U.S. Department of Transportation

Federal Railroad Administration

Demonstration of Commercial-Off-The-Shelf Change Detection on Railway Images

Office of Research, Development
and Technology
Washington, DC 20590

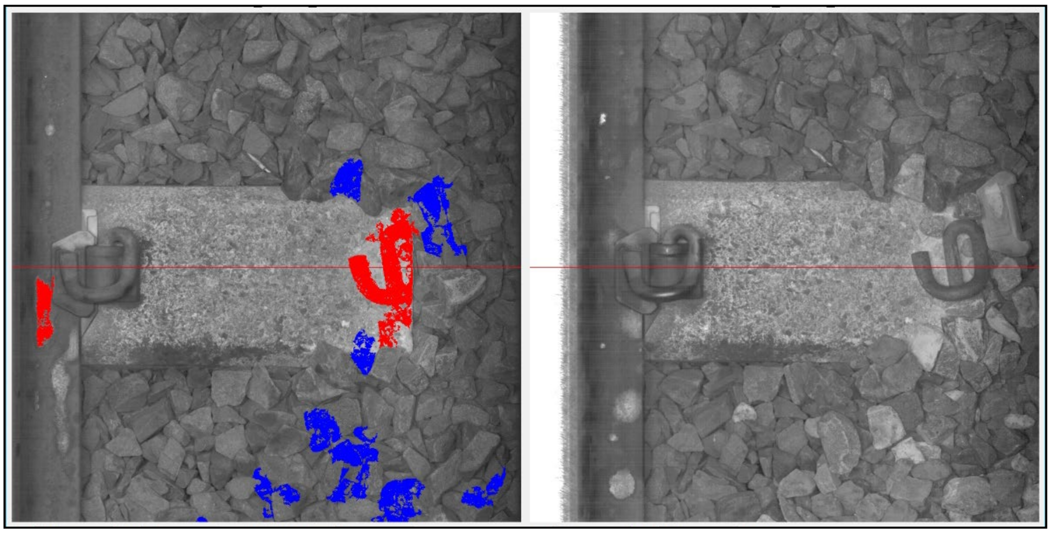

DOT/FRA/ORD-19/22

Final Report
July 2019

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	July 2019	Technical Report, 2017

4. TITLE AND SUBTITLE
Demonstration of Commercial-Off-The-Shelf Change Detection on Railway Images

5. FUNDING NUMBERS
DTFR53-17-C-00001

6. AUTHOR(S)
Herbert Henderson,[1] Atle Borsholm[2]

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
ENSCO Inc.[1]
5400 Port Royal Road
Springfield, VA 22151

Harris Corporation[2]
385 Interlocken Crescent, Suite 300
Broomfield, CO 80021

8. PERFORMING ORGANIZATION REPORT NUMBER
SERV-REPT-0001686

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-19/22

11. SUPPLEMENTARY NOTES
COR: Jay Baillargeon

12a. DISTRIBUTION/AVAILABILITY STATEMENT
This document is available to the public through the FRA website.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

This study, conducted jointly by ENSCO, Inc. and Harris Corporation, assessed the potential benefits of applying commercially available change detection software to railway machine vision images. Change detection involves registering "Before" and "After" images on a pixel-by-pixel basis to isolate changed areas. Two modes of change detection were assessed: intensity-based detection that identifies change where pixel intensity differs and thematic-based detection that identifies change where pre-determined pixel themes (e.g., concrete, ballast, etc.) occur. The intensity-based mode has shown potential for use on rail-based images; the thematic-based mode has shown potential to help fill a supporting role of filtering non-relevant changes.

14. SUBJECT TERMS
Track inspection, change detection, machine vision inspection

15. NUMBER OF PAGES
123

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	

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

Inches	0	1	2	3	4	5

Centimeters	0	1	2	3	4	5	6	7	8	9	10	11	12	13

QUICK FAHRENHEIT - CELSIUS TEMPERATURE CONVERSIO

°F	-40°	-22°	-4°	14°	32°	50°	68°	86°	104°	122°	140°	158°	176°	194°	212°
°C	-40°	-30°	-20°	-10°	0°	10°	20°	30°	40°	50°	60°	70°	80°	90°	100°

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

Contents

Executive Summary 1

1. Introduction 2
   - 1.1 Background 2
   - 1.2 Objectives 3
   - 1.3 Overall Approach 3
   - 1.4 Scope 3
   - 1.5 Organization of the Report 4
2. Change Detection Software 5
   - 2.1 Definition and Scope of Change-Based Processing 5
   - 2.2 Change Detection Processing 5
   - 2.3 Image Review Tool 9
3. Technical Approach 10
   - 3.1 Description of Analyzed Datasets 10
4. Results 12
   - 4.1 Macroscopic Overview of Results 12
   - 4.2 Relevant Changes 14
   - 4.3 Missed Changes 27
   - 4.4 Non-Relevant Changes 30
5. Managing Non-Relevant Changes 48
   - 5.1 Address Root Cause of Non-Relevant, False Changes 48
   - 5.2 Classify and Ignore Common Non-Relevant Changes 52
   - 5.3 Region-Based Change Detection to Filter Non-Relevant Changes 55
6. Ideal Change-Based Track Inspection Process 62
   - 6.1 Description of Change-Based Track Inspection 62
   - 6.2 Ideal Manual Review Process 64
   - 6.3 Ideal, Closed-Loop Maintenance Tracking 65
7. Conclusion 66
   - 7.1 Perceived Strengths and Weaknesses 68
   - 7.2 Identified Areas for Additional Development 69
   - 7.3 Technology Assessment 69
   - 7.4 Recommendations 69

Appendix A. Change Detection Software Assessment 71

Abbreviations and Acronyms 115

Illustrations

Figure 2-1. Example of a Detected Intensity-Based Change (Frame 332 in 2012/2013 Data Set) 7

Figure 2-2. Example of Detected Thematic-Based Change (Frame 368 in 2012/2013 Data Set).. 8

Figure 2-3. Image Review Tool to Accommodate Rail-Based Images 9

Figure 4-1. Template Used to Present Change Detection Results 12

Figure 4-2. Macroscopic Overview of Results 13

Figure 4-3. Graphical Representation of Unchanged Frame and Unchanged Pixel Percentages. 14

Figure 4-4. Intensity-Based Change Response for a Missing Fastener—Example 1 15

Figure 4-5. Intensity-Based Change Response for a Missing Fastener—Example 2 16

Figure 4-6. Intensity-Based Change Response for a Rotated Fastener.. 17

Figure 4-7. Intensity-Based Change Response for a Rotated Fastener Plate Retainer Clip 18

Figure 4-8. Intensity-Based Change Response at a Crumbled Tie Shoulder............................... 19

Figure 4-9. Thematic-Based Change Response at a Crumbled Tie 20

Figure 4-10. Intensity-Based Change Response for a Scenario that Resembles Fouled Ballast .. 21

Figure 4-11. Intensity-Based Change Response for a Rail Surface Anomaly 22

Figure 4-12. Intensity-Based Change Response for Standing Water 23

Figure 4-13. Intensity-Based Change Response for a Replaced Concrete Tie 24

Figure 4-14. Intensity-Based Change Response for a Replaced Fastener at a Replaced Concrete Tie 25

Figure 4-15. Intensity-Based Change Response for a Replaced Third Rail Stand 26

Figure 4-16. Intensity-Based Change Response for a New Third Rail Retainer Clip 27

Figure 4-17. Example Used to Clarify Missed Intensity-Based Changes 29

Figure 4-18. Example of a Relevant Intensity-Based Change Missed by the Evaluated Software 30

Figure 4-19. Example 1 of Non-Relevant, True Changes Associated with Ballast...................... 34

Figure 4-20. Example 2 of Non-Relevant, True Changes Associated with Ballast...................... 35

Figure 4-21. Example of a Non-relevant, True Change Caused by Rail Grease 36

Figure 4-22. Example of a Non-Relevant, True Change Caused by Grease and Rail Dust 37

Figure 4-23. Example of a Non-Relevant, True Change Caused by Moisture 38

Figure 4-24. Example of a Non-Relevant, True Change Caused by Trash (a Bottle) 40

Figure 4-25. Close-up View of a Non-Relevant, True Change Caused by Trash (a Bottle) 41

Figure 4-26. Example of a Non-Relevant, True Change Resulting from a Left-Over Track Component 42
Figure 4-27. Example of a Non-Relevant, True Change Caused by Pine Needles 43
Figure 4-28. Example of a Non-Relevant, True Change Caused by a Stick 44
Figure 4-29. Example of a Non-Relevant, False Change Caused by Non-Perfect Intensity Normalization 45
Figure 4-30. Example of a Non-Relevant, False Change Caused by Non-Perfect Intensity Normalization 45
Figure 4-31. Example of a Non-Relevant, False Change Caused by Non-Perfect Co-Registration 46
Figure 4-32. Examples of False, Thematic-Based Changes 47
Figure 5-1. Close-up View of a Non-Relevant, False Change—Non-Perfect Intensity Normalization 49
Figure 5-2. Close-up View of False Change Caused by Non-Perfect Co-Registration........ 50
Figure 5-3. Close-up View of a Cluster of False, Thematic-Based Changes 51
Figure 5-4. Demonstration of Improvements from Re-Training the Thematic Classifier........ 52
Figure 5-5. Intensity-Based Result—Isolated Changes in Ballast........ 53
Figure 5-6. Thematic-Based Result—Isolated Changes in Ballast........ 54
Figure 5-7. Example Confusion Matrix Used to Ignore Non-Relevant Changes........ 54
Figure 5-8. Example Four-Step Process Used to Apply Region-Based Change Detection........ 56
Figure 5-9. Demonstration of Improvement from Region-Based Change Detection........ 57
Figure 5-10. Non-Relevant, True Change After Applying Region-Based Change Detection........ 57
Figure 5-11. Example of a Rotated Base Plate Retainer Clip........ 58
Figure 5-12. A Region-Based Change Detection Template Customized for a Specific Objective 59
Figure 5-13. Frame used to Estimate Detection Probability (Frame Size = 2048 x 7214)........ 60
Figure 5-14. Region-Based Change Detection Applied to a Base Plate Retainer Clip 60
Figure 6-1. Envisioned Process for Ideal, Change-Based Track Inspection........ 63

Tables

Table 3-1. Data Set Summary Statistics 11
Table 4-1. Static-Frame and Static-Pixel Results for Each Evaluated Data Set 31
Table 4-2. Percentage of Frames with at Least One, Non-Relevant, True Change 32
Table 4-3. Strategies for Managing Isolated Changes in Ballast 33
Table 4-4. Strategies for Managing Non-Relevant Changes from Surface Discoloration 36
Table 4-5. Strategies for Managing Non-Relevant Changes from Trash and Natural Debris 39

Executive Summary

In 2017, the Federal Railroad Administration (FRA) sponsored an ongoing a study to assess the potential benefits of applying commercially available change detection software (ENVI, provided by Exelis VIS, a subsidiary of Harris Corporation) to railway machine vision images; this work was conducted by ENSCO, Inc. and Harris Corporation at their facilities. Change detection involves registering "Before" and "After" images on a pixel-by-pixel basis to isolate changed areas. Two modes of change detection were assessed: intensity-based and thematic-based. The intensity-based mode identifies change where pixel intensity differs and the thematic-based mode identifies change where pixel theme differs. Pixel theme is determined ahead of time using an algorithm to assign a theme (e.g., concrete, ballast, etc.) to each pixel. The intensity-based mode showed eminent promise on rail-based images. The thematic-based mode showed potential to help fill a supporting role of filtering non-relevant changes.

Machine vision offers the potential to improve track inspection thoroughness while reducing disruption to railroad operations, however, today's algorithm development process is expensive, and the corresponding performance is not sufficient to significantly avoid manual image review. What is needed is a machine vision approach capable of isolating a broad number of conditions based on a low up-front development cost. The corresponding performance must be sufficient to substantially reduce the need for manual image review so that the total cost of machine-vision-based track inspection is in line with that of prevailing track inspection methods.

Based on results from this study, change detection demonstrated the potential to realize the objectives stated above. Change detection could identify changes corresponding to many conditions while providing a low probability of missing a relevant change. Low miss probabilities are expected to justify a no-change-no-review policy if change-based processing is deployed in the rail sector. Under such a policy, manual review is applied only where change is detected. Such a policy leads to a cost-effective track inspection method only if the number of reported, non-relevant changes is not overwhelming. This study has shown that non-relevant changes have the potential to be overwhelming in rail-based images, but viable approaches for managing non-relevant changes exist. These approaches may involve machine vision algorithms, but their development cost is estimated to be a third of that associated with conventional machine-vision algorithms. It is easier to filter and manage non-relevant changes than to detect many different conditions using conventional algorithms.

The change detection software evaluated in this study was originally developed for aerial-based imagery. While core aspects of the software are applicable to rail-based images, other aspects require modification to establish a commercial capability suitable for the rail sector. The following minimum development needs were identified during this study: 1) automating the process of applying co-registration and intensity normalization to rail-based images, 2) adding real-time compensation for light exposure to existing rail-based imaging software, and 3) integrating change detection into a pre-existing rail-based image review tool.

The study concluded that change-based processing is currently better-suited for office-based application as opposed to implementation on an inspection vehicle in real-time. Although not assessed during this study, change-based track inspection is expected to benefit from the use of three-dimensional image data.

1. Introduction

The purpose of this study was to evaluate the potential benefits of leveraging image-based change detection in a rail environment. In this study, a leading commercial change detection software package (ENVI, provided by Exelis VIS, a subsidiary of Harris Corporation) was used to apply change detection to rail-based images. This report will detail the change-based processing that was applied; present corresponding results and in-depth analysis; identify areas where improvements are deemed necessary; and provide final conclusions and recommended next steps.

1.1 Background

Visual track inspection is still largely addressed by traditional walking and hi-rail-based inspections, but trends to add machine vision capabilities are accelerating. Machine vision offers improved safety, less disruption to railroad operations, and the potential to improve inspection thoroughness. However, within a rail environment, today's machine vision approaches are burdened by two important realities.

1) A need exists for detecting more than 30 visually observable conditions-of-interest in a rail environment.[1] Today, each condition-of-interest requires its own machine vision algorithm development cycle. Based on current-day algorithm development approaches, the cost to develop a matured algorithm for one condition-of-interest is on the order of $125,000.[2]
2) Imaging conditions in a rail environment are uncontrolled and challenging to accommodate in a robust way using traditional machine vision algorithms. The net result is that, even after considerable expense, corresponding algorithm performance is not sufficient to significantly minimize the need for manual image review.

Due primarily to high costs associated with machine vision algorithm development, a need exists for an algorithm framework that is better-suited to accommodate the uncontrolled nature of rail environments. The desired framework needs to be adaptable to accommodate a wide range of detection tasks without requiring an expensive algorithm development cycle for each task. Additionally, the framework needs to continue providing useful results even if conditions in the environment change somewhat after the framework was deployed. A primary motivation for this study is to explore the potential for image-based change detection to act as a cornerstone within such a framework.

Comprehensive, autonomous inspection is a goal for machine vision technology in the rail sector. While current-day utilization of the technology is well short of this goal, machine vision is beginning to improve inspection efficiency in a variety of ways. One way involves using machine vision algorithms to automatically locate assets such as ties, joint bars, and rail clips

[1] Based on railroad customer enquiries made between 2012 and 2016 in the form of requests for information (RFIs), requests for proposals (RFPs), or directly a result of ongoing customer-supplier relationships.

[2] Cost estimate for algorithm development is based on known costs to develop a machine vision algorithm for detecting missing fasteners under FRA report, Robust Anomaly Detection for Vision-Based Inspection of Railway Components, as well as other relevant examples known to ENSCO.

that need to be inspected. Once such assets have been located in an image file, image review software enables manual review to proceed efficiently from one asset to the next without the traditional need to traverse the physical distance between assets.

In situations where safety is compromised only when anomalous conditions emerge in clusters (e.g., missing fasteners), it is possible to apply business rules to achieve a low enough missed detection rate to justify reviewing only detected conditions rather than all candidate images. Other than this scenario, missed detection rates are not low enough to avoid the need for 100 percent manual screening.

At the current state of development, machine vision algorithms do not detect all required track conditions. As a result, traditional modes of track inspection are still prevalent, and manual image review in an office setting is emerging as a viable alternative. An office-based review offers improved safety and less disruption to railroad operations.

An important aspect of visual track inspection is converting identified conditions into actionable results as part of a track maintenance program. Today's image-based track inspection logs location and description information for each condition-of-interest in a database format. This provides several benefits including:

1) Automatic generation of track condition reports
2) Images showing each detected condition
3) Outputs needed to support efficient work order creation
4) Records needed for long-term asset tracking

Demand for more and better machine vision algorithms is ever-present in the rail sector—finding cost-effective ways to meet this demand is key to enabling machine-vision-based track inspection to experience broader adoption.

1.2 Objectives

The goal of this project was to assess the potential benefits of applying commercially available change detection software to railway machine vision images.

1.3 Overall Approach

This study evaluated the application of commercially available change detection software, ENVI, to railway machine vision images. Change detection involves registering "Before" and "After" images on a pixel-by-pixel basis to isolate changed areas. Two modes of change detection were assessed: intensity-based and thematic-based. The intensity-based mode identifies change where pixel intensity differs and the thematic-based mode identifies change where pixel theme differs. Pixel theme is determined ahead of time using an algorithm to assign a theme (e.g., concrete, ballast, etc.) to each pixel.

1.4 Scope

This investigation was limited to the evaluation of the viability of change detection applied to rail images using existing software. The effort did not focus on the customization of software to establish an optimized change detection approach. Recommendations for further development required to enable this approach to be used in an efficient manner were outputs of this study.

1.5 Organization of the Report

This report provides an in-depth analysis of a study that applied commercial, off-the-shelf change detection software to railway images.

Section 2 provides an overview of the commercial change detection software evaluated during this study, including a high-level summary of modifications made to the software to accommodate rail-based images.

Section 3 describes the technical approach applied during this study and includes a description of the data sets (images) processed during the study.

Section 4 presents results obtained from applying change detection to rail-based images under this study. Results include examples of relevant changes, missed changes, and non-relevant changes.

Section 5 provides discussion and in-depth analysis related to three strategies assessed during this study for managing non-relevant changes.

Section 6 describes an envisioned, ideal, change-based, track inspection process. The ideal process establishes a basis for identifying areas where further development is recommended.

Section 7 provides a summary of the findings obtained during this study, including an assessment of whether image-based change detection is expected to be useful within the rail sector. A list of perceived strengths and weaknesses associated with applying change detection to railway images is included, as well as a list of identified areas where further development is deemed necessary.

Appendix A includes a report written by Exelis VIS, the providers of the change detection software evaluated during this study. The report further describes the change detection software assessed during this study, details the processing steps applied to rail-based images during this study, and includes additional example results.

2. Change Detection Software

This chapter defines change detection, describes intensity-based and thematic-based modes of change detection, and summarizes the change detection software assessed during this study.

2.1 Definition and Scope of Change-Based Processing

This section defines change detection, including the broader scope of change-based processing, and clarifies aspects within the broader scope that were evaluated during this study. Change detection is the process of automatically identifying regions in a dataset where new information (change) exists relative to a baseline state recorded at an earlier time. Within the rail sector, change detection is generally applicable to one-dimension (e.g., track geometry), two-dimension (e.g., rail-based imagery), and three-dimensional (e.g., rail and track profile) data. While the software evaluated herein accommodates both two and three-dimensional data, this study only evaluated the software on continuous, two-dimensional track images.

This study evaluated the following aspects of image-based change detection:

1) **Image warping** – A process of stretching or shrinking an image along one or more axes to improve co-registration prior to change detection
2) **Intensity normalization** – A process associated with intensity-based change detection used to compensate for global lighting differences between two images
3) **Co-registration** – A process of aligning two images on a pixel-by-pixel basis
4) **Thematic classification** – A deep-learning-based classifier used to assign a texture or material type (a theme) to each pixel in an image, as needed to support thematic-based change detection
5) **Image differencing** – A process of subtracting one image from another on a pixel-by-pixel basis and comparing the results to a threshold to determine where change exists
6) **Cluster processing** – A process of identifying contiguous groups of pixels where change has been detected and using the identified groups to remove small regions of change—intended primarily to help manage non-relevant changes
7) **Intensity-based change detection** – A change detection mode where reported changes are based on differences in pixel intensity values between image pairs
8) **Thematic-based change detection** – A change detection mode where reported changes are based on differences in pixel theme (e.g., material type) between a "Before" and "After" image pair

2.2 Change Detection Processing

This section presents an overview of the steps applied to a pair of images to achieve change detection results. See Appendix A for an in-depth analysis of how change detection processing was applied under this project.

2.2.1 Processing

To achieve change detection results that correspond well with changes perceived by observation, specific preprocessing steps are necessary. These steps include co-registration and intensity normalization.

Co-Registration

Co-registration precisely aligns two images prior to applying change detection. Co-registration is achieved using an algorithm that relies on a set of one or more anchor points. An anchor point is a matching location in both images (a ground-truth registration point) used by a co-registration algorithm to pin a “Before” image to a corresponding “After” image. At least one anchor point is required to achieve co-registration; however, in general, using more than one anchor point improves registration accuracy. When more than one anchor point is provided, the co-registration algorithm warps (stretches or shrinks) all unpinned pixels as needed to fit between the specified set of anchor points.

The software evaluated during this study includes feature extraction algorithms sufficient to automatically select anchor points as needed to fully automate the co-registration process. However, this algorithm is not currently adapted to work with rail-based, line scan image formats. Therefore, this study used manually-entered anchor points for co-registration.

Intensity Normalization

Intensity normalization removes any global intensity differences between two images prior to applying intensity-based change detection. Global intensity differences can be caused by the environment or by the imaging equipment. Changes in ambient lighting (e.g., sunny versus overcast) is an example of an environmentally-induced difference. Swapping a new camera for an old camera is an example where imaging equipment has potential to cause a global intensity difference. That is, if the sensor response in a new camera is different from that of an old camera, a global intensity difference may exist. Without intensity normalization, normally unnoticeable, global intensity differences would tend to dominate the reported change response.

In the simplest case, intensity normalization is achieved by first computing the difference between the average intensity in both images being compared. The difference is then added to the darker of the two images to bring its average brightness up to the same level as the other image. In practice, intensity normalization is more complex than this because it must also attempt to properly accommodate additional factors. Additional factors include boundary conditions between consecutive frames and imaging artifacts (e.g., caused by a dirty lens). Appendix A clarifies details associated with intensity normalization steps applied during this project.

2.2.2 Intensity-Based Change Detection

When applying intensity-based change detection, pixel intensity serves as the compared value. Change is detected at each pixel where the normalized intensity difference is greater than a configurable threshold. Areas where change is detected are then subjected to cluster processing. Cluster processing filters out small changes and reports large changes based on a configurable, size threshold. In this report, intensity-based change detection results are presented using two color channels, one corresponding to changes from dark to light (blue areas in Figure 2-1); the other from light to dark (red areas in Figure 2-1).

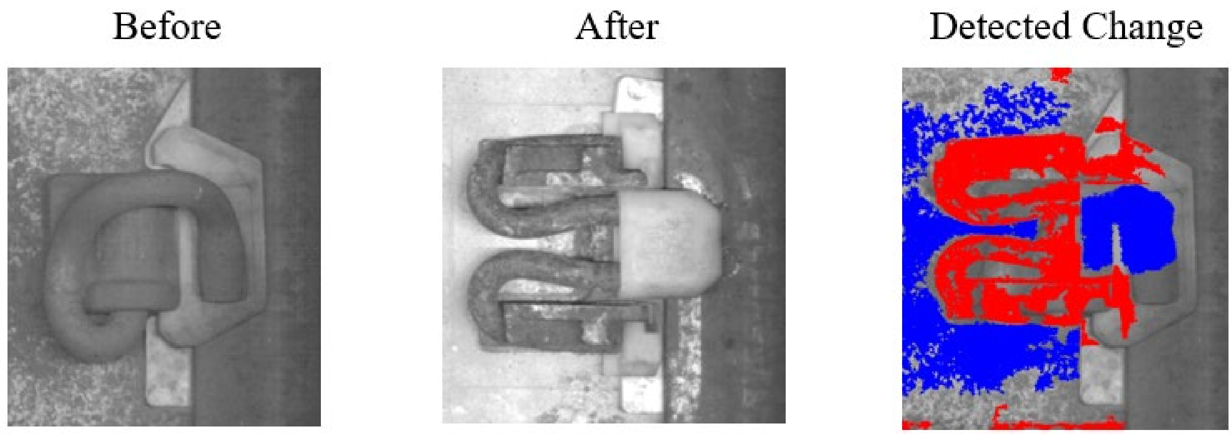

Figure 2-1. Example of a Detected Intensity-Based Change (Frame 332 in 2012/2013 Data Set)

2.2.3 Thematic-Based Change Detection

When applying thematic-based change detection, a theme (e.g., a material type) is first assigned to each pixel in both the "Before" and "After" images by a machine vision algorithm. Change is then reported at locations where pixel theme does not agree between the two images being compared.

Figure 2-2 shows an example of thematic-based change detection results associated with ballast covering a tie shoulder. As with the intensity-based mode, areas where a thematic change has occurred are presented using color-coded overlays. Pixel themes and corresponding color codes presented in this report are as follows:

1) Ballast (yellow)
2) Concrete (dark blue)
3) Wood (light blue—not present in Figure 2-2)
4) Rail surface and rail base (green)
5) Rail fastener (red)

By convention, the color used to report a detected thematic change is based on the theme present in the "Before" image.

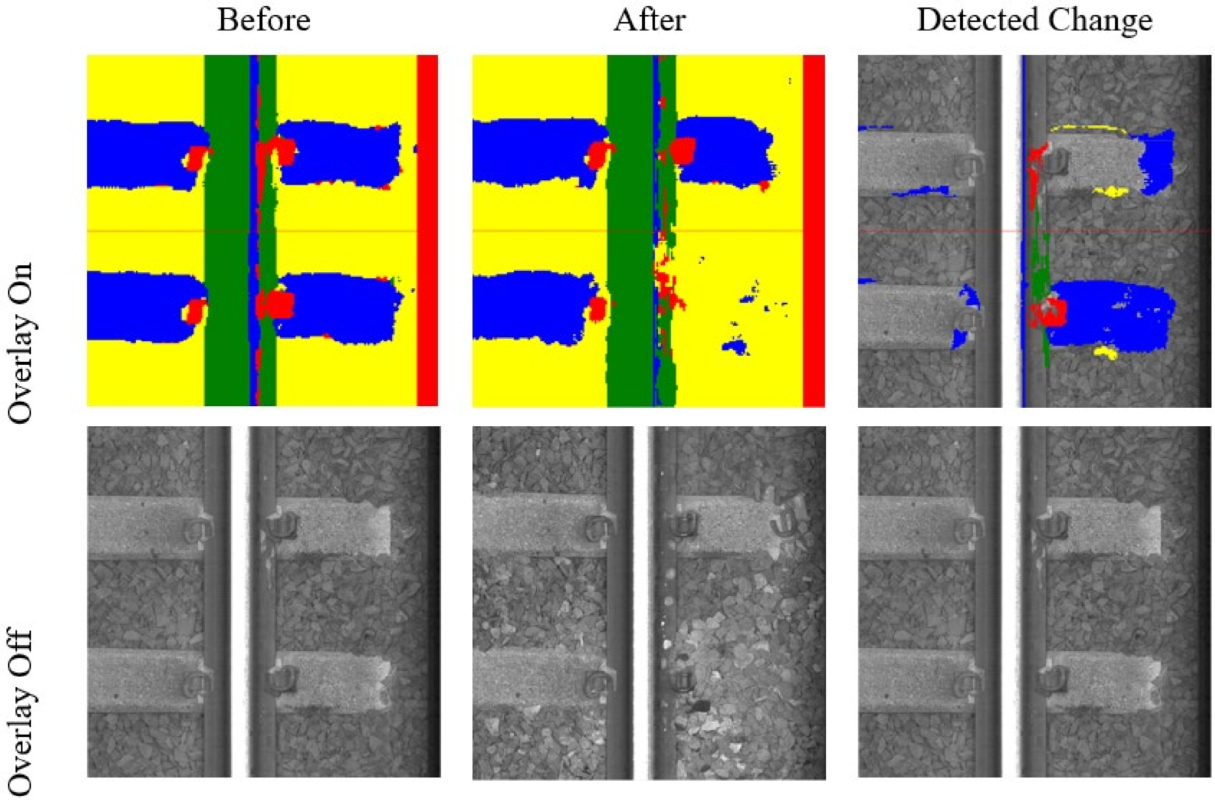

Figure 2-2. Example of Detected Thematic-Based Change (Frame 368 in 2012/2013 Data Set)

Under the thematic-based change detection approach, pixel theme values are established by a machine vision algorithm before applying change detection. The algorithm used in this study is based on deep learning technology. Deep learning uses raw intensity values from an image as parallel inputs to a network that then returns a corresponding classification result indicating the most likely theme (material type) for each pixel. The network must be trained ahead of time for each theme of interest. During training, images of a given theme (e.g., square patches of concrete tie, ballast, etc.) are used as inputs to the network. During training, numerical weights within the network automatically evolve as required to ultimately classify the material type, texture, or object in the training images. The set of numerical weights established during training for a given theme are then reusable (as a model) to identify the theme in arbitrary images.

There are no known limitations on the materials, textures, or component types that the thematic classifier can be trained to accommodate if training images are available. Training the thematic classifier requires engineering effort and is normally handled on a need-driven basis. Two important clarifications are:

1) Training is a one-time process that only needs to be performed when a new thematic class is added.

2) A significant number of training images is needed for each thematic class (e.g., several hundred for coarse results to several thousand for refined results).

Although training the thematic classifier requires engineering effort, the process is inexpensive compared to developing a separate algorithm for each theme of interest.

2.3 Image Review Tool

An image review tool was added to the pre-existing change detection software evaluated under this study. The image viewer associated with the pre-existing software was developed to accommodate rectangular, aerial imagery and corresponding mosaics. In contrast, the rail-based images assessed during this project are continuous, line scan images.

Figure 2-3 shows the image review tool developed under this project to accommodate rail-based images. In this example, the tool is displaying intensity-based change detection results associated with a replaced tie, including corresponding disturbances in the surrounding ballast. Within the viewer, change detection results are overlaid on top of the "Before" image on the left. The "After" image is on the right. The review tool for rail-based images supports synchronized, side-by-side viewing of continuous track bed images. It includes the following capabilities: standard play-forward, play-reverse, pause, goto-frame-number, zoom, and pan capabilities. See Appendix A for more details related to the rail-based image review tool.

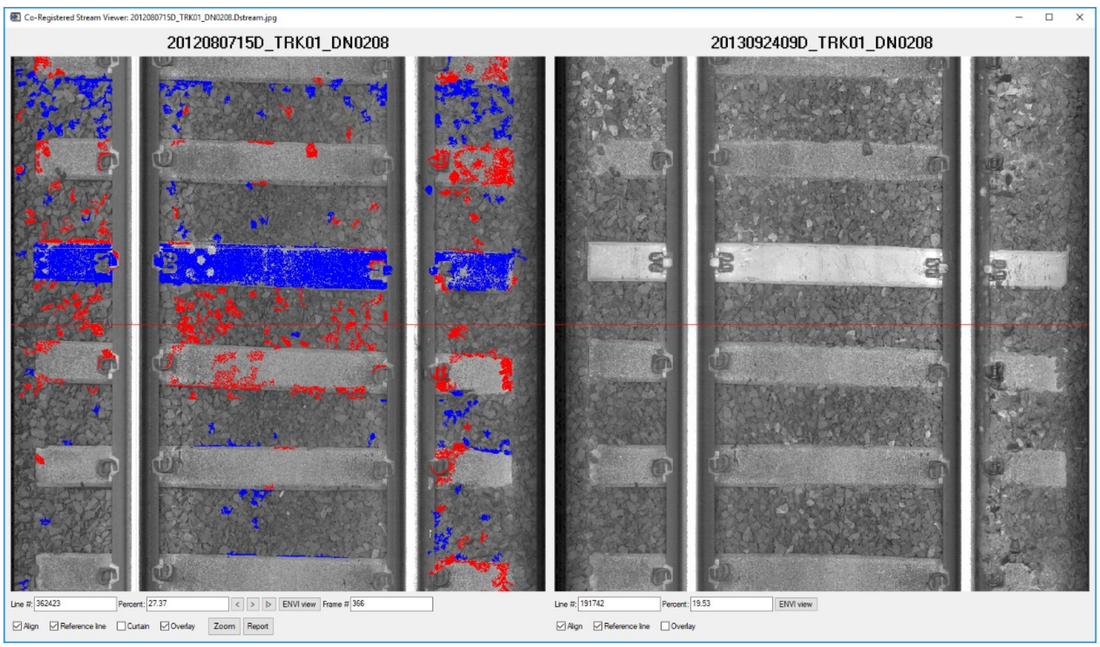

Figure 2-3. Image Review Tool to Accommodate Rail-Based Images

3. Technical Approach

ENSCO provided Exelis VIS with track-based images from an image archive. The images were recorded on two separate railroads and included both a “Before” and “After” file for each railroad. The “Before” and “After” files were spaced in time by 8 months in one case and 13.5 months in the other case.

Exelis VIS made adaptations to their existing ENVI software to accommodate the rail-based images and then applied intensity-based and thematic-based change detection to the provided image sets. Exelis VIS prepared and submitted a final report to ENSCO (see Appendix A). The Exelis VIS report details how change detection was applied to rail-based images and includes additional examples of the achieved change detection results.

Exelis VIS presented its final report to ENSCO, provided them with a copy of the ENVI software adapted to accommodate rail-based images, and briefed them on how to use the software. ENSCO then used the adapted software to analyze the provided change detection results.

A primary goal during ENSCO’s analysis was to determine whether applying image-based change detection to rail-based images might add value to a machine vision-based track inspection process. To help make this determination, regions of change were considered across three categories:

1) Relevant Changes – Changes detected by the software that are deemed to be useful within the context of a change-based track inspection process.
2) Non-relevant Changes – Changes detected by the software that are deemed to be non-useful within the context of a change-based track inspection process.
3) Missed Changes – Changes not detected by the software that are deemed to be useful within the context of a change-based track inspection process.

The categorized results were then evaluated to establish areas for improvement. Additional analysis was applied to assess strategies for managing non-relevant changes. A list of perceived strengths and weaknesses associated with applying image-based change detection to rail-based images was identified. The technology readiness level of the as-is state of the assessed software was determined and reported. Overall findings from the study were then reviewed to establish final conclusions and a list of recommended next steps.

3.1 Description of Analyzed Datasets

Change detection processing was applied to pre-existing track images pulled from archives maintained by ENSCO. The images used in this study are high-resolution, continuous images that span the full width of the track (from tie-end to tie-end). To cover the full width of the track, images are recorded using an array of four line scan cameras. The width of the image produced by each camera is 2,048 pixels (as measured perpendicular to the direction of travel). Images from each of the four cameras are stitched together during the image-capture process to achieve a composite image spanning the full width of the track. The width of the composite image is 4 x 2,048 = 8,192 pixels. High-intensity line scan lighting based on LED technology is used as an illumination source during track imaging. The pixel size used in this study, as projected onto the nominal tie surface, is 0.38 mm.

Track constructions represented in the image sets used in this study include:

- Concrete tie
- Wood tie
- Direct fixation
- Special trackwork at turnouts

Conditions of potential interest present in the image sets include:

- Crumbled ties
- Missing fasteners
- Rotated base plate retainer clips
- Disturbed ballast (used as an approximate surrogate for fouled ballast)
- Standing water
- Maintenance activity

Only image data was processed and analyzed in this study, while metadata (geolocation information) was not utilized. Additional statistics for each processed image set are included in Table 3-1 for reference.

Table 3-1. Data Set Summary Statistics

Description	Data Set 1	Data Set 2
Track Construction	Direct fixation	Concrete tie
Collection Date 1	June 29, 2016	August 7, 2012
Collection Date 2	March 8, 2017	September 24, 2013
Time Between Collections	~ 8 months	~ 13.5 months
Length of Processed Track Imagery	3.2 miles	1 mile
Image File Name 1	2016062902_TCIS.stream.jpg	2012080715D_TRK01_DN0208.pgm
Image File Name 2	2017030801_TCIS.stream.jpg	2013092409D_TRK01_DN0208.pgm

4. Results

This section presents and discusses representative examples of change detection results achieved during this study.

Change detection results in this report are typically presented as a set of four images arranged in a standard format. The standard format is based on a template shown in Figure 4-1. The top row of the template is used to show a “Before” and “After” image with no color overlay applied. The bottom row of the template is used to show the same “Before” and “After” images, except with a color overlay applied to the “Before” image to indicate areas where change is reported.

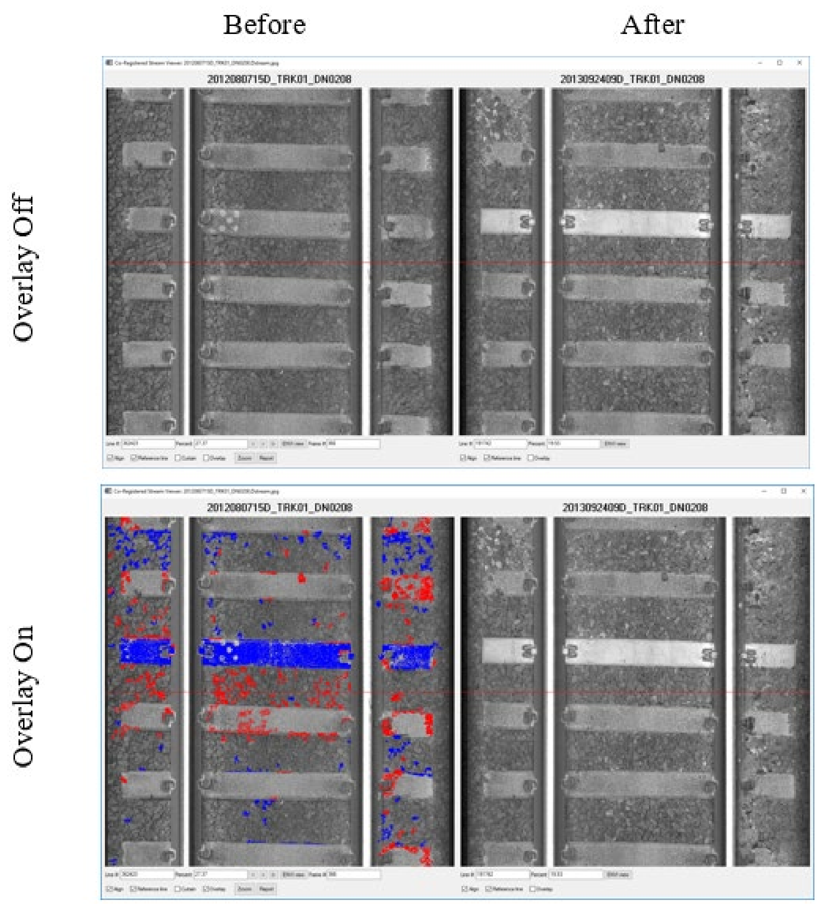

Figure 4-1. Template Used to Present Change Detection Results

4.1 Macroscopic Overview of Results

Figure 4-2 shows a graph intended to provide a macroscopic overview of the intensity-based change detection results. The horizontal axis of Figure 4-2 shows frame numbers corresponding to approximately1 mile of track image data. The vertical axis of Figure 4-2 indicates the number of changed pixels in each frame (in millions of pixels). Peaks in the graph of Figure 4-2 that extend above the green line correspond to tie replacement activity. Peaks between the red and green lines correspond to false changes caused by a correctable co-registration error (discussed in more detail later). Finally, all other changes observed during this study fall below the red line of Figure 4-2.

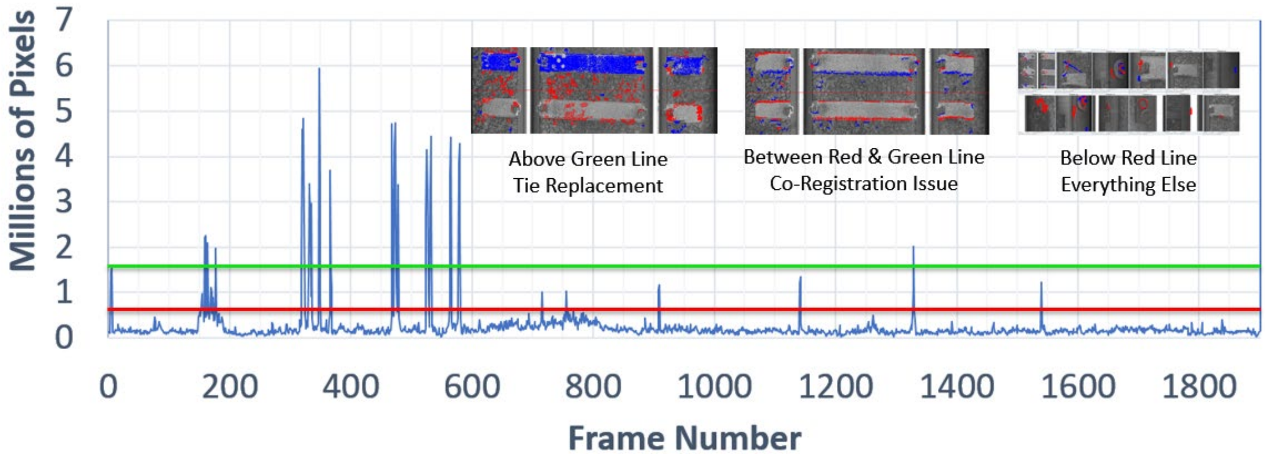

Figure 4-2. Macroscopic Overview of Results

Figure 4-3 shows the same curve plotted in Figure 4-2, except the axes are not labeled, and the plot region was scaled so that the area inside the border represents the total area within the corresponding mile-long track image. Figure 4-3 helps highlight that the total image area where change was reported (i.e., the area below the plotted curve) is significantly less than the total image area where no change was reported (i.e., the area above the plotted curve). Another point highlighted by Figure 4-3 is that the percentage of unchanged frames in the corresponding data set is zero, while the percentage of unchanged pixels is 98.4 percent. The low unchanged frame percentage suggests that dividing track images into full-width frames is not expected to be a productive way to present image-based change detection results during manual image review. In contrast, the high unchanged pixel percentage leaves open the possibility for other, more efficient approaches.

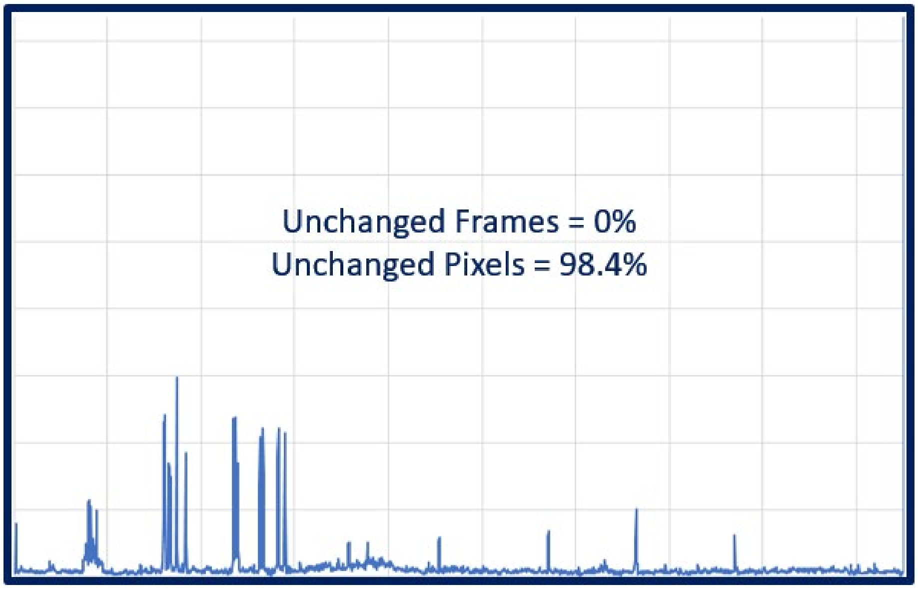

Figure 4-3. Graphical Representation of Unchanged Frame and Unchanged Pixel Percentages

4.2 Relevant Changes

This section presents examples where the evaluated software detected a relevant change. A relevant change is a region in an image that is important with respect to achieving a specified objective. In the case of machine-vision-based track inspection, the primary objective is to identify potentially unsafe track conditions.

4.2.1 Missing and Rotated Fastening Elements

Figure 4-4 shows an example of a missing fastener identified using intensity-based change detection. The fastener is shown in an installed state in the "Before" image and in an uninstalled state in the "After" image. The reported difference between the two states is then shown in red. Figure 4-4 shows the only example of a missing fastener found in the evaluated datasets, however, many other changes related to fasteners were found and properly identified by the software. A few of these are included in this section. Importantly, there were no observed incidences of relevant changes to a rail fastener that the software failed to identify.

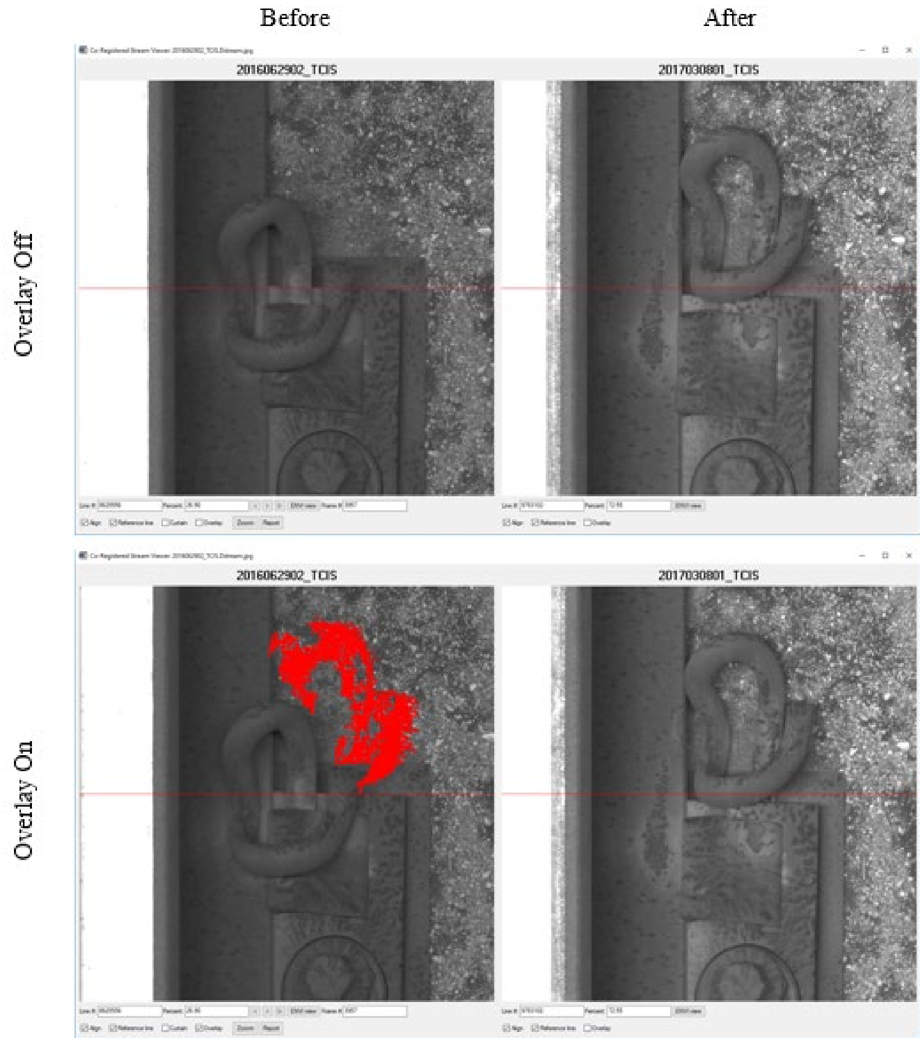

Figure 4-4. Intensity-Based Change Response for a Missing Fastener—Example 1

Figure 4-5 shows an example of detected changes corresponding to a rail fastener that was added to the imaged scene. Here, a corresponding ability to detect a fastener removed from a scene (i.e., a missing fastener) is implied.

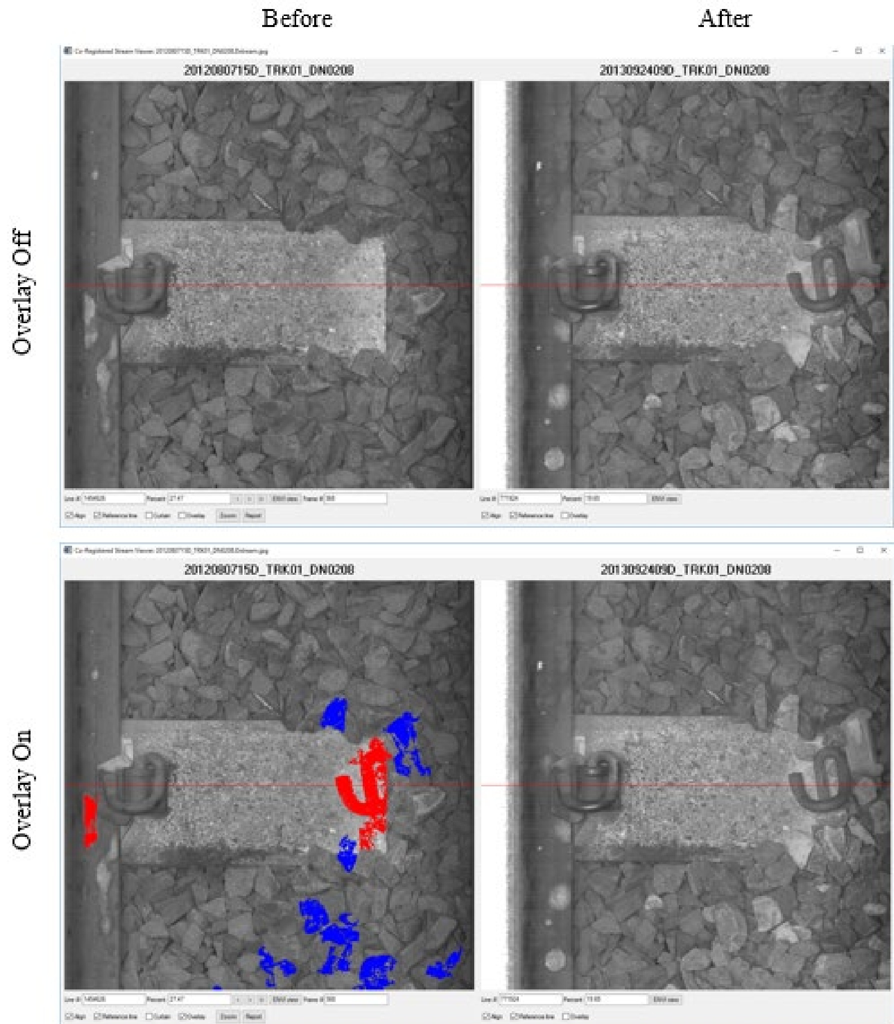

Figure 4-5. Intensity-Based Change Response for a Missing Fastener—Example 2

Figure 4-6 shows detected change of a fastening element that is slightly rotated in the “After” image with respect to its state in the “Before” image. While the fastening element in this example is already in a missing state in the “Before” image, this example serves to demonstrate that the evaluated software can detect relatively small changes in fastening element orientation. This is a preliminary indication that intensity-based change detection is suitable for locating rotated fasteners.

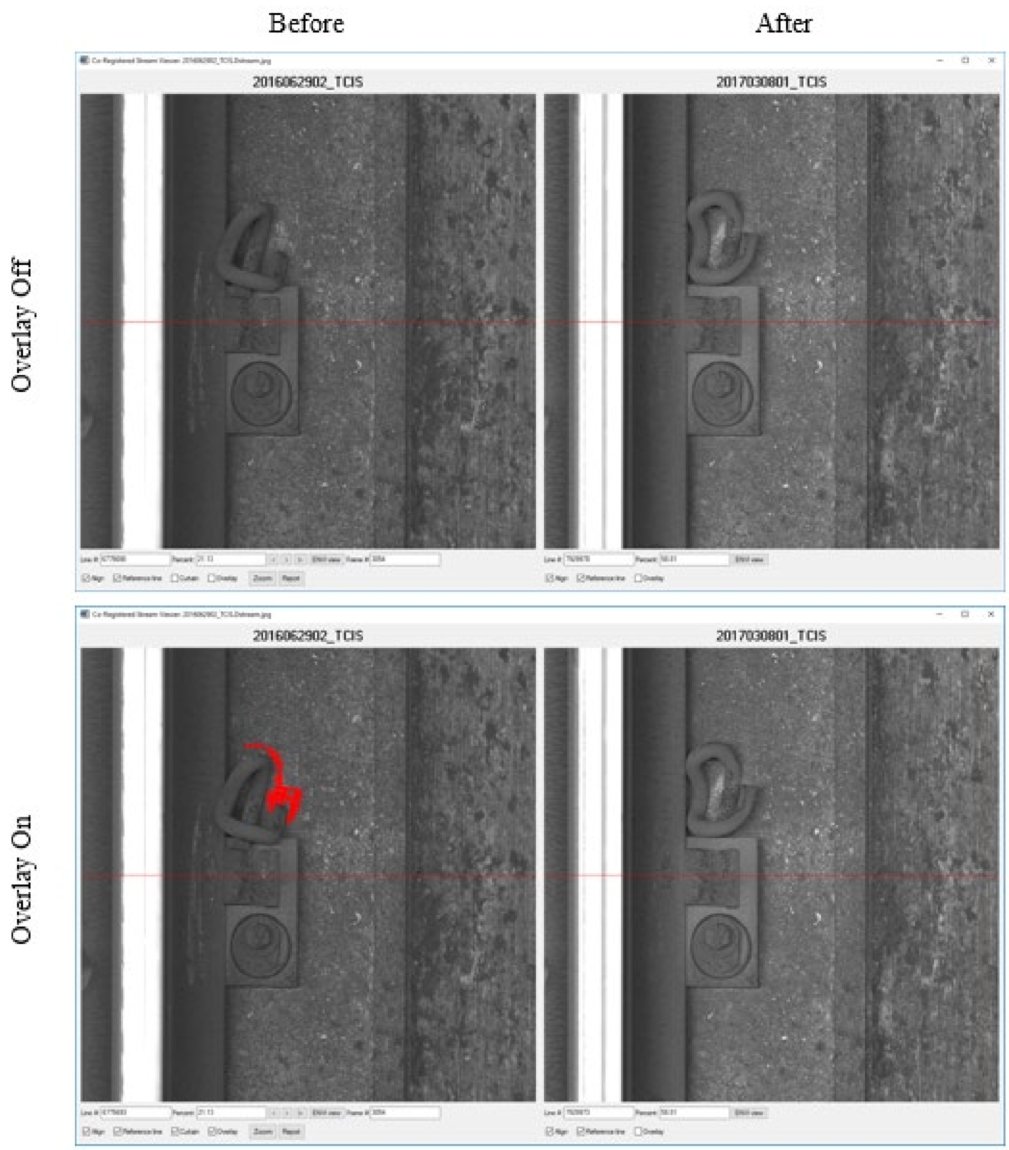

Figure 4-6. Intensity-Based Change Response for a Rotated Fastener

Figure 4-7 shows an example of a case where the evaluated software identified a rotated retainer clip used to hold a fastener plate to the underlying structure. This condition indicates that the corresponding bolt is loose—a state believed to precede a more serious missing bolt condition.

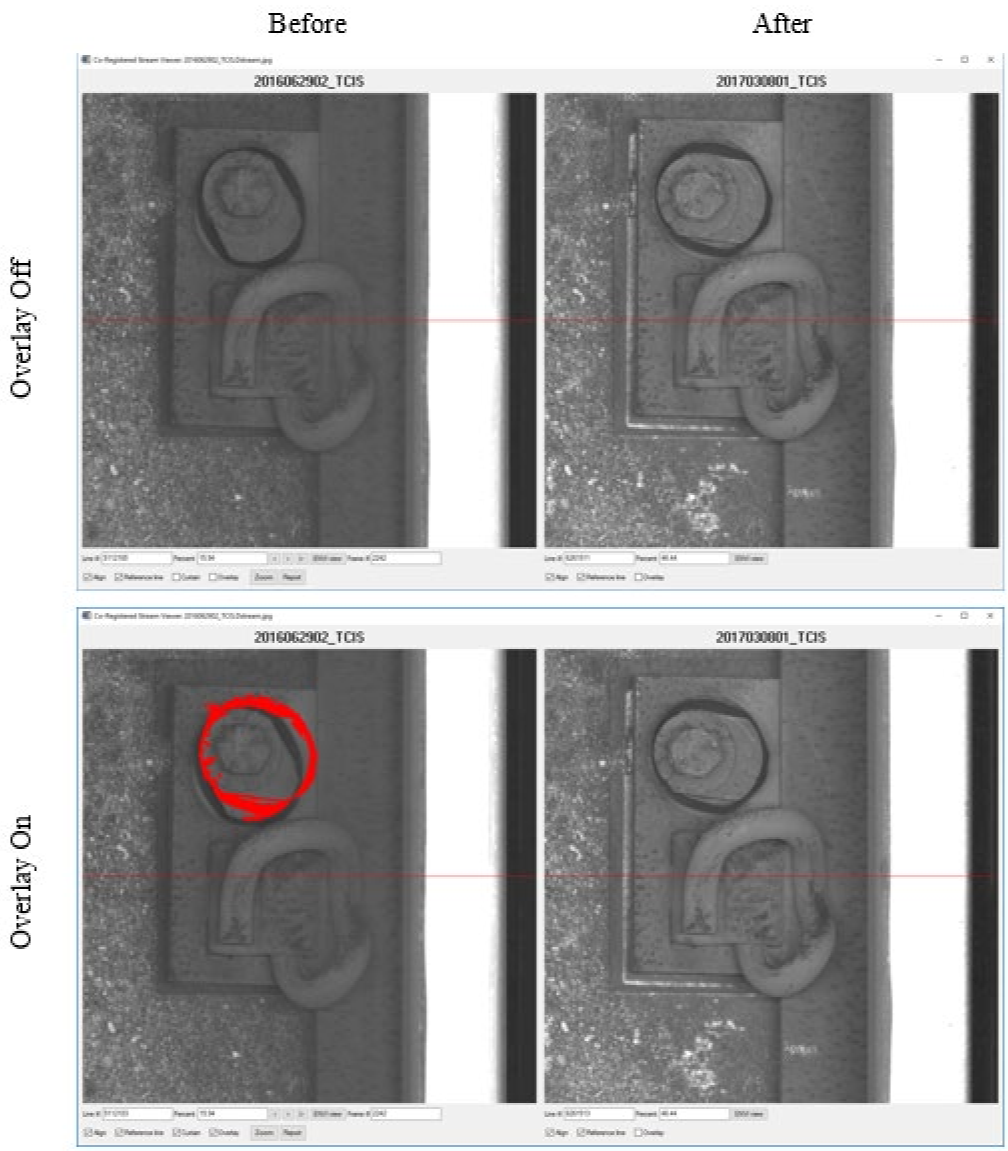

Figure 4-7. Intensity-Based Change Response for a Rotated Fastener Plate Retainer Clip

4.2.2 Crumbled Tie

The image sets did not include ties in a non-crumbled state in the "Before" image and a crumbled state in the "After" images. The change detection system successfully identified many less-pronounced changes in the state of an already-crumbled tie. Figure 4-8 and Figure 4-9 show an example of such a case. In Figure 4-8, the results are from applying intensity-based change detection, while in Figure 4-9, the results are from applying thematic-based change detection to the same crumbled tie incident.

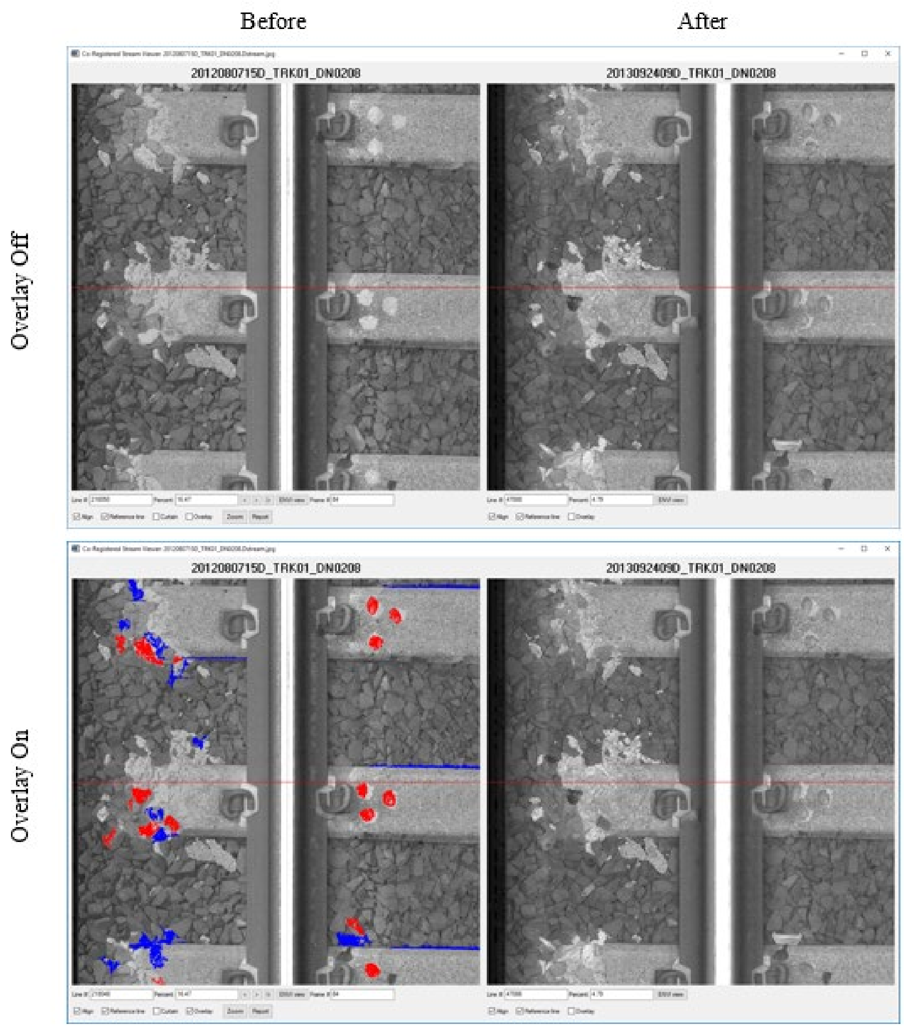

Figure 4-8. Intensity-Based Change Response at a Crumbled Tie Shoulder

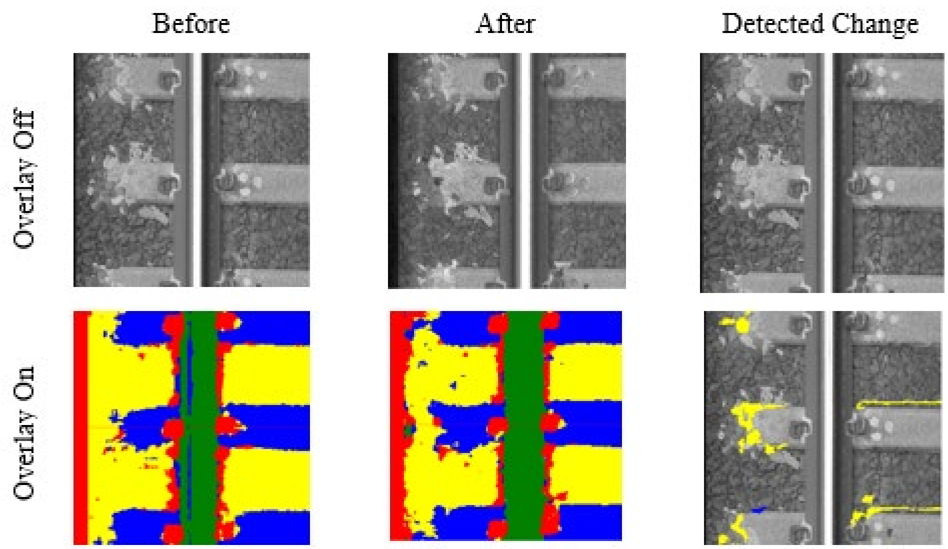

Figure 4-9. Thematic-Based Change Response at a Crumbled Tie

4.2.3 Fouled Ballast

The image sets did not include an example of fouled ballast. Instead, an example that looks roughly like some instances of fouled ballast is shown in Figure 4-10. In this example, new (brighter-shaded) ballast was added to the track near a region where a tie was replaced. The lighter-shaded ballast in this example resembles mud-stained ballast commonly associated with fouled ballast. While this example does not show actual fouled ballast, the indication is that fouled ballast would manifest as a detected change when applying intensity-based change detection.

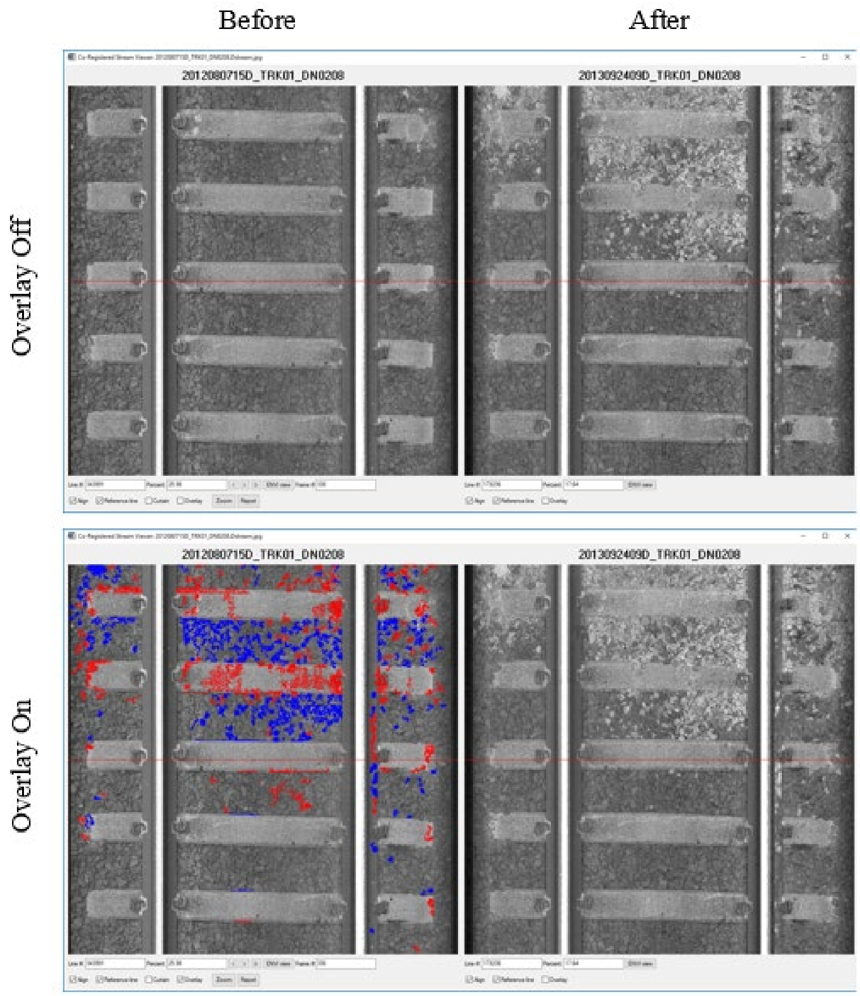

Figure 4-10. Intensity-Based Change Response for a Scenario that Resembles Fouled Ballast

4.2.4 Rail Surface Anomaly

Figure 4-11 shows an example of a detected change on the running surface of the rail. This example is intended to demonstrate the potential for change-based processing of actual rail surface images. Actual rail surface images depict the rail surface with significantly better image quality than that shown in this example.

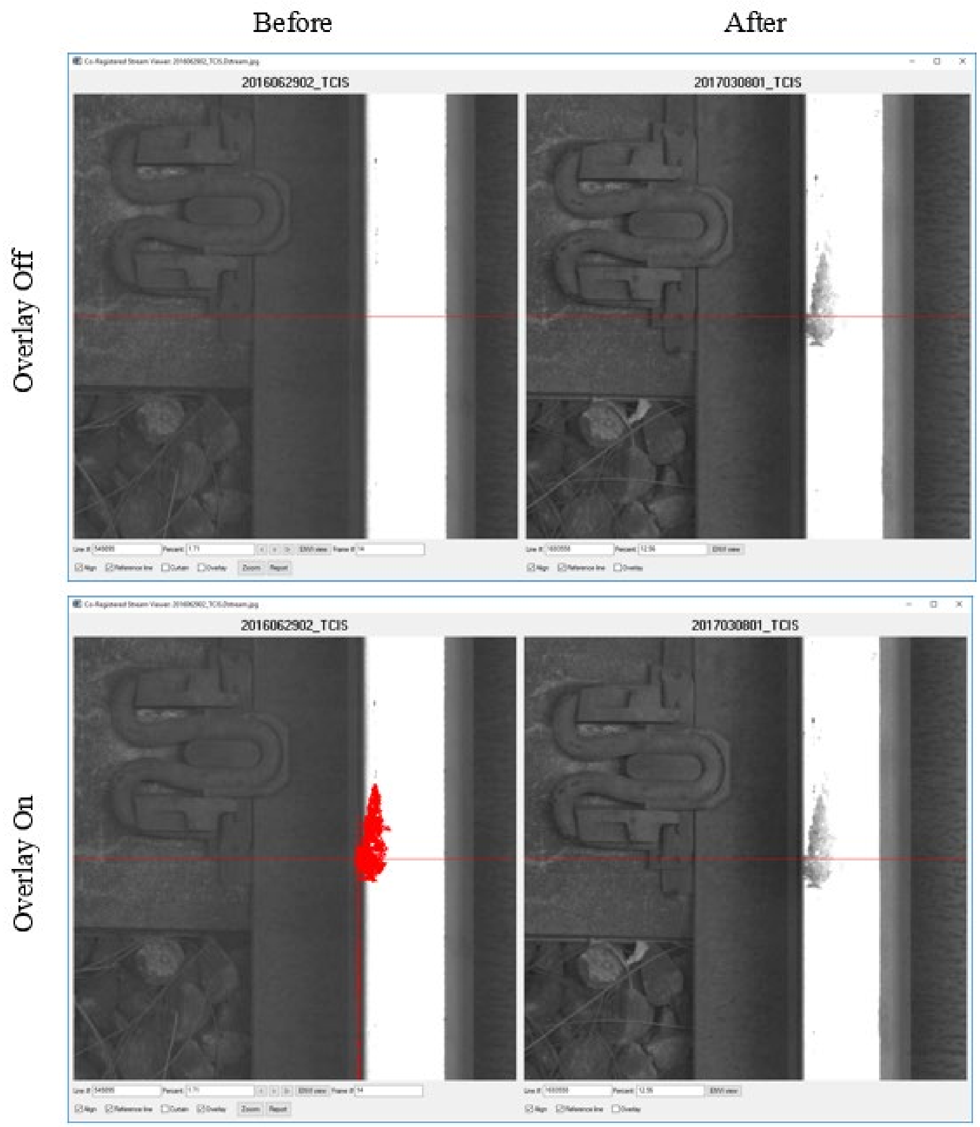

Figure 4-11. Intensity-Based Change Response for a Rail Surface Anomaly

In general, railroads are not interested in identifying all changed rail surface conditions. However, tracking the degradation rate of various infrastructure assets, such as the running rails, is a key component of proactive maintenance planning. Traditional machine vision approaches have been shown to be useful in identifying areas of rail surface degradation, however, tracking the rate of degradation means that an algorithm must also be able to differentiate between old and new degradation. Due to change-based processing identifying new areas of change, it is inherently well-suited for degradation studies.

Applying change detection to rail surface images may also automatically identify freshly ground rail locations. Combining image-based change detection with a machine vision algorithm that can identify rail grinding (already in existence today) can confirm the grinding plan.

4.2.5 Standing Water

Standing water in a track environment is of interest to railroads because it can accelerate deterioration of infrastructure components, particularly the rail base. Figure 4-12 shows an example of standing water identified as a change using intensity-based change detection. In general, railroads are interested in cases involving more significant amounts of standing water, whereas Figure 4-12 is intended to demonstrate that standing water is detectable using intensity-based change detection.

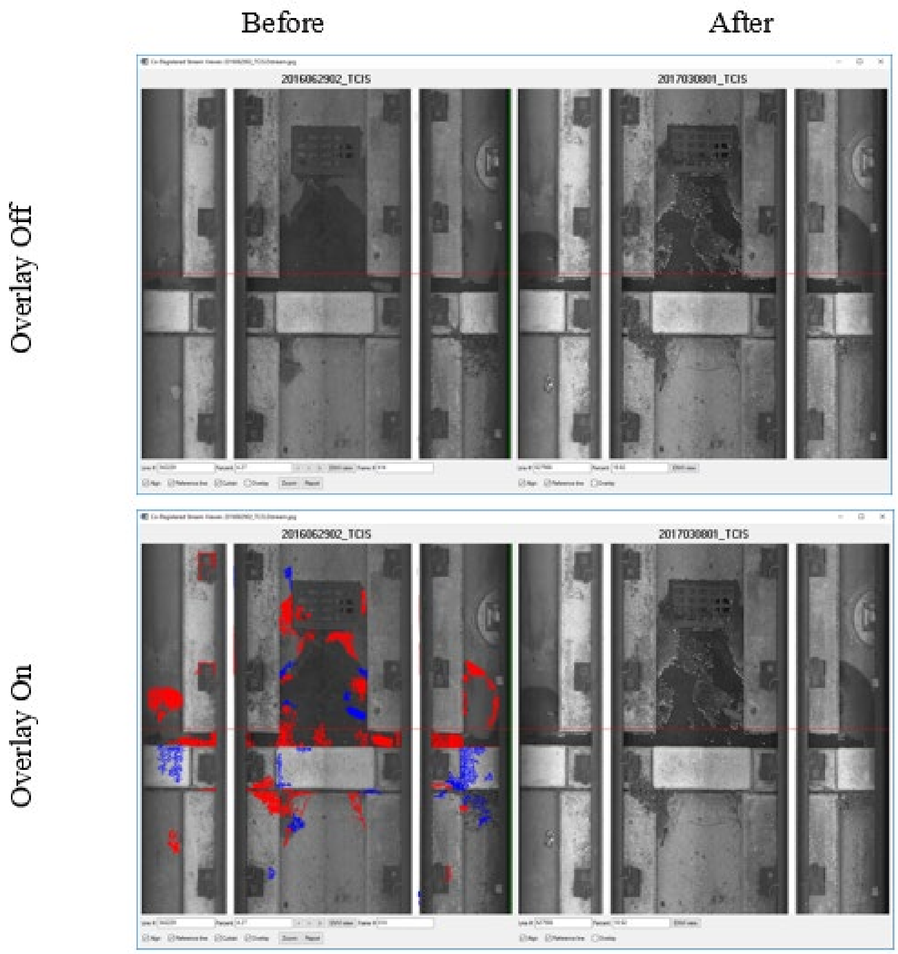

Figure 4-12. Intensity-Based Change Response for Standing Water

4.2.6 Maintenance Activity

The software detected small and large maintenance activities. This section includes examples of a few of the many detected changes. In general, identifying potentially unsafe conditions is regarded to be more important than identifying maintenance activity. Many of the maintenance activity examples in this section suggest that inverse changes corresponding to potentially unsafe conditions would also be detectable using change-based processing. For example, demonstrating an ability to detect a newly-installed third rail retainer clip also demonstrates an ability to detect the inverse scenario in which a third rail retainer clip is missing.

Localized Tie Replacement

Figure 4-13 shows an example of an intensity-based change response associated with replacing a concrete tie. In this study, the change response associated with localized tie replacement was consistently several orders of magnitude larger than other detected changes, as measured by the total number of changed pixels per frame.

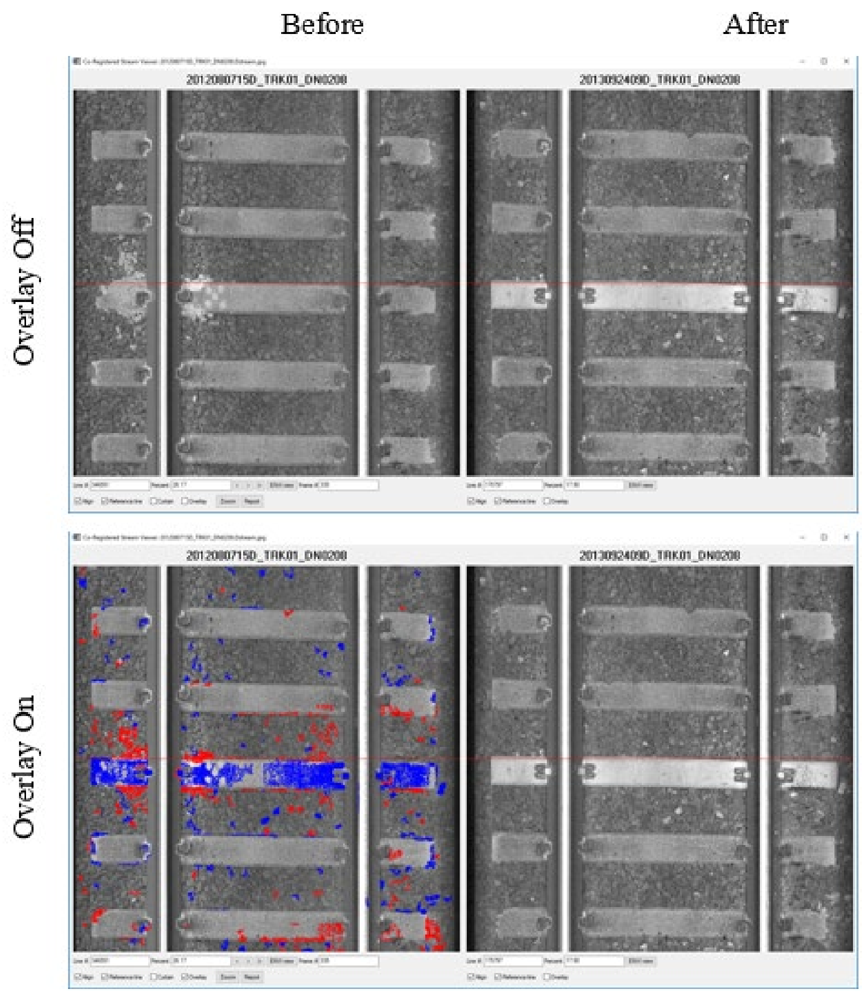

Figure 4-13. Intensity-Based Change Response for a Replaced Concrete Tie

Replaced Fastening Element

Figure 4-14 shows an example of a change response associated with replacing a rail fastener in conjunction with the underlying tie. Many examples like this were found in the evaluated data—all were properly detected. Section 5.3 describes a region-based approach to change detection that would allow changes such as replaced fasteners to be automatically labeled with a corresponding assed type (i.e., fastener). Among other capabilities, automatically labeling changes with a corresponding asset type would help enable automated maintenance confirmation.

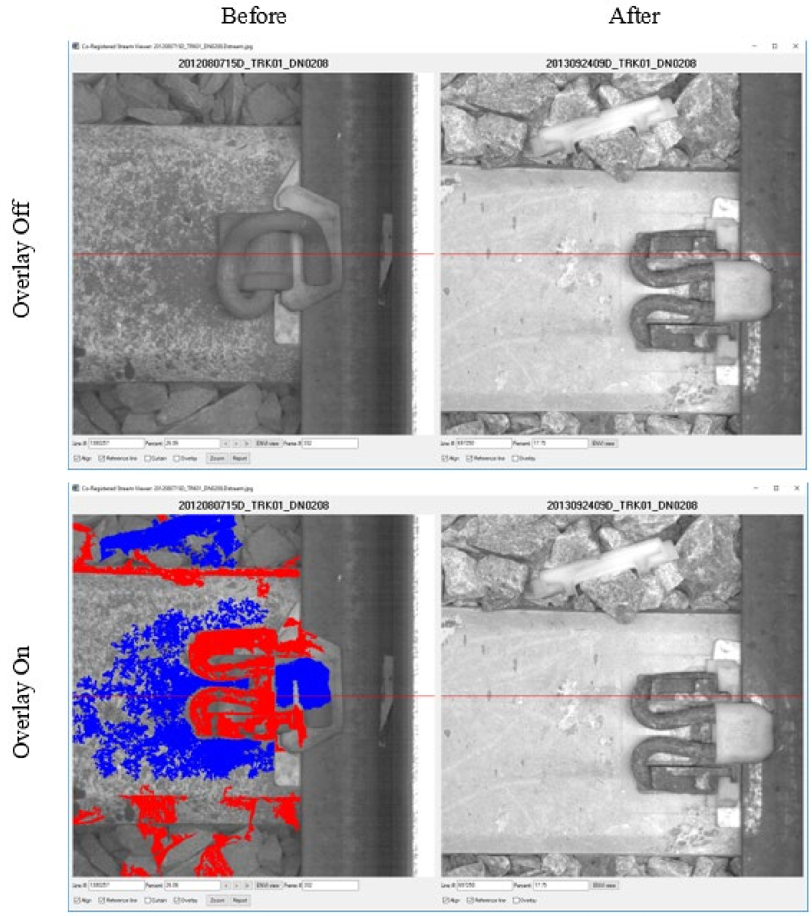

Figure 4-14. Intensity-Based Change Response for a Replaced Fastener at a Replaced Concrete Tie

Replaced Third Rail Stand

Figure 4-15 shows an example of a change response associated with replacing a third rail pot and stand. Many examples like this were properly detected by the evaluated software. In this example, several left over track components are also reported as a change.

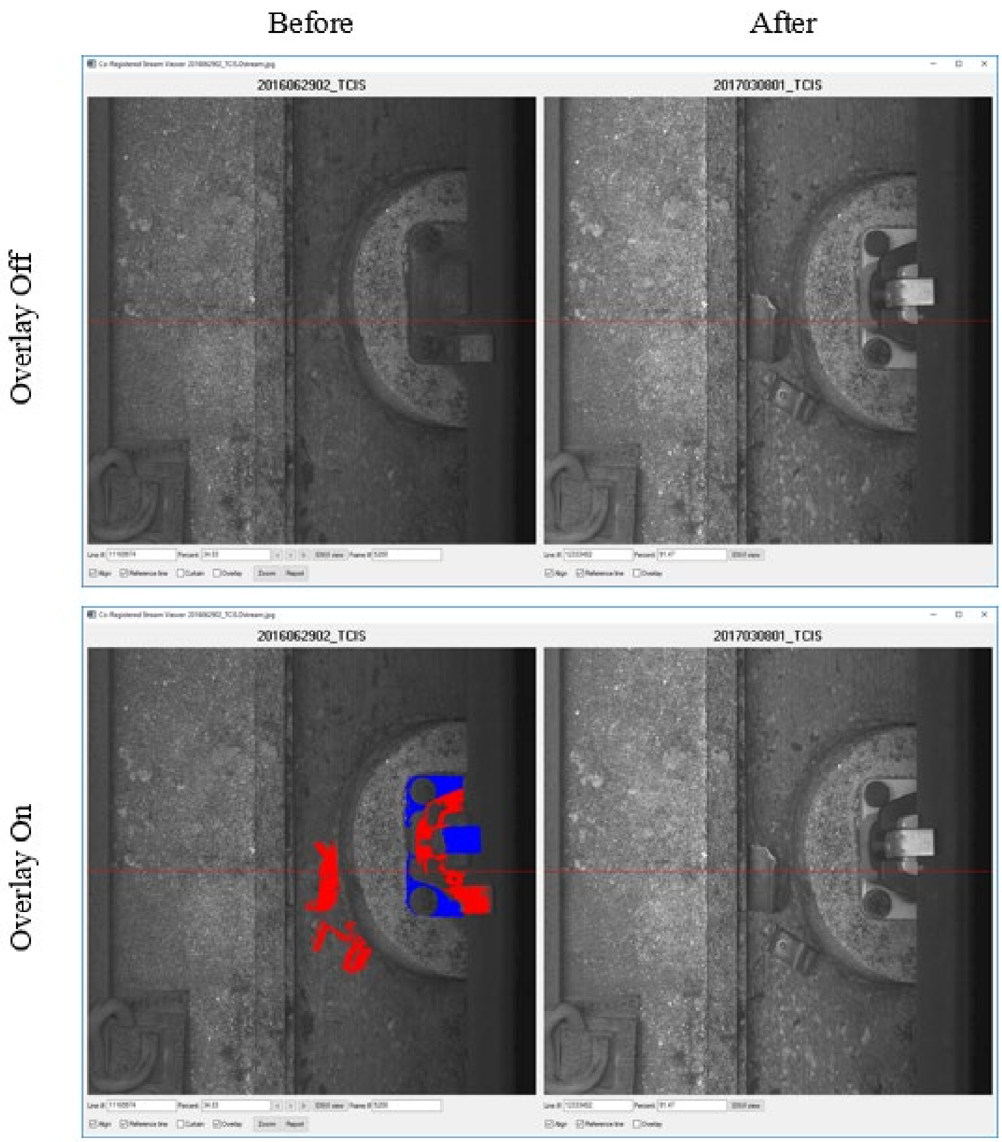

Figure 4-15. Intensity-Based Change Response for a Replaced Third Rail Stand

New Third Rail Retainer Clip

Figure 4-16 shows one of many examples of a change associated with installing a new retainer clip on a third rail. In this example, the base of the third rail is shown running vertically along the right edge of the image. The region corresponding to the newly-installed retainer clip is highlighted in blue in the "Before" image of the "Overlay On" row.

As with the previous three examples, this example demonstrates the potential to use change detection to track maintenance activity. This example also demonstrates that a corresponding inverse change (a missing third rail retainer) would also be detectable.

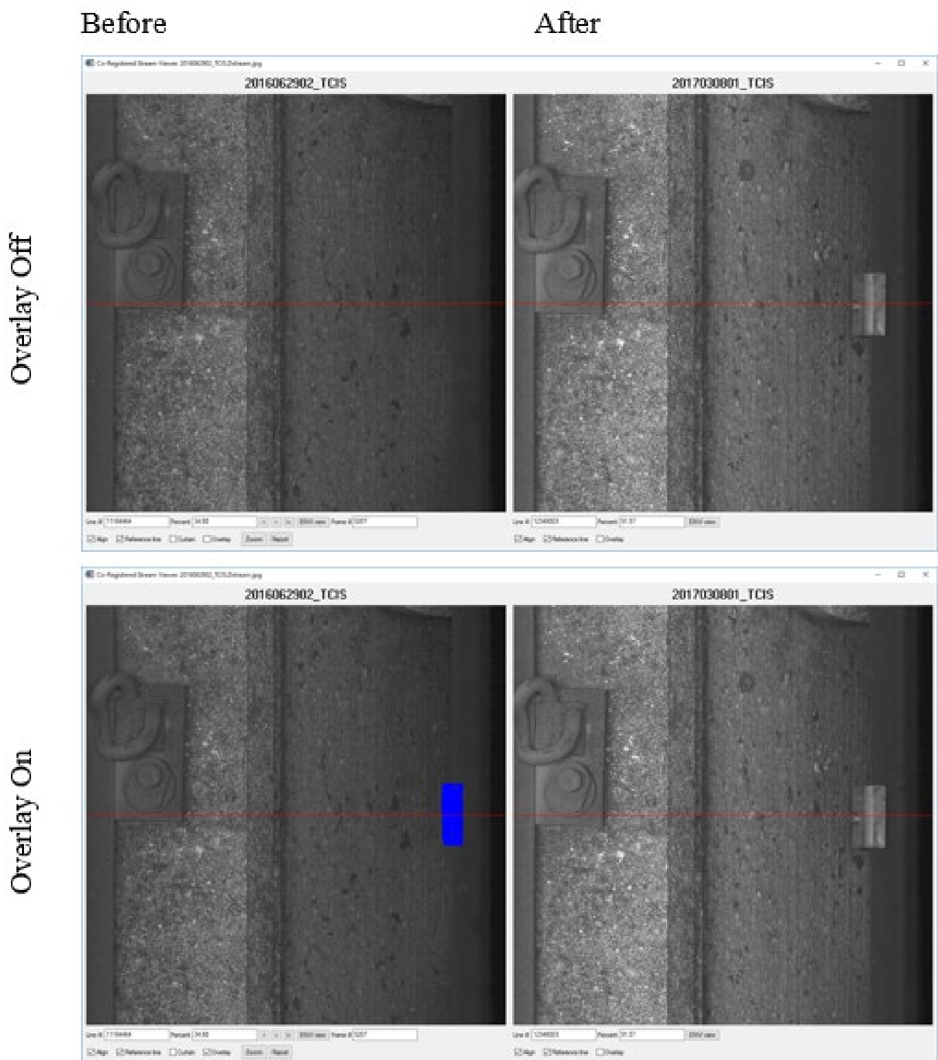

Figure 4-16. Intensity-Based Change Response for a New Third Rail Retainer Clip

4.3 Missed Changes

A missed change is a region in an image where change detection is unsuccessful at highlighting a difference between two compared images, given that the difference is perceptible, based on normal (unaided) observation.

4.3.1 Reasons for Missed Changes

This section explains why intensity-based change detection misses some changes that are otherwise perceptible based on observation and clarifies an inherent relationship between missed changes and non-relevant changes.

The intensity-based change detection algorithm applied in this study detects changes as small as a single pixel; however, it only reports changes that impact a contiguous region larger than a configurable number of pixels. The detected changes in this study are based on a minimum size threshold of 10,000 pixels, which corresponds to an area of roughly 15 square centimeters (15 cm^2 or 2.3 in^2). Similarly, the assessed change detection algorithm can detect changes in intensity as small as a fraction of a percent, but it only counts a pixel as a change when the corresponding intensity difference is greater than a configurable percentage (e.g., 10 percent). How these configurable limits lead to missed changes is further discussed in conjunction with sample results.

In Figure 4-17, several grease-like deposits are present on both the rail base and along one edge of the depicted tie shoulder. The largest grease-like deposit on the rail base in the "Before" image is not present in the "After" image and is, therefore, detected as a change. The smaller spot enclosed in a green box in the "After" image of Figure 4-17 is not present in the "Before" image. Although the smaller spot corresponds to a perceptible change, it is not detected as a change by the software. This is because the size of the changed region inside the green box is below the 10,000-pixel size threshold selected for this study.

The grease-like deposit on the edge of the tie shoulder enclosed in a yellow box in Figure 4-17 is an observably large change based on perception. In this case, the change would appear to be large enough to trigger a reported change; however, because intensity and size thresholds are applied concurrently by the software, no change is reported. To clarify, the overall changed region (based on perception) becomes divided into many sub-regions when intensity thresholding is applied by the software. In this example, no resulting sub-region is large enough to trigger a reported change.

Due to the changes described above associated with grease are relatively unimportant, it is beneficial that the software did not report all the changes. In cases like these, size and intensity thresholds help regulate the number of non-relevant changes reported by the software. The next section clarifies how size and intensity thresholds can also lead to missed, relevant changes.

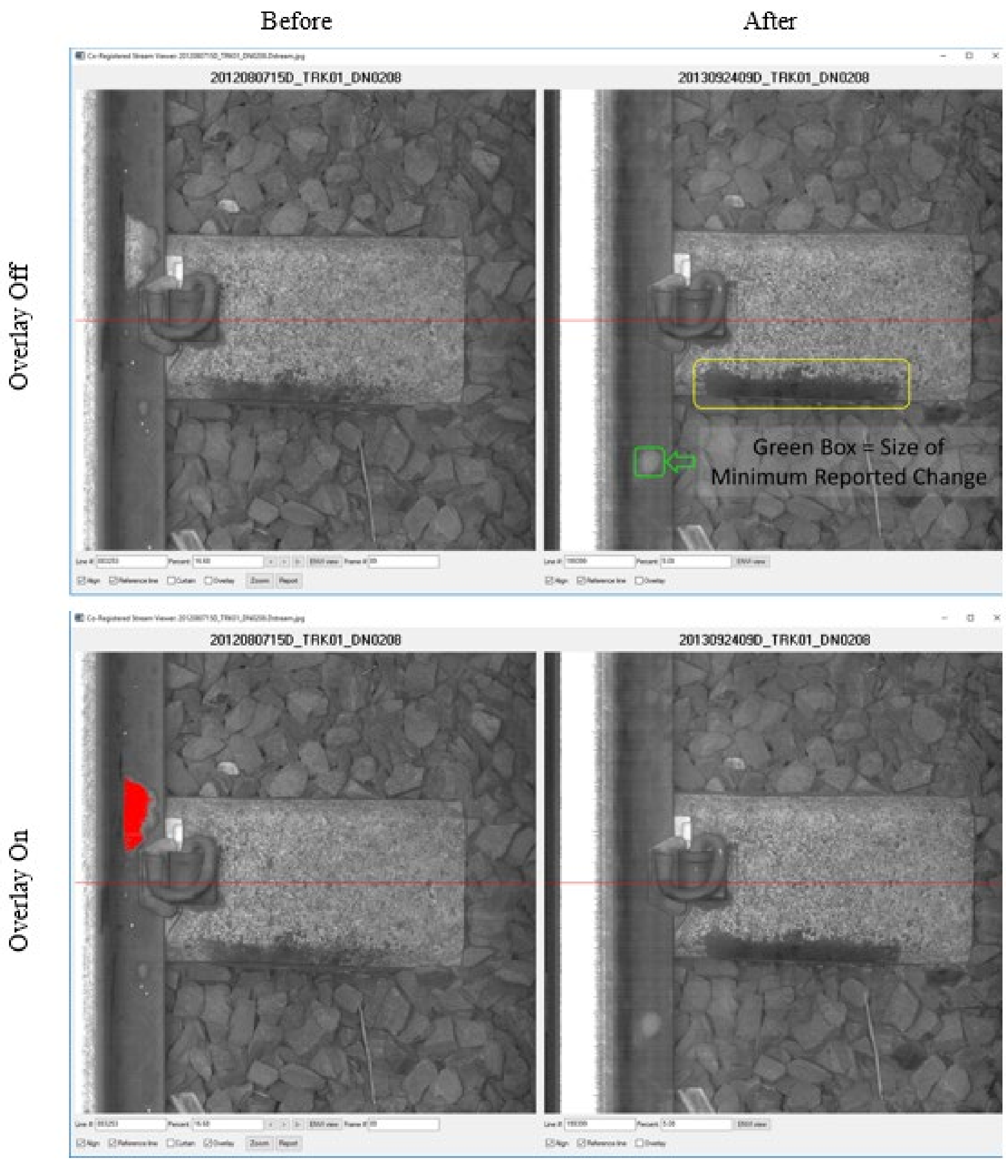

Figure 4-17. Example Used to Clarify Missed Intensity-Based Changes

4.3.2 Example of a Relevant Missed Change

The size threshold used (10,000 pixels) to reject small changes was selected to achieve a reasonable balance between detecting too many non-relevant changes and missing too many larger, relevant changes. Some relevant changes were not detected. An example of one is shown in Figure 4-18. In Figure 4-18, the rotated retainer clip inside the green box is reported as a change (see red area in the “Overlay On” image). In contrast, the rotated retainer clip inside the yellow box corresponds to a missed, relevant change (note the corresponding absence of red pixels in the “Overlay On” image).

The change inside the green box was detected because it manifests as a contiguous region larger than the minimum size threshold (10,000 pixels) described above. Conversely, the missed change inside the yellow box did not manifest as a large enough contiguous region and was, therefore, not reported as a change. Thus, although the underlying condition-of-interest (a rotated retainer clip) is the same in both cases, one case was detected and the other was not. Section 5.3.3 provides a more in-depth analysis of missed changes and demonstrates how applying a region-based change detection approach improves system performance.

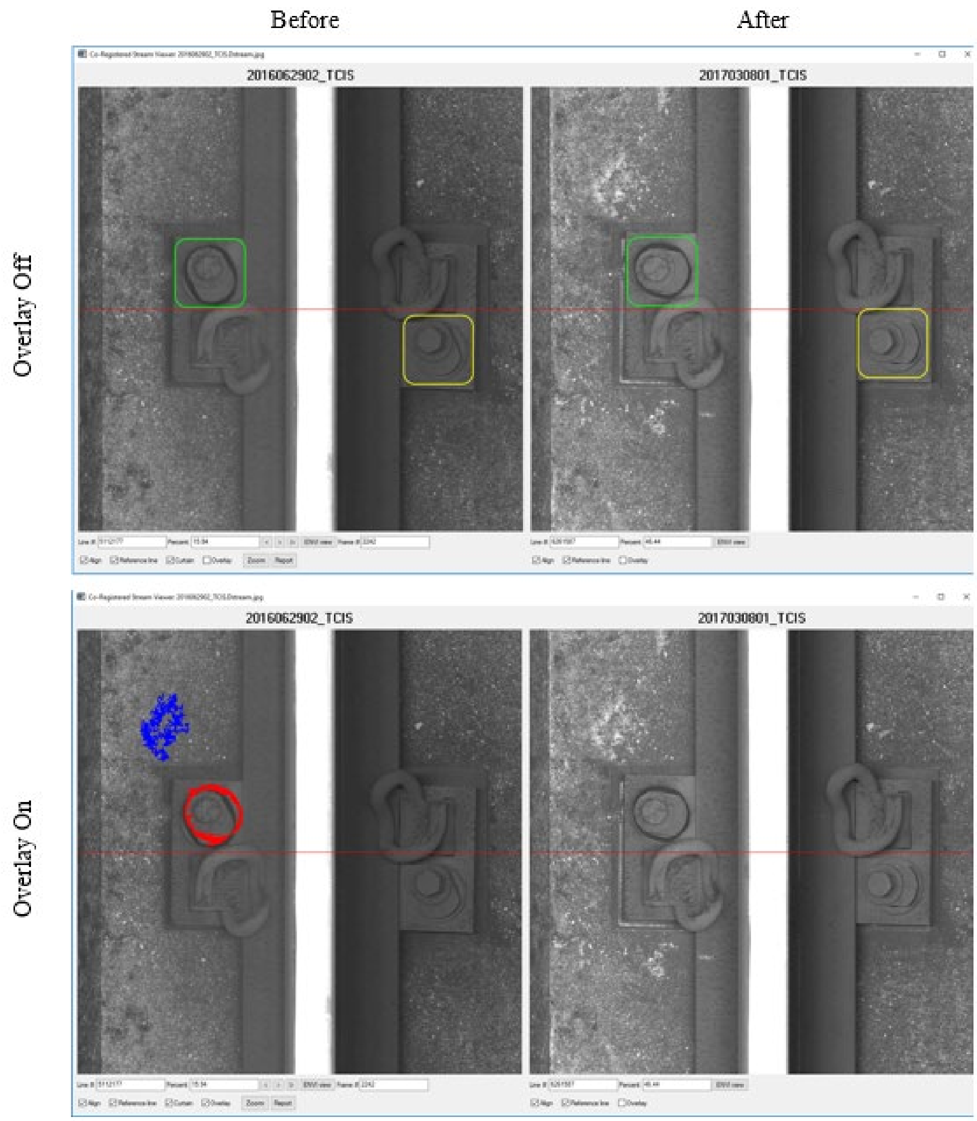

Figure 4-18. Example of a Relevant Intensity-Based Change Missed by the Evaluated Software

4.4 Non-Relevant Changes

A non-relevant change is a region where a difference is reported that is not considered important with respect to achieving a specified objective.

When applying image-based change detection to rail based images, non-relevant changes have the same undesirable impact as false alarms associated with applying conventional machine vision algorithms. That is, both non-relevant changes and false alarms increase the need for manual image review.

The following two parameters help provide insight when assessing the potential impact of non-relevant changes:

1) Static-frame percentage – The percentage of the total number of frames of data where no change was detected

2) Static-pixel percentage – The percentage of the total number of pixels in a data set where no change was detected

The static-frame percentage provides insight into whether detected changes are uniformly or non-uniformly distributed throughout a data set. The static-pixel percentage provides insight into how dense detected changes are within an average frame. Table 4-1 summarizes the static-frame and static-pixel percentages corresponding to the image sets evaluated in this study.

Table 4-1. Static-Frame and Static-Pixel Results for Each Evaluated Data Set

			Static-Frame Percentage		Static-Pixel Percentage	
Data Set	**Number of Frames in Data Set**	**Number of Pixels in Data Set ($x10^9$)**	**Intensity-Based Processing**	**Thematic-Based Processing**	**Intensity-Based Processing**	**Thematic-Based Processing[3]**
2012–2013	1902	31.9%	0%	0%	98.4%	N/A
2016–2017	5763	96.7%	16%	0%	98.0%	N/A

The low static-frame percentages shown in Table 4-1 indicate a high probability that any given frame will include a reported region of change. As nearly every image frame includes at least one detected region of change, managing non-relevant changes is expected to be an important aspect associated with any deployment of image-based change detection in the rail sector.

The high static-pixel percentages shown in Table 4-1 indicate a low probability that any given pixel within a frame will exhibit change. A high static-pixel percentage is favorable in the sense that it means spatial filtering is a viable option to help manage non-relevant changes.

Non-relevant changes can be either true or false. Non-relevant, true changes correspond to changes that are real but not important. Non-relevant, false changes show up in the results, but they do not correspond to a change in the real world; they are introduced by non-perfect image processing. The following two subsections present examples of true and false non-relevant changes encountered during this study.

4.4.1 Non-Relevant True Changes

A non-relevant true change is a detected change where the reported difference is real (true) but unimportant within the context of achieving a designated objective. Non-relevant, true changes associated with these data sets generally fall into three categories:

1) Surface discoloration (e.g., typically caused by grease, rail dust, or moisture)

2) Isolated changes in ballast
3) Trash and natural debris

Non-relevant, true changes associated with thematic-based change detection were not assessed during this study because they are currently intermixed with a significant number of false, thematic-based changes. Thus, until the number of false, thematic-based changes is reduced, there is no benefit in assessing non-relevant, true changes associated with thematic-based change detection. For reference, Section 5.1.3 presents results demonstrating how false thematic-based changes can be reduced.

Table 4-2 summarizes the percentage of frames found to include at least one, non-relevant, true change within the intensity-based results. The high percentages in Table 4-2 are another reason that managing non-relevant, true changes is expected to be important if change-based processing is used in the rail sector.

Table 4-2. Percentage of Frames with at Least One, Non-Relevant, True Change

Data Set	Track Construction	Percentage of Frames with at Least One, Non-Relevant, True Change[3]
2012–2013	Concrete tie	93%
2016–2017	Concrete slab	68%

The remainder of this section presents representative examples of non-relevant, true changes observed during this study. Related discussion clarifies potential options for managing each case.

Isolated Changes in Ballast

Figure 4-19 and Figure 4-20 show typical examples of isolated changes in ballast. Such changes are typically characterized by a relatively small region where ballast stones were disturbed. Isolated changes in ballast were the most common source of non-relevant, true change noted in this study. Although common within the evaluated data sets, isolated changes in ballast are not ultimately expected to be problematic. Table 4-3 summarizes strategies expected to be useful in managing isolated changes in ballast.

[3] Percentages shown here are based on evaluating 500 consecutive image frames from each image set, starting from a randomly chosen frame.

Table 4-3. Strategies for Managing Isolated Changes in Ballast

Applicable Management Strategy	Justification
Classify and Ignore – Section 5.2	Ballast is a common texture found in rail-based images making it possible to classify and automatically ignore changes in this category
Region-Based Change Detection – Section 5.3	Isolated changes in ballast can be spatially filtered relative to many other relevant changes

In addition to the management strategies listed above, it is anticipated that three-dimensional track profile data would be useful in filtering isolated changes in ballast when the corresponding ballast is sitting on a tie. The capability to isolate (and ignore) changes in ballast on a tie is an important aspect of any change-based processing intended to isolate degradation associated with ties.

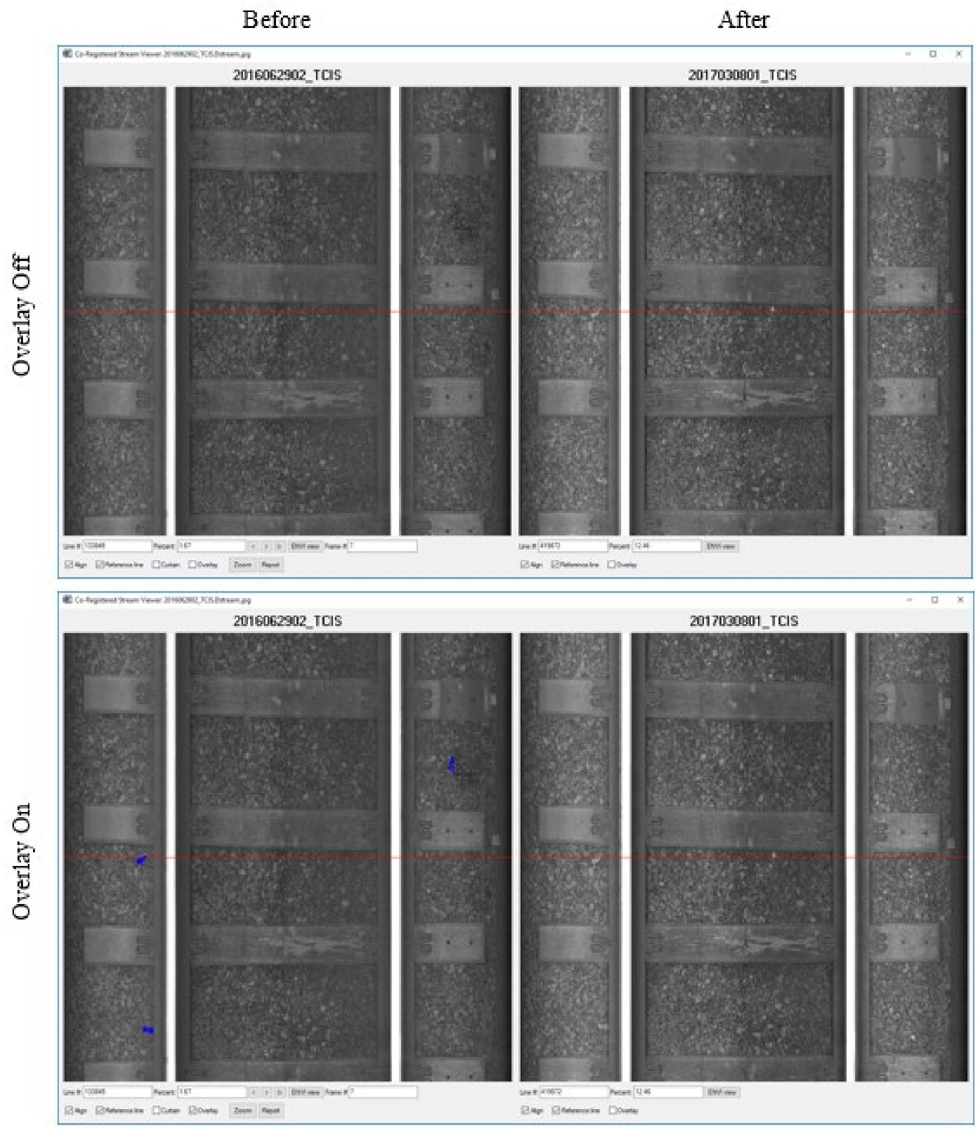

Figure 4-19. Example 1 of Non-Relevant, True Changes Associated with Ballast

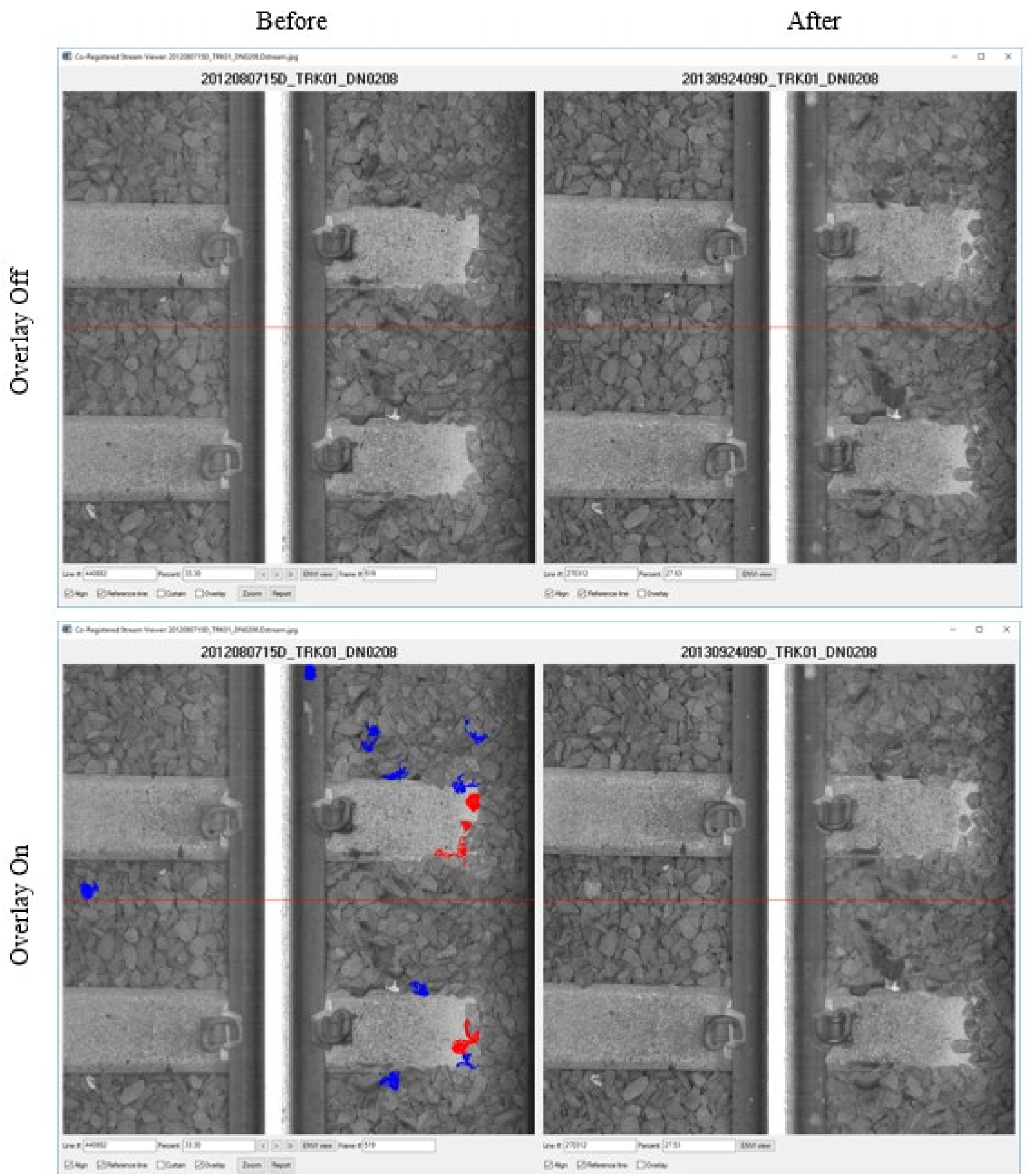

Figure 4-20. Example 2 of Non-Relevant, True Changes Associated with Ballast

Surface Discoloration

Figure 4-21 through Figure 4-23 show examples of surface discoloration caused by grease, rail dust, and moisture, in that order. Changes in this category were the second-most prevalent source of non-relevant, true changes. Table 4-4 summarizes strategies expected to be useful in managing non-relevant changes related to surface discoloration.

Table 4-4. Strategies for Managing Non-Relevant Changes from Surface Discoloration

Applicable Management Strategy	Justification
Classify and Ignore – [Section 5.2]	Some changes in this category are expected to be common enough to enable automatic classification so that the changes can be ignored
Region-Based Change Detection – [Section 5.3]	Many changes in this category can be spatially filtered to reduce their impact on detecting other relevant changes

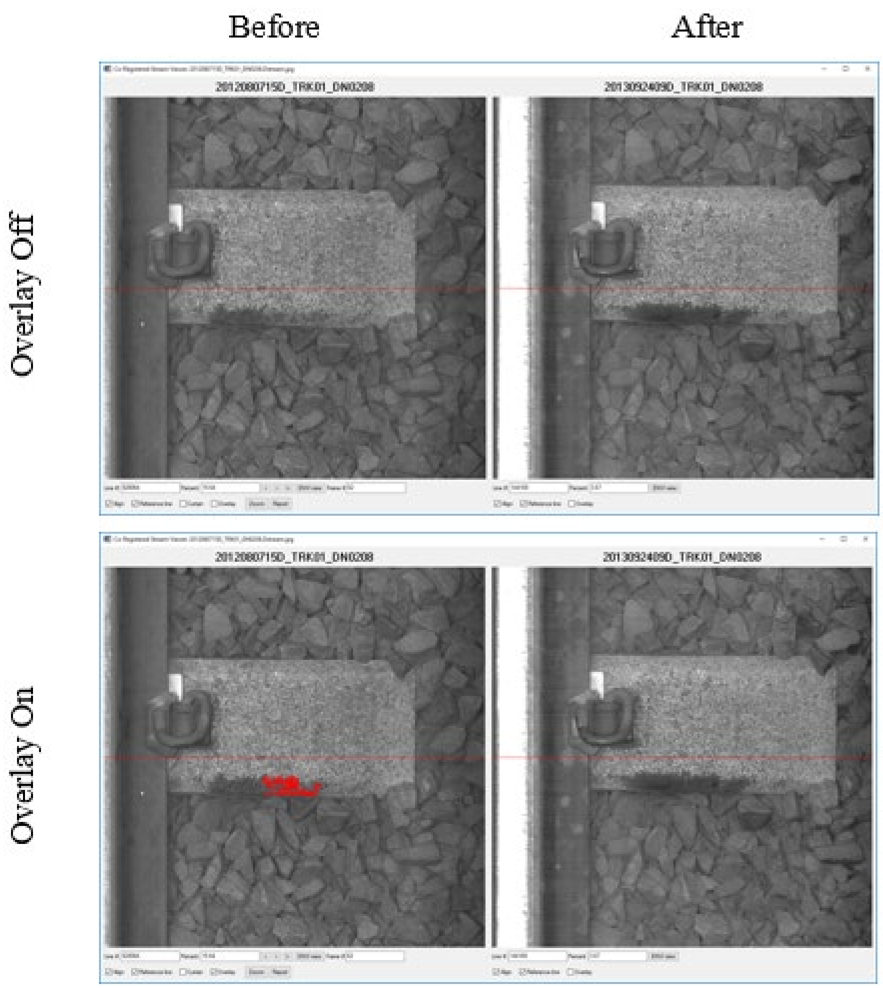

Figure 4-21. Example of a Non-relevant, True Change Caused by Rail Grease

Figure 4-22 shows cases where a grease spot on the rail base and build up from rail dust on a tie clip cause non-relevant, true changes. Although the white spot on the rail base in this example resembles white paint, the spot is more likely fresh grease. Fresh grease is known to make the rail base shiny, which tends to manifest as a bright spot in the line scan image due to high-intensity lighting.

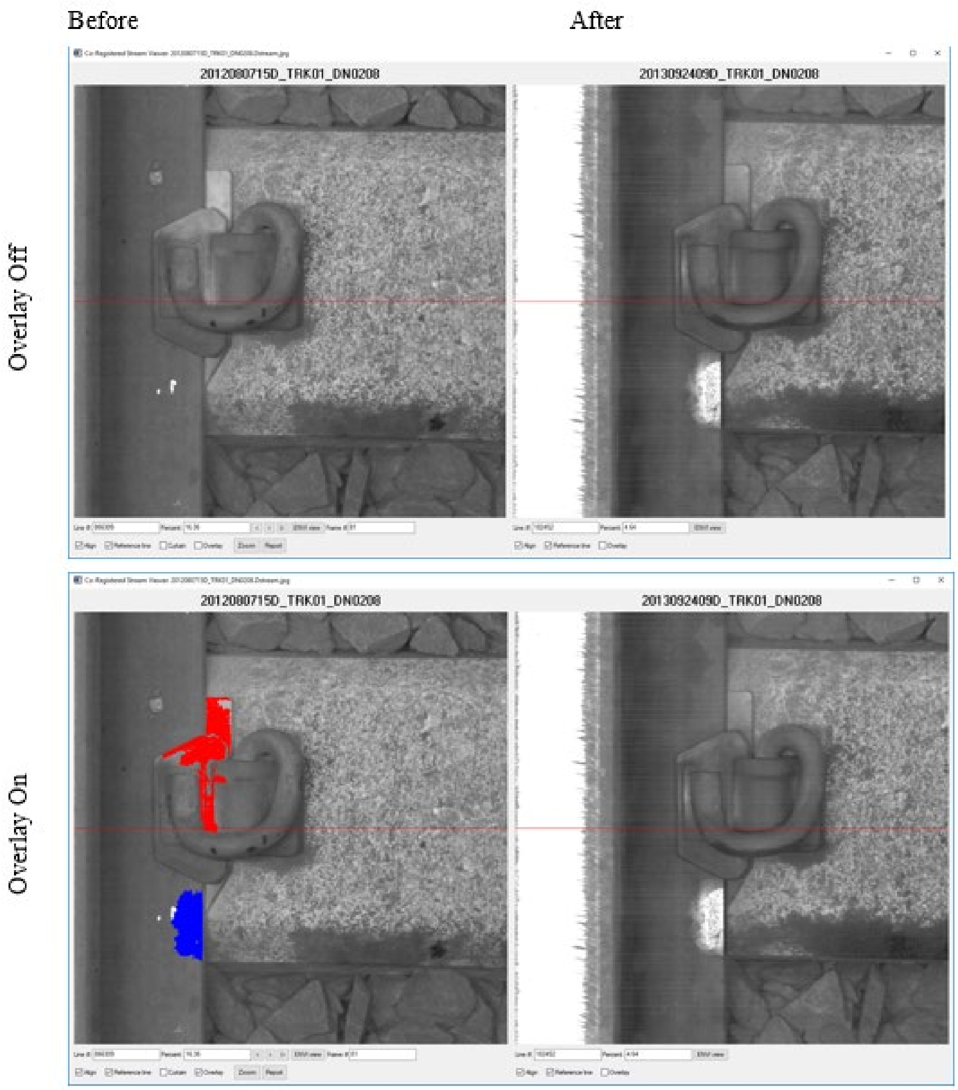

Figure 4-22. Example of a Non-Relevant, True Change Caused by Grease and Rail Dust

Figure 4-23 shows an example where moisture caused a non-relevant, true change. Non-relevant, true changes caused by moisture are not expected to be common under dry track conditions but could, potentially be problematic following rain. It is expected that image-based change detection would not be useful if wet track is compared to dry track; however, it remains to be seen to what degree wet/damp track could be compared to wet/damp track.

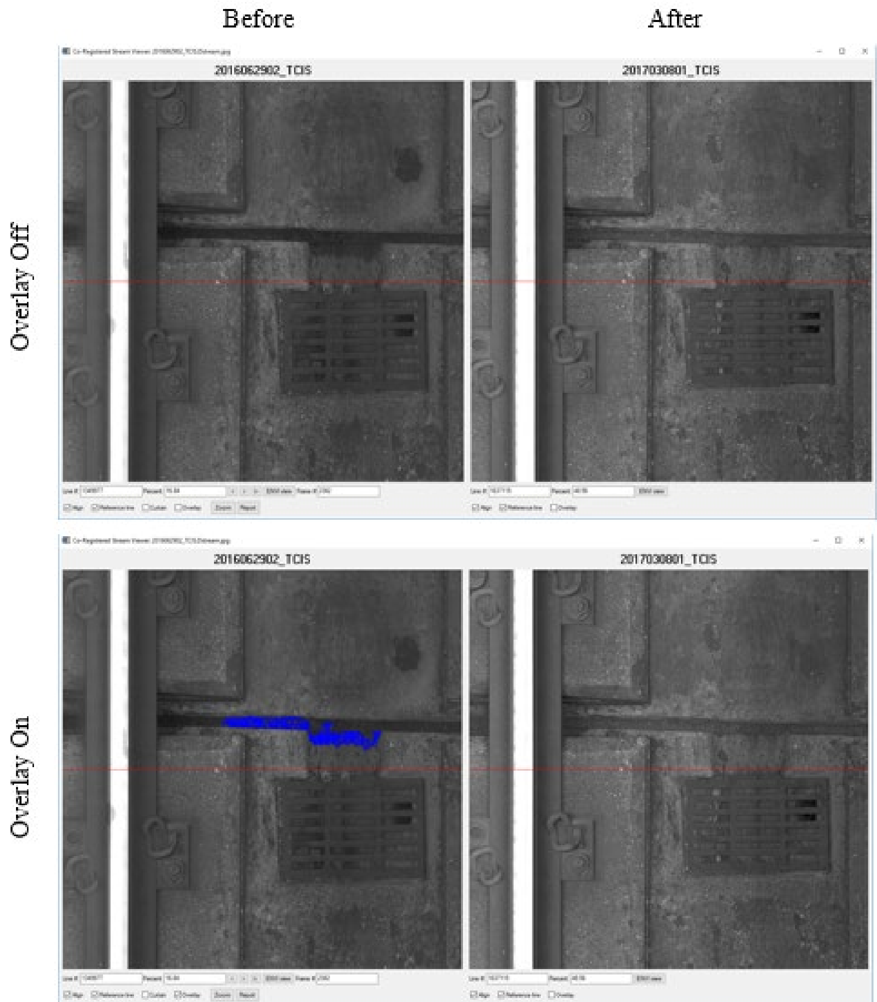

Figure 4-23. Example of a Non-Relevant, True Change Caused by Moisture

Trash and Natural Debris

Trash and natural debris were the third most prevalent source of non-relevant, true changes. Examples include bottles, left over track components, pine needles, and sticks, as shown in Figure 4-24 through Figure 4-28. Table 4-5 summarizes strategies expected to be useful in managing non-relevant changes related to trash and natural debris.

Table 4-5. Strategies for Managing Non-Relevant Changes from Trash and Natural Debris

Applicable Management Strategy	Justification
Classify and Ignore – Section 5.2	Some changes in this category (e.g., pine needles, leaves, etc.) are expected to be common enough to enable automatic classification so that the changes can be ignored
Region-Based Change Detection – Section 5.3	Most changes in this category can be spatially filtered to reduce their impact on detecting other relevant changes

In addition to the management strategies listed above, it is anticipated that three-dimensional track profile data would be useful in filtering some non-relevant changes in this category (e.g., cans, bottles, and other sizeable, three-dimensional forms of trash and natural debris).

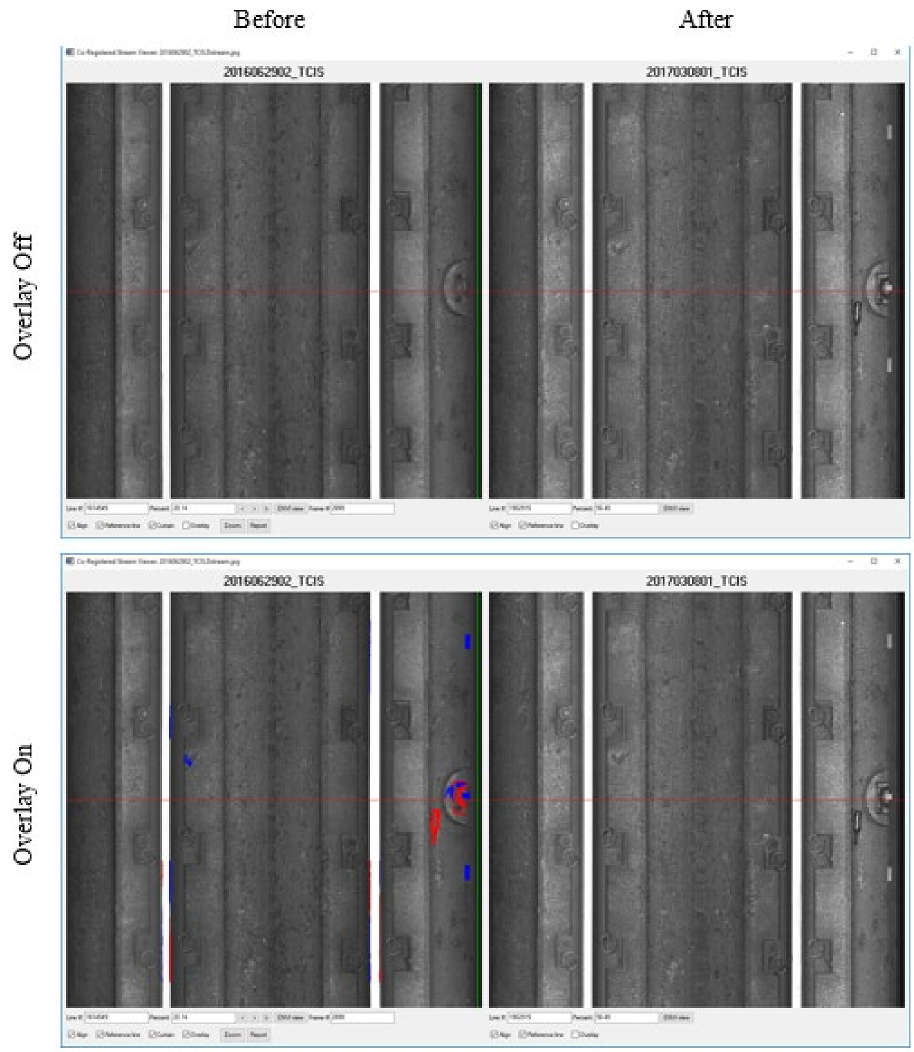

Figure 4-24. Example of a Non-Relevant, True Change Caused by Trash (a Bottle)

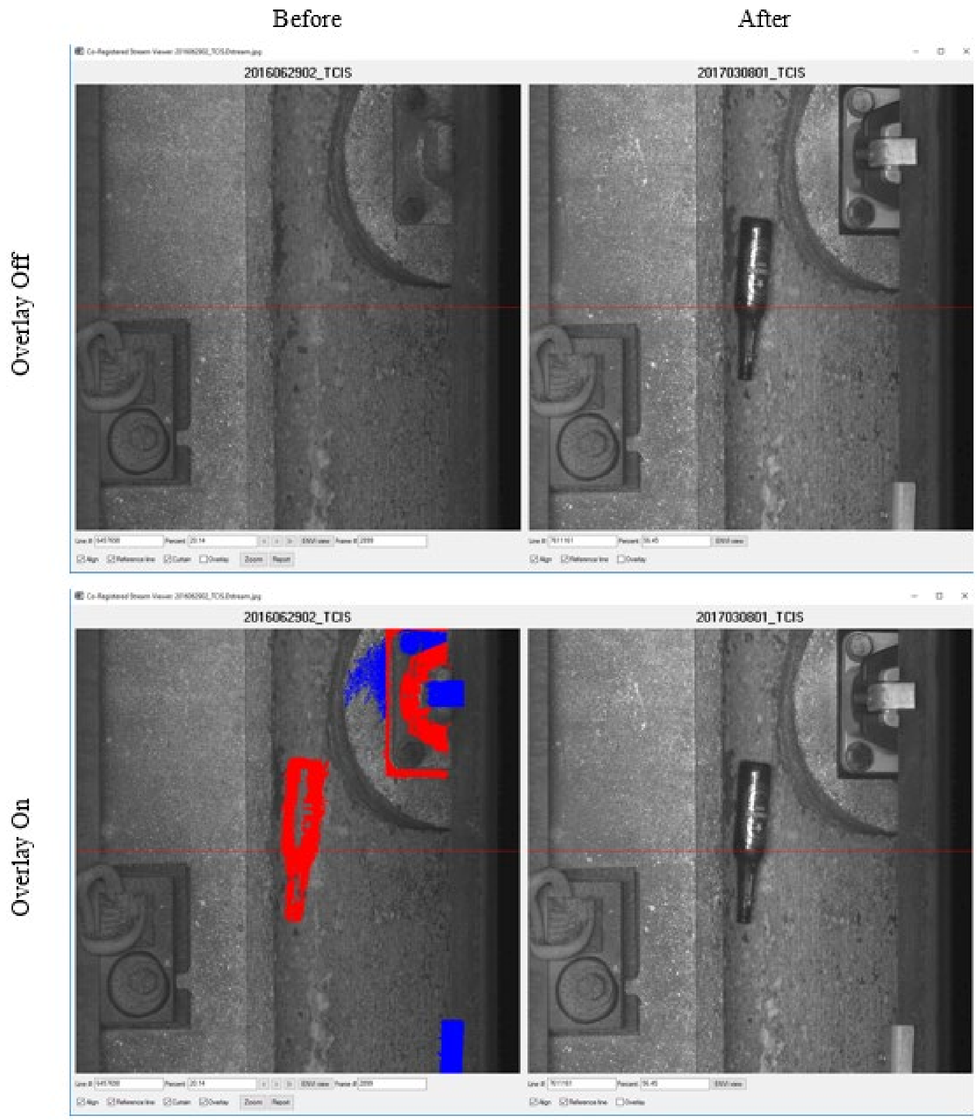

Figure 4-25. Close-up View of a Non-Relevant, True Change Caused by Trash (a Bottle)

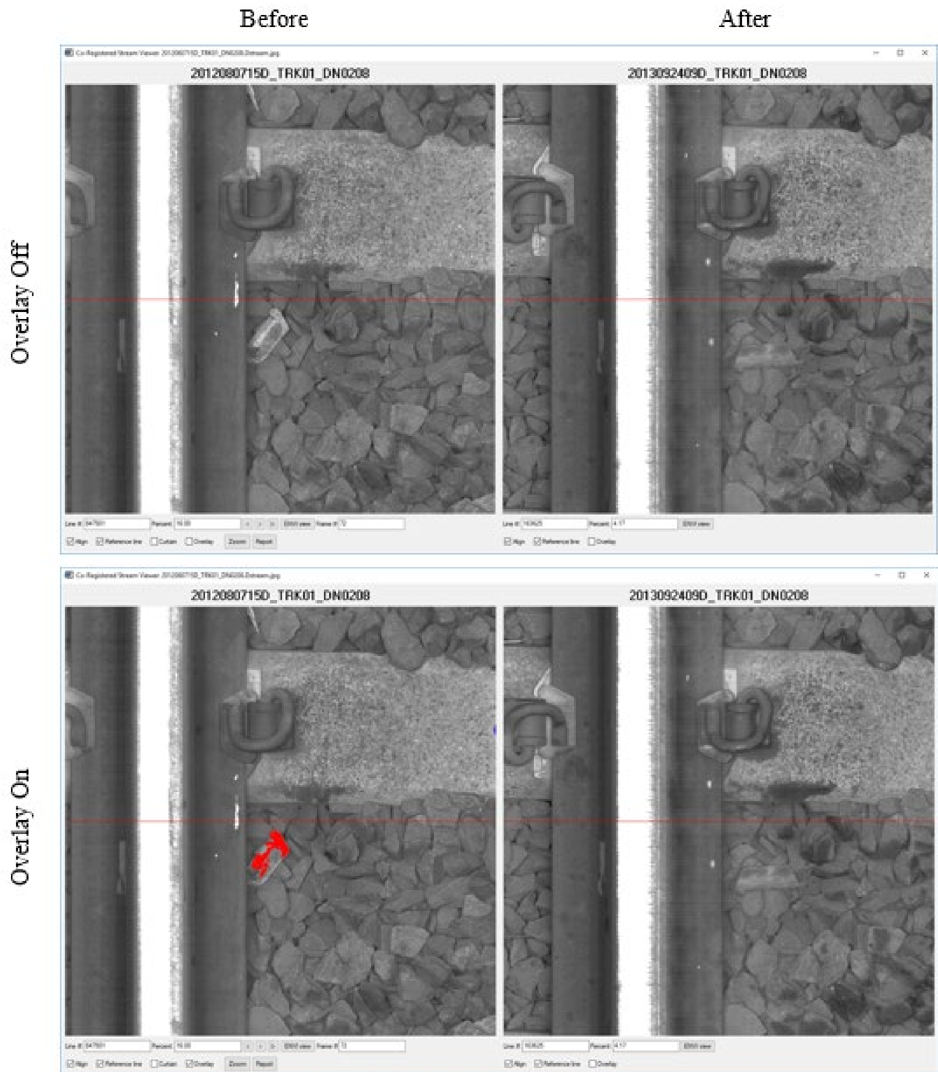

Figure 4-26. Example of a Non-Relevant, True Change Resulting from a Left-Over Track Component

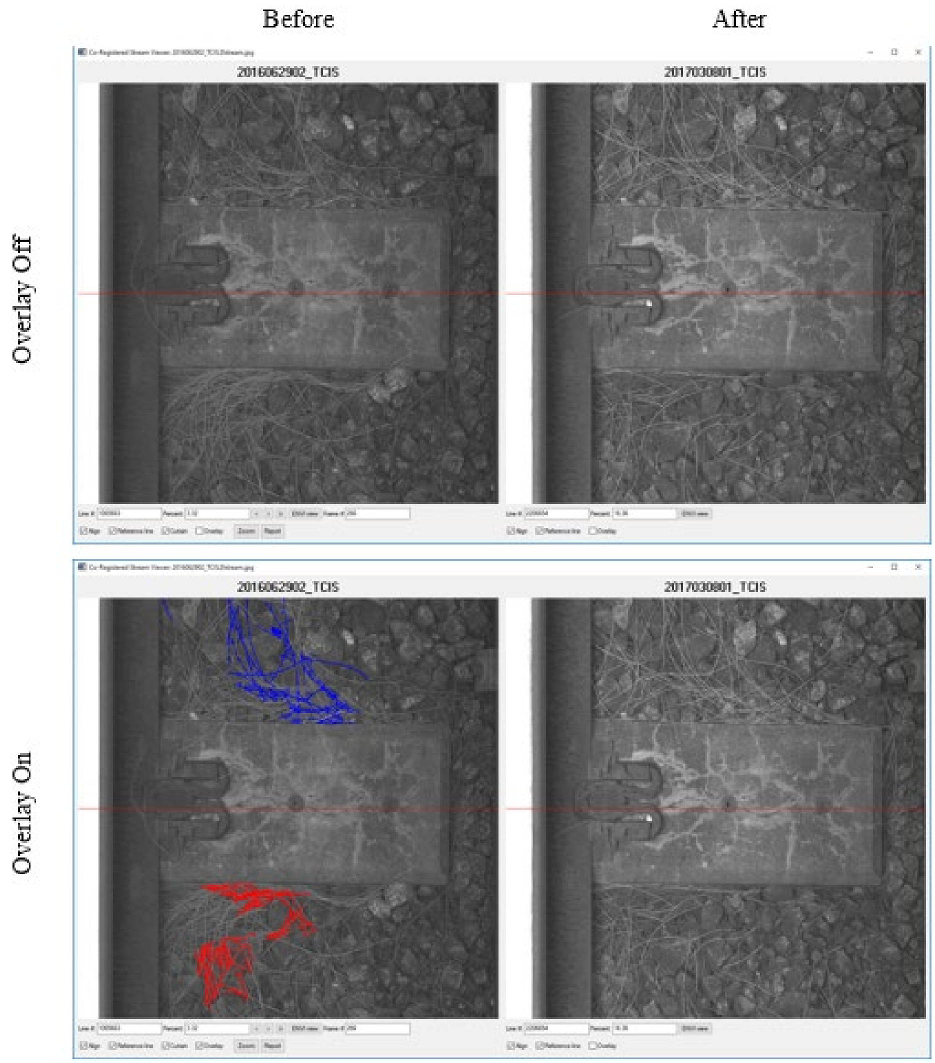

Figure 4-27. Example of a Non-Relevant, True Change Caused by Pine Needles

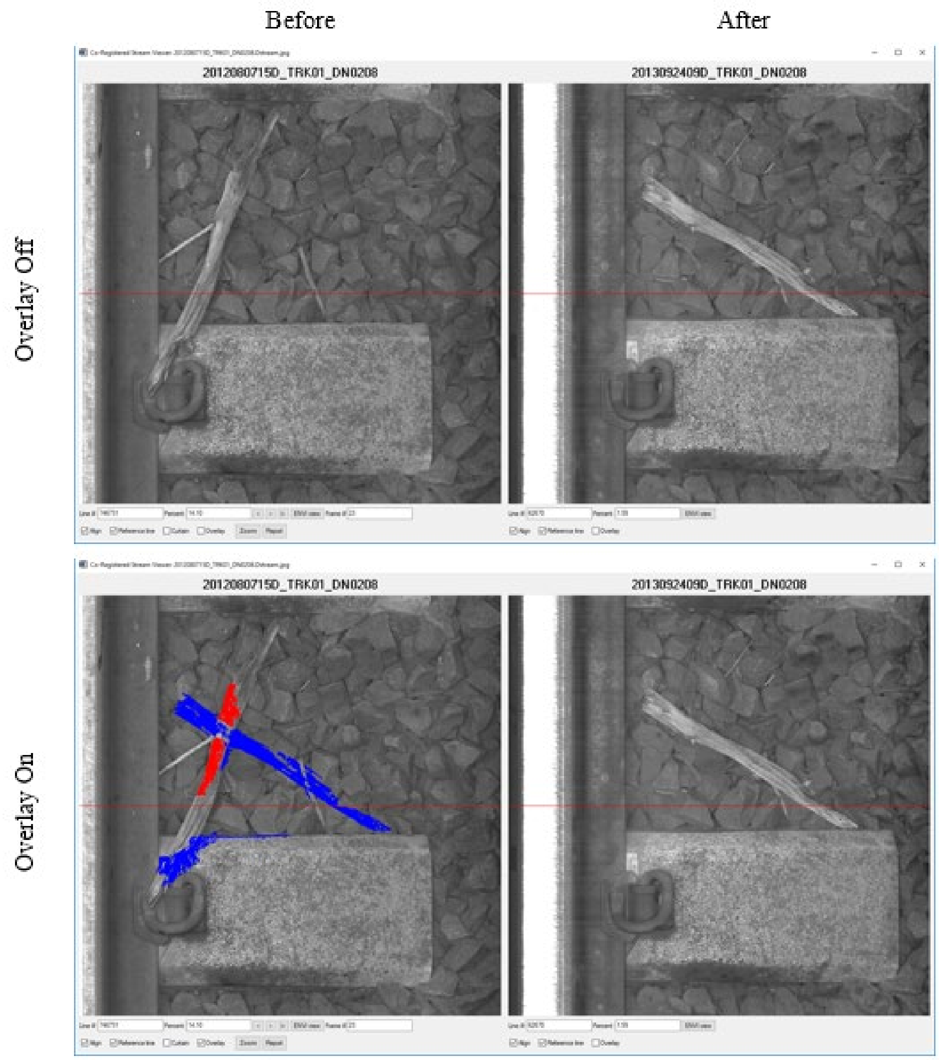

Figure 4-28. Example of a Non-Relevant, True Change Caused by a Stick

4.4.2 Non-Relevant False Changes

A non-relevant false change is a detected change where the reported difference is not real (false) and, therefore, unimportant within the context of achieving a specified objective. False changes associated with intensity-based change detection observed in this study were caused by either non-perfect intensity normalization or non-perfect co-registration. False changes associated with thematic-based change detection are the result of incorrectly classifying a material type in one or both images being compared. This section presents examples of false changes observed during this study. Section 5.1 describes how non-relevant, false changes could be addressed if change detection is deployed in the rail sector.

False Change from Non-Perfect Intensity Normalization

Figure 4-29 shows many false changes corresponding to fastener plates that are the result of non-perfect intensity normalization. False changes such as these were common in the 2016–2017 transit data set (direct fixation track). Although problems associated with intensity normalization were less common in the 2012–2013 data set (concrete tie), they were still present. Figure 4-30 shows an example of a false change related to non-perfect intensity normalization from the 2012–2013 data set.

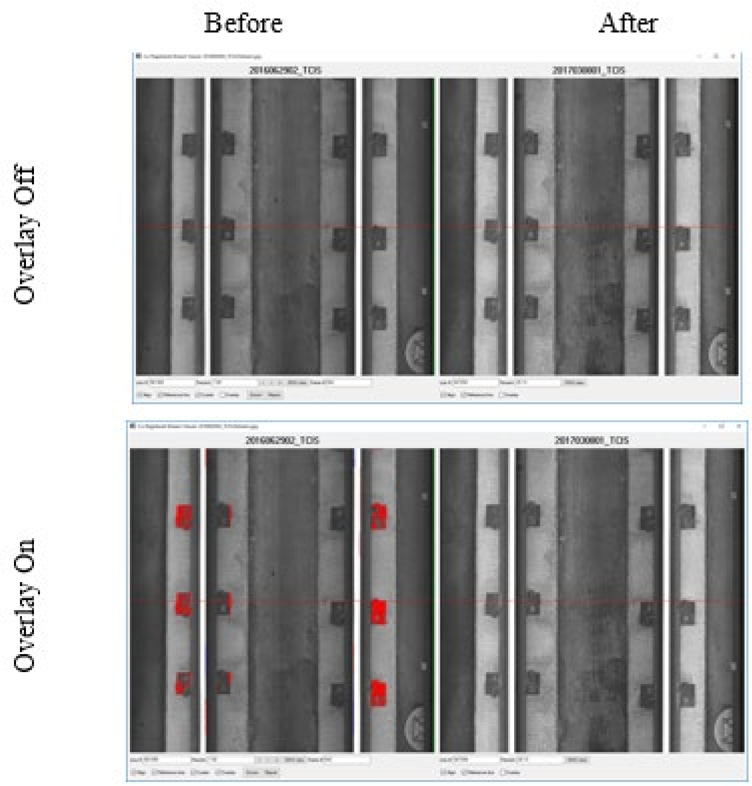

Figure 4-29. Example of a Non-Relevant, False Change Caused by Non-Perfect Intensity Normalization

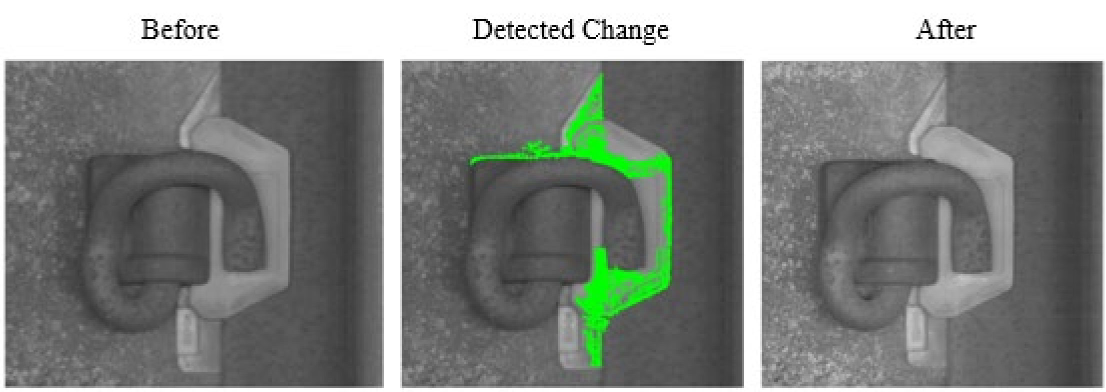

Figure 4-30. Example of a Non-Relevant, False Change Caused by Non-Perfect Intensity Normalization

False Change from Non-Perfect Co-Registration

Figure 4-31 shows an example of false changes caused by non-ideal registration between the "Before" and "After" images when applying intensity-based change detection. In this example, the false changes manifest as narrow, horizontally-disposed regions above and below many ties as shown in the "Overlay On" portion of Figure 4-31.

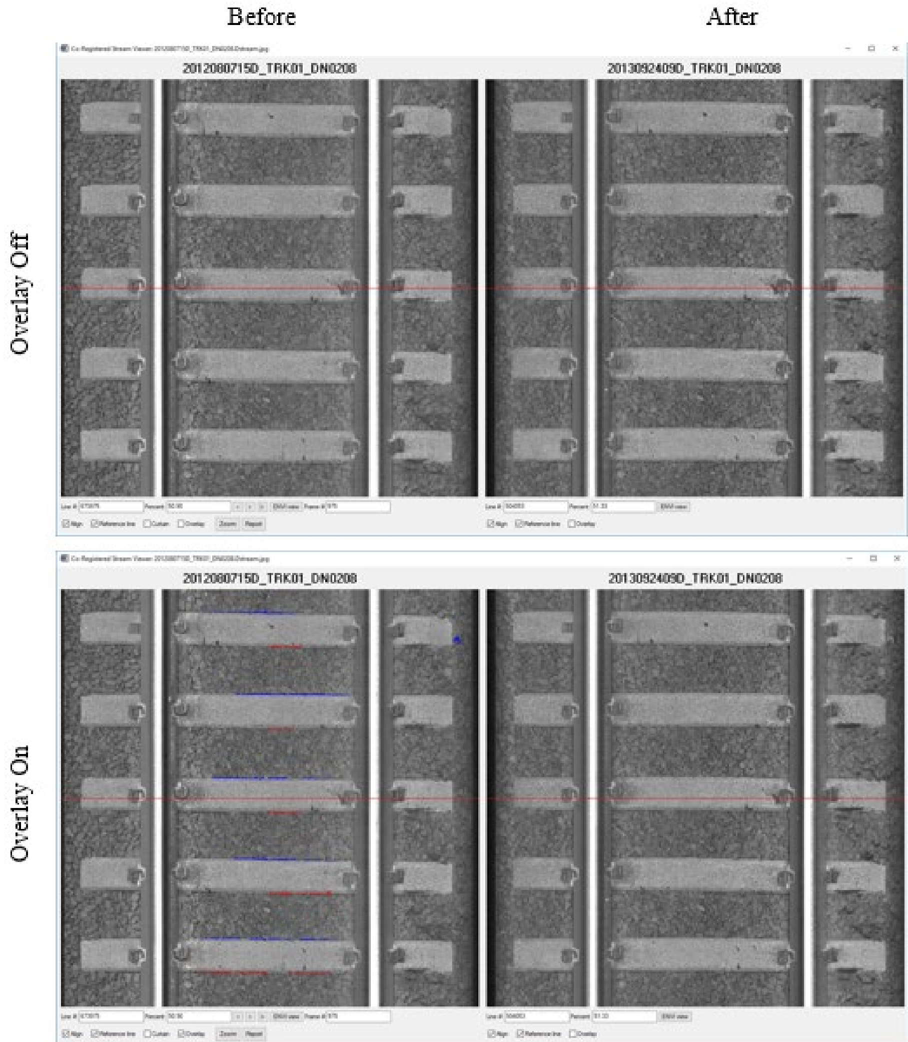

Figure 4-31. Example of a Non-Relevant, False Change Caused by Non-Perfect Co-Registration

False Thematic Changes

The thematic-based change detection approach evaluated in this study produced too many false changes to evaluate the approach for change-based track inspection. Figure 4-32 shows examples of many false, thematic changes typical of those observed throughout the data sets evaluated during this study. Although the number of false, thematic changes is currently too high, the technique has potential for improvement (see Section 5.1.3 for related discussion).

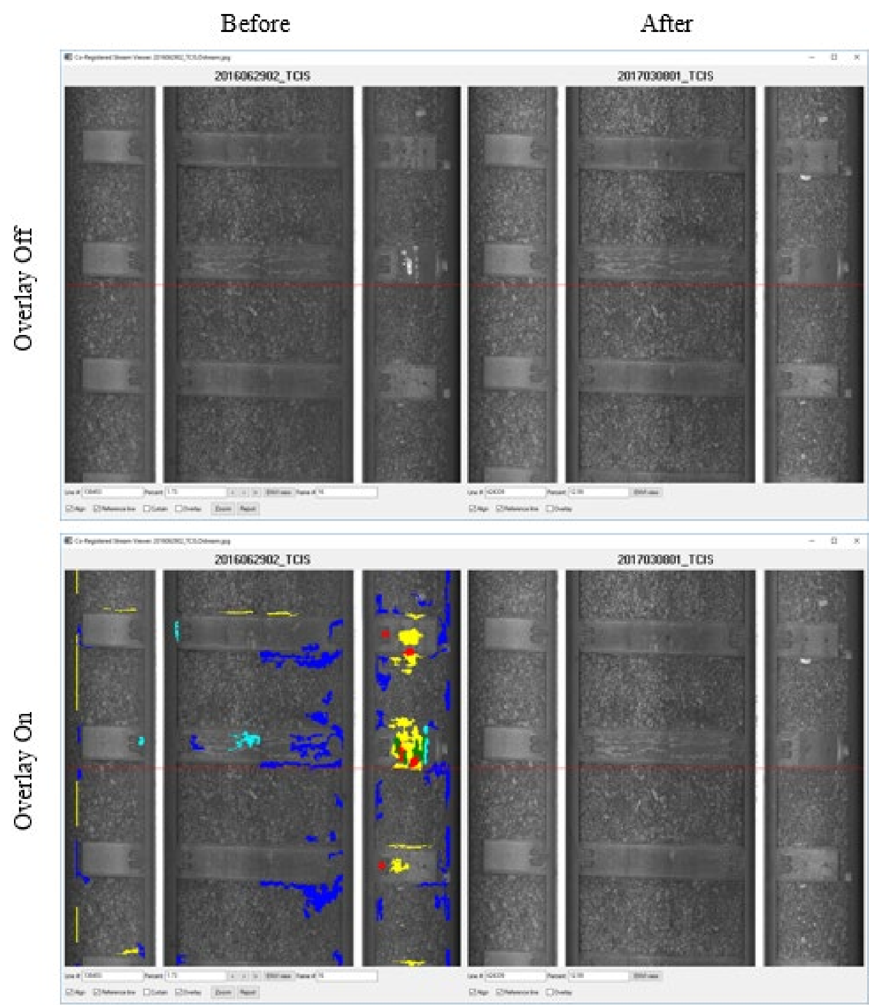

Figure 4-32. Examples of False, Thematic-Based Changes

5. Managing Non-Relevant Changes

The analysis results indicate that non-relevant changes need to be managed if change detection is used in a rail environment. This section describes options for eliminating, reducing, and managing sources of non-relevant changes. The following four strategies are discussed:

- Section 5.1 – Address Root Cause of Non-Relevant, False Changes
- Section 5.2 – Classify and Ignore Common Non-Relevant Changes
- Section 5.3 – Region-Based Change Detection to Filter Non-Relevant Changes

5.1 Address Root Cause of Non-Relevant, False Changes

Subsections under this section describe how each of the following sources of non-relevant, false changes could be eliminated or reduced if change detection is deployed in the rail sector:

- Section 5.1.1 – False Changes from Non-Perfect Intensity Normalization
- Section 5.1.2 – False Change from Non-Perfect Co-Registration
- Section 5.1.3 – False Change from Non-Perfect Thematic Classification

5.1.1 False Changes from Non-Perfect Intensity Normalization

Figure 5-1 shows an enlarged view of one of the fastener plates from Figure 4-29 to help clarify the underlying cause of the false changes. In Figure 5-1, the "Before" image is slightly darker than the "After" image. Intensity normalization typically compensates for overall differences in intensity, but in this case, intensity normalization was ineffective because of an image quality issue. In the "After" image of Figure 5-1, a percentage of the pixels associated with the concrete region are saturated due to excess light. The saturated pixels produce a non-linear brightness response in the concrete region but not in the fastener plate region (the fastener plate region is darker, so pixels in that region are not saturated, or are less saturated). Saturated pixels can induce an error during intensity normalization that makes non-saturated pixels manifest as false, changed pixels.

A solution to address the root cause of this problem involves adding a real-time capability that automatically optimizes light exposure during rail imaging. This would significantly eliminate saturation seen in Figure 5-1. Automatic optimization for light exposure is already built into today's area scan cameras—a similar approach could be leveraged during real time line scan imaging. Applying real-time compensation for light exposure during line scan imaging may only require a modification to real-time imaging software. This means no hardware changes would be necessary for existing imaging systems.

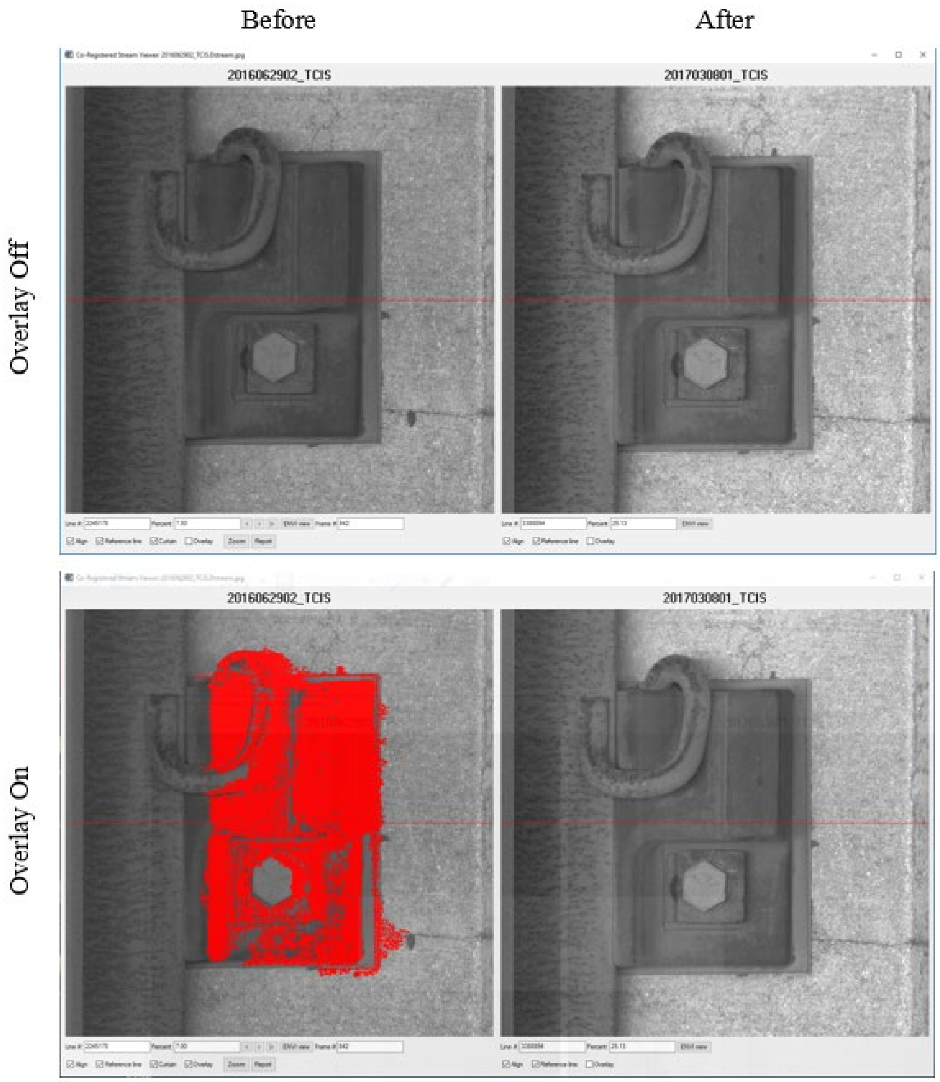

Figure 5-1. Close-up View of a Non-Relevant, False Change—Non-Perfect Intensity Normalization

5.1.2 False Change from Non-Perfect Co-Registration

Figure 5-1 shows an enlarged view of one of the false change regions shown in Figure 4-31. In this case, false change is reported because of an image registration error along the direction of travel. The noted registration error causes thin sections of concrete in the "Before" image to be registered with ballast in the "After" image. Since ballast and concrete are different shades of grey, the net result is that an intensity difference is detected and reported as change along tie edges.

The co-registration errors responsible for the corresponding false changes are believed to be caused by small differences in tachometer slippage between "Before" and "After" imaging runs. Regardless of the root cause, automating the co-registration process is expected to handle false changes along tie edges. Anchor points manually entered on image sets were used to facilitate this study. Automating the co-registration process would allow an algorithm to use a larger number of anchor points so that co-registration accuracy remains high throughout an entire imaging run.

The software evaluated in this study includes automated co-registration, but it is configured to work with aerial images. Modifications would be needed to adapt the existing capability for rail-based images.

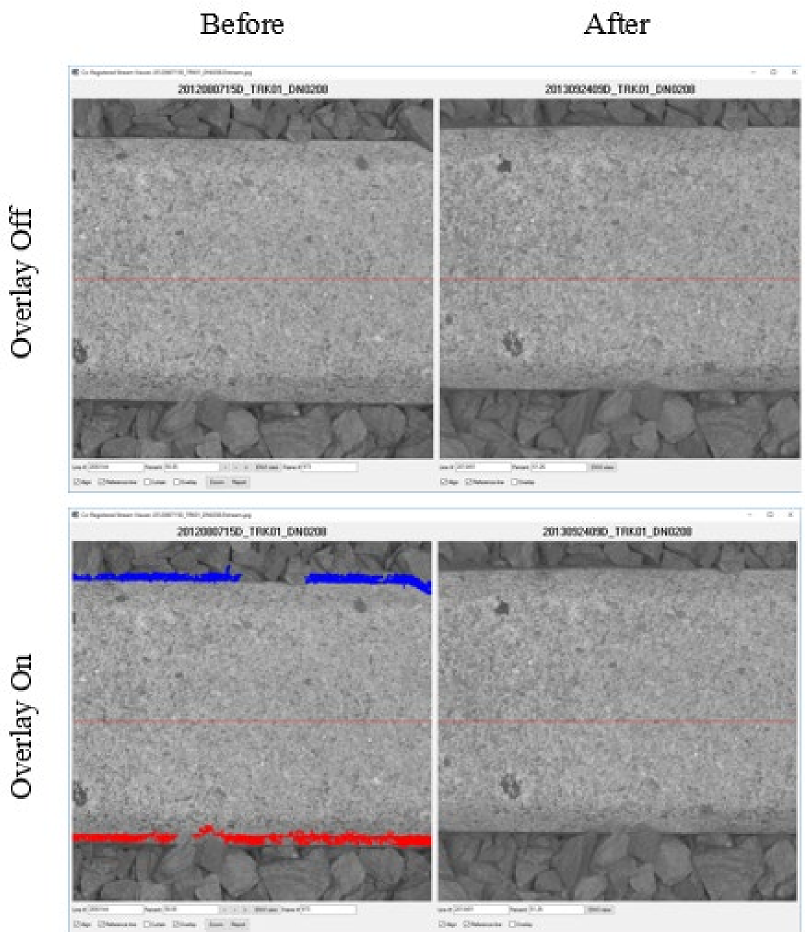

Figure 5-2. Close-up View of False Change Caused by Non-Perfect Co-Registration

5.1.3 False Change from Non-Perfect Thematic Classification

Figure 5-3 provides a close-up view of a portion of the false changes shown in Figure 4-32. Each colored region in Figure 5-3 indicates a location where the software is reporting a thematic change. In the "Overlay Off" portion of Figure 5-3, several grease spots present in the "Before" image are gone in the "After" image. Regions in Figure 5-3 where grease is present on the depicted metal plate correspond to true detections (i.e., something changed, and the software detected the change). However, many surrounding regions in Figure 5-3 where change is reported correspond to false changes.

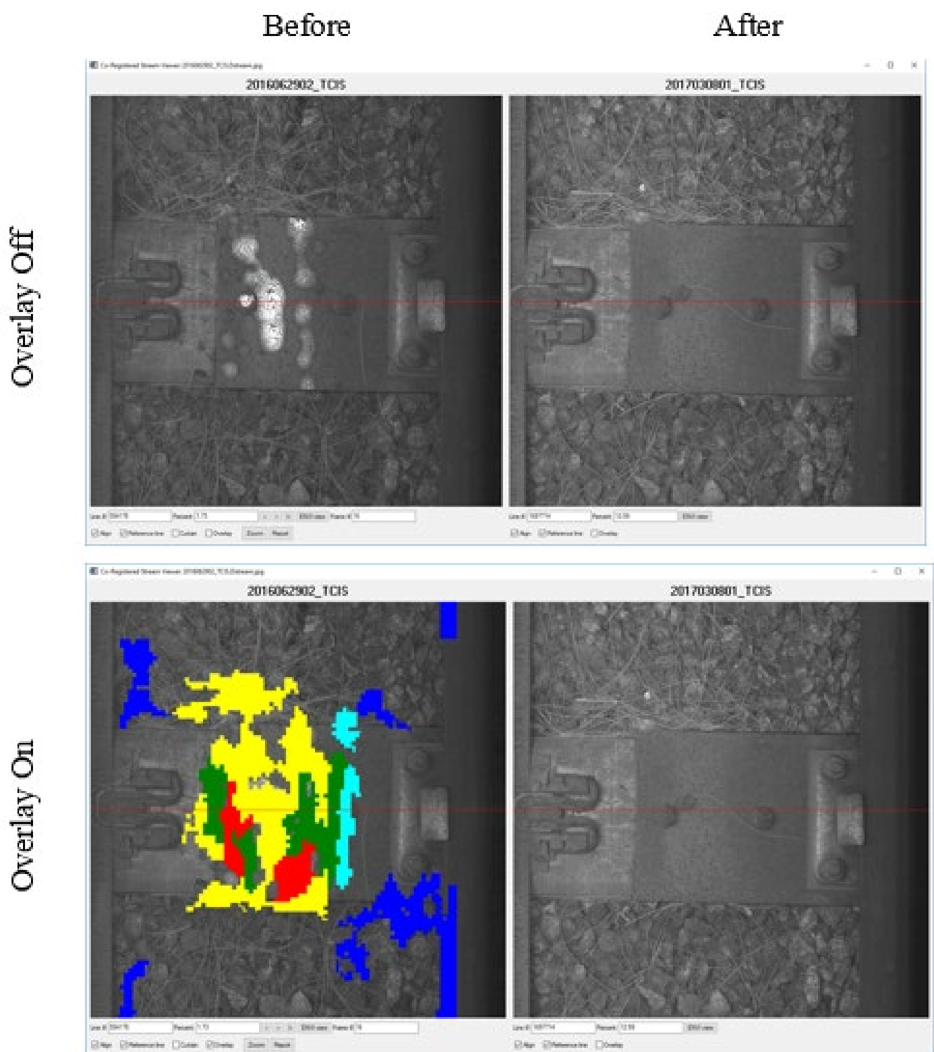

Figure 5-3. Close-up View of a Cluster of False, Thematic-Based Changes

This example serves to clarify that thematic-based change detection can detect changes; however, its ability to accurately localize the underlying changed regions as well as its ability to correctly categorize changes needs further improvement. As discussed next, additional classifier training has been shown to significantly reduce the prevalence of false thematic changes.

The thematic classifier used in this study is based on deep learning technology. One strength of deep learning is that it can learn to classify almost anything. A weakness is that the technique's accuracy significantly depends on the number of sample images used during training. In this study, hundreds of sample images were used to train each of the five thematic classes (rail, ballast, concrete, wood, and fasteners). Performance would be better if thousands of training samples are used per class.

Figure 5-4 shows an example of thematic-based change detection results achieved before and after additional fastener training samples were used. Additional training samples significantly lowered the number of false changes (red regions) associated with fasteners. In this example, additional training was only applied to the fastener class; however, similar improvements to all classes are possible if additional training is applied to all classes.

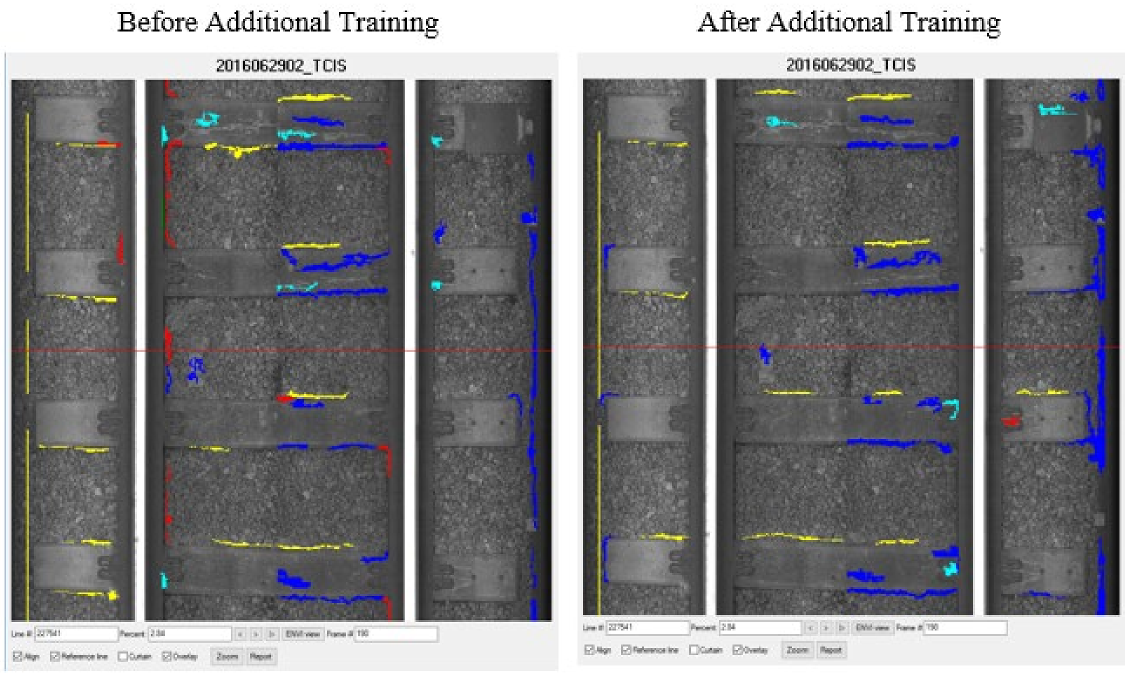

Figure 5-4. Demonstration of Improvements from Re-Training the Thematic Classifier

An important point is that training a thematic classifier based on deep learning technology is not labor intensive if classes to be learned are selected based on the following guidelines:

1) Inter-class variation remains high meaning that each class looks visually different from all other classes (e.g., classifying wood and concrete is easier than classifying different types of concrete).

2) The total number of thematic classes remains relatively low (e.g., 20 or so).

3) Training images are readily available in the underlying data (e.g., rail infrastructure components are typically prevalent enough whereas many conditions-of-interest are not).

5.2 Classify and Ignore Common Non-Relevant Changes

The thematic classifier used in this study is based on deep learning technology, which is well-suited for classifying objects and image textures when training images are available in large quantities (e.g., thousands). As such, thematic-based classification is well-suited for automatically classifying common, non-relevant changes so that they can be ignored, as clarified next by example.

Isolated changes in ballast accounted for approximately 75 percent of the non-relevant, intensity-based changes.[4] Figure 5-5 through Figure 5-7 help clarify how thematic-based classification can

[4] These statistics are based on non-relevant, true changes only (non-relevant, false changes are not represented)

be used to classify and ignore non-relevant changes associated with ballast. Figure 5-5 shows an example of the intensity-based change response corresponding to a non-relevant, isolated change in ballast. Figure 5-6 shows the corresponding thematic-based classes and change response. Finally, Figure 5-7 shows a confusion matrix that could be used to instruct a decision engine when, and when not to, ignore intensity-based changes.

In the example introduced above, each detected intensity-based change (Figure 5-5) would be used to retrieve a corresponding thematic class change (Figure 5-6) at the same location. The confusion matrix of Figure 5-7 would then use the thematic class change to automatically determine whether to report the corresponding intensity-based change. In this example, the thematic-based class results would successfully filter all the non-relevant changes in ballast shown in Figure 5-5.

The approach described above would eliminate approximately 80 percent of the detected, non-relevant changes in ballast. If additional training is used to improve the thematic classifier, it is expected that the approach would further reduce and potentially eliminate non-relevant changes associated with ballast.

It is anticipated that leaves during the fall season may result in a significant number of non-relevant changes. As training samples associated with leaves would be prevalent, it is expected that the classify-and-ignore approach described above would be suitable for filtering leaves. In a broader sense, it is expected that the classify-and-ignore approach can handle any situation involving many non-relevant changes.

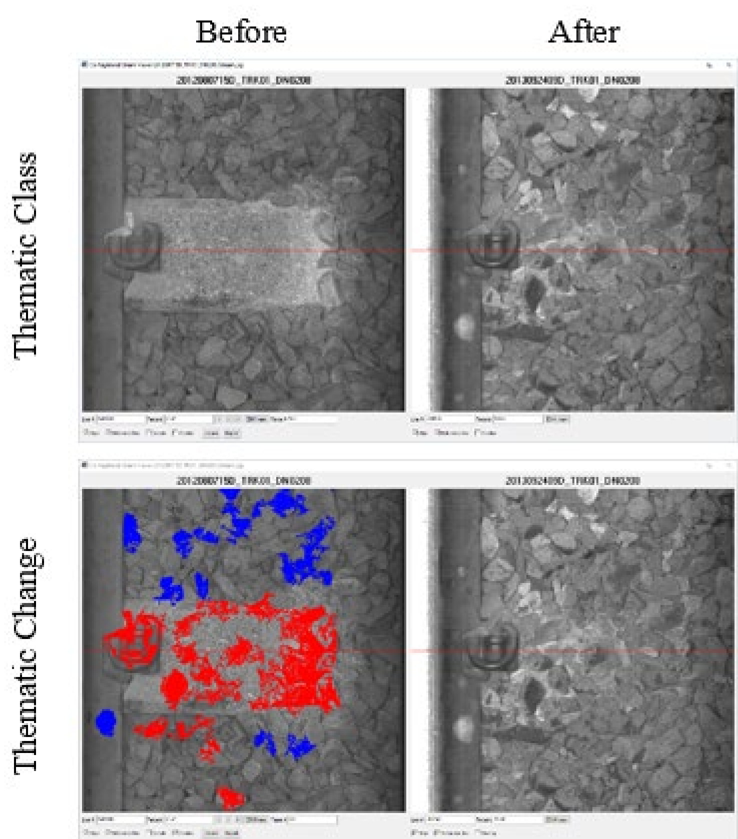

Figure 5-5. Intensity-Based Result—Isolated Changes in Ballast

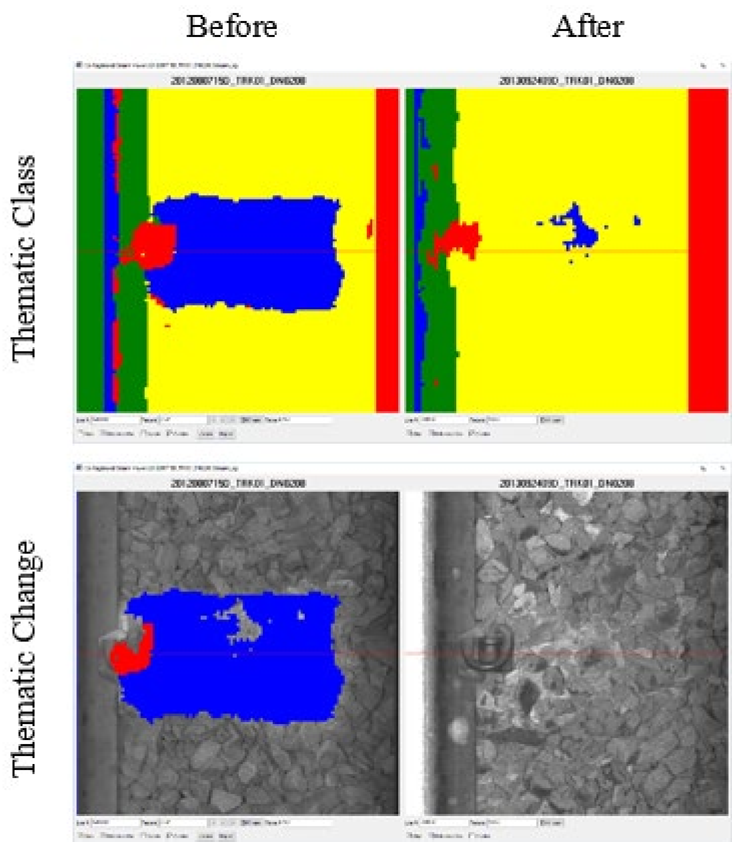

Figure 5-6. Thematic-Based Result—Isolated Changes in Ballast

After State

Confusion Matrix (Before State)	Ballast	Concrete	Wood	Rail	Fastener
Ballast	✗	✗	✗	✗	✗
Concrete	✗	✗	✓	✓	✓
Wood	✗	✓	✗	✓	✓
Rail	✗	✓	✓	✗	✓
Fastener	✗	✓	✓	✓	✗

✗ Do Not Report Change
✓ Report Change

Figure 5-7. Example Confusion Matrix Used to Ignore Non-Relevant Changes

5.3 Region-Based Change Detection to Filter Non-Relevant Changes

This section describes and assesses a region-based change detection approach intended to provide another means of managing non-relevant changes. In addition to helping reduce non-relevant changes, the region-based approach described here has also demonstrated abilities to:

- Simultaneously lower the probability of missing relevant changes
- Enable relevant changes to be labeled with a corresponding asset type

5.3.1 Overview of Region-Based Change Detection

Region-based change detection involves using two-dimensional and/or three-dimensional spatial templates to indicate specific regions in an image (or within a three-dimensional point cloud) where change detection results are desired, while ignoring other regions. When applying region-based change detection, co-registration and intensity normalization are applied to an entire track image and change detection is applied only within specific regions. This leads to many benefits further clarified in later sections.

The spatial templates referenced above are specific areas in an overall track image where change detection results are needed. Templates (or regions of interest) can be specified manually by a user (e.g., by dragging a box around a region of interest in an image of the track) or generated automatically by a machine vision algorithm. Whether generated manually or automatically, templates only need to be created one time and then updated only if the underlying track infrastructure changes.

When the underlying change detection mode is intensity-based (rather than thematic-based), size and intensity thresholds used with region-based change detection are optimized to address specific detection goals associated with specific region types. For example, if templates correspond to regions where fasteners exist, then size and intensity thresholds optimized for detecting missing and rotated fasteners are used.

5.3.2 Region-Based Change Detection Applied to Missing Fasteners

A case study was performed using region-based change detection to detect missing and rotated fasteners. In the case study, a MATLAB script was written to automatically generate a set of templates required by region-based change detection. Referencing Figure 5-8, the template-creation process involved four steps.

- In Step 1, thematic classification results were used to locate the centroid of each tie's central section and both of its shoulders.
- In Step 2, tie-region centroids were used to estimate the center line of the underlying tie, including any skew angle.
- Step 3 used tie center lines and known track dimensions to position horizontal and vertical lanes where fasteners are expected.
- In Step 4, lanes were used to position a set of four rectangular fastener templates for each tie.

The MATLAB script used to implement the four steps described above was applied to each tie in the evaluated data set. The right side of Figure 5-8 shows a corresponding set of fastener templates superimposed on a sample of the processed image file.

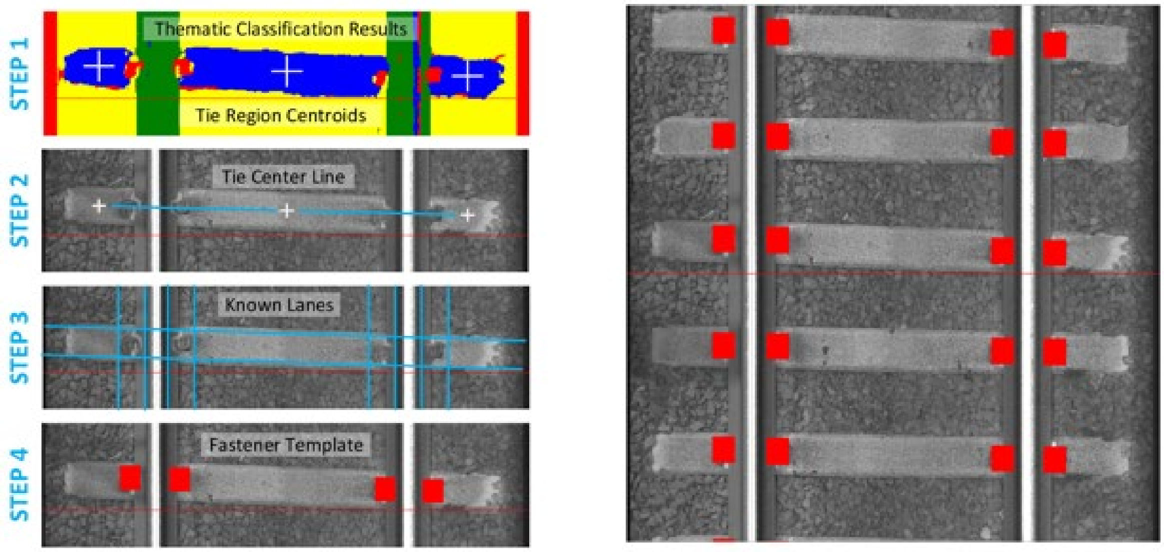

Figure 5-8. Example Four-Step Process Used to Apply Region-Based Change Detection

The region-based change detection approach described above was applied to approximately 1 mile of concrete tie track. The corresponding line scan image consisted of 1,902 frames and included approximately 14,512 fasteners. Prior to applying the fastener templates to the image set, all 1,902 frames in the data set included at least one reported region of change. This means it would have been necessary to manually review all 14,512 fastener regions in the data set to fully screen for missing fasteners. After region-based change detection was applied, only 218 out of 14,512 fastener regions included a detected change. Based on this example, region-based change detection reduced the need for manual image review by 98.5 percent, computed as follows: (1 – 218/14,512) * 100 = 98.5 percent.

The left side of Figure 5-9 shows a sample of the raw, intensity-based change detection results achieved without using region-based change detection. The right side of Figure 5-9 shows change detection results achieved using region-based change detection. Note that in Figure 5-9, most of the non-relevant changes shown on the left are not present on the right after region-based change detection was applied.

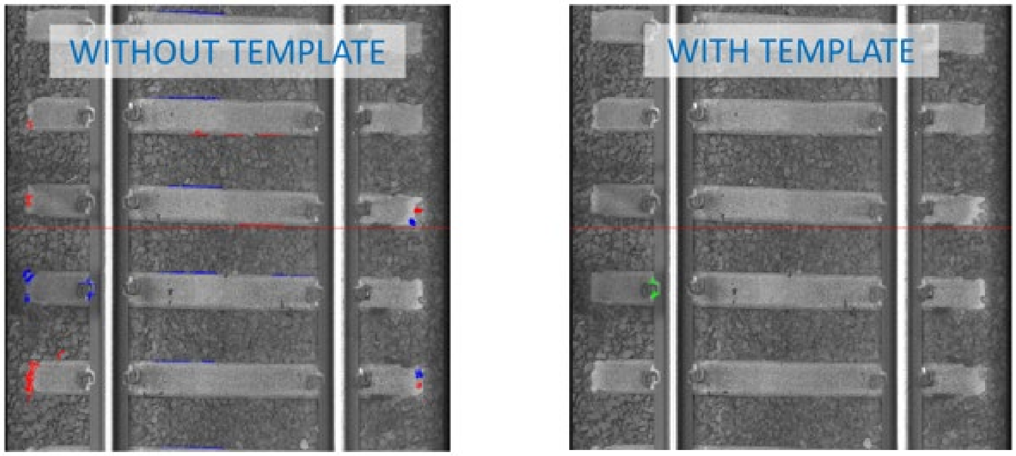

Figure 5-9. Demonstration of Improvement from Region-Based Change Detection

For reference, Figure 5-10 shows a close-up view of the fastener from Figure 5-9 where region-based change detection produced a non-relevant change. A likely cause of this change is non-perfect intensity normalization resulting from improper light exposure during imaging. Provisions for improving light exposure during imaging, as discussed in Section 5.1.1, are expected to further reduce non-relevant changes such as that of Figure 5-10.

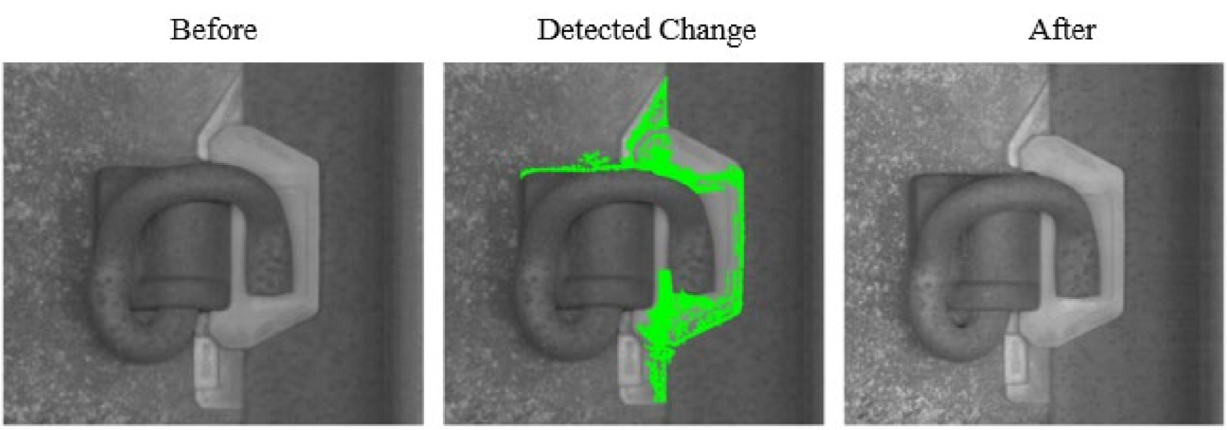

Figure 5-10. Non-Relevant, True Change After Applying Region-Based Change Detection

In this case study, no relevant changes associated with fasteners were missed out of more than 100 possible opportunities. The data included cases of missing, rotated, moved, and replaced fasteners.

5.3.3 Region-Based Change Detection Applied to Rotated Base Plate Retainer Clips

To help clarify broader potential for region-based change detection, this section discusses a second example based on using a custom template intended to detect rotated base plate retainer clips. The discussion in this section also clarifies how region-based change detection improves the probability of detecting relevant changes while simultaneously reducing the probability of detecting non-relevant changes.

In the example discussed here, the goal is to detect rotated retainer clips such as the one shown in Figure 5-11. Such clips are used to hold fastener baseplates to the underlying substructure. A rotated clip is a sign that the corresponding bolt is loose, a condition believed to precede a more serious missing bolt condition.

Figure 5-11. Example of a Rotated Base Plate Retainer Clip

In Figure 5-12, a donut-shaped template (shown in orange) is approximately matched to the shape of the expected region of change associated with the rotated retainer clip introduced in Figure 5-11. Matching a template to the shape of an expected region of change increases the probability of detecting a relevant change while simultaneously decreasing the probability of detecting a non-relevant change. This expected win-win scenario is a result of the following axiomatic assumptions:

> **Axiom 1** - The probability of detecting a change of random size is inversely proportional to the size threshold used when declaring that change exists, all else equal.
>
> **Axiom 2** - Given an arbitrary number of uniformly distributed, non-relevant changes of random size, the probability of detecting a non-relevant change based on a given size threshold is proportional to the total image area over which change detection is applied, all else equal.
>
> **Axiom 3** - The probability of detecting a relevant change is unaffected by the size of the search area so long as the relevant change is fully enclosed within the search area.

How these axioms establish the win-win scenario stated above is clarified by example. Figure 5-13 shows four donut-shaped templates superimposed on one frame of image data. Based on dimensions provided in Figure 5-13 the area of the corresponding image frame is 101 times greater than the per-frame area inside the four donut templates. Since the area inside the donuts is 101 times less than that of a frame, the donut templates reduce the total area over which change detection is applied by a ratio of 101:1. Based on Axiom 2 stated above, this decreases the probability of detecting a non-relevant change inside the donuts by a factor of 101 (all other factors equal). The large decrease in probability arises because not as much room for a non-relevant change to occur exists inside the donuts compared to outside the donuts. The calculation presented here is based on a single frame; however, if a reasonable assumption is made that non-relevant changes are distributed uniformly throughout a complete image set, the result is applicable to the full image set.

In this study, a 10,000-pixel size threshold was used when intensity-based change detection was applied to a full data set. When applying region-based change detection to the task of detecting rotated base plate retainer clips it is preliminarily estimated that a 1,000-pixel threshold would be appropriate. This corresponds to a 10:1 reduction in threshold size. Based on Axiom 1 stated above, a 10:1 reduction in threshold size increases the probability of detecting a rotated base plate retainer clip by a factor of 10, all else equal. Unfortunately, it also increases the probability of detecting a non-relevant change by a factor of 10; however, the area reduction from the donut-shaped template offsets this 10:1 increase by a 101:1 decrease. The net result is a 10:1 decrease in the probability of detecting a non-relevant change inside a donut area. Meanwhile, based on Axiom 3, the donut-shaped template does not lower the probability of detecting a relevant change because the template fully covers the zone where a relevant change would emerge. This example analysis has demonstrated how region-based change detection is expected to provide order-of-magnitude improvements in detection probability simultaneous with order-of-magnitude reduction in non-relevant changes.

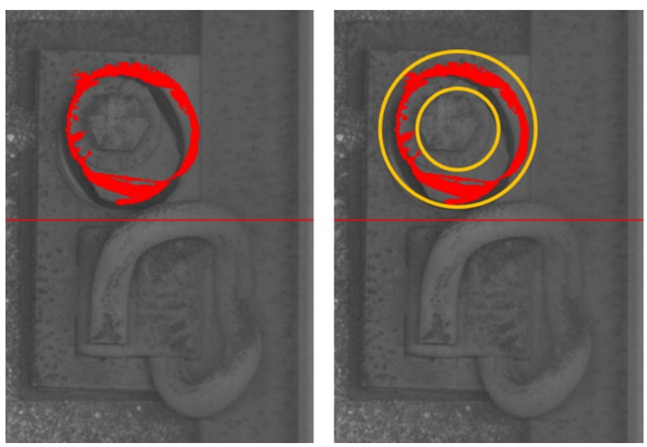

Figure 5-12. A Region-Based Change Detection Template Customized for a Specific Objective

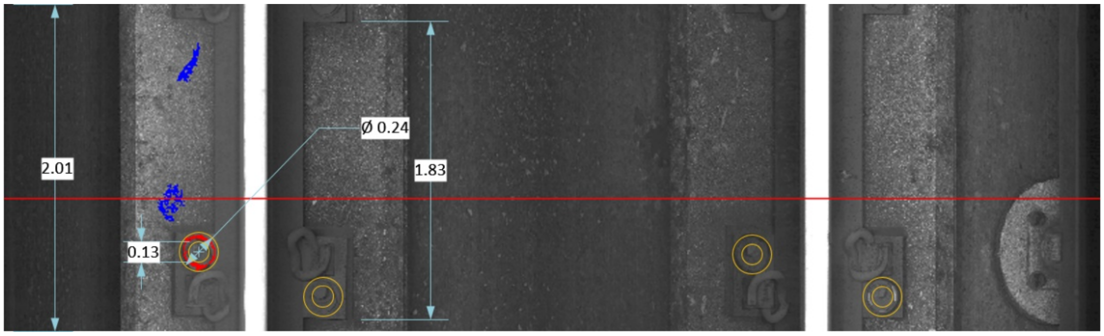

Figure 5-13. Frame used to Estimate Detection Probability (Frame Size = 2048 x 7214)

To help confirm that the assertions related to the win-win scenario described above are reasonable, region-based change detection was applied to the missed change shown earlier in Figure 4-18. In doing so, the corresponding size threshold was reduced from 10,000 pixels to approximately 1,000 pixels. If doing so causes the previously missed clip to be reported as a change then the theory is supported; otherwise it is not. The corresponding results shown in Figure 5-14 indicate that the previously missed clip is now robustly reported as a change. Based on this result, the theory is supported, which provides a preliminary basis of confidence that the win-win scenario outlined above is achievable on a broader scale.

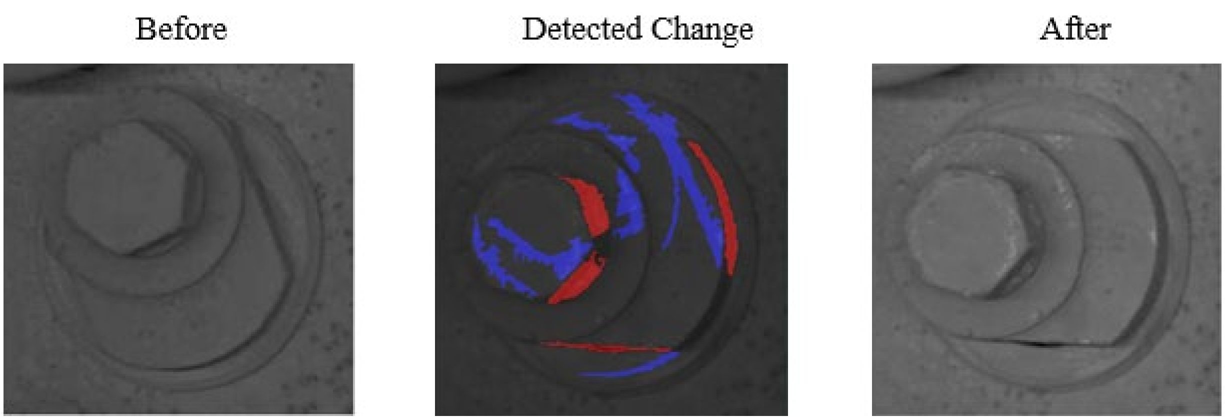

Figure 5-14. Region-Based Change Detection Applied to a Base Plate Retainer Clip

5.3.4 Algorithm Development Associated with Region-Based Change Detection

Using region-based change detection to detect rotated base plate retainer clips throughout a rail network would require hundreds-of-thousands of donut-shaped templates to be overlaid on top of the corresponding track image (a one-time process). While a machine vision algorithm would be needed for template placement, developing the algorithm is expected to be easier than developing an algorithm to detect the rotated clips.

Developing a template-placement algorithm suitable for detecting rotated base plate retainer clips using region-based change detection would typically involve isolating one or more base plate corners (or edges). Isolated corners (or edges) would then provide reference lines needed for template placement. Detecting a condition such as rotated clips using a conventional algorithm development process would typically involve isolating base plate corners (or edges) as an initial step. Additionally, conventional algorithm development would typically involve some form of image segmentation, feature extraction, and training a detector (or classifier). Each of these additional steps requires considerable effort. Thus, while machine vision algorithm development is still needed when using region-based change detection, significant savings potentially exists. Specifically, a three-fold (3:1) reduction in algorithm development cost is conservatively expected in typical situations where region-based change detection can be substituted for conventional algorithm development.

Although the discussion presented above centers around a specific example (detection of rotated base plate retainer clips), the underlying principals are expected to apply to many examples relevant to machine-vision-based track inspection.

6. Ideal Change-Based Track Inspection Process

This section describes an ideal, change-based, track inspection process. The ideal process described here, contrasted with the as-is state of the evaluated software, establishes a basis for the identified areas for additional development listed in Section 7.3.

6.1 Description of Change-Based Track Inspection

Figure 6-1 shows an ideal, change-based track inspection process. This process would be capable of automatically identifying changes associated with many conditions-of-interest. The corresponding probability of missing a relevant change would be low enough in all cases to justify a no-change-no-review policy (e.g., significantly less than 1 percent). Although optimal survey intervals for change-based processing are not yet known, an optimal range is expected to be between 2 and 26 weeks, depending on specific circumstances. Under the ideal scenario of Figure 6-1, change-based processing will be compatible with images captured from a variety of platforms, including high-rail vehicles, dedicated measurement cars, and unmanned aerial vehicles (UAVs).

With continued reference to Figure 6-1, change-based track inspection (top) is shown running in parallel with a track maintenance process (bottom). Based on this ideal scenario, the track inspection process generates a list of new (unrepaired) track conditions at known locations. The list is provided to the track maintenance process following each track survey (Figure 6-1, Item 13) and used to generate work orders. As work orders are completed, the track maintenance process feeds a list of repaired conditions at known locations back to the track inspection process (Figure 6-1, Item 12). The track inspection process then uses the list to separate new (unrepaired) track conditions from repaired track conditions during the next change-based track survey (Figure 6-1, Item 9). Here, a key point is that, without track maintenance feedback, change-based track inspection has no inherent way to isolate changes associated with maintenance activity from changes related to potentially unsafe track conditions.

Following each imaging survey (Figure 6-1, Item 1) a decision will be made based on whether the survey is the first survey (Figure 6-1, Item 2). If the current survey is the first survey, track images will be manually reviewed (Figure 6-1, Item 4), and automated template-placement algorithms will be applied (Figure 6-1, Item 5). The purpose of manually reviewing track images following a first survey is to identify any pre-existing track conditions. The purpose of automated template placement (Figure 6-1, Item 5) is to designate regions where changes-of-interest are to be reported. This step supports automated filtering of non-relevant changes (Figure 6-1, Item 10) and also enables detected changes to be grouped by asset type when reports are generated (Figure 6-1, Item 11).

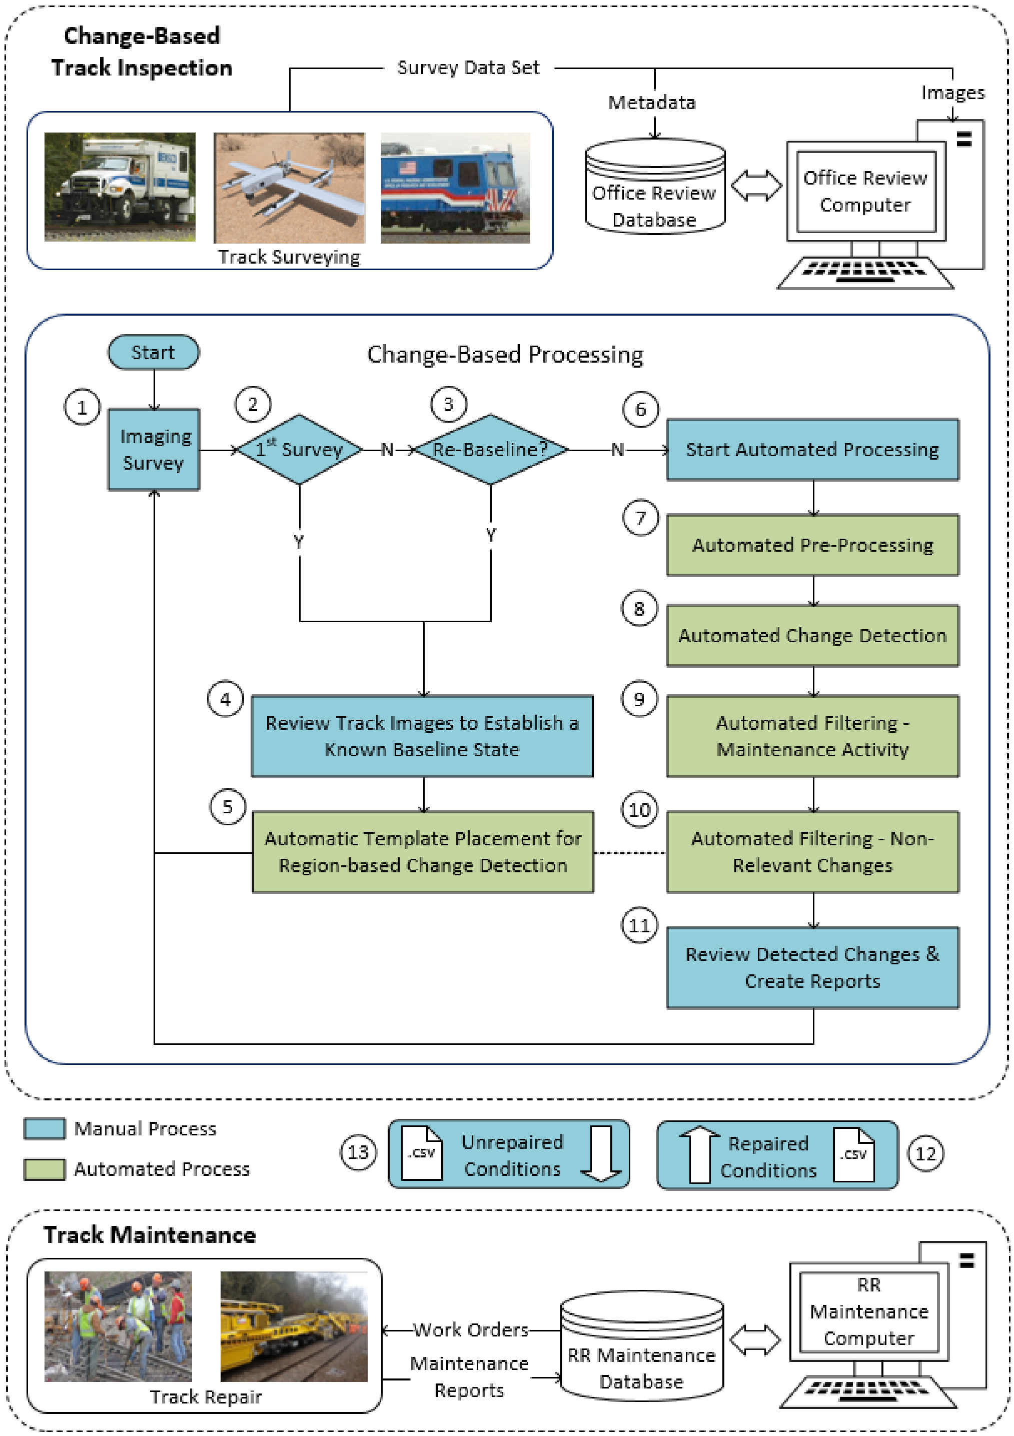

Figure 6-1. Envisioned Process for Ideal, Change-Based Track Inspection

If the current imaging survey is not the first survey, a decision will be made regarding whether to establish a new baseline for change-based processing (Figure 6-1, Item 3). Establishing a new baseline involves manually reviewing track images to identify any potentially unsafe track conditions that might have been missed by automated processing. Optimal re-baselining intervals are not yet known, but estimated at between 6 and 24 months, depending on circumstances that may be specific to each railroad.

During the normal course of change-based track inspection, images and corresponding metadata from each survey are hand-carried from the inspection platform(s) to an office setting. Due to large file sizes associated with machine vision data, this data offloading step is expected to remain a manual process into the foreseeable future.[5] Once images and metadata are moved from the field to the office, change detection processing is initiated manually (Figure 6-1, Item 6). The manual initiation process is expected to involve the following steps:

1) Loading images and metadata from the latest imaging survey into an office-based computer
2) Designating "Before" and "After" survey's to be compared
3) Starting the automated comparison process

Loading image files from the latest imaging survey typically involves inserting a disk drive from the field directly into the office computer. Metadata will typically reside on the same drive as the images and will automatically be imported into a database that resides on the office computer. When performing change-based processing, the "Before" survey will typically be the prior survey and the "After" survey will typically be the most recent survey; however, in general, any two surveys can be compared.

Once initiated, change-based processing will take place in a fully-automated manner. As indicated in Figure 6-1 (Items 7–10), the automated processing steps are expected to include:

1) Pre-processing – intensity normalization and co-registration
2) Change detection – identifying regions where differences exist
3) Filtering maintenance activity – separating maintenance activity from relevant changes
4) Filtering non-relevant changes – identifying and removing unimportant changes.

Following automated processing, areas where relevant change is reported will be manually reviewed (Figure 6-1, Item 11). Manually-reviewed results will then be used to convey track repair needs to a parallel track maintenance process (Figure 6-1, Item 13). Additional details associated with the envisioned, ideal, manual review process are further described in the next section.

6.2 Ideal Manual Review Process

Following each survey, manual review of relevant changes would take place in an office setting. Except for a first survey and during periodic re-baselining, it would only be necessary to manually review areas where change is detected. Non-relevant changes and maintenance activity would automatically be filtered from the result set.

During manual review, a track inspector would have access to capabilities including:

---

[5] Based on current, 4G wireless technology operating at 20 Mbits/second, approximately a thousand-fold increase in wireless bandwidth would be needed to support automated (wireless) data offloading sufficient to support robust, change-based track inspection.

- Show Changes Relative to a Specified Date
- Show Only Changes that Meet a Specific Business Rule
- Review Historical Changes at a Given Location
- Show/Hide Changes for Specific Asset Types
- Show/Hide Changes at Specific Locations
- Show/Hide Changes Associated with Maintenance Activity
- Add/Delete Locations where Change Monitoring is to Occur
- Confirm/Audit Work Order Status

Following the ideal manual review process described above, confirmed changes-of-interest would be exportable in the form of an electronic report. Among other optional information, the report would list detected changes (by asset type) and a corresponding location for each change. Reports would include hot links allowing corresponding images to be manually reviewed without the need for specialized software. Information conveyed in such a report would be filterable, allowing report content to be targeted to a specific audience or used for various other purposes.

6.3 Ideal, Closed-Loop Maintenance Tracking

Unrepaired conditions confirmed during the ideal manual review process described above would be exportable in a widely-used file format (e.g., .csv) to support track maintenance (Figure 6-1, bubble 13). Unrepaired conditions exported from the track inspection process would be imported into an external work-order tracking system. It is envisioned that the work-order tracking system would be based on software and a database specific to each railroad. It is assumed that the work order tracking process would track repair needs identified by an arbitrary number of sources in addition to the ideal change-based track inspection process.

Following each imaging survey, the status of repair activity in the work-order tracking system would be transferred to the change-based track inspection system via a common file format (e.g., .csv) (Figure 6-1, bubble 12). Such a transfer establishes a closed loop allowing the track inspection process to be aware of repairs made by the track repair process. A closed loop helps streamline the ideal manual review process described above by allowing changes associated with repair activity to be separated from changes corresponding to potentially unsafe, new conditions. The closed loop also allows changes exported from the track inspection process to be automatically filtered to remove any open, unrepaired changes left over from all prior imaging surveys. In addition to streamlining the amount of manual image review, this step also helps avoid duplicate work order creation.

7. Conclusion

Based on results from this study, image-based change detection has demonstrated eminent potential to add value within the rail sector. This study has demonstrated that change detection can identify changes corresponding to many conditions of interest. Abilities to detect the following relevant changes were demonstrated during this study:

1) Missing rail fasteners
2) Rotated rail fasteners
3) Rotated base plate retainer clips
4) Changes in crumbled tie state
5) Changes consistent with fouled ballast
6) Many forms of maintenance activities
7) Rail surface anomalies
8) Standing water

While not directly confirmed during this study, observed results indicate that image-based change detection would be able to detect changes that result from many of the following, additional conditions (organized by imaging system):

1) Rail Surface Imaging
    a. Completely broken rail
    b. Wheel burns
    c. Rail grinding activity
    d. Other significant rail surface anomalies
2) Rail Web Imaging
    a. Completely broken joint bars
    b. Completely broken rail
    c. Missing joint bar nuts and bolts
    d. Rotated joint bar nuts
    e. Excessive metal flow at rail welds on heavy haul routes
    f. Excessive rail gaps
3) Full Width Track Imaging
    a. Missing base plate bolts
    b. Missing tie spikes
    c. Completely broken base plates
    d. Completely broken rail

  e. Changes in concrete tie crack growth [6]
  f. Skewed ties
  g. Significant rail base deterioration [7]
  h. Land-slide debris encroachment
  i. Erosion of track foundation
4) Power Rail Imaging
  a. Sagging third rail cover boards
  b. Missing or broken power rail retainer clips (third rail anchors)
  c. Completely broken, moved, or missing third rail pots
  d. Significant, visible third rail surface anomalies

If change-based processing is deployed in the rail sector, the probability of missing relevant changes is expected to be low enough to justify a no-change-no-review policy. Under such a policy, manual review is applied only where change is detected. Such a policy leads to a cost-effective track inspection method only if the number of non-relevant detected changes is manageable. Although a need for managing non-relevant changes was identified during this study, three approaches outlined below were assessed and may address this need.

Approach 1 - Automate Co-Registration

a. False changes along tie edges were a leading source of non-relevant, false changes in this study.
b. When “Before” and “After” images are properly co-registered using automated, co-registration, false changes along tie edges are expected to be significantly eliminated.

Approach 2 – Automate Compensation for Light Exposure During Line Scan Imaging

a. False changes from over exposure were the second-most prevalent cause of non-relevant, false change within the assessed data sets.
b. Automated compensation for light exposure during rail-based, line scan imaging is expected to be achievable and would significantly assist in managing false changes caused by improper light exposure.

Approach 3 – Apply Techniques to Filter Non-Relevant Changes

a. Region-Based Change Detection - A technique described in Section 5.3, referred to as region-based change detection, can provide order-of-magnitude reductions in non-relevant changes while simultaneously raising the probability of detecting relevant

---

[6] Detecting small changes in concrete tie crack growth is expected to be possible by applying region-based change detection to individual ties. In doing so, size and intensity thresholds optimized for detecting small crack growth would be used. In addition, it is expected that thematic ballast classification would be able to minimize false changes caused by differences in ballast on and around the perimeter of each tie.

[7] Detecting rail base deterioration is expected to be possible by comparing “Before” and “After” image sets spaced by a significant amount of tie (e.g., several years).

changes (see results presented in Section 5.3.3 for details). When applied to a specific task of detecting missing and rotated fasteners, region-based change detection eliminated the need to review 98.5 percent of the total fasteners without missing any relevant changes. While region-based change detection is expected to involve machine vision algorithm development, the cost of such development has been conservatively estimated to be one-third that of current-day algorithm development. The expected savings is based on findings indicating that it is easier to filter and manage non-relevant changes than it is to detect many different conditions using conventional algorithms.

b. Thematic-Based Change Detection - The performance of the deep-learning-based, thematic classification algorithm assessed during this study was shown to improve significantly in response to increasing the number of images used during training. Based on this finding, the approach can reasonably be expected to ultimately add value during image-based change detection. Thematic-based change detection is expected to add value primarily by facilitating the process of identifying and filtering non-relevant intensity-based changes.

7.1 Perceived Strengths and Weaknesses

Machine-vision-based track inspection is a complex, multifaceted problem. No system capable of addressing all related aspects of the problem is known to exist. Thus, like other machine vision approaches available today, image-based change detection has strengths and weaknesses, as further clarified below.

7.1.1 Strengths

1) Image-based change detection can detect many conditions relevant to both track inspection and maintenance tracking without a need for expensive algorithm development for each condition.
2) Missed relevant changes are expected to be low enough to justify a no-change-no-review policy, thereby significantly reducing current needs for manual image review associated with machine-vision-based track inspection.
3) The image alignment process associated with image-based change detection facilitates other capabilities including:
    a. Cradle-to-grave asset tracking
    b. Asset degradation rate analysis

Image-based change detection can both augment and benefit from traditional machine vision algorithms.

7.1.2 Weaknesses

1) Although real-time implementation is possible, image-based change detection is better-suited in the near term for office-based processing.

2) Manual image review is needed to locate pre-existing conditions when change-based processing is first deployed and is expected to be needed at periodic intervals (preliminarily once every 6–24 months) thereafter.
3) At least two imaging surveys are needed before change-based processing can be applied.
4) Change-based track inspection may not be possible during rain or snow.

7.2 Identified Areas for Additional Development

This section lists areas where additional development is deemed necessary to realize the ideal change-based track inspection process described in Section 6. Development efforts listed here pertain to:

1) The change detection software evaluated during this study
2) The real-time imaging software used to collect the images evaluated during this study
3) A commercially-available, rail-based image review tool

Identified areas for additional development include:

1) Integrating intensity-based change detection into a pre-existing, rail-based, image review tool:
    a. Include a thematic classification engine such as the one evaluated during this study
    b. Automate the process of co-registering “Before” and “After” images
    c. Automate the process of normalizing intensity in the “Before” and “After” images
    d. Include provisions for accommodating region-based change detection (Section 5.3)
2) Adding automatic, real-time control of gain and light exposure during line scan imaging
3) Adding provisions during change-based processing to allow thematic classification results to be used to automatically