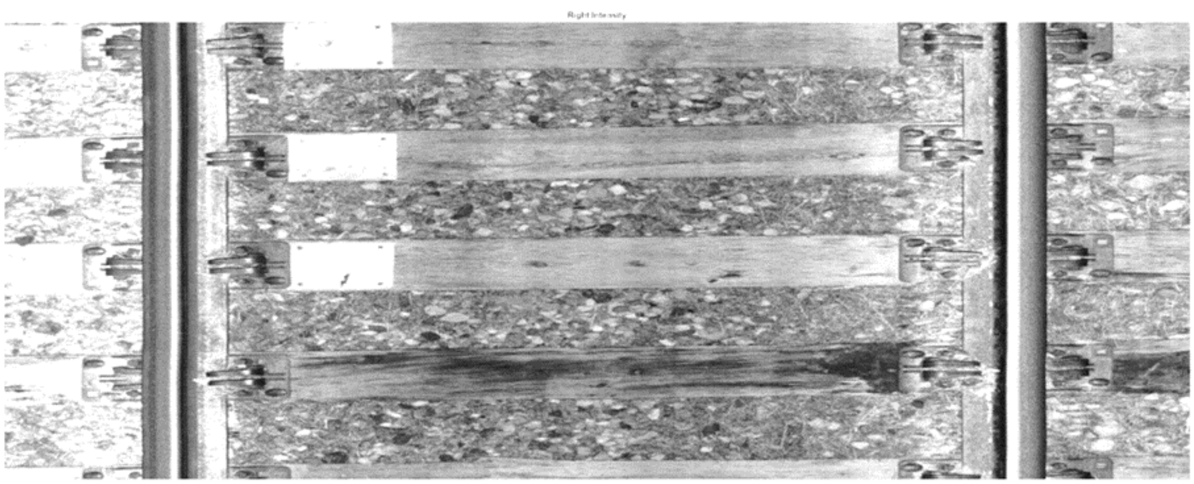

Figure 5 – Intensity Image

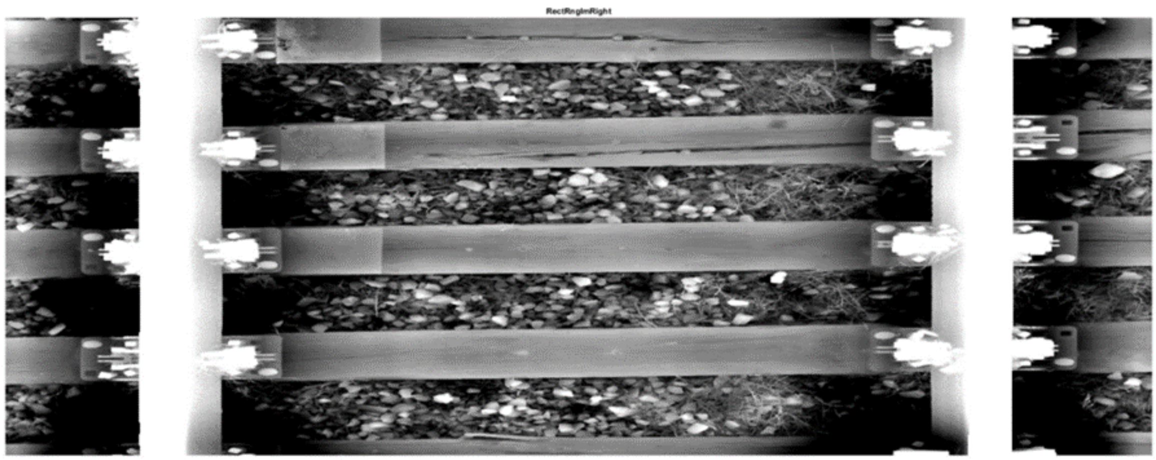

Figure 6 – Range Image

The resulting images are approximately 4,000 pixels wide with a lateral resolution of one point per millimeter and a longitudinal scan interval (as the inspection vehicle travels along the track) of one scan per millimeter. The result is a 1 mm by 1 mm scan of the railway (Figure 7).

Figure 7 – 1 mm by 1 mm Scan of Railway

2.2.1 Motion Compensation

An inertial measurement unit is enclosed in each sensor head and used to track minute changes in the pitch, roll, heading, and acceleration of each sensor. These orientation data are used by a computer algorithm to remove vehicle motion artifacts from 2-dimensional and 3-dimensional data (Figure 8).

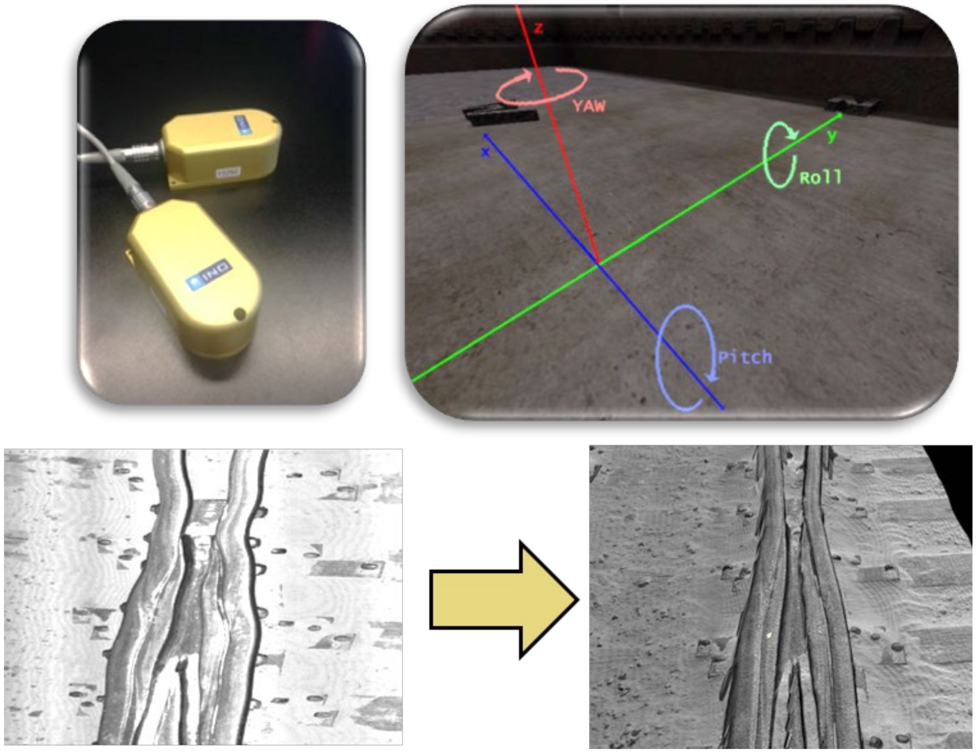

Figure 8 – Correction for Vehicle Motion

2.3 Test Site

For this project, data was collected on Amtrak starting at approximately MP 25 on the south leg of Wye-Landlith until approximately MP 17 at the Hook Interlocking (Figure 9, the data collection "loop"). The test zone runs north of Wilmington, DE to Marcus Hook, PA, through the towns of Bellefonte, DE and Claymont, DE.

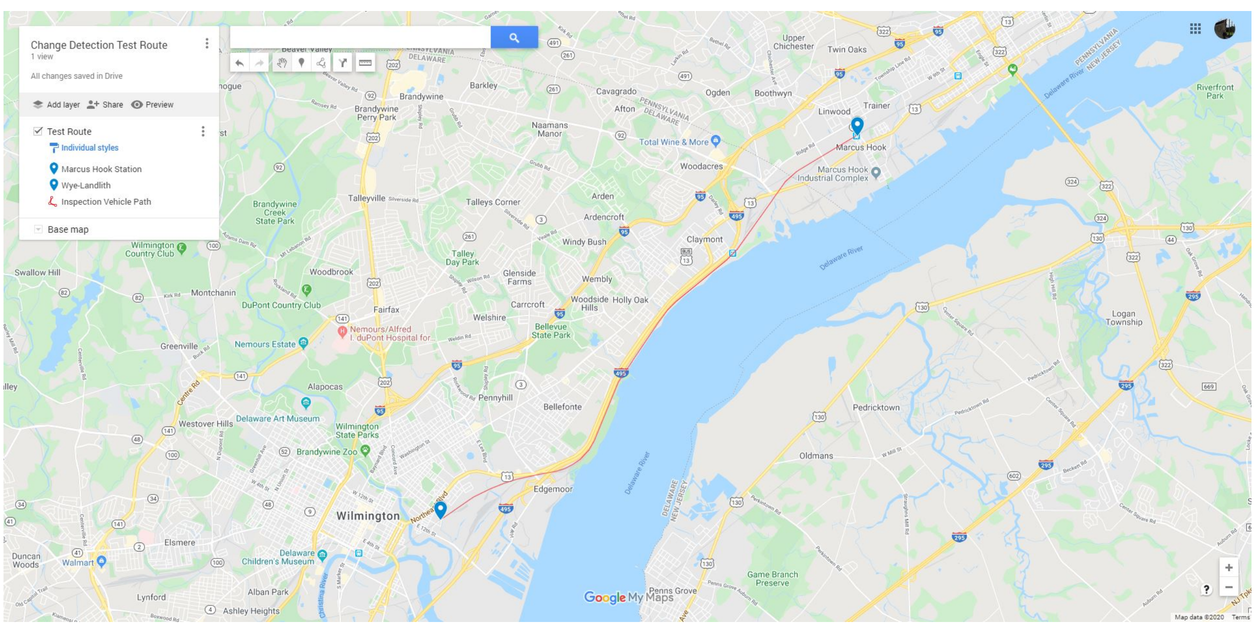

Figure 9 – Test Loop (MP 25 at Wye-Landlith to MP 17 at the Hook Interlocking)

This route was selected for a variety of reasons, including:

- It is an active length of track where organic changes are expected to develop during the test period due to wear as well as the performance of corrective maintenance,thus supporting a number of the identified use cases.
- It will permit the inspection of the identified regions of interest, including examples of both wooden and concrete ties as well as special trackwork such as switches and crossings.
- This length of track is currently inspected twice weekly during the daytime by an inspector in a hi-rail vehicle, providing an opportunity to validate detected changes against manual field inspection observations.

2.4 Test Scheduling and Methodology

Field testing for this project involved three field deployments, spaced approximately 2 months apart, with an initial deployment in May 2019, a second in July 2019, and the final deployment in September 2019.

Each field deployment involved the use of repeat runs (i.e., driving the inspection route in the same direction multiple times) as well as multiple reverse runs (i.e., driving the inspection route in the opposite direction) to model the noise floor and establish repeatability of automated change detection.

The field data collection plan for each deployment was as follows:

Deployment 1 (May 2019)

- Day 1 night-time between approximately 11:00pm on May 13 and 1:00am on May 14: One primary direction inspection run (decreasing chainage from MP 25 to MP 17) and one secondary direction inspection run
- Day 1 daytime on May 14: The introduction of manual changes
- Day 2 night-time between 10:30pm and just before midnight on May 15: One primary and one secondary direction inspection run
- Day 2 daytime on May 15: Manual changes are reversed (to some degree; for example, it would be impossible to restore ballast to its exact original state)
- Day 3 night-time between approximately 11:00pm on May 15 to 1:00am on May 16: One primary and one secondary direction inspection run

Deployment 2 (June 2019)

- Day 4 night-time between approximately 11:00pm on June 26 to 1:00am on June 27: One primary and one secondary direction inspection run
- Day 5 night-time between approximately 11:00pm on June 28 to 1:00am on June 28: One primary and one secondary direction inspection run

Deployment 3 (September 2019)

- Day 6 night-time between approximately 11:00pm on September 4 to 1:00am on September 5: One primary and one secondary direction inspection run

- Day 7 night-time between approximately 11:00pm on September 5 to 1:00am on September 6: One primary and one secondary direction inspection run

In general, the data collection went according to plan, but there were some occassions when portions of the test loop were unavailable due to ongoing maintenance. There were also occassions when it was not possible for the inspection vehicle to leave the track in order to drive the loop in the reverse direction, which required the inspection vehicle to simply travel in reverse to the starting point.

2.5 Introduction of Deliberate Changes

A series of deliberate changes (Figure 10) were made during the first field deployment between test runs to safeguard against insufficient organic changes to track conditions over the project timeframe. Changes were made to the ballast level in the crib area as well as at the end of ties, tie skew, fastener installation, and joint bar bolting.

Figure 10 – Deliberate Changes to Ballast Level

Detailed field notes and pictures were recorded for each change (Figure 11), and this information was subsequently compiled into a single log of track conditions and deliberate changes (Figure 12).

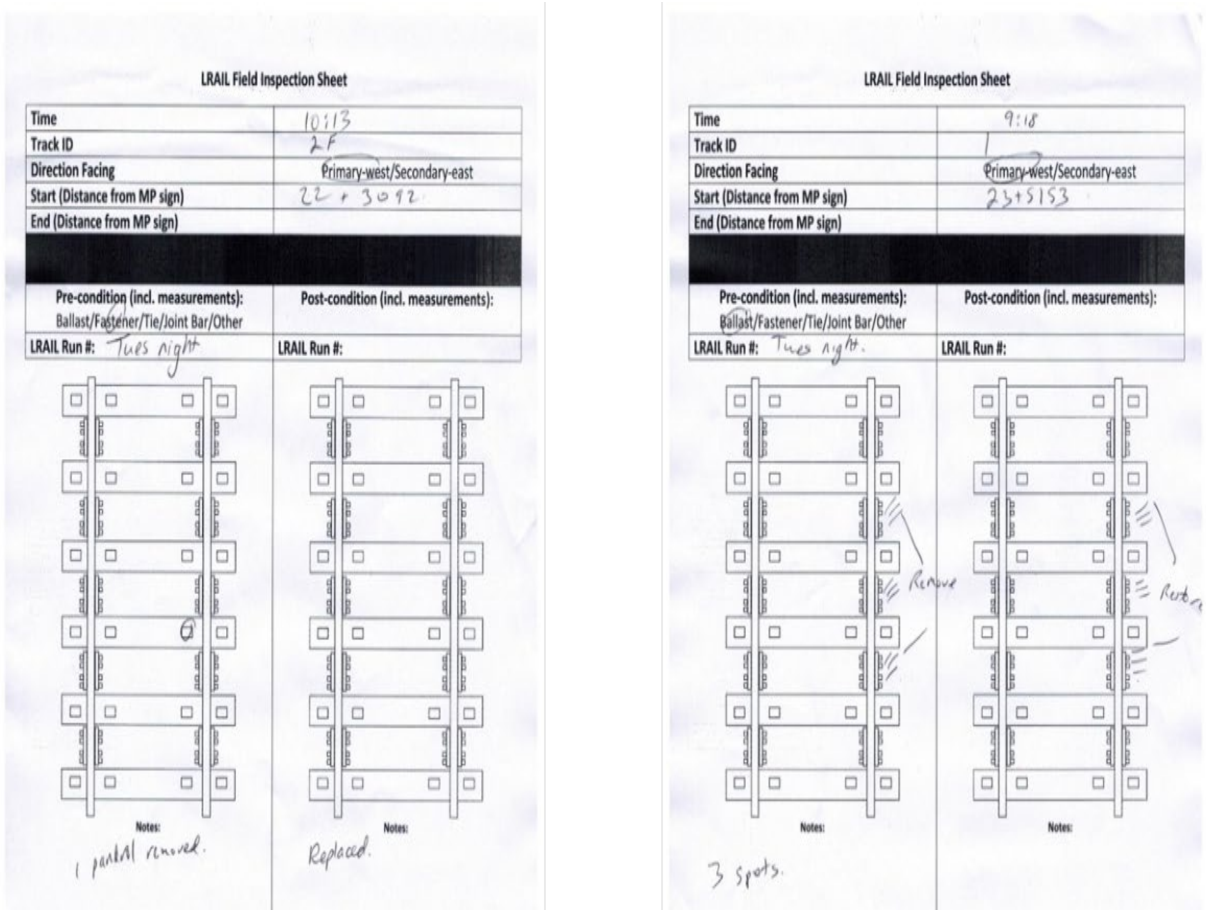

Figure 11 – Field Notes to Record Deliberate Changes

Amtrak Change Log

*All changes were made during the daytime on Tuesday May 14 so they will be captured in the FIS files from Tuesday May 14 PM and early monrning on the 15th (Day2). The track was to be restored to the original state (to the degree that is is possible; ballast would not be exactly the same) during the day of May the 15th. So then the conditions for May 13 (Day1) should be different from the 14th (Day2) and from the 15th (Day3).

*There was a problem with acquisition software during the decrease mileage run of Day3, so the decrease mileage run of Day3 contains 2 runs (run1 and run2). Due to restart of acquisition software, we did not scan about 74 sections (around 150m) near milepost 23 (milepost 23 - 400m)

*Day1, we could not make a u-turn so the run2 is in reverse direction. We need to use a backward calibration file to process the data of Day1/run2

* Concrete ties are from milepost 17 to milepost 20 + 558m; wooden ties are from milepost 20 + 558m to milepost 25

* Day4 (Deploy2/Day1), there's problem with the antenna so there's no GPS signal for both runs of Day4

										Section ID		
Track	Location (MP)	Time	Direction Facing	Change Type	Notes	Run1 (Deploy1/Day1)	Run2 (Deploy1/Day1)	Run1 (Deploy1/Day2)	Run2 (Deploy1/Day2)	Run1 (Deploy1/Day3) Decrease Mileage	Run2 (Deploy1/Day3) Decrease Mileage	Run3 (Deploy1/Day3) Increase Mileage
1	24	9:03am	Primary	Ballast	Ballast was removed from 4 end of tie locations on field side	1108	5622	1085	5646	1210		5923
1	23+5344	9:06am	Primary	Ballast	Ballast was removed from 2-3 crib locations (2 very empty, 1 partially)	1118	5613	1095	5637	1219		5913
1	23+5274	9:12am	Primary	Ballast	Ballast was removed from 4 end of tie locations on field side	1127	5603	1104	5627	1228		5904
1	23+5192	9:14am	Primary	Ballast	Ballast was removed from 2-3 crib locations (2 very empty, 1 partially)	1141	5590	1118	5614	1243		5890
1	23+5153	9:18am	Primary	Ballast	Ballast was removed from 3 end of tie locations on field side	1147	5584	1124	5608	1248		5885
1	23+5100	9:20am	Primary	Ballast	Ballast was removed from 2-3 crib locations (2 very empty, 1 partially)	1154	5577	1131	5601	1255		5877
1	23+5012 - 23+5000	9:36am	Primary	Tie skew	1 tie was skewed and 7 ties later a second was. The plate is still on the tie. The ballast at the end of the tie on the field side was also removed.	1167, 1169	5561, 5563	1144, 1146	5585, 5587	1269, 1271		5862, 5864
1	23+4916	9:41am	Primary	Tie skew	1 tie was skewed and ballast at the end of the tie on the field side was removed. The plate is still on the tie.	1181	5549	1158	5573	1282		5850
2F	22+3092	10:13	Primary	Fastener	1 pandrol removed; right hand rail	2266	4463	2244	4488		32	4766
2F	22+3100	10:14	Primary	Fastener	1 pandrol removed; right hand rail	2265	4464	2243	4490		30	4767

Figure 12 – Deliberate Change Log

3. Change Detection Process

3.1 3-Dimensional Data Analysis

Intensity and range data scans captured in the field are analyzed by the LRAIL's proprietary processing library which combines artificial intelligence (a deep neural network) and traditional image processing to automatically detect railway features and assess their condition (Figure 13). These results, including detect fasteners, joint bar bolts, ballast fouling levels, and tie skew angles, serve as inputs to a proprietary change detection algorithm described in Section 3.3.

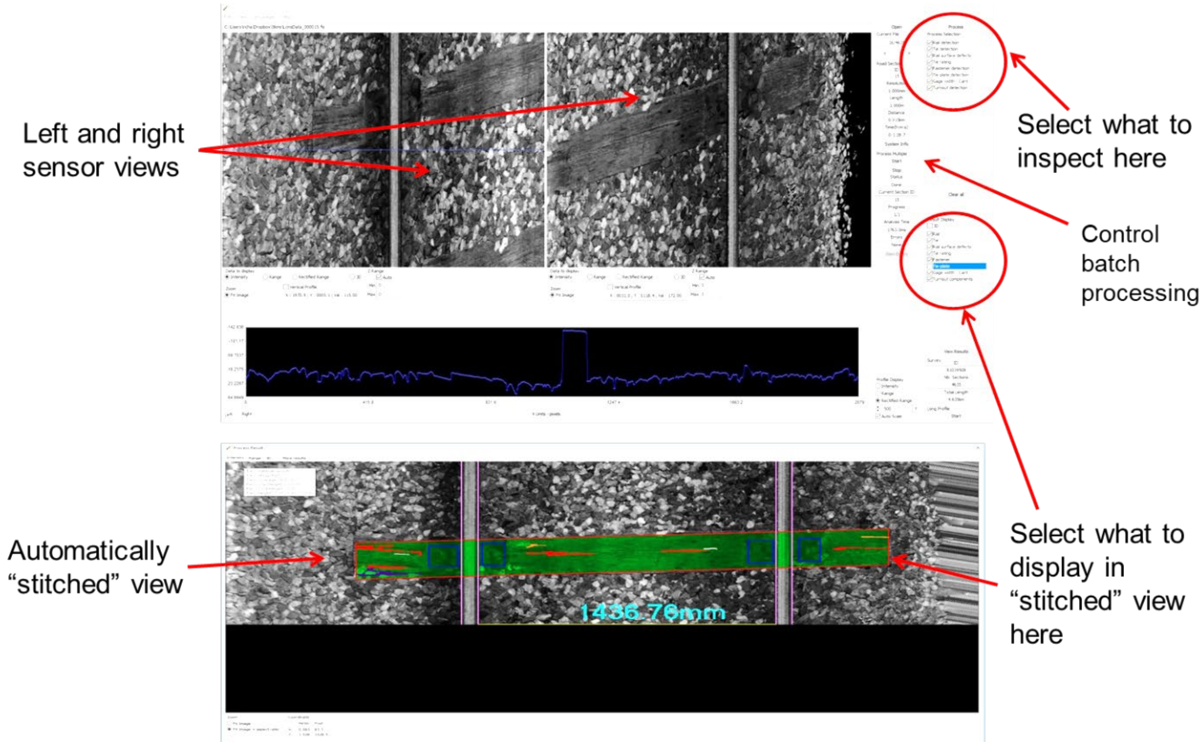

Figure 13 – LRAIL Inspection Software Interface

3.2 Change Detection Process Flow

The process flow for change detection is presented in Figure 14.

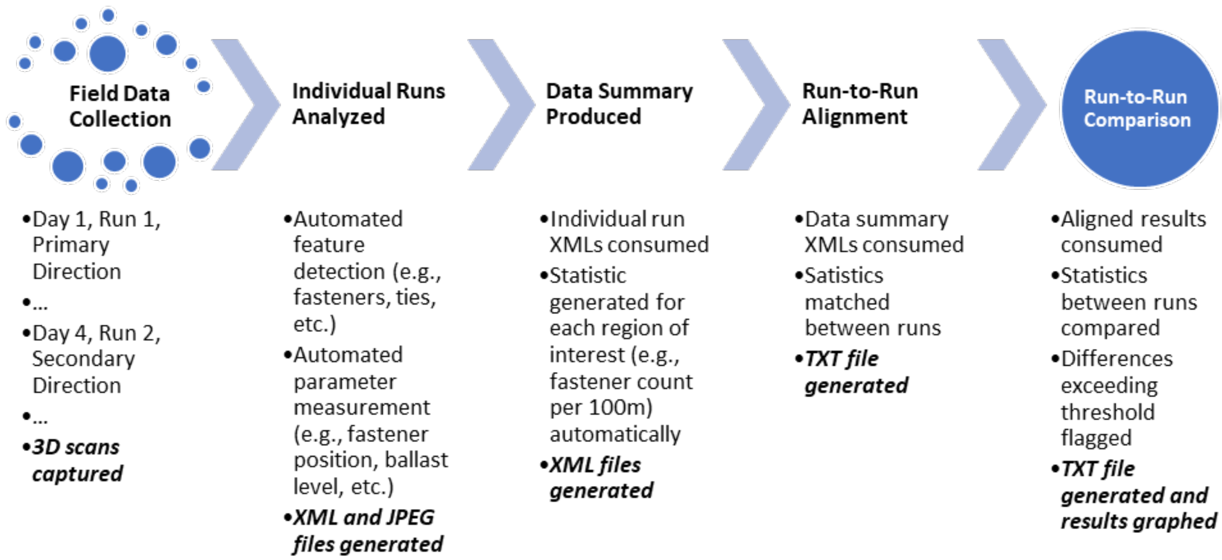

Figure 14 – Change Detection Process Flow Summary

3.3 Run-to-Run Alignment

To model the noise floor of the system as well as to detect change between repeat runs, it is essential to first have a method to align the two runs. The alignment process involves matching data streams from each run using location information to ensure that comparisons are performed at the exact same location. Errors in alignment result in the detection of false changes, failed detection of real changes, and false modeling of the noise floor.

For this project both linear position (chainage via wheel encoder) as well as geospatial position (latitude, longitude, and elevation via GNSS receiver) were captured to align positions between separate data collection runs to detect changes.

However, there are inherent challenges in aligning data due to a variety of factors related to positional information:

- Stop and start locations are never the same between two runs.
- Inspection runs can be captured in different travel directions (increasing vs. decreasing chainage).
- Wheel encoder positional accuracy tends to decrease as the driving distance increases, effectively resulting in the accumulation of error.
- GNSS positional data varies due to satellite visibility, multipath error, and tropospheric conditions.

Two methods were developed in order to align data; the first was a semi-automated method and the second an automated method.

3.3.1 Semi-automated Alignment Method

This method is useful when GPS data are not available (e.g., in remote locations or tunnel environments). The semi-automated method utilizes the linear position of a single physical reflector target (Figure 15) at the start of the test run to align the two runs. To develop this approach, reflector targets were installed at the start, end, and midpoint of the test route. Although, ultimately, only a single target was required.

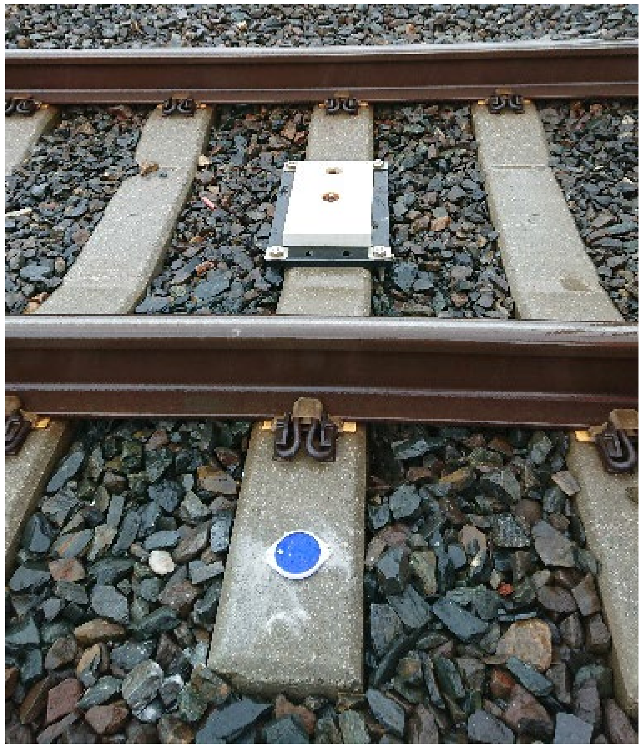

Figure 15 – Reflector Target

With semi-automated alignment, the linear location (as measured by the wheel encoder) of a single target was determined by selecting the target in the 3-dimensional scan from each of the two runs. An automated algorithm then worked outward from the matched target location in each run to align the data on a tie-by-tie basis. This method is referred to as semi-automated, as there is a single manual step at the start of the process while the remaining steps are all computer-driven.

Automated tie-by-tie matching is performed throughout the length of the runs to compensate for wheel encoder errors which naturally accumulate as the track is driven. For example, even with a relatively low wheel encoder error of approximately 0.2 percent, the positional error after driving 1 km of track (0.62 mile) will be about 2 m (2.18 yd). Without correction, this error would result in the reported position of the same tie between runs to vary by 2 m after 1 km, making it quite difficult to match between runs.

3.3.2 Automated Alignment Method

A second, fully automated, method was also developed for alignment (Figure 16). This method does not require the installation of physical targets and provides superior alignment. This method relies on using the automatically determined GPS position (i.e., latitude, longitude and elevation) of each tie to align two runs.

This algorithm compares the GPS position of all ties in each run and selects a single tie location in each run closest to a tie in the other run. The algorithm then moves outward to adjacent ties and compares automatically reported tie parameters, such as skew angle, fastener count, and ballast cover, to validate the initial match. Once validated, data streams from each run are aligned with one another on a tie-by-tie basis.

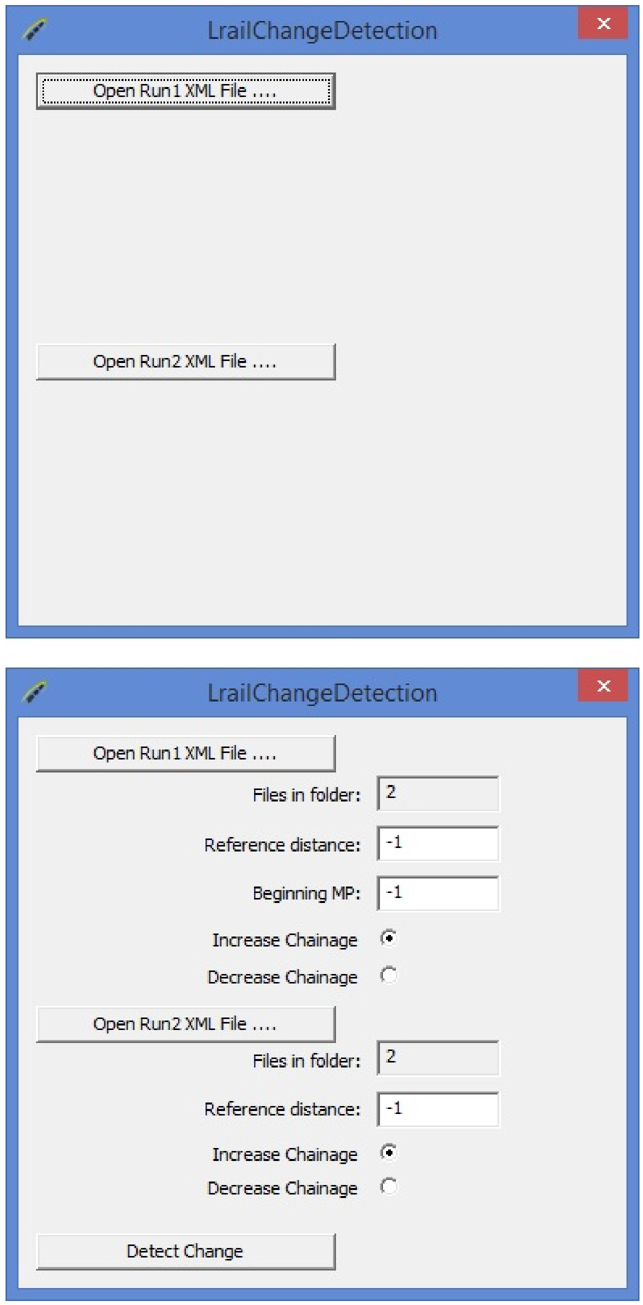

Figure 16 – Automated Run-to-Run Alignment Software

3.4 Repeatability, Mean, and Standard Deviation

Back-to-back runs, collected in opposite directions (one in increasing chainage and the other in decreasing chainage) and without deliberate changes, were analyzed to determine the repeatability, mean, and standard deviation of each measurement.

3.5 Modeling the Noise Floor

One of the project objectives was to determine a practical limit for the magnitude of change that can be detected (i.e., the noise floor of the system). To accomplish this objective, data were compared between repeat and reverse direction inspection runs of the unchanged test loop.

To determine a noise floor for each parameter, measurement results were graphed for each set of forward and reverse-direction back-to-back runs (therefore without changes) as well as for runs over the life of the project (therefore with changes).

Engineers reviewed graphs to determine the natural error in measurement encountered for repeat runs of unchanged track. As well, the magnitude of known changes were noted. The magnitude of the largest outliers was used to set the minimum threshold for reporting a valid change between runs. This approach was conservative, with the potential to ignore some small changes, but is the best method to eliminate false positives from the results.

4. Change Detection Performance Results

The following sections describe the method used to evaluate the performance of the change detection tool and presents the results for each criterion studied.

A table of key areas of interest for change detection (Table 1) was developed through consultation with FRA and Amtrak. Draft thresholds for each region of interest were also developed based on the magnitude of change that a track inspector would deem relevant to detect.

Table 1 – Areas of Interest for Change Detection

Area of Interest	**Reporting Metric (including draft thresholds)**
Ballast Height in Crib Area	Significant changes in average level relative to top of rail (**a change >= 10% in volume or area in a single crib)**
Ballast Height within Gage and End of Tie Areas	Significant changes in average level relative to top of rail within the gage or end of tie areas along a length of track (**a change >= 30% in average height (less ballast compared to the last run) for an individual crib or >= 10% in area for >= 3 m (9.84 ft))**
Ballast Fouling	Changes in fouling (presence of non-ballast materials), changes in the presence of water or moisture **(a change in affected area >=0.15 m^2 (1.61 ft^2))**
Tie Skew Angle	Significant changes in the angle of multiple ties **(at least 3 consecutive ties with an increase >= 2 degrees per 1 km (0.62 mile))**
Tie Condition	Significant changes to individual ratings of ties per km **(>= 10% change in any tie rating category)**
Joint Bar Fasteners	Significant changes to bolting **(one or more missing bolts for any given bar)**
Rail End Gap	Significant changes in joint gaps **(an increase of at least 10 mm (0.39 in) for any given joint)**
Rail Fastener Inventory	Significant changes in present fastener counts per km (0.62 mi) as a percentage (**+/-1%**), significant changes in the position of multiple fasteners **(>= 10 mm (0.39in) for three consecutive fasteners in the same position along the same rail)**
Rail Head Surface	Significant changes in surface defects **(at least a 10% increase in total surface area of defects per 1km (0.62mi) of track)**

4.1 Ballast-in-Crib Height Change Detection

This task involved run-to-run comparisons regarding the mean level of ballast material present in the crib area of the same crib between repeat runs (Figure 17).

Figure 17 – Ballast-in-Crib Height Measurements

Ballast level is measured as the absolute distance between the planer surface of the top-most point on each rail and the mean height of the ballast surface (Figure 18).

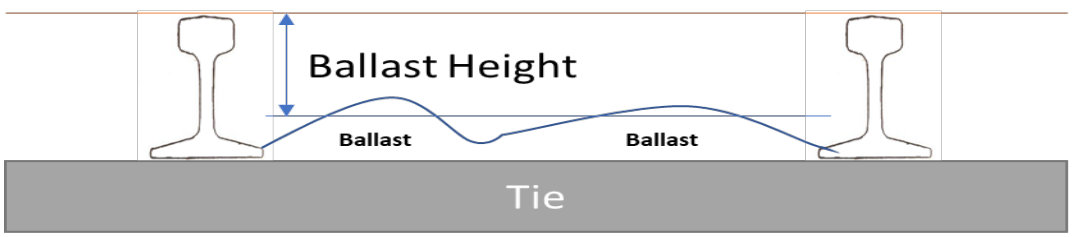

Figure 18 – Ballast Height Measurement

A total of approximately 22,000 measurements of ballast-in-crib level were made in each run.

Overall repeatability for ballast-in-crib measurements was excellent, with repeat runs achieving 99.93 percent agreement. Overall mean for ballast-in-crib level differences was -0.055 mm, meaning that, on average, measurements between runs varied by 0.055 mm. Standard deviation of reported ballast-in-crib difference between runs was 0.751 mm.

Table 2 – Ballast-in-Crib Repeatability, Mean, and Standard Deviation

Ballast-in-Crib (mm) 22,000 Instances			
D2R2 vs D2R1		D7R2 vs D7R1	
Repeatability	99.70%	Repeatability	99.97%
Mean	-0.107	Mean	-0.032
Std	0.717	Std	0.934
D3R2 vs D3R1		D3R2 vs D3R1a	
Repeatability	99.97%	Repeatability	99.95%
Mean	-0.034	Mean	-0.137
Std	0.728	Std	0.743
D4R2 vs D4R1		D3R2 vs D3R1b	
Repeatability	99.97%	Repeatability	99.98%
Mean	-0.109	Mean	0.018
Std	0.724	Std	0.721
D5R2 vs D5R1		Overall	
Repeatability	99.97%	Repeatability	99.93%
Mean	0.013	Mean	-0.055
Std	0.689	Std	0.751

Upon review of the back-to-back runs, the comparison between the second and first runs of Day 2 (D2R2 and D2R1) contained the largest outliers. No deliberate changes were made to track conditions between these runs, and they were captured within approximately two hours of one another.

Figure 19 presents the run-to-run comparison with the difference in measurements plotted along the Y-axis and the location for each comparison plotted along the X-axis. Each circle in the plot represents a height comparison between runs for a given crib. Note that a negative value implies an increase in ballast level, since the absolute distance between the rails and the ballast surface decreased while a positive value implies a decrease in ballast height. Note that special trackwork and bridge installation locations were deliberately omitted from analysis. These areas are located by arrows in the figure.

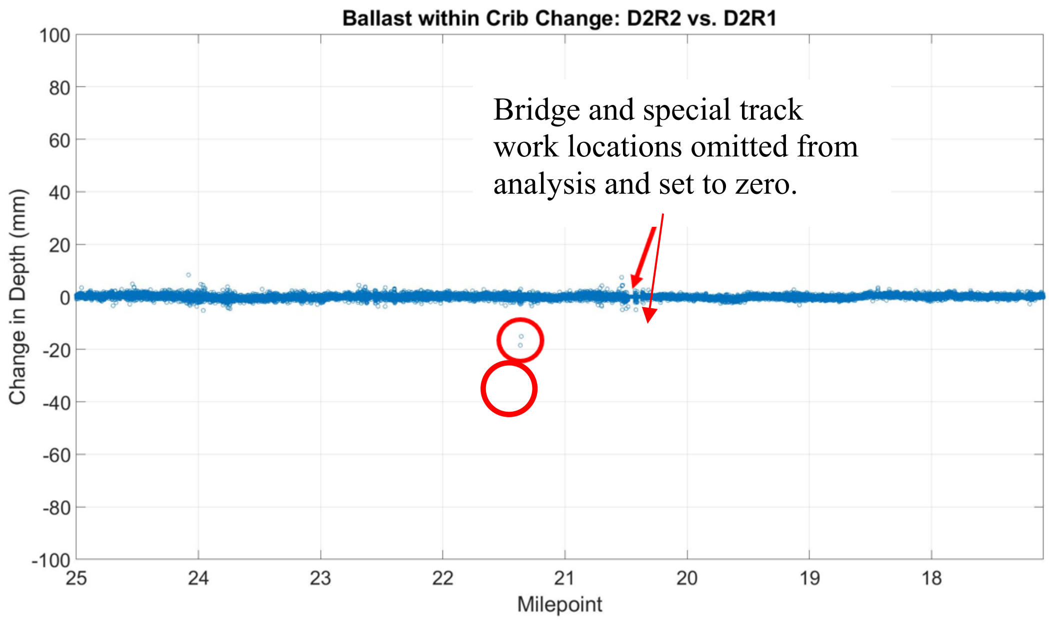

Figure 19 – Modeling Ballast-in-Crib Height Noise Floor (D2R2 vs. D2R1)

As can be seen in Figure 19, there were two significant outliers near MP 21.25 (circled in red) that approached 20 mm in magnitude. Upon investigation of the 3-dimensional scans (Figure 20) a utility box mounted in the crib area was discovered as the source of error.

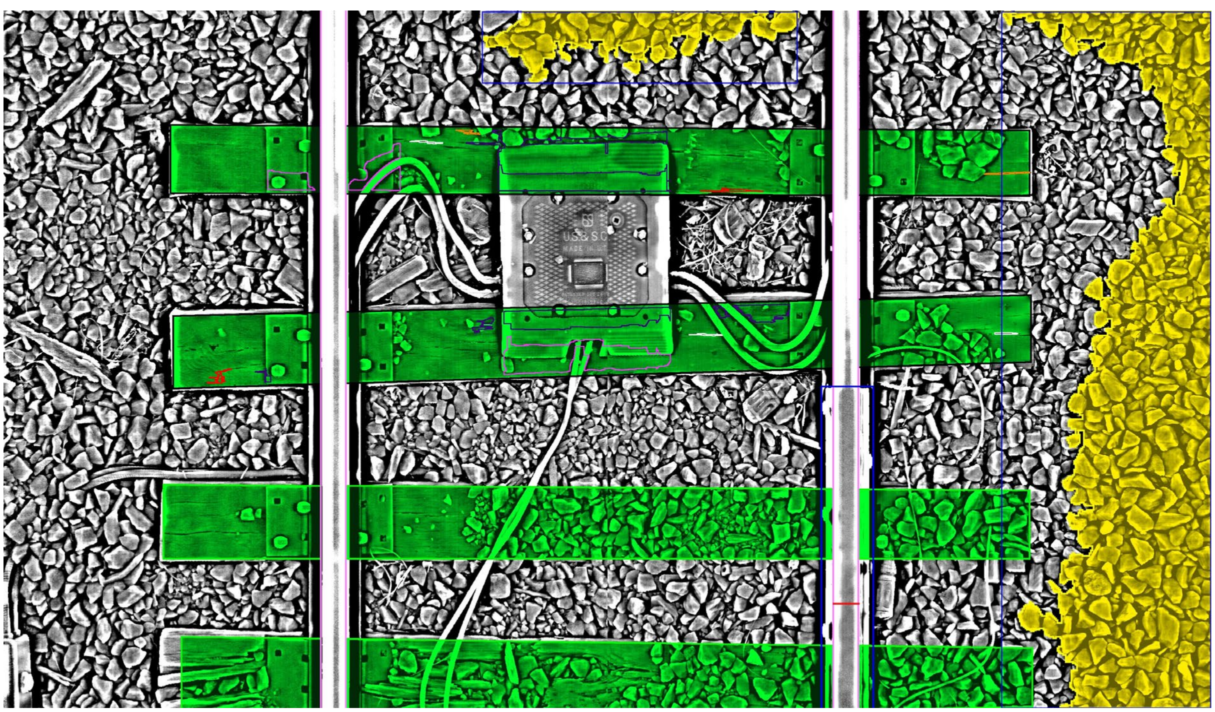

Figure 20 – Utility Box that Led to False Change Detection

When comparing the 3-dimensional data from the two runs, small variations can be seen which resulted from different laser light travel paths between the primary and secondary direction.

In future development, the researchers believe that this error can be eliminated (e.g., through the deliberate detection of such boxes); however, for the time being a noise floor of 20 mm is recommended to avoid change detection false positives.

Deliberate changes were made to ballast-in-crib height (Figure 21) to validate the performance of this algorithm.

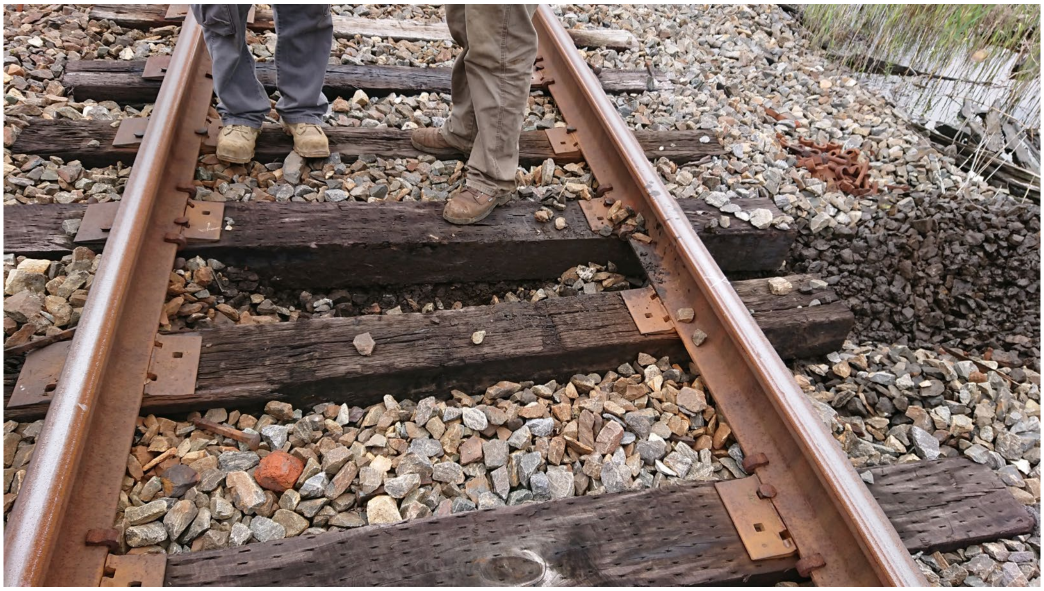

Figure 21 – Deliberate Ballast-in-Crib Height Change

Deliberate changes to the ballast height were made at the following locations between Day 1 and Day 2:

- MP 23+5344; ballast was removed from three cribs.
- MP 23+5192; ballast was removed from three cribs.
- MP 23+5100; ballast was removed from three cribs.

Consequently, when plotting the difference in ballast-in-crib height from any of the scans captured on Day 1 to those captured on Day 2, researchers expected that a number of data points around MP24 would exceed the 20-mm (0.87 inches) noise floor. Figure 22 presents this comparison (D1R1 vs. D2R1).

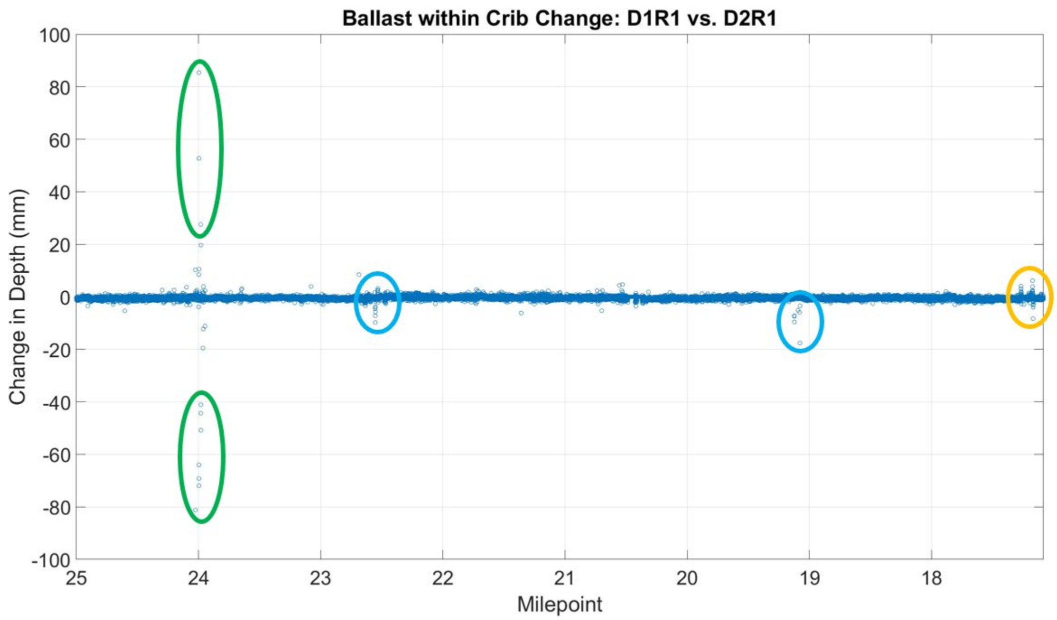

Figure 22 – Detecting Deliberate Ballast-in-Crib Change (D1R1 and D2R1)

As expected, in Figure 22 there is a clear trend of data points exceeding the 20-mm noise floor at the expected position of MP24 (circled in green). Note that there is a total of 10 reported change locations, as opposed to the expected 9. This was the unintended result of deliberate changes made to the tie skew angle in the same general area which also produced a change to ballast-in-crib height adjacent to the skewed tie.

Also of note in Figure 22 were some small responses near MP 22.5 and MP 19 (circled in blue) and two near MP 17 (circled in orange). Although these signals were below the 20-mm noise floor, upon investigation of the 3-dimensional scans there are some interesting things learned. The points near MP 22.5 and MP 19 were simply due to excessive changes in surface moisture between D1R1 and D2R1; the ballast material was clearly saturated on D1R1 (which was collected following very heavy rain). By D2R1, the surface was much drier here. However, the two responses near MP 17 told a different story. When the 3-dimensional scans (Figure 23 and Figure 24) were reviewed for this location, small, unplanned, and unrecorded changes (about 10 mm) to a track sensor installation could be found.

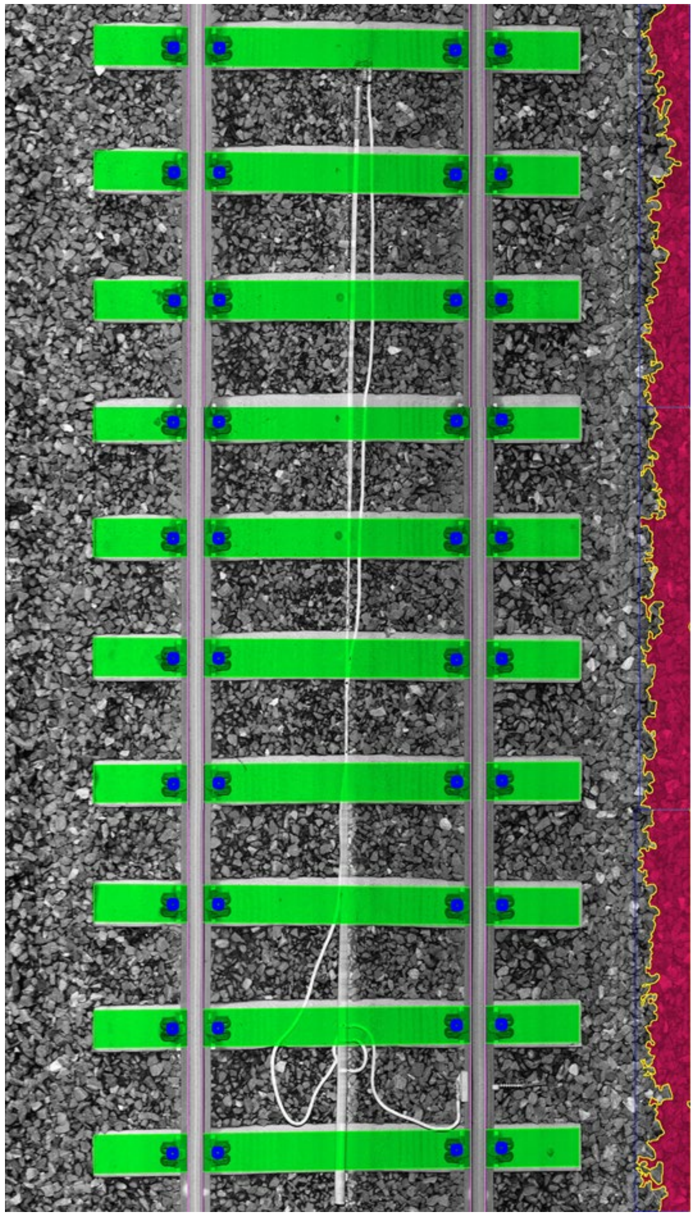

Figure 23 – Unknown Sensor Installation in Crib Area (D1R1)

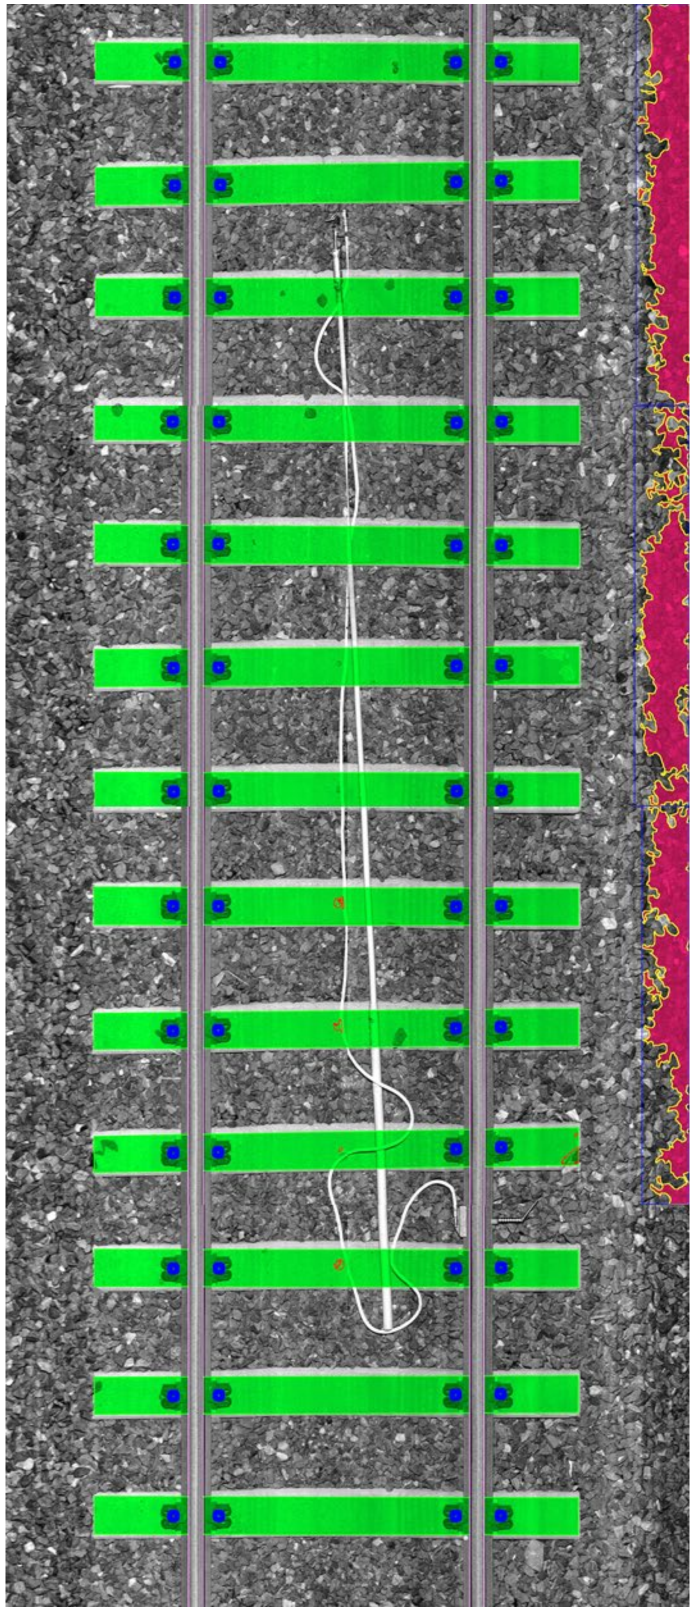

Figure 24 – Unknown Sensor Installation in Crib Area (D2R1)

The magnitude of change for this feature was quite small, as the position of the sensor installation components (a cylindrical tube with some cabling) were moved only a few centimeters, but the change was accurately detected by the system.

Organic changes to ballast-in-crib height occurred over the test period and were automatically detected. An example of this can be seen when comparing ballast-in-crib heights measured in September 2019 (D7R1) versus those measured in May (D2R1), presented in Figure 25.

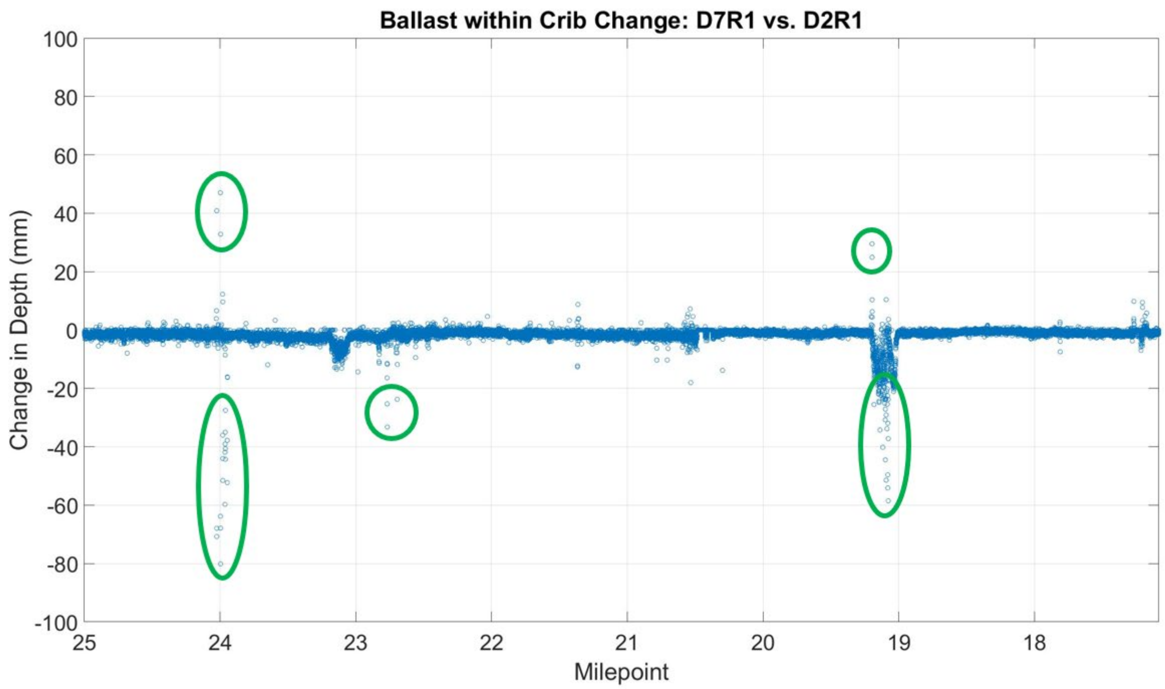

Figure 25 – Ballast-in-Crib Changes between May and September

The figure shows a number of track locations with significant changes to ballast-in-crib levels, including changes at MP 24, MP 22.75, and MP 19 (circled in green). The majority of the changes between May and September involved the reduction of ballast-in-crib levels, although there were a few increases as well.

4.2 Within Gauge and End-of-Tie Ballast Level Change Detection

This task involved run-to-run comparisons of ballast height at the end of each tie and in the gauge area along a track segments (Figure 26).

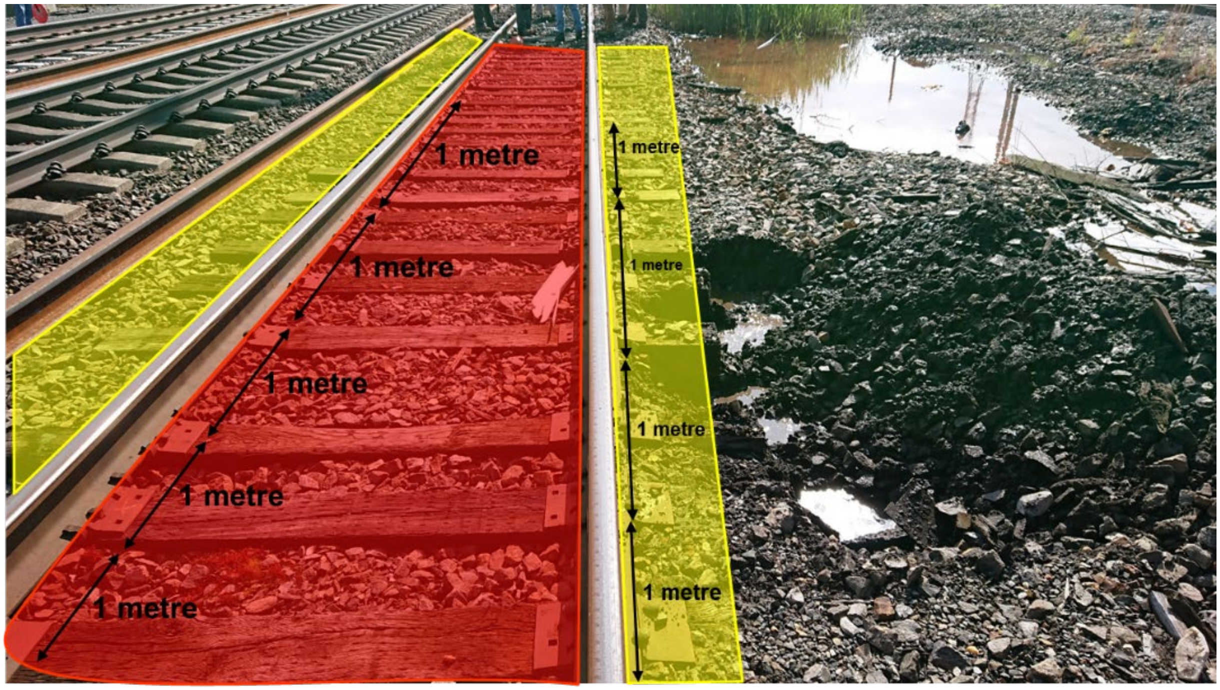

Figure 26 – Running Meter Ballast Height Measurements

An average height of ballast per meter of travel was computed for three track zones – left rail, gauge area, and right rail – and compared between runs. A total of approximately 12,700 measurements of ballast height per meter of travel (in each zone) were made in each run.

Overall repeatability for measurements of ballast height per meter of travel was excellent, with repeat runs achieving 99.99 percent agreement. The overall mean for measurements of ballast height per meter of travel was -0.059 mm, meaning that they varied by 0.059 mm between runs. The standard deviation of reported ballast height per meter of travel between runs was 2.938 mm.

Table 3 – Ballast-Height-per-Meter-of-Travel Repeatability, Mean, and Standard Deviation

Ballast Level in 3 Zones: Left Rail, Gauge, Right Rail (mm per 1m length) 12.7k instances							
D2R2 vs D2R1				D7R2 vs D7R1			
	Left	Gauge	Right		Left	Gauge	Right
Repeatability	100.00%	100.00%	99.90%	Repeatability	100%	100%	100%
Mean	0.843	-0.004	-0.782	Mean	0.521	-0.039	-0.707
Std	0.9176	0.9055	2.4336	Std	0.885	0.61	0.871
D3R2 vs D3R1				D3R2 vs D3R1a			
Repeatability	100.00%	100.00%	100.00%	Repeatability	100.00%	100.00%	100.00%
Mean	0.407	-0.036	-0.433	Mean	0.272	-0.15	-0.357
Std	0.689	0.658	0.753	Std	0.899	0.571	0.854
D4R2 vs D4R1				D3R2 vs D3R1b			
Repeatability	100.00%	100.00%	100.00%	Repeatability	100.00%	100.00%	100.00%
Mean	-1.146	-0.105	0.7181	Mean	0.4752	0.021	-0.471
Std	0.726	0.588	0.7553	Std	0.584	0.701	0.703
D5R2 vs D5R1				Overall Left, Gauge, Right			
Repeatability	100.00%	100.00%	99.99%	Repeatability	100.00%	100.00%	99.98%
Mean	-0.552	0.051	0.224	Mean	0.117	-0.037	-0.258
Std	1.071	1.329	3.061	Std	0.825	0.766	1.347
Overall All Zones							
Repeatability	99.99%						
Mean	-0.059						
Std	2.938						

Upon review of the back-to-back runs, the comparison between the second and first runs of Day 2 (D2R2 and D2R1) were determined to contain the largest outliers for this parameter. No deliberate changes were made to track conditions between these runs, and they were captured within approximately 2 hours of one another. Figure 27 presents the run-to-run comparison with the difference in measurements plotted along the y-axis and the location for each comparison plotted along the x-axis. A negative value indicates a decrease in ballast level between scans, and a positive value denotes an increase. The three graphs present changes located in the left field side (when scanned in decreasing chainage), gauge side, and the right field side. Each circle in the plot represents a comparison between runs of average ballast height per meter of track.

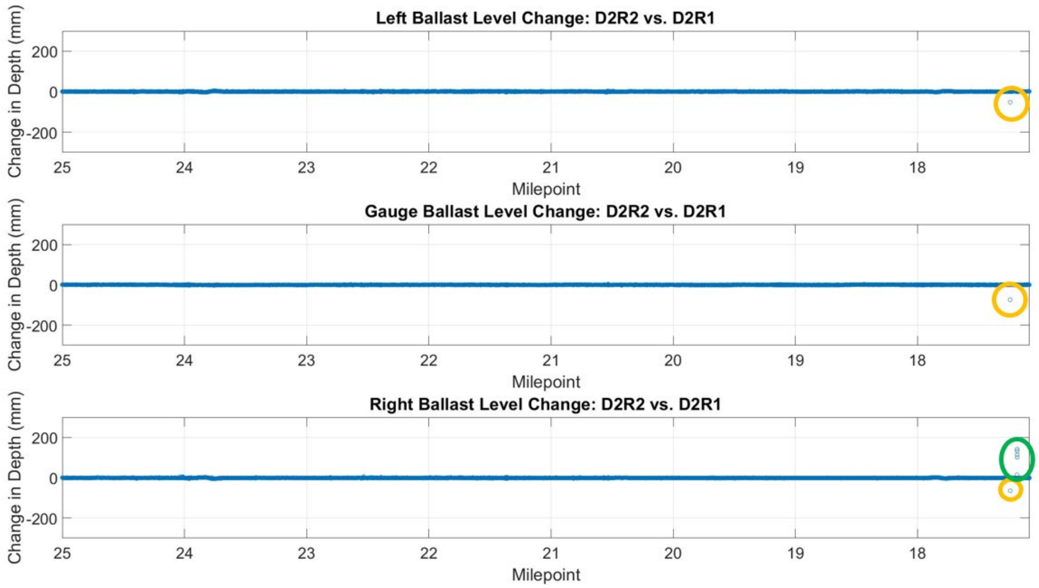

Figure 27 – Modeling Ballast Height at Gauge and End-of-Tie Noise Floor (D2R2 vs. D2R1)

As shown in Figure 27, the typical noise level sat around zero to 5 mm, but there were multiple outliers at magnitudes as high as 150 mm near MP 17 (circled in green and orange). Compared to the average mean difference between runs of -0.059 mm, these values seemed far too large to use as the noise floor. Upon review of the 3-dimensional scans, researchers learned that there were in fact a number of real changes at that location which were unplanned and unrecorded:

- The removal of wooden planks adjacent to the right rail (circled in green on middle graph in Figure 27 and in 3-dimensional scans in Figure 28)
- The removal of a long cylindrical object in the gauge area (circled in green on bottom graph in Figure 27 and in 3-dimensional scans in Figure 29)

Figure 28 – Unplanned Removal of Wooden Planks (D2R1 on left and D2R2 on right)

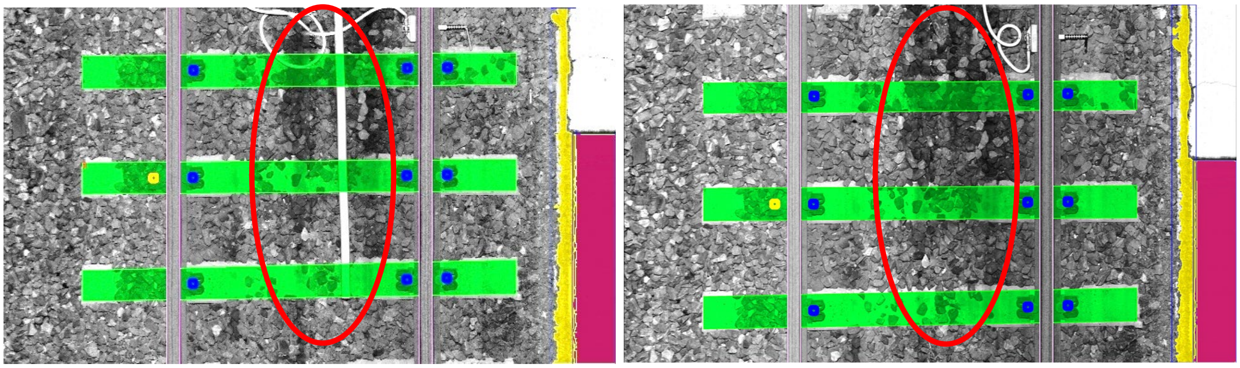

Figure 29 – Unplanned Removal of Pole (D2R1 left and D3R1 right)

Consequently, these legitimate changes were omitted when considering appropriate values for the noise floor.

There were also changes detected in the ballast adjacent to the left rail, ballast in the crib, and ballast adjcent to the right rail – all at the same location near MP 17 (circled in orange in the top graph of Figure 27). However, upon review of the 3-dimensional scans (Figure 30) these reported changes were determined to be the result of variability between scans related to top-of-rail detection (used as a point of reference for ballast height measurements). In one run, the highest points on the two rails (green circles in Figure 30) were used to define the frame of reference for ballast height measurements on the left side, while in the second run an adjacent wooden plank was selected (orange circle in Figure 30) in error by the algorithm and combined with the other rail to form the reference plane.

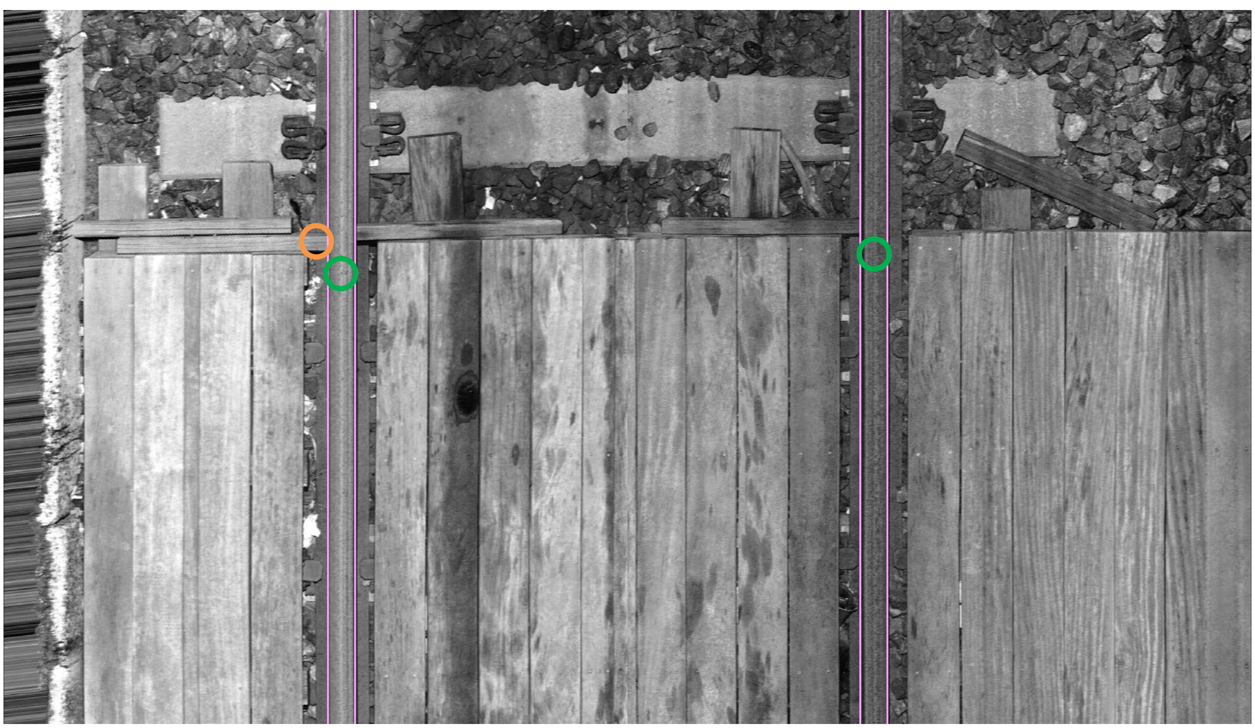

Figure 30 – Road Crossing Leads to False Change

Researchers reviewed additional data and found this error to be intermittent; this can be seen in Figure 31, below, where the false changes do not appear in the left rail ballast graph (top graph) nor in the crib ballast graph (middle graph).

Researchers believe this error can be easily eliminated through algorithm enhancement. As this error is infrequent, this outlier was ignored for the purposes of setting the noise floor. Ultimately, a noise floor of 5 mm was determined as the best fit with the data, resulting in detected changes under 5 mm being considered non-relevant.

Deliberate changes in ballast height to validate algorithm performance in the field were made at the following locations between Day 1 and Day 2:

- MP 24; ballast was removed from four end-of-tie locations on the field side.
- MP 23+5344; ballast was removed from three cribs.
- MP 23+5274; ballast was removed from four end-of-tie locations on the field side.
- MP 23+5192; ballast was removed from three cribs.
- MP 23+5153; ballast was removed from three end-of-tie locations on the field side.
- MP 23+5100; ballast was removed from three cribs.

Consequently, when comparing the average ballast height per meter for any of the scans captured on Day 1 to those captured on Day 2, a number of data points that exceeded the 5 mm noise floor were expected near MP 24 (and MP17 due to the legitimate changes noted prior). Figure 31 presents this comparison (D1R1 vs. D2R1).

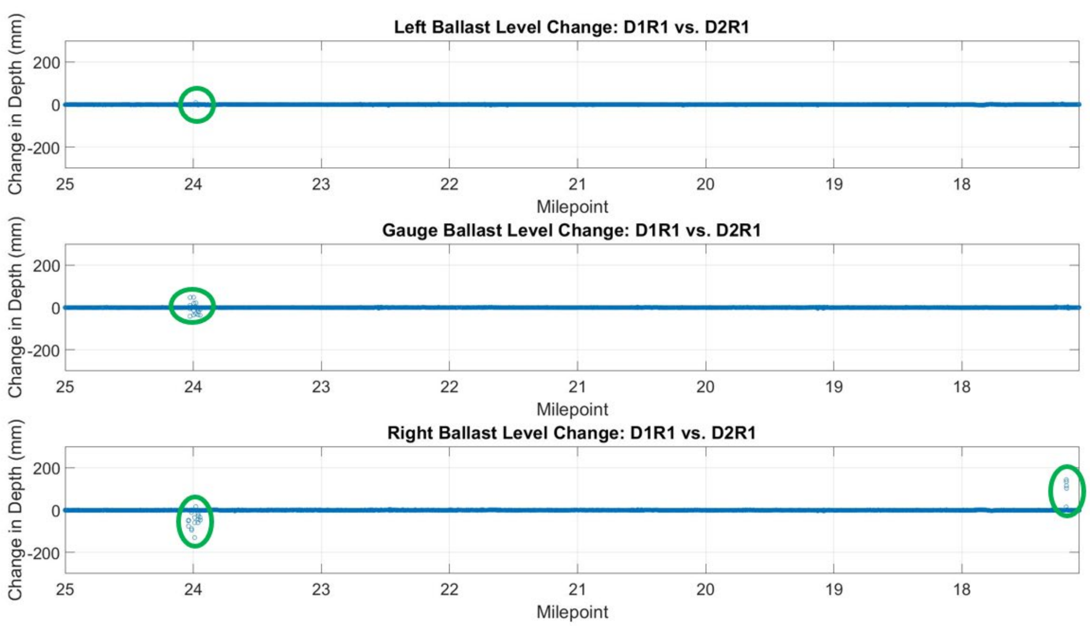

Figure 31 – Detecting Deliberate Ballast in Gauge and End-of-Tie Change (D1R1 and D2R1)

As expected, in Figure 31 there is a clear trend of data points exceeding the 5 mm noise floor at the expected positions of MP 24 and MP 17. As with the crib-to-crib height comparison method, there was an excess of reported change locations due to the impact of tie skew changes made at the same location as well as the imprecise nature of moving ballast.

Organic changes to ballast-height-per-meter-of-travel occurred over the test period and were automatically detected. An example of this can be seen when comparing ballast-in-crib heights measured in September 2019 (D7R1) versus those measured in May 2019 (D2R1) and presented in Figure 25.

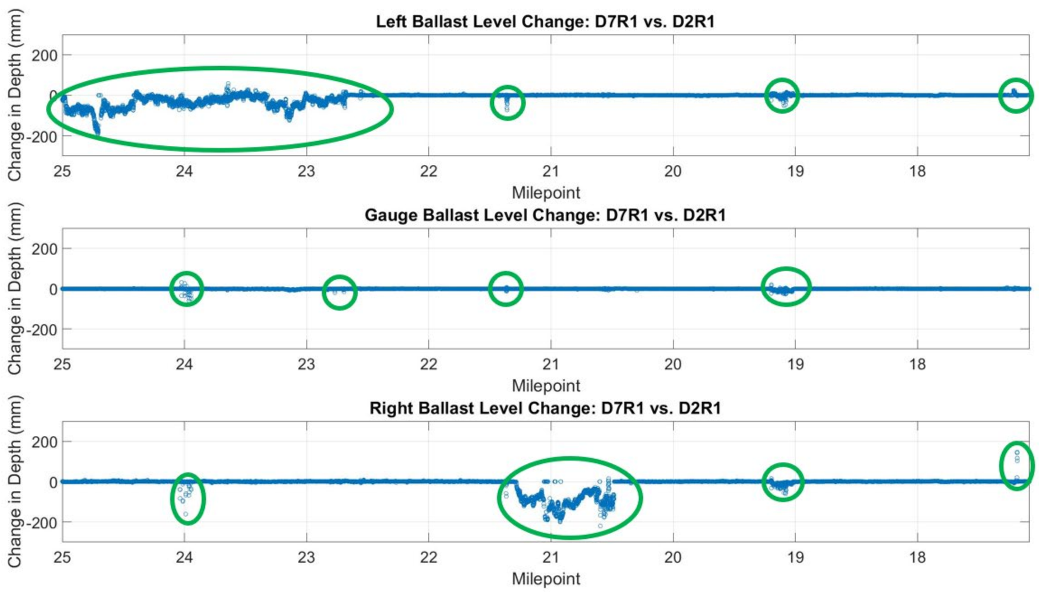

Figure 32 – Height of Changes in Ballast per Meter of Travel between May and September

Figure 32 shows numerous track locations, in particular between MP 25 and MP 22.75, with significant changes in the order of 100–200 mm (3.93–7.87 in) (circled in green). There were also much smaller and more localized changes near MP 19 and MP 17. The predominant change between May and September 2019 was the increase in ballast level at specific locations, but there were some instances of small reductions as well.

4.3 Ballast Fouling Change Detection

This task involved run-to-run comparisons regarding the fouling status of ballast material present in three zones: the crib area, the end-of-tie near the left rail, and the end-of-tie near the right rail. The total area of fouled ballast per 20 m of track length was used as a statistic for change detection. A total of approximately 640 measurements of 20 m track lengths were made in each run.

Overall repeatability for ballast fouling area measurements was excellent with repeat runs returning values with a 98.80 percent agreement. The overall mean for ballast fouling measurements was 0.004 m^2 – meaning that on average, ballast fouling measurements varied by

0.004 m^2 between repeat runs. Standard deviation of reported ballast fouling between runs was 0.074 m^2.

Table 4 – Ballast Fouling Mean and Standard Deviation

Fouled Ballast in 3 Zones: Left Rail, Gauge, Right Rail (m2 per 20m track length) 640 instances							
D2R2 vs D2R1				D7R2 vs D7R1			
	Left	Gauge	Right		Left	Gauge	Right
Repeatability	97.56%	99.66%	99.25%	Repeatability	98.94%	99.52%	99.09%
Mean	-0.003	0.013	0.015	Mean	-0.004	0.004	0.008
Std	0.039	0.125	0.069	Std	0.026	0.13	0.056
D3R2 vs D3R1				D3R2 vs D3R1a			
Repeatability	98.45%	99.52%	98.72%	Repeatability	98.46%	99.11%	98.28%
Mean	-0.005	0.003	0.012	Mean	-0.0016	-0.004	0.003
Std	0.03	0.117	0.073	Std	0.015	0.09	0.036
D4R2 vs D4R1				D3R2 vs D3R1b			
Repeatability	97.90%	99.28%	99.29%	Repeatability	98.44%	99.72%	98.94%
Mean	0.005	0.003	0.01	Mean	-0.006	0.006	0.017
Std	0.036	0.15	0.053	Std	0.037	0.13	0.091
D5R2 vs D5R1				Overall Left, Gauge, Right			
Repeatability	96.91%	99.63%	98.21%	Repeatability	98.09%	99.49%	98.83%
Mean	-0.006	0.009	0.009	Mean	-0.003	0.005	0.011
Std	0.04	0.113	0.083	Std	0.034	0.122	0.066
Overall All Zones							
Repeatability	98.80%						
Mean	0.004						
Std	0.074						

Upon review of the back-to-back runs, the comparison between the second and first runs of Day 2 (D2R2 and D2R1) contained the largest outliers. No deliberate changes were made to track conditions between these runs, and they were captured within approximately 2 hours of one another. Figure 33 presents the measurement results from each run overlaid on top of one another (D2R1 in blue and D2R2 in orange).

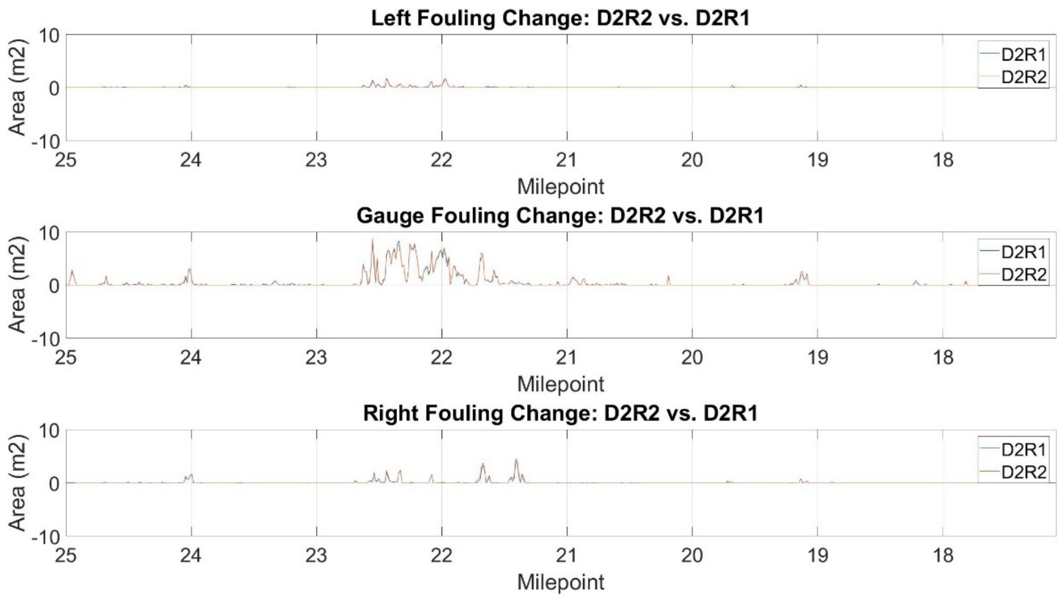

Figure 33 – Modeling Ballast Fouling Noise Floor (D2R2 vs. D2R1)

Figure 33 shows that these measurements are highly repeatable. As well, if the mathematical difference in these measurements is plotted (Figure 34), one can see that the typical variation per measurement is quite low as it is less than 1 m^2 over a 20-meter length of track.

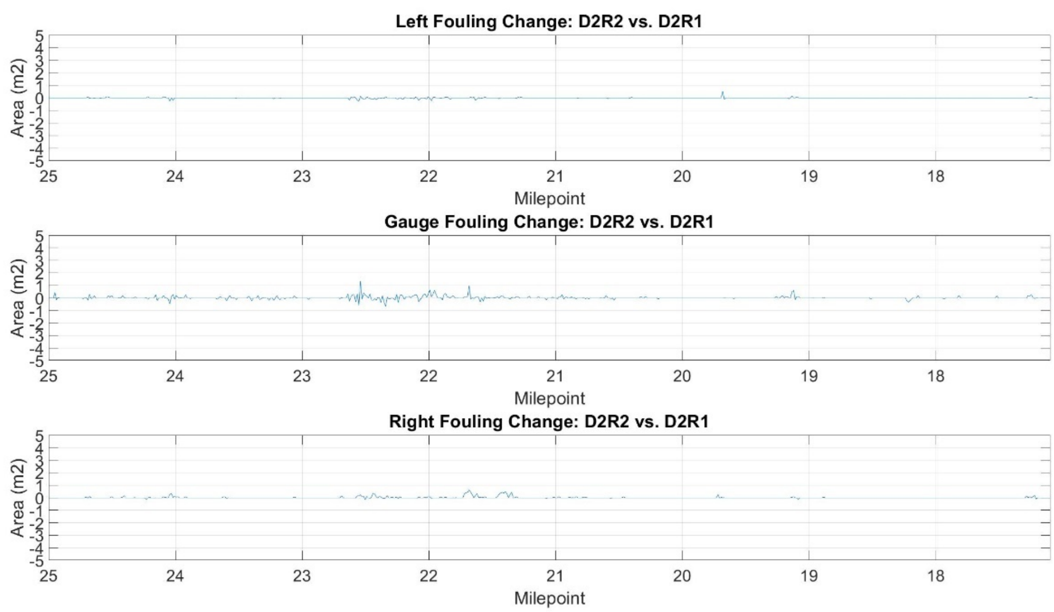

Figure 34 – Difference in Area of Fouling between Runs (D2R2 vs. D2R1)

Figure 33 shows that the calculated difference in fouled area between these runs was typically less than 1 m^2. Thus, it can be inferred that the noise floor for this parameter was approximately

1 m^2. In other words, detected changes in ballast fouling (both positive and negative) that are 1 m^2 and under can be considered as non-relevant.

Deliberate changes were not made to ballast fouling. Organic changes to fouling occurred over the test period and were automatically detected. For example, the reduction in surface fouling between May and June is visible in the range data (Figure 35).

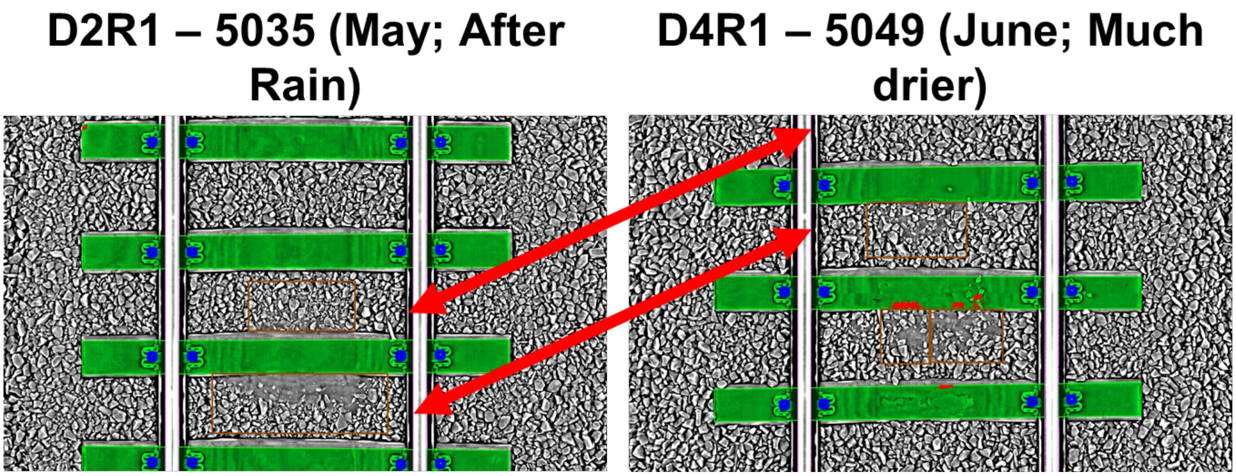

Figure 35 – Reduction in Surface Fouling between May and June

This same trend can be seen when plotting the detected changes in surface fouling between September 2019 (D7R2) and May (D2R1) (Figure 36, circled in green).

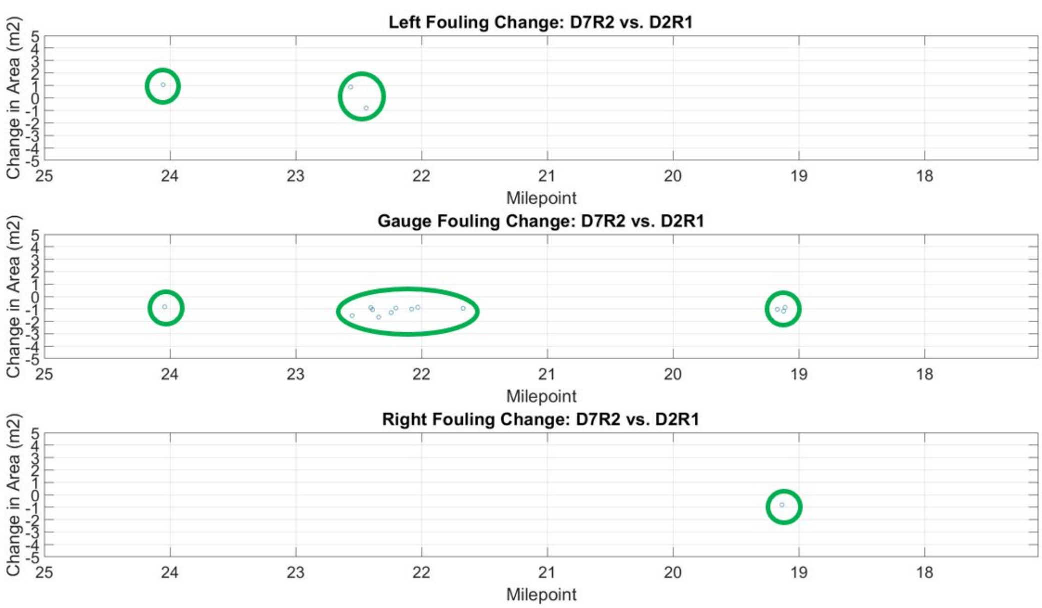

Figure 36 – Organic Changes in Ballast Fouling over the Test Period (D7R2 vs. D2R1)

The majority of organic change observed between May and September was related to a decrease in the level of fouling. This was likely because the test conditions in May were very wet (i.e., standing water) with significant surface fouling which was absorbed over time. There are, however, two instances of fouling increase at MP 24 and MP 22.5.

4.4 Tie Skew Change Detection

This task involved run-to-run comparisons regarding the degree of tie skew (Figure 37).

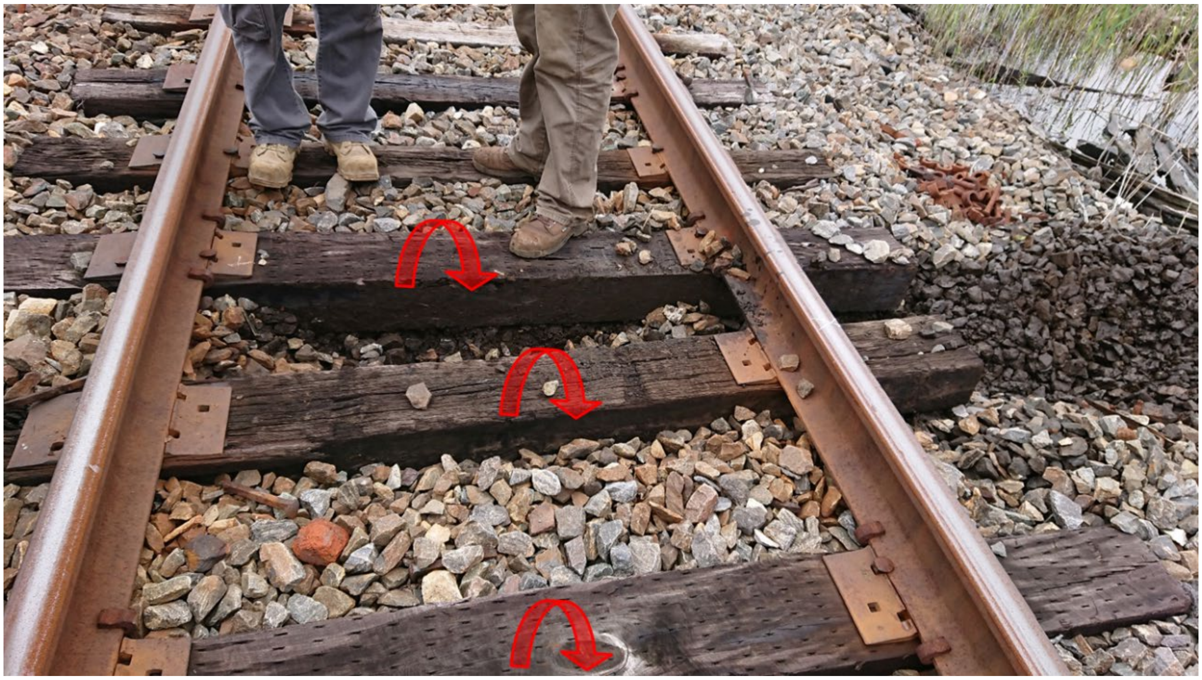

Figure 37 – Tie Skew Change Detection

A total of approximately 22,000 measurements of tie (containing a mixture of concrete and timber ties) skew were made in each run.

4.4.1 Repeatability, Mean, and Standard Deviation

First, note that the accuracy of concrete tie skew angle measurement was superior to those for wooden ties due to the less accurate edge detection of wooden ties owing to their non-uniform surface. Overall repeatability for skew measurements was reasonable, with repeat runs returning values with a 95.72 percent agreement. Overall mean for skew angle measurements was 0.022, meaning that on average, the reported skew angle for a given tie varied by 0.022 degrees between repeat runs. Standard deviation of reported skew angle between runs was 0.285 degrees.

Table 5 – Skew Angle Measurement Repeatability, Mean, and Standard Deviation

Skew Angle (degree) 22,000 Instances			
D2R2 vs D2R1		**D7R2 vs D7R1**	
Repeatability	95.66%	**Repeatability**	96.12%
Mean	-0.005	**Mean**	0.008
Std	0.26	**Std**	0.305
D3R2 vs D3R1		**D3R2 vs D3R1a**	
Repeatability	95.04%	**Repeatability**	96.77%
Mean	0.007	**Mean**	0.01
Std	0.274	**Std**	0.277
D4R2 vs D4R1		**D3R2 vs D3R1b**	
Repeatability	96.92%	**Repeatability**	94.17%
Mean	0.005	**Mean**	0.006
Std	0.275	**Std**	0.273
D5R2 vs D5R1		**Overall**	
Repeatability	95.39%	**Repeatability**	95.72%
Mean	0.125	**Mean**	0.022
Std	0.329	**Std**	0.285

Upon review of the back-to-back runs, the comparison between the second and first runs of Day 2 (D2R2 and D2R1) were determined to contain the largest outliers. No deliberate changes were made to track conditions between these runs and they were captured within approximately 2 hours of one another.

Figure 38 presents the run-to-run comparison with the difference in measurements plotted along the Y-axis and the location for comparison along the X-axis. Each circle in the plot represents a skew angle comparison between runs for an individual tie, with a total of approximately 22,000 ties being analyzed. Note that track locations with bridges or special track work were deliberately omitted from the analysis and are shown as a zero value.

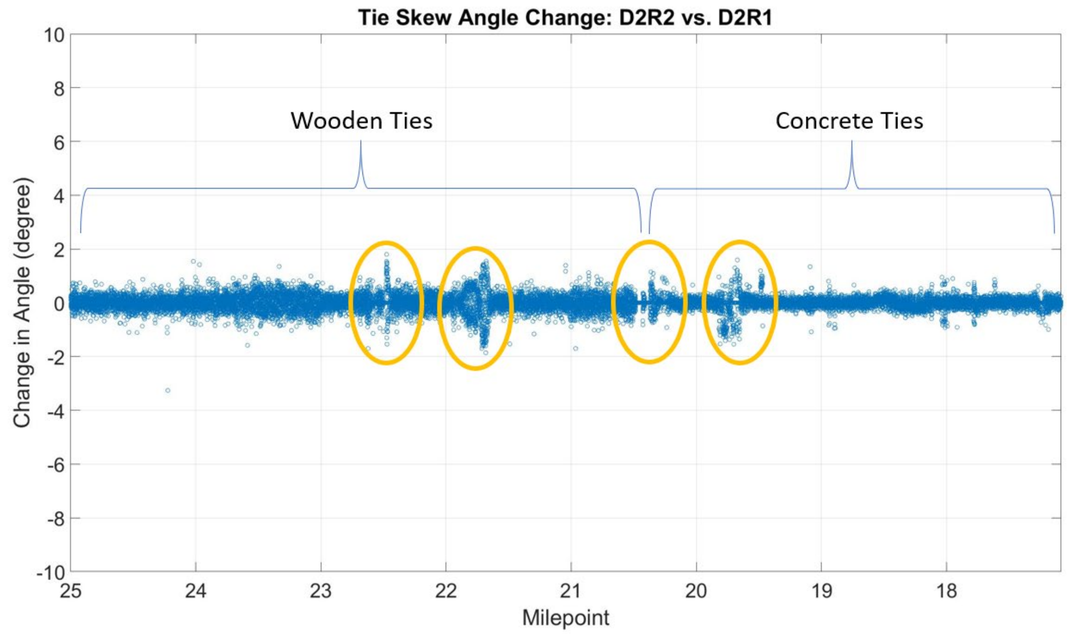

Figure 38 – Modeling Tie Skew Noise Floor (D2R2 vs. D2R1)

Figure 38 shows that the calculated difference in tie skew between these runs (effectively the noise) is typically less than 1 degree. However, also note that the differences between MP 25 and MP 20.5 were noticeably noisier than between MP 20.5 and MP 17. This was because MP 25 to MP 20.5 consists of wooden ties and MP 20.5 to MP 17 consists of concrete ties.

Lastly, and of particular note, there were a number of 2-degree changes at MP 22.5, MP 21.75 and MP 19.75 (yellow circles). Initially, researchers believed that these high-noise measurements were simply due to small differences (a few millimeters) in the edge detection of ties between runs, which, in turn, resulted in differences in computed tie-skew values.

However, this did not fully explain why this noise was clustered at specific locations along the test route, as opposed to being uniform along its length. Upon further investigation, a different hypothesis emerged. If one analyzes skew angle for two runs that were captured in the same direction (but separated by days or weeks, so there could be some organic change), the plot of skew-angle differences would be much flatter than an immediate back-to-back run captured in opposite directions.

Consider Figure 39, which compares skew angle differences between June 26 (Day 4) and June 27 (Day 5). While there was a small increase in signal near MP 22.7, the general plot of differences is significantly flatter than the previous opposite-direction run comparison and does not exhibit spiking at MP 22.5, MP 21.75, or MP 19.75.

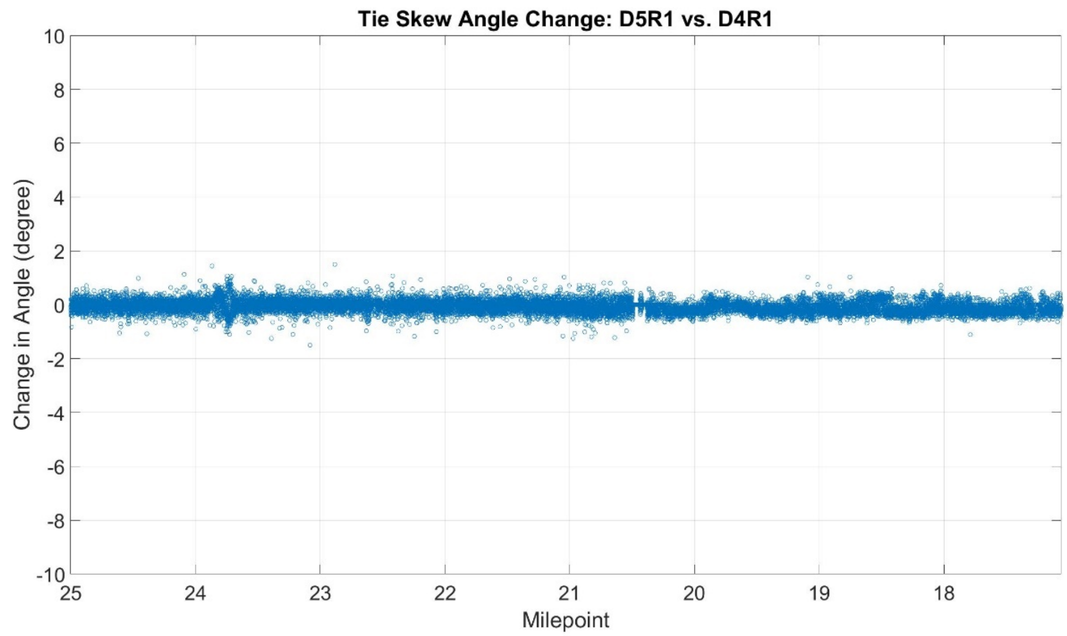

Figure 39 – Same Direction Runs Produce More Consistent Skew Angle Measurements

Initially, there was no logical explanation why skew angle measurements from opposite-direction repeat runs would be noisier than same-direction runs at specific track locations. To better understand potentially unique conditions encountered at these locations, additional LRAIL algorithms were used to calculate geometry parameters for the test loop (Figure 40).

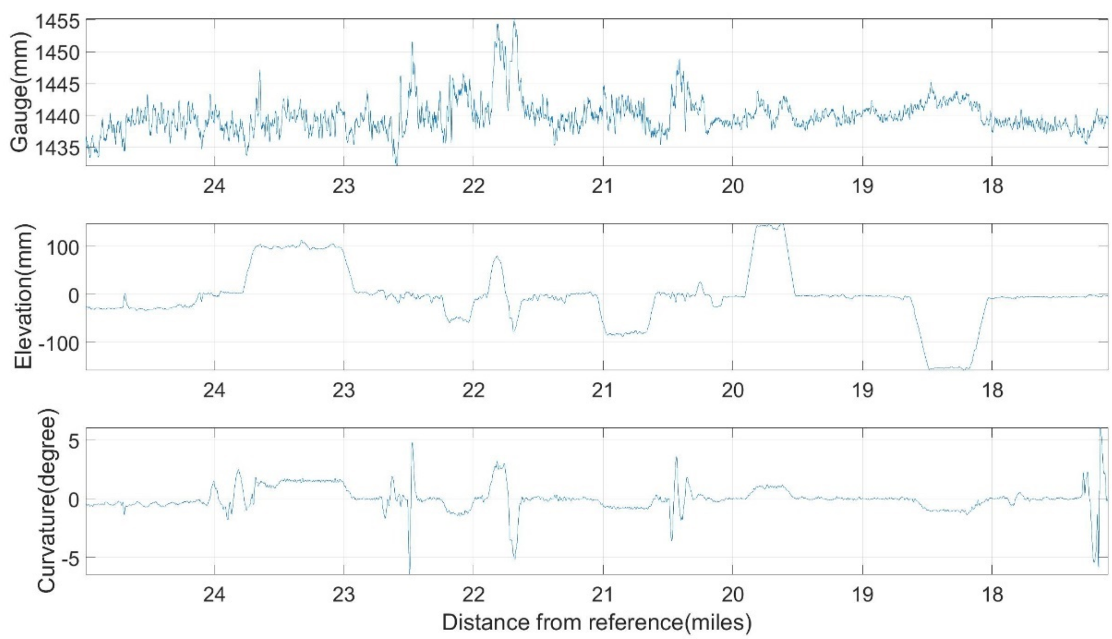

Figure 40 – Track Geometry for Test Section

Track alignment was then overlaid on top of the skew angle plot to look for patterns of note. Interestingly, the track locations with the highest skew angle measurement noise tended to corresponded to the track locations where there was a rapid change in horizontal track alignment and in crosslevel (yellow circles in Figure 41). Adding the locations of special track work to Figure 41 (red squares) helps to further corroborate this hypothesis. Multiple turnouts can be found at MP 22.5 and MP 20.46, two bridges at MP 24.69 and MP 21.92, and multiple road crossings at MP 19.68 and MP 17.23.

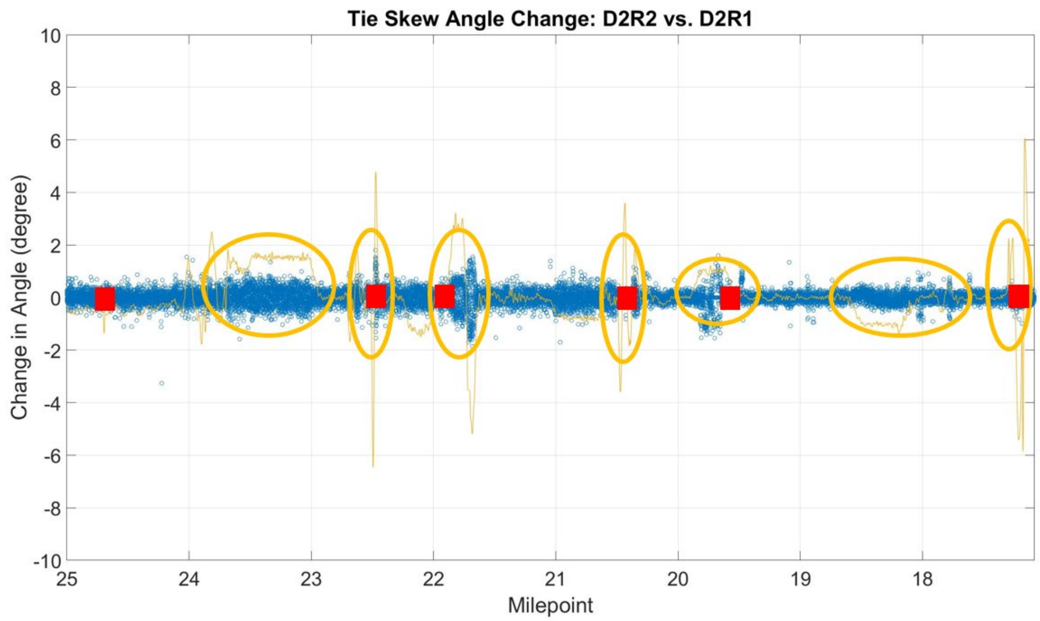

Figure 41 – Track Alignment (yellow line) Overlaid on Plot of Tie Skew Angle Change

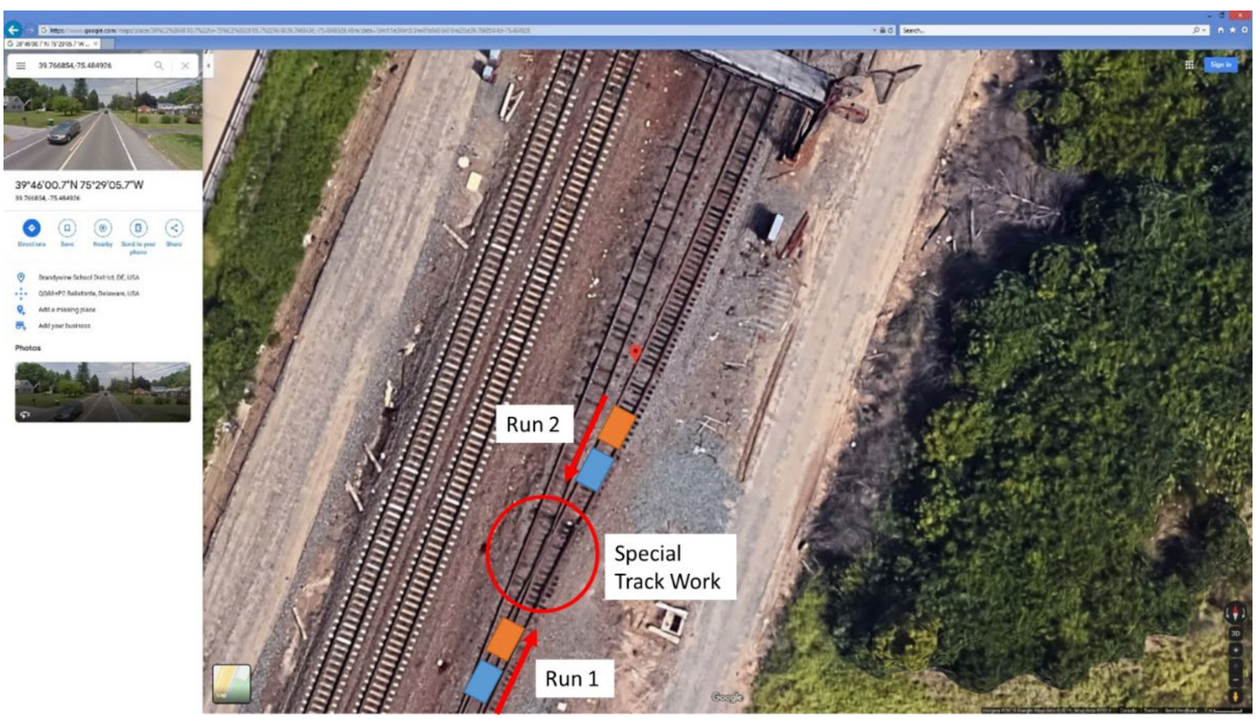

Figure 42 – Google Aerial Image Showing Location with High Skew Angle Noise

Looking at these same locations using aerial images (Figure 42), additional details are visible. Each area of high noise was near special trackwork at switches and crossings. These features explain the rapid alignment and crosslevel changes indicated by the geometry data. Differences in skew-angle measurements were believed to be related to these features by way of the tow vehicle's response to special trackwork as it crossed. Researchers believe that extreme tow-vehicle vibration was translated to the inspection trailer which temporarily pulled the inspection trailer out of its geometrically calibrated position with regard to the track.

Interestingly, the track position where this condition occurred varied based on the direction of travel (i.e., up-chain versus down-chain). Since the tow vehicle entered the special trackwork from a different side in each travel direction, the phenomenon also occurred in a different location (Figure 27; see Run 1 and Run 2 illustrations).

A solution has not yet been developed to address this scenario, but a number of options are being considered, including using in-sensor accelerometer data to automatically flag these locations. As well, note that this issue would likely not be present (or would be minimized) in a scenario where the LRAIL inspection system is installed in the bed of the hi-rail as opposed to in a separate trailer.

As such, for the time being, a noise floor of 2.5 degrees is recommended to account for typical outliers and to avoid incorrectly flagging track locations where this phenomenon occurs.

Deliberate changes were made to tie skew to validate the ability to detect changes for this parameter. Deliberate changes were made at the following locations between Day 1 and Day 2:

- MP 23+5012 to 23+5000; one tie was skewed at the starting MP and approximately seven ties later a second tie was skewed.
- MP 23+4916; one tie was skewed.

Consequently, when comparing the tie skew for any of the scans captured on Day 1 to those captured on Day 2, one should expect to see three data points in close proximity that exceed the 2.5-degree noise floor at approximately MP24. Figure 43 presents this comparison (D1R1 vs. D2R1).

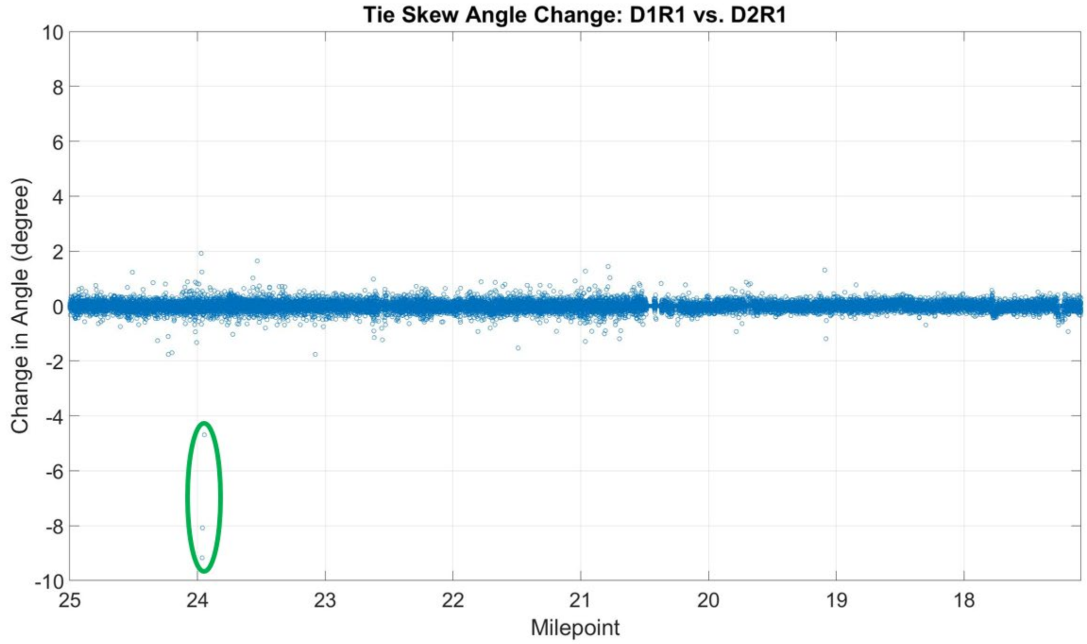

Figure 43 – Detecting Deliberate Tie Skew Change (D1R1 and D2R1)

As expected, in Figure 43 there is a clear cluster (circled in green) of wooden ties exceeding the 2.5-degree noise floor at the expected position of MP 24.

There were no organic changes (changes above 2.5 degrees) to tie skew detected between May 2019 and September 2019. This can be observed by comparing tie skew results captured on Day 7 (Run 1) in September 2019 versus those measured on Day 4 (Run 1) in May 2019 (Figure 44).

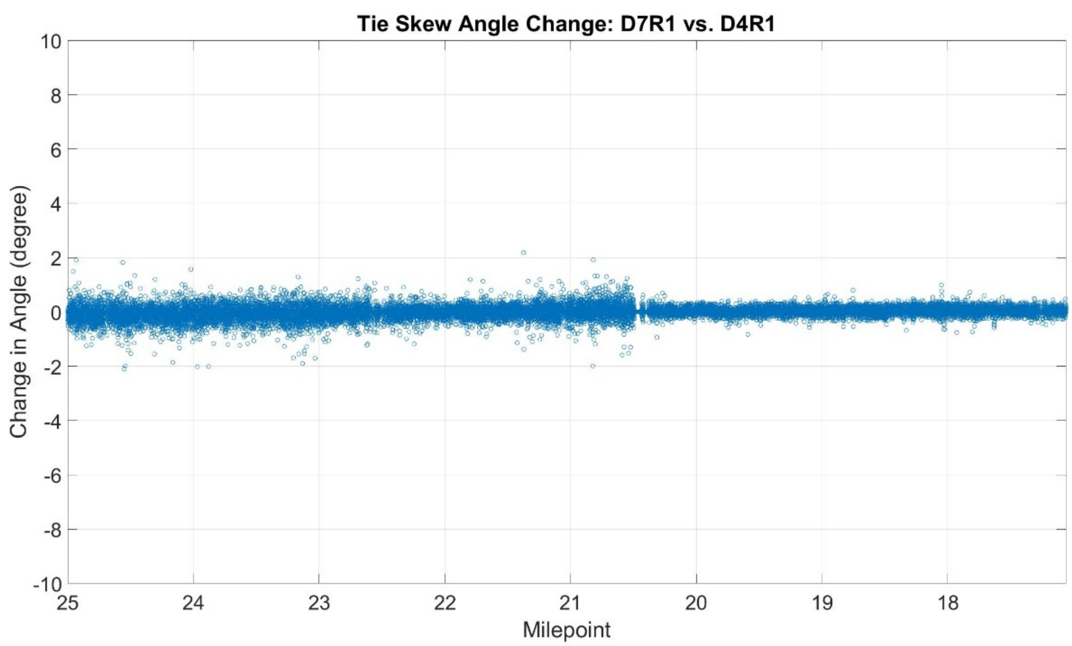

Figure 44 – Organic Changes to Tie Skew (D7R1 vs. D2R1)

4.5 Tie Condition Rating Change Detection

This task involved run-to-run comparisons regarding the condition rating of both wooden and concrete ties between runs. Condition was evaluated in terms of the percentage of "acceptable" ties per mile.

Wooden tie rating involves the detection of cracking, splits, and holes in the surface of the tie with individual defects grouped according to severity and displayed using color-coding. (Figure 45 and Figure 46).

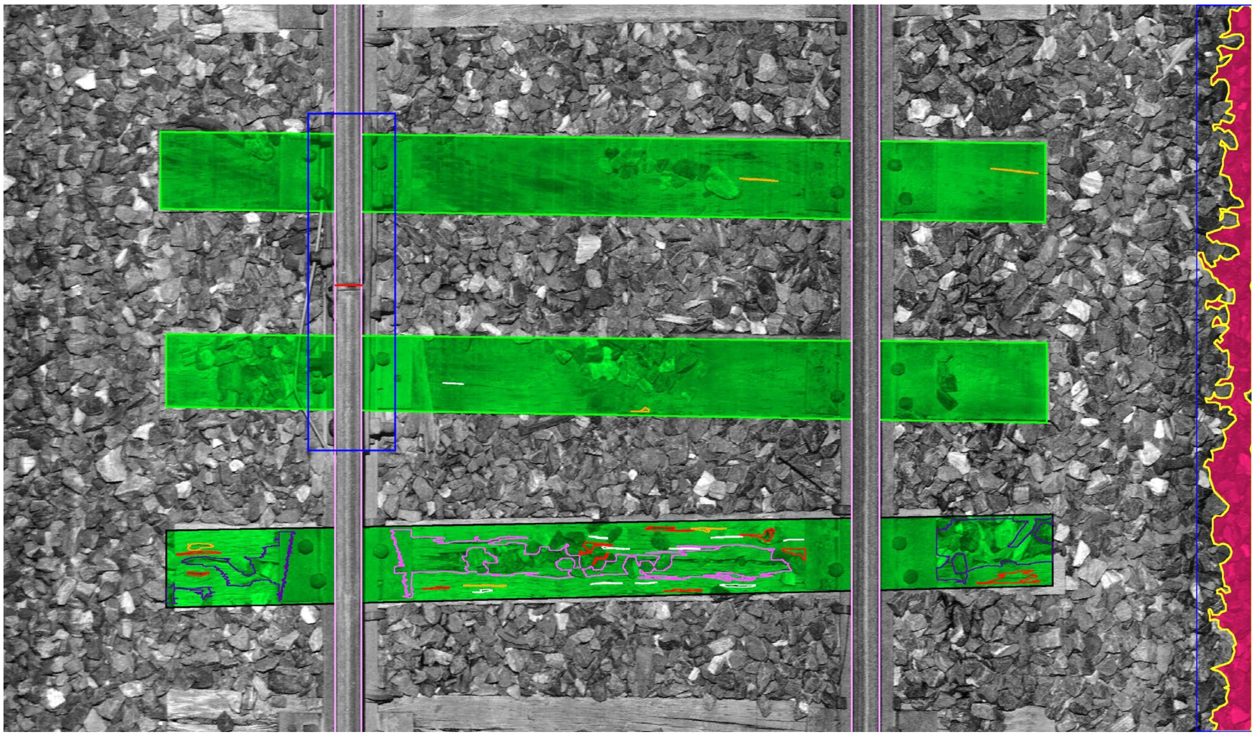

Figure 45 – Wooden Tie Showing Color-coded Defects

Defect Colour	Depth	Width	Length
Very Severe	Defects in ties which contain ballast materials		
Severe	2 cm+	5 cm+	60 cm+
Moderate	1-2 cm	3-5 cm	15-60 cm
Light	Not considered	1-3 cm	10-15 cm
Very Light	Not considered	0.5-1 cm	Not considered
Unmarked	Not considered	Under 0.5 cm	Not considered

Figure 46 – Severity Rating for Individual Defects

Then an overall condition was determined for the individual tie based on the aggregate defects (Figure 47).

Tie Box Colour	Overall Tie Rating
D: Failed Tie	More than 3.7% of tie surface area contains ***Very Severe, Severe, Moderate or Light*** defects.
C: Near Failure Tie	Between 3.7% and 3.1% of tie surface area contains ***Very Severe, Severe, Moderate or Light*** defects.
B: Fair Condition Tie	Between 3.1% and 2.6% of tie surface area contains ***Very Severe, Severe, Moderate or Light*** defects.
A: Good Condition Tie	Any tie which does not fall into the 3 above categories.

Figure 47 – Overall Tie Condition Score Based on Individual Defects

For concrete ties, chips and cracks were first detected (Figure 48) and then an overall condition was determined based on a combination of defect size as well as location on the tie:

- A: Acceptable (Green): Small defects that do not fall in the fastener area
- B: Poor (Red): Larger defects or an accumulation of multiple small defects (however, not in the fastener area)
- C: Failed (Black): Defects in the fastener area or near the rail seat, ties broken in half, defects which cover a large surface area

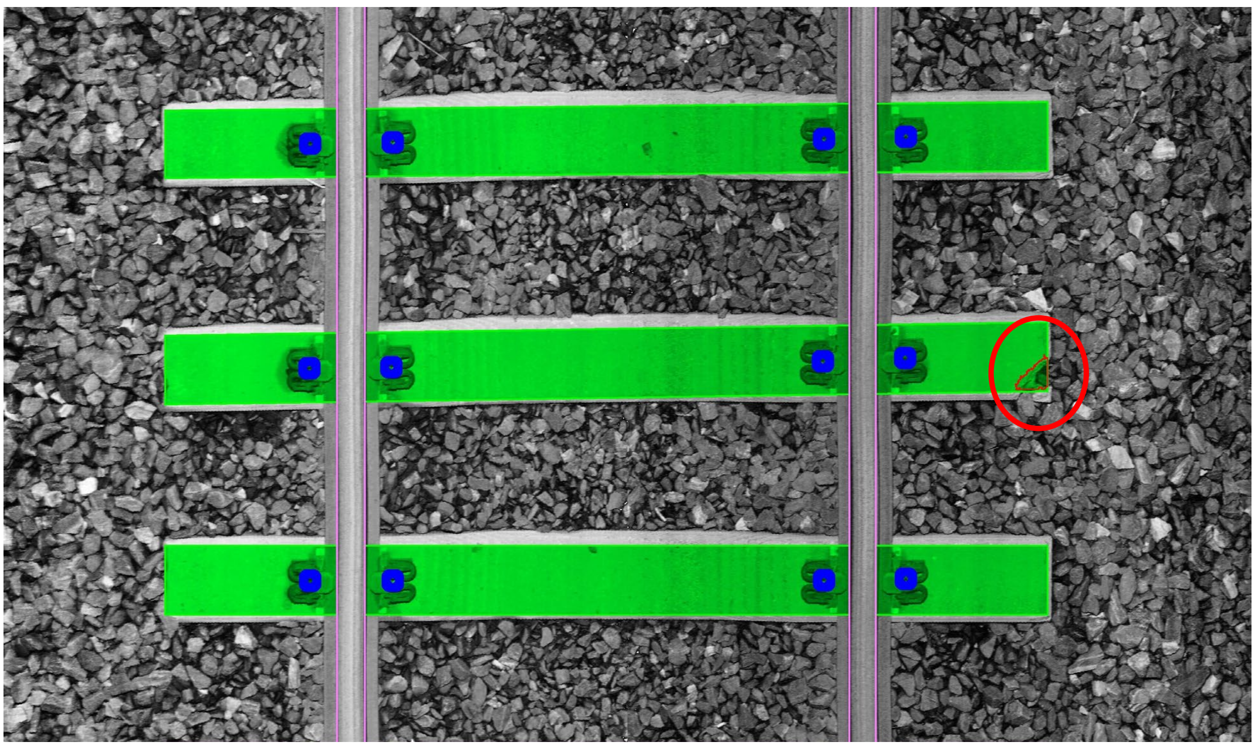

Figure 48 – Intensity Image Showing Automatically Detected Chipping on Concrete Tie

Change detection was focused on detecting changes to the percentage of ties falling into the “acceptable” tie category which included ties rated as “green” and “yellow” and excluded “red” and “black” ties. A total of 22,000 ties were inspected with the percentage of “acceptable” ties per mile being reported.

Overall repeatability for tie condition rating was excellent, with repeat runs achieving 100 percent repeatability. The overall mean for tie condition rating was 4.511 units, meaning that on average the reported number of failed ties per mile varied by 4.511 between repeat runs. Standard deviation of reported failed ties per mile between runs was 7.640.

Table 6 – Acceptable Tie Per Mile Measurement Repeatability, Mean, and Standard Deviation

Acceptable Ties (Units/Mi) 22,000 Ties Across 9Mi			
D2R2 vs D2R1		**D7R2 vs D7R1**	
Repeatability	100.00%	**Repeatability**	100%
Mean	-1.5	**Mean**	6.917
Std	6.515	**Std**	10.629
D3R2 vs D3R1		**D3R2 vs D3R1a**	
Repeatability	100.00%	**Repeatability**	100.00%
Mean	5.333	**Mean**	15
Std	5.437	**Std**	10.53
D4R2 vs D4R1		**D3R2 vs D3R1b**	
Repeatability	100.00%	**Repeatability**	100.00%
Mean	7.33	**Mean**	0.5
Std	9.14	**Std**	2.89
D5R2 vs D5R1		**Overall**	
Repeatability	100.00%	**Repeatability**	100.00%
Mean	-2	**Mean**	4.511
Std	8.34	**Std**	7.64

Upon review of the back-to-back runs, the comparison between the second and first runs of Day 2 (D2R2 and D2R1) were determined to contain the largest outliers. No deliberate changes were made to track conditions between these runs, and they were captured within approximately 2 hours of one another. The percentage of acceptable ties per mile between runs were compared with the difference plotted along the Y-axis and the location for each comparison plotted along the X-axis (Figure 49). The percentage acceptable values for D2R1 are indicated using a "o" symbol and an "x" symbol for D2R2.

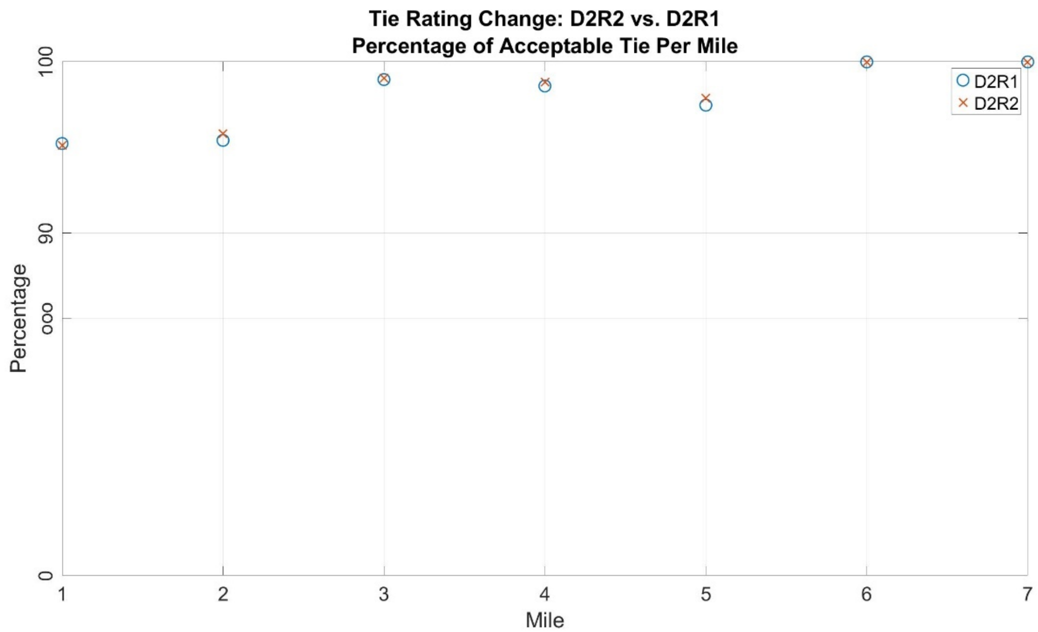

Figure 49 – Modeling Tie Condition Rating Noise Floor (D2R2 vs. D2R1)

Figure 49 shows that the difference in reported percentage of acceptable ties per mile between repeat runs was typically less than 0.25 percent. Thus, it can be inferred that the noise floor for this parameter was approximately 0.25 percent. In other words, detected changes in percentage acceptable tiles per mile that were 0.25 percent and under can be considered as non-relevant.

When applying a noise floor of 0.25 percent, the percentage of failed ties per mile appeared to organically increase over the test period. This can be observed when comparing the percentage of bad ties per mile in September 2019 (D7R2) to May 2019 (D2R1) in Figure 50. The percentage bad tie values for D2R1 are indicated using a "o" symbol and an "x" symbol for D7R2.

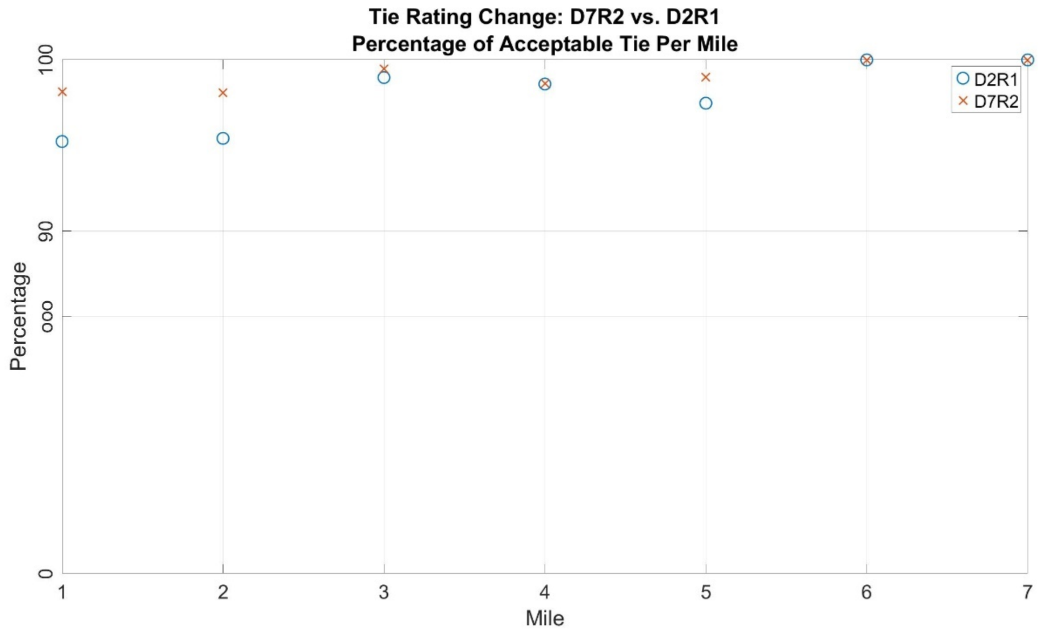

Figure 50 – Organic Changes to Tie Condition Between May 2019 and September 2019 (D7R2 vs. D2R1)

Figure 50 shows that the percentage of acceptable ties increased by approximately 2 percent for the first kilometer and second kilometer and 1 percent at the fifth mile of the test loop. The other locations effectively reported the same condition between runs.

4.6 Joint Bar (and Welding Strip) Bolting

This task involved run-to-run comparisons regarding the number of bolts installed on joint bars and welding strips. A total of 108 joints were inspected with bolt count per joint being reported.

Overall performance of joint bar change detection was outstanding, with 100 percent repeatability, a mean of zero, and a standard deviation of zero.

Table 7 – Joint Bar Bolt Count Repeatability, Mean, and Standard Deviation

Joint Bar Bolting (Count) 108 instances			
D2R2 vs D2R1		D7R2 vs D7R1	
Repeatability	100.00%	Repeatability	100.00%
Mean	0	Mean	0
Std	0	Std	0
D3R2 vs D3R1		D3R2 vs D3R1a	
Repeatability	100.00%	Repeatability	100.00%
Mean	0	Mean	0
Std	0	Std	0
D4R2 vs D4R1		D3R2 vs D3R1b	
Repeatability	100.00%	Repeatability	100.00%
Mean	0	Mean	0
Std	0	Std	0
D5R2 vs D5R1		Overall	
Repeatability	100.00%	Repeatability	100.00%
Mean	0	Mean	0
Std	0	Std	0

To model the noise floor of this parameter, joint bar bolt counts from the second and first runs of Day 2 (D2R1 and D2R2) were compared, with the difference plotted along the Y-axis and the location of each comparison plotted along the X-axis in Figure 51. These runs were selected for comparison as there were no deliberate changes made to track conditions between these runs, and they were captured within approximately 2 hours of one another. Additionally, note that each run was collected in the opposite driving direction with respect to the other (e.g., one was collected in increasing chainage and the other in decreasing chainage).

Each circle in the plot represents a single joint with a related bolt count comparison between runs with a total of approximately 108 joints being analyzed. The two graphs present changes related to bolting for the left and right rails, respectively.

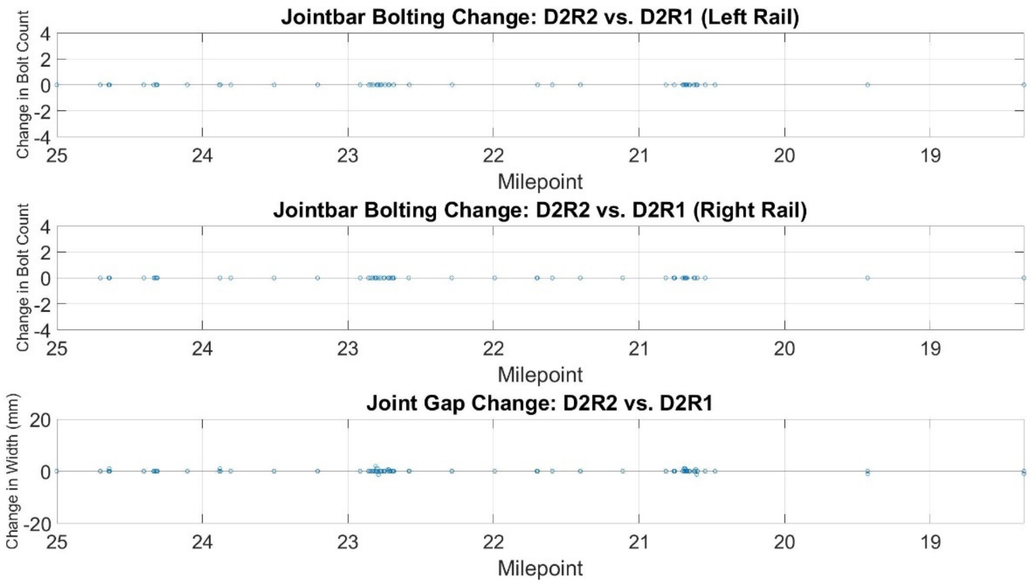

Figure 51 – Modeling Joint Bar Bolting Noise Floor (D2R2 vs. D2R1)

Figure 51 shows that the noise floor for joint bar counts was nil. In other words, there were no instances of a change in count being reported between unchanged runs. Thus, any report of joint bar bolting change should be treated as legitimate.

Deliberate changes were made to welding strip and joint bar bolt counts to validate the ability to detect changes for this parameter (Figure 52).

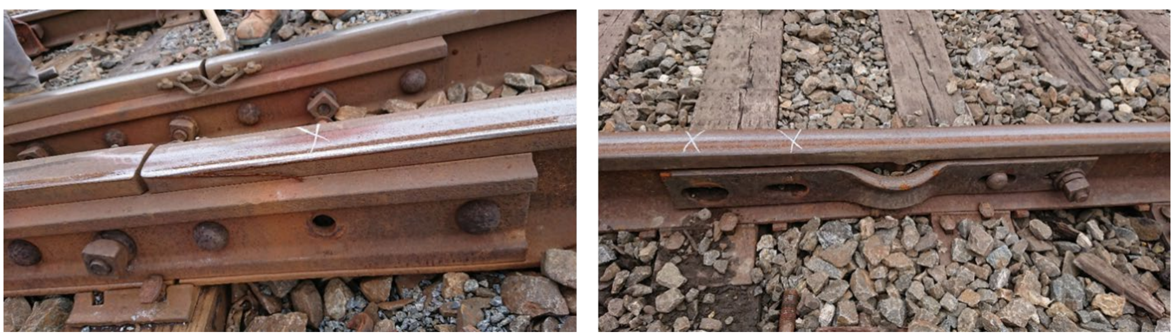

Figure 52 – Deliberate Changes to Joint Bar Bolting

Deliberate changes to welding strip and joint bar bolt counts were made at the following locations between Day 1 and Day 2:

- MP 22+3102; two bolts on the left rail's welding strip were removed.
- MP 22+3076; two bolts on the right rail's welding strip were removed.
- MP 22+2998; one bolt on the right rail's joint bar was removed.

Consequently, when comparing joint bar and welding strip bolt count between scans captured on Day 1 to those captured on Day 2, there should have been five data points that exceeded the nil noise floor at approximately MP 23. However, note that the bolts were removed just prior to D1R1 and replaced just before D2R1, so when comparing these two runs, the result was graphed as a positive change, with five bolts being added as opposed to five bolts being removed. Figure 53 presents this comparison (D1R1 vs. D2R1).

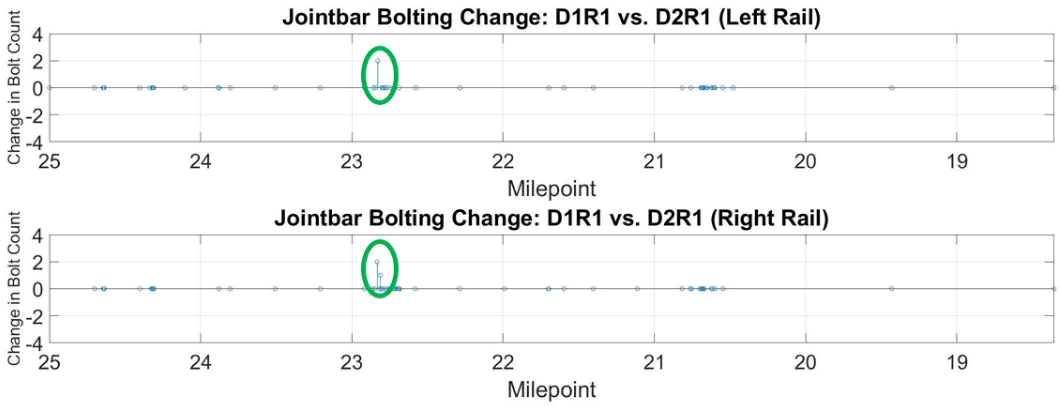

Figure 53 – Detecting Deliberate Joint Bar Bolting Change (D1R1 and D2R1)

As expected, in Figure 53, there is a single instance of a joint bar on the left rail which had two bolts added to it (corresponding to the deliberate change at MP 22+3102) and two instances of change on the right rail, with one instance involving the addition of single bolt (corresponding to the deliberate change at MP 22+2998), and the other shows the addition of two bolts (corresponding to the deliberate change at MP 22+3076). The corresponding 3-dimensional scan images for the change at MP 22+3076 are presented in Figure 54 and Figure 55.

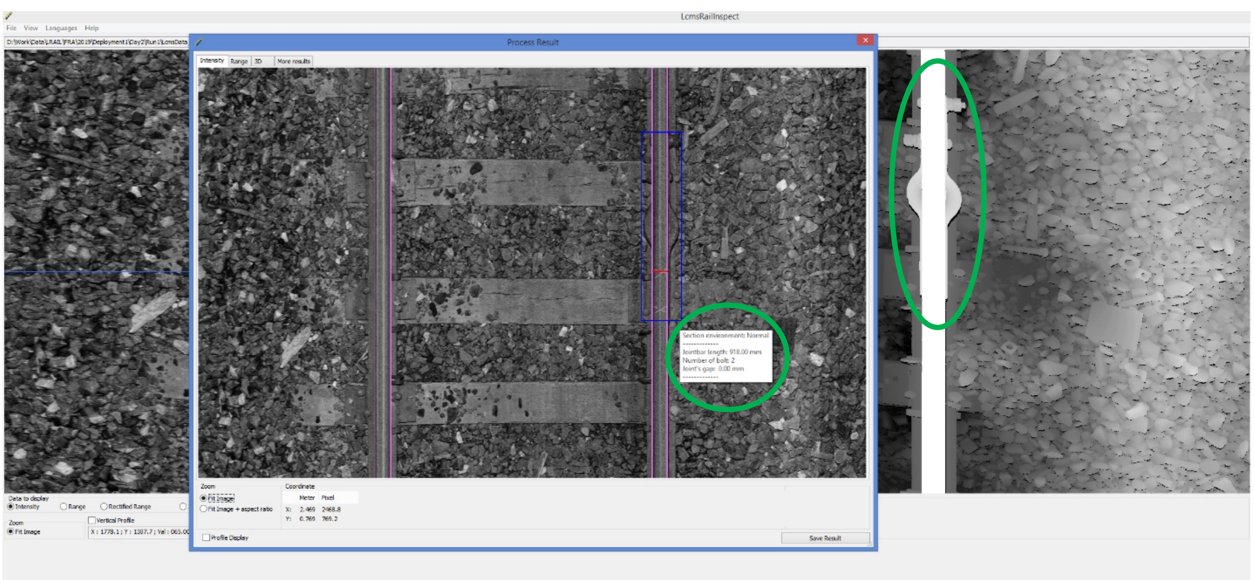

Figure 54 – Right Rail at MP 22+3076 with Two Bolts (D1R1)

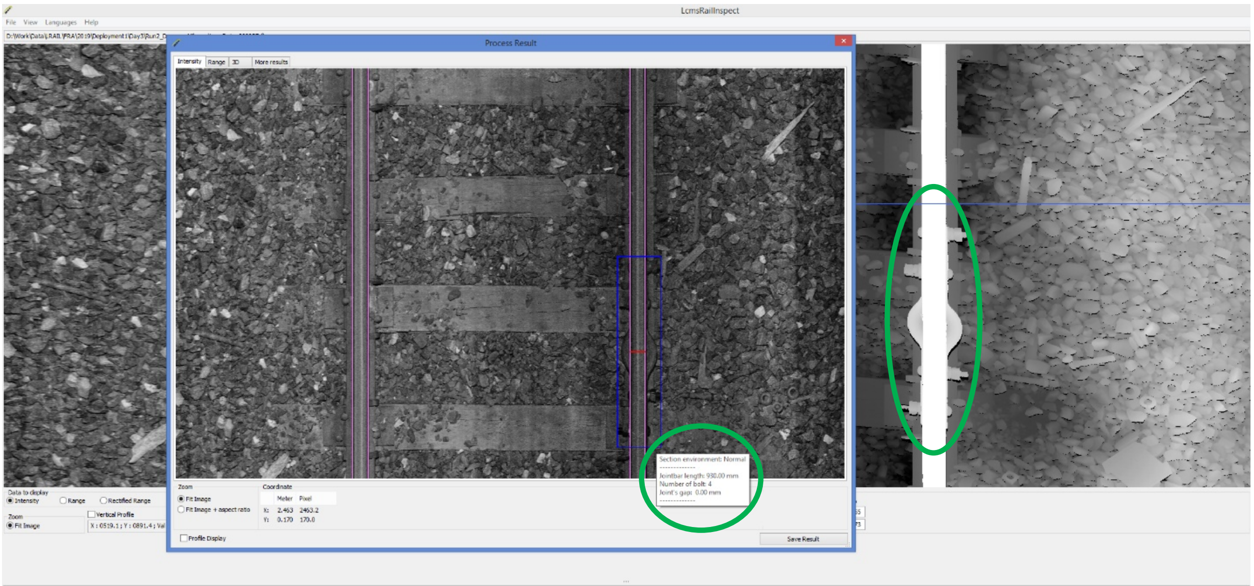

Figure 55 – Right Rail at MP 22+3076 with Four Bolts (D2R1)

There were three organic changes to joint bar bolt counts over the test period. The first two changes occurred between May (D2R1) and June (D4R1). Figure 56 presents these changes (circled in green) along with the known deliberate changes (circled in orange).

Both organic changes can be found at MP 24.6. In May, each joint bar was equipped with a total of six bolts (Figure 57 and Figure 57), and in June each bar had 8 bolts (Figure 57 and Figure 59). The bolting in May also had all of the bolt heads on one side of the rail while the bolting in June alternated bolt heads on either side of the rail.

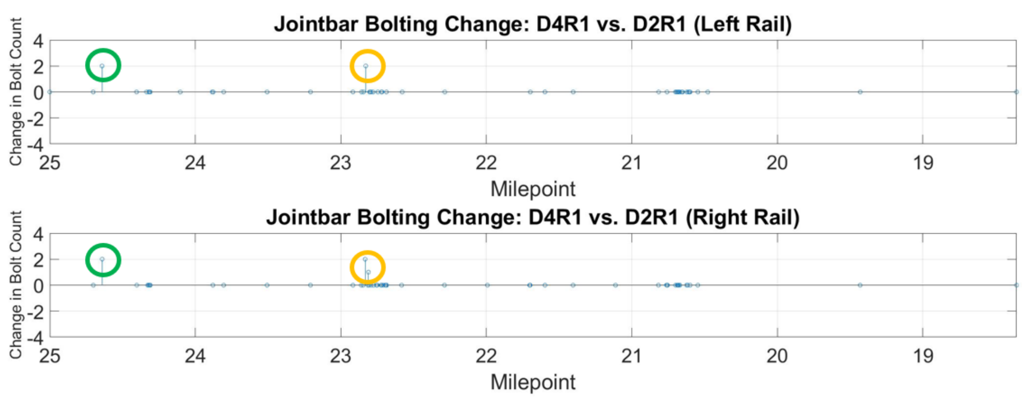

Figure 56 – Joint Bar Bolting Changes between May and June (D4R1 vs. D2R1)

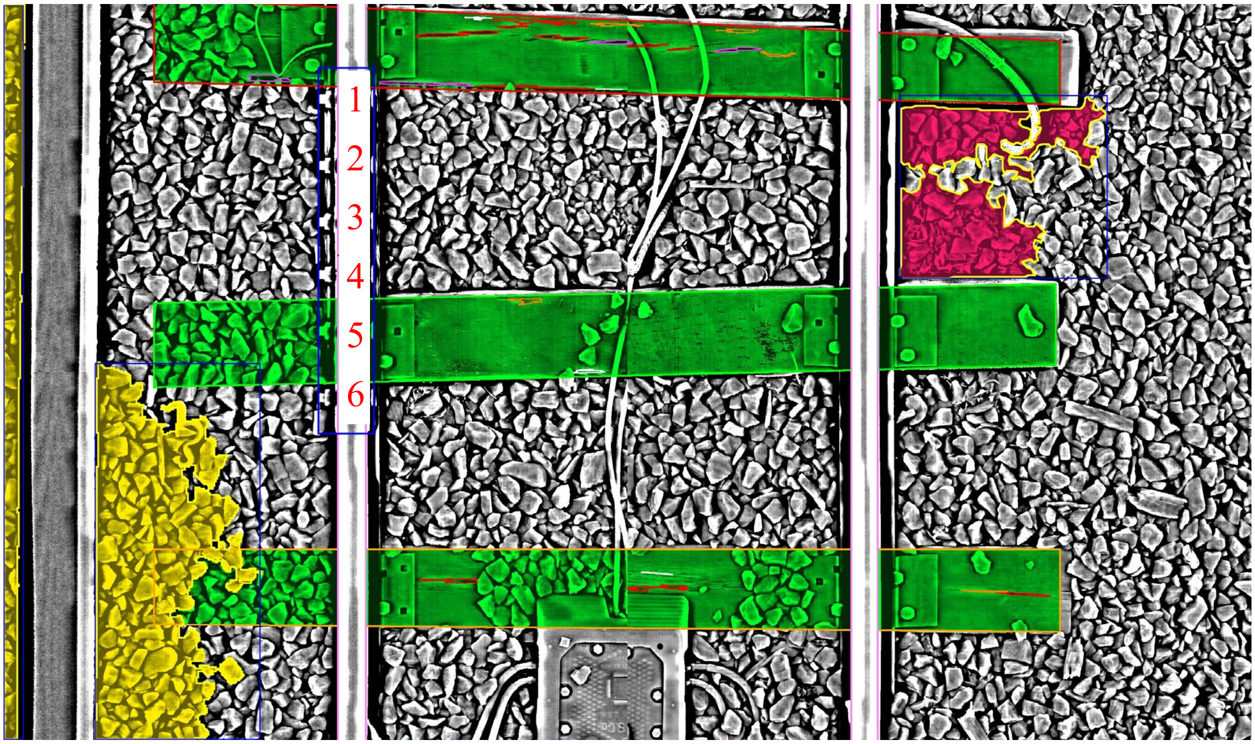

Figure 57 – Left Rail Joint Bar in May Showing Six Bolts with Heads on Gauge Side

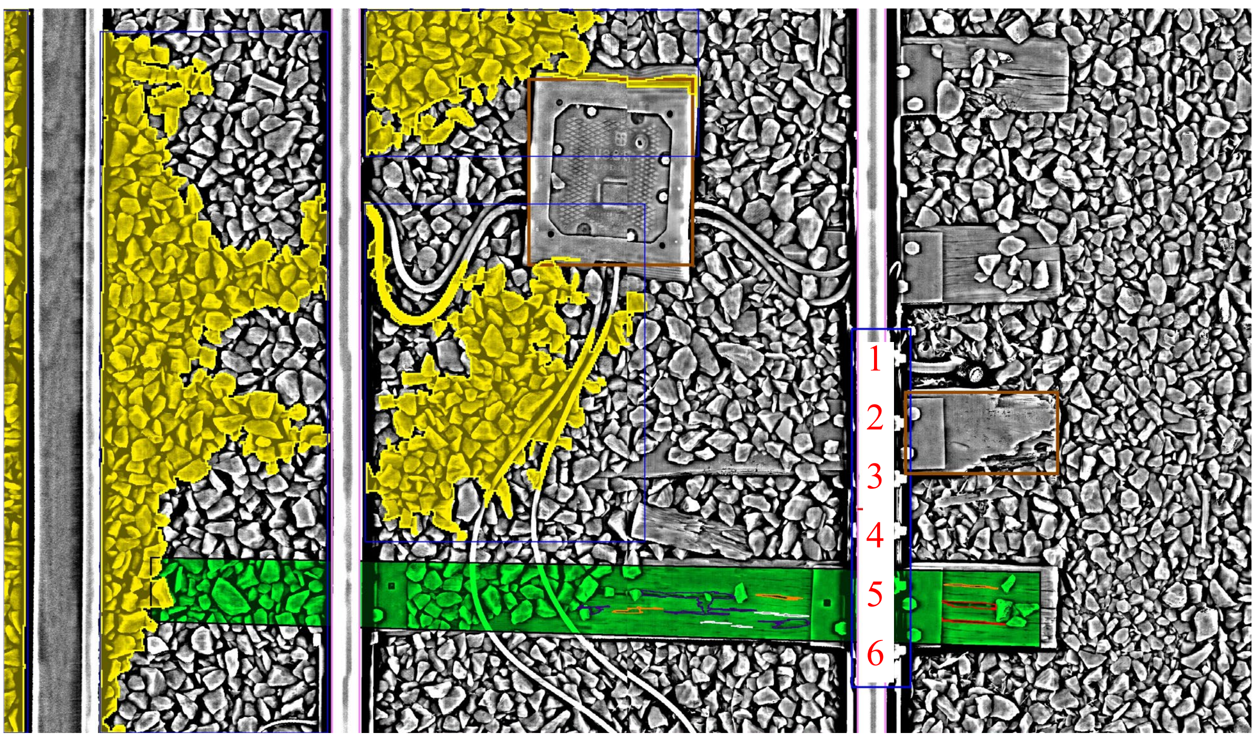

Figure 58 – Right Rail Joint Bar in May Showing Six Bolts with 100 percent of Heads on the Gauge Side of the Rail

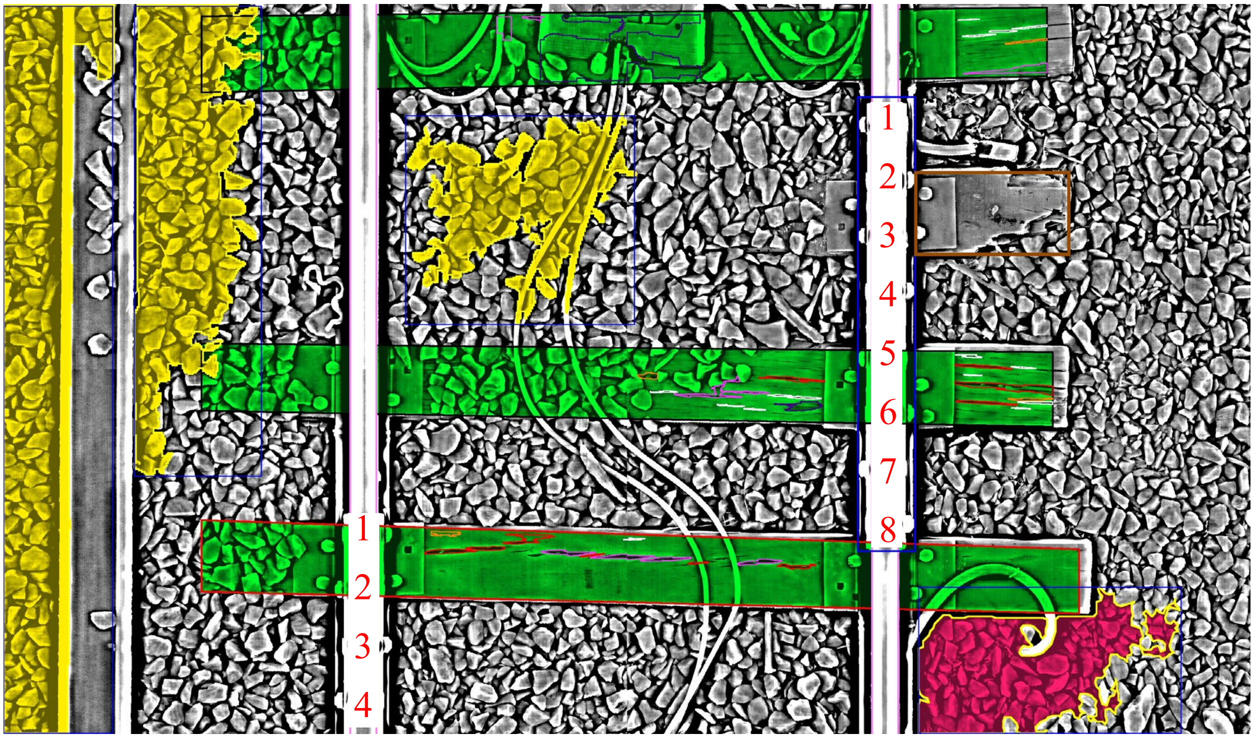

Figure 59 – Part 1 of Left and Right Rail Joints in June Showing Eight Bolts (four on either side of the joint) with Staggered Heads

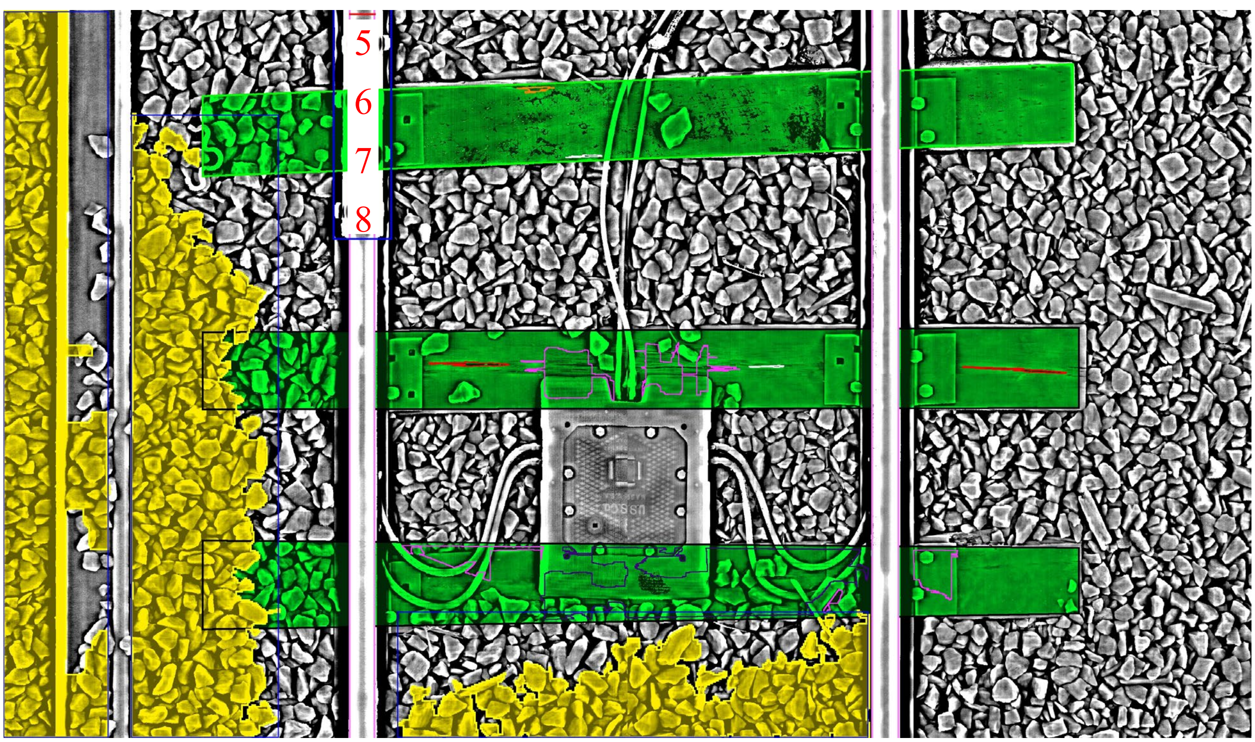

Figure 60 – Part 2 of Left and Right Rail Joints in June Showing Eight Bolts (four on either side of the joint) with Heads Alternating between Gauge and Field Sides

In addition, a third change occurred between June and the final inspection in September. The bolt count for the left rail at the same location changed from eight to four. Upon review of the 3-dimensional scans one can see that the change is caused by ballast material which is now covering bolt ends (it is 25 mm above the bolt heads on average) on the left side of the rail and preventing their detection (Figure 61 and Figure 62). Thus, a change is detected at this location.

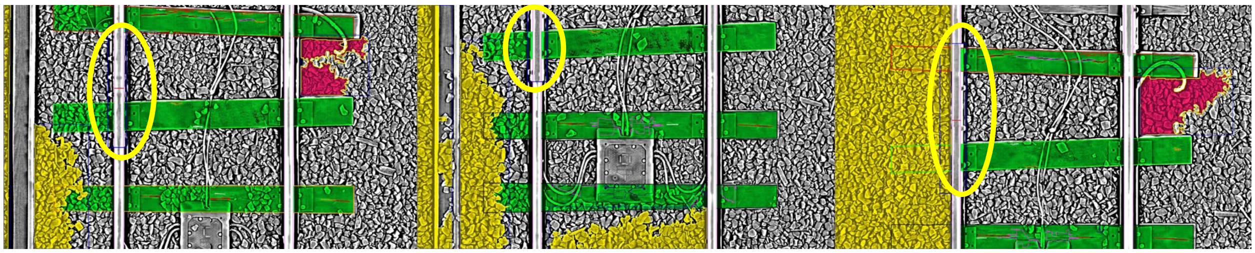

Figure 61 – Ends of Joint Bar Bolts Buried by Ballast (from left to right: May, June, and September)

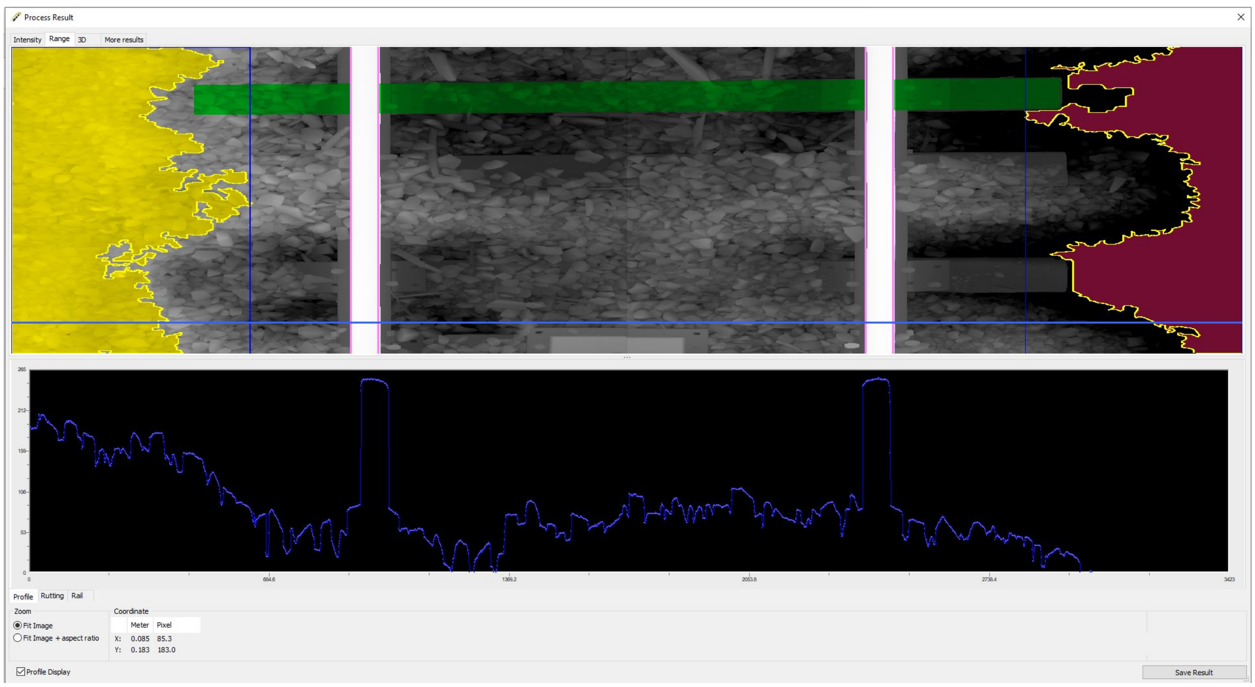

Figure 62 – 3-Dimensional Range Data Showing Height Profile of Ballast

4.7 Joint Gap Change Detection Results

This task involved run-to-run comparisons of the gap measurement at each joint along the test route. A total of 108 joints were inspected with joint gap magnitude being reported.

Overall performance of joint gap measurement was very good, with a run-to-run repeatability of 97.98 percent, a mean of 0.008, and a standard deviation of 0.722.

Table 8 – Joint Gap Repeatability, Mean, and Standard Deviation

Joint Gap Measurement (mm) 108 instances			
D2R2 vs D2R1		D7R2 vs D7R1	
Repeatability	99.75%	Repeatability	96.98%
Mean	0.052	Mean	0.014
Std	0.389	Std	0.88
D3R2 vs D3R1		D3R2 vs D3R1a	
Repeatability	97.86%	Repeatability	100.00%
Mean	-0.035	Mean	0
Std	0.97	Std	0
D4R2 vs D4R1		D3R2 vs D3R1b	
Repeatability	94.69%	Repeatability	97.01%
Mean	0.047	Mean	-0.049
Std	1.134	Std	1.359
D5R2 vs D5R1		Overall	
Repeatability	99.54%	Repeatability	97.98%
Mean	0.029	Mean	0.008
Std	0.325	Std	0.722

To model the noise floor of this parameter, joint gaps from the second and first runs of Day 2 (D2R2 and D1R2) were compared, with the difference plotted along the Y-axis and the location of each comparison plotted along the X-axis (Figure 63). Each circle in the plot represents a report of the change in joint gap between runs, for a total of 108 joints.

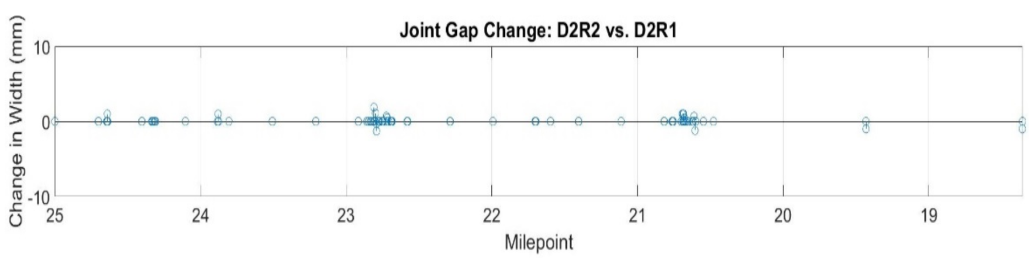

Figure 63 – Modeling Joint Gap Noise Floor (D2R2 vs. D2R1)

Figure 63 shows that the difference between measured joint gaps between runs was typically less than 2 mm. Thus, a 2-mm noise floor was sensible for this parameter.

Deliberate changes were not made to joint gaps.

Joint gaps experienced organic change over the test period in the range of 2 to 16 mm. For example, as shown in Figure 64, a joint with a measured gap of 11 mm on May 15 (at 11:36 PM–

11:51 PM) decreased to just 1 mm (Figure 65) when measured on June 27 (at 11:09 PM–11:55 PM).

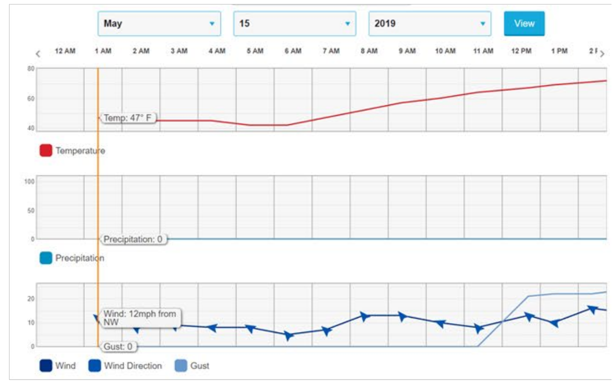

47ºF

https://www.wunderground.com/history/daily/us/de/new-castle/KILG/date/2019-5-15

Figure 64 – Organic Joint Gap Change; Joint 1 at 11 mm in May

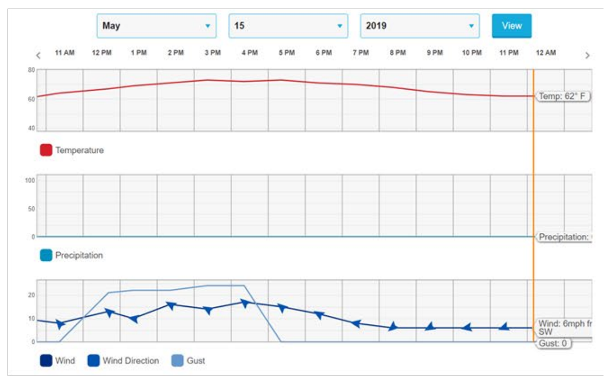

62ºF

https://www.wunderground.com/history/daily/us/de/new-castle/KILG/date/2019-5-15

Figure 65 – Organic Joint Gap Change; Joint 1 at 1 mm in June

While actual rail temperature information was not available for these tests, ambient air temperature information was provided for the day proceeding the gap measurement as an indicator of likely rail temperatures.

4.8 Fastener Inventory Change Detection Results

This task involved run-to-run comparisons regarding the number of acceptable (present and properly seated/installed) fasteners versus unacceptable fasteners (missing or loose) on 22,000 ties (containing roughly 88,000 fasteners) along the test route.

Repeatability, mean, and standard deviation for fastener counts was exceptional, with 100 percent repeatability, a mean of zero, and a standard deviation of zero.

Table 9 – Fastener Inspection Repeatability, Mean, and Standard Deviation

Fastener Inspection (Count) 22,000 instances			
D2R2 vs D2R1		D7R2 vs D7R1	
Repeatability	100.00%	Repeatability	100.00%
Mean	0	Mean	0
Std	0	Std	0
D3R2 vs D3R1		D3R2 vs D3R1a	
Repeatability	100.00%	Repeatability	100.00%
Mean	0	Mean	0
Std	0	Std	0
D4R2 vs D4R1		D3R2 vs D3R1b	
Repeatability	100.00%	Repeatability	100.00%
Mean	0	Mean	0
Std	0	Std	0
D5R2 vs D5R1		Overall	
Repeatability	100.00%	Repeatability	100.00%
Mean	0	Mean	0
Std	0	Std	0

To model the noise floor of this parameter, fastener counts from the second and first runs of Day 2 (D2R2 and D1R2) were compared, with the difference plotted along the Y-axis and the location of each comparison along the X-axis in Figure 66.

Each circle in the plot represents a report of the change in fastener status (number of acceptable and defective fasteners) for an individual tie between runs, for a total of 22,000 ties.

These runs were selected for comparison as there were no deliberate changes made to track conditions between these runs, and they were captured within approximately 2 hours of one another. Additionally, note that each run was collected in the opposite driving direction with respect to the other (e.g., one was collected in increasing chainage and the other in decreasing chainage).

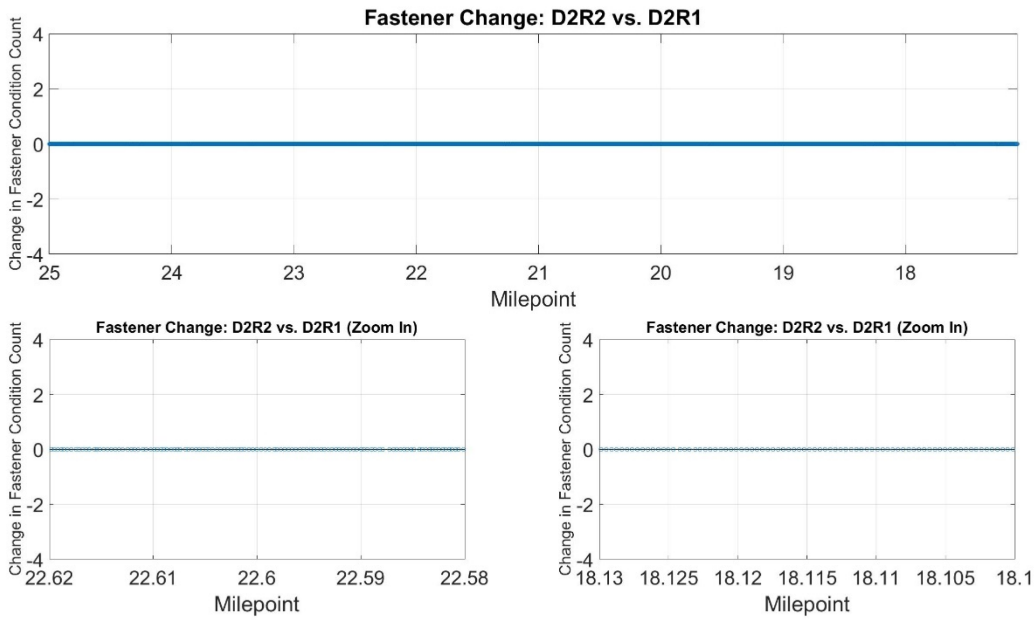

Figure 66 – Modeling Fastener Status Noise Floor (D2R2 vs. D2R1)

Upon initial review, the results shown in Figure 66 suggested a noise floor of zero. However, as run-to-run comparison was expanded beyond this dataset, researchers determined that there was indeed a small amount of variance from run-to-run, as can be seen in Figure 67 and Figure 68, which compare conditions from D2R1 to conditions in D4R1 and D4R2. The conditions between D4R1 and D4R2 should have been identical, yet there was a reported difference at MP 18 in D4R2 (circled in red) that was not present in D4R1.

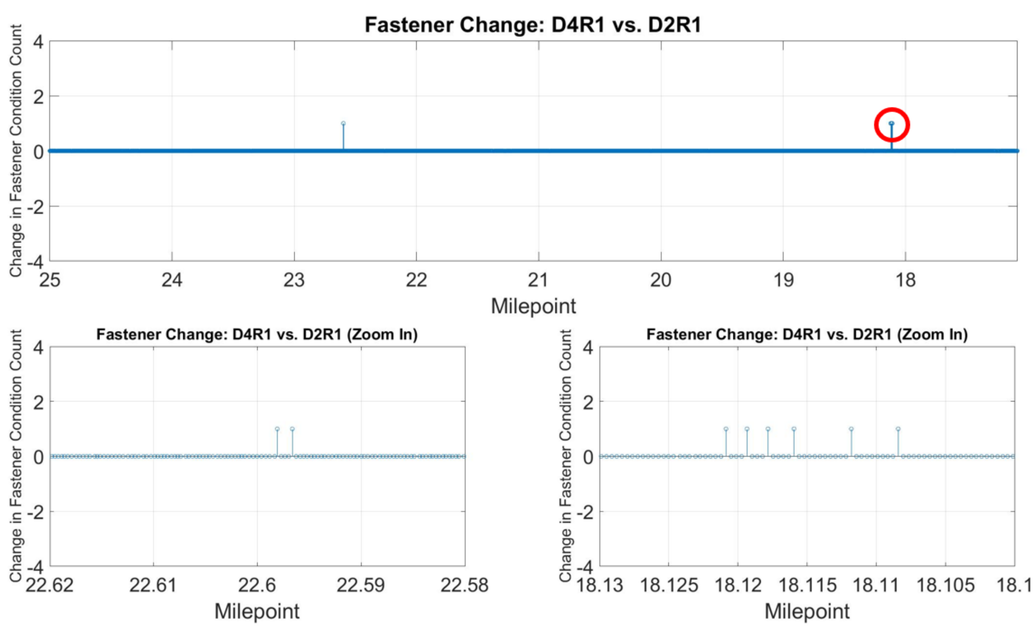

Figure 67 – Modeling Fastener Status Noise Floor (D4R1 vs. D2R1)

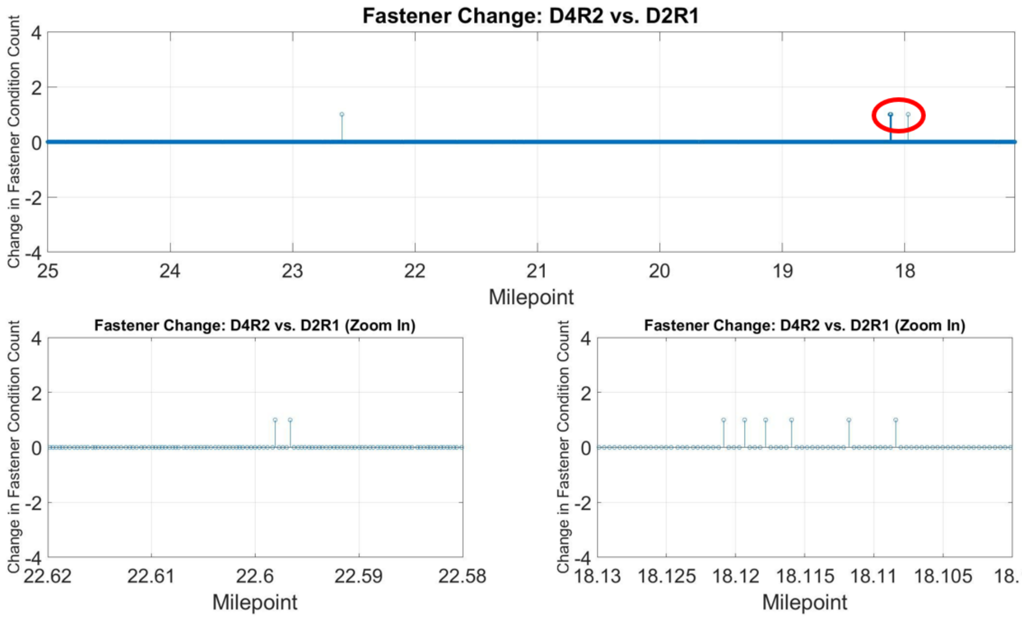

Figure 68 – Modeling Fastener Status Noise Floor (D4R2 vs. D2R1)

Upon further analysis, this difference was due to very small changes in the presentation of intensity and range data (as a result of differences in factors such moisture levels and vehicle orientation) between runs for the fastener in question. These small differences resulted in the fastener in question narrowly being rated as unacceptable in one run (due to it being borderline loose) and not in the other run, producing a kind of "data bin hopping" for a single fastener.

From a practical perspective, a change of fastener status for a single tie would not be of note to a field inspector. Thus, to compensate for this natural variance and to make the measurement meaningful, a threshold of more than one change per 10 m of track was devised.

Deliberate changes were made to fastener status (i.e., acceptable vs. defective) to validate the ability to detect changes for this parameter.

Figure 69 – Deliberate Changes to Fastener Conditions in the Field (missing on the left and loose on the right)

Deliberate changes to the fastener status were made at the following locations between Day 1 and Day 2:

- MP 22+3092; one clip was removed on the right-hand rail.
- MP 22+3100; one clip was removed on the right-hand rail.
- MP 18+10 to MP 18+90; three loosened and three missing clips along this stretch

Consequently, when comparing fastener acceptability status along the test loop of any of the scans captured on Day 1 to those captured on Day 2, one should expect to see two data points at approximately MP 22 and six data points at approximately MP 18. Figure 70 presents this comparison.

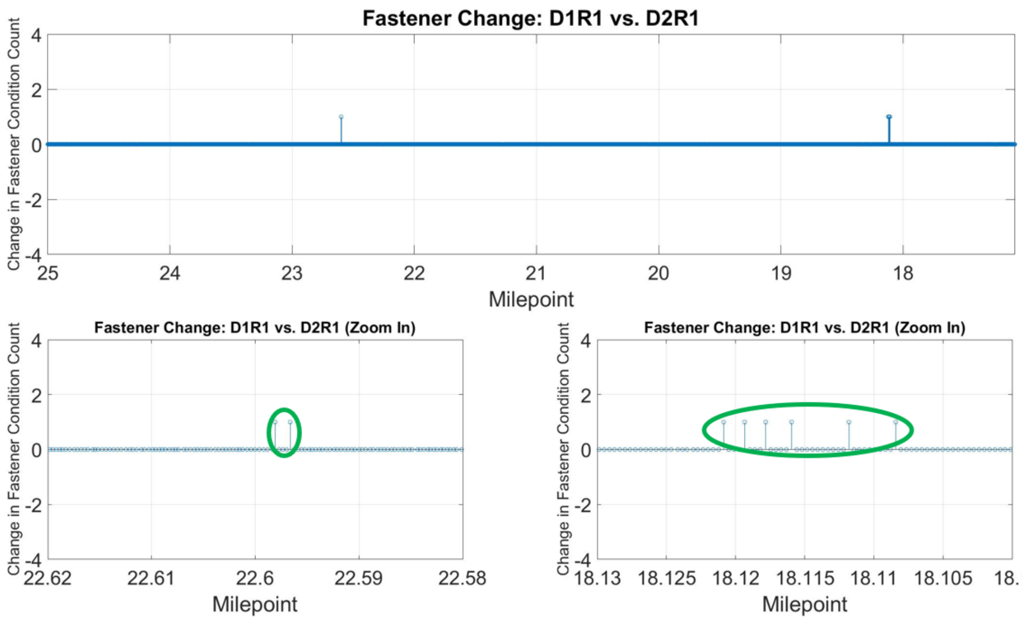

Figure 70 – Detecting Deliberate Fastener Change (D1R1 and D2R1)

Figure 70 shows that there were two instances of fastener status change at approximately MP 22 and six instances of fastener status change at approximately MP 18, as expected.

There were no organic changes to fastener counts over the test period.

4.9 Rail Surface Change Detection Results

This task involved run-to-run comparisons regarding the quantity of rail surface damage on the head of the rail between runs, with results being reported as surface area of defect (mm^2) per mile of track.

Upon review of the data, very little evidence of actual rail surface damage could be found. Two scenarios were considered. The first utilized a minimum defect size of 4 mm long by 5 mm wide by 4 mm deep which resulted in no defects being reported along the test loop.

A stricter criteria of 4 mm long by 5 mm wide by 2 mm deep was then applied, with only 20 defects detected along the 8-mile test loop. In aggregate, these defects amounted to just 50–200 mm^2 (2–8 square inches) per mile of track which is minute in comparison to a 1-mile length of track. Upon review of the 3-dimensional data for these candidate defects, most presented only a faint appearance of a defect in the intensity images, with no corroborating result in the 3-dimensional range data. Consequently, the research team believes that these defects were not legitimate and were in fact simply noise.

Repeatability, mean, and standard deviation for rail surface defect area per mile of track were not calculated due to the very low number of reported instances and the belief that the candidate defects were in fact not legitimate.

Insufficient data were available to properly model the noise floor of change in rail surface defect area per mile. Organic changes in area of real surface defects per mile were not detected over the course of the project.

5. Conclusion

The extended field trial proved to be a success. Repeatability, mean, and standard deviation of change measurements were determined and noise floors for each measured parameter were established.

LRAIL change measurements were determined to be highly repeatable, ranging between 95.72 percent (for skew angle) to 100 percent (for fasteners, joint bar bolting, tie rating, ballast level in gauge, and ballast level on left field side), with an overall average of 99.28 percent.

A convenient trailer-based inspection platform and a real-time user interface was developed to streamline the field data collection process.

The change detection process was also streamlined and further automated through the development of an automated run-to-run alignment tool.

Most importantly, the ability to detect organic changes above the noise floor was validated with numerous examples, including ballast level, ballast fouling, tie condition, joint bar bolting, and joint gap being detected.

6. Recommendations

Based on the results of this project, the research team recommends this technology be deployed on an autonomous test platform. This would present an opportunity to evaluate the impact of higher data collection speeds (above hi-rail limits) on collected data accuracy and repeatability.

Doing so would directly support the FRA Track Research Division's strategic priority of developing track inspection technologies that detect defects before they become failures in service by providing continuous change-detection capabilities across the network.

The team recommends this technology be installed on the FRA autonomous inspection vehicle, DOTX-225, if possible. This approach has the advantages of collocating LRAIL technology along with existing ATGMS sensors to permit the simultaneous measurement and leveraging of existing location referencing hardware (DGPS and encoder), the existing onboard power system, and existing cellular data communication capabilities.

Abbreviations and Acronyms

FRA	Federal Railroad Administration
LRAIL	Laser Railway Inspection System
MP	Milepoint
LCMS	Laser Crack Measurement System
GPS	Global Positioning System
GNSS	Global Navigational Satellite System
IMU	Inertial Measurement Unit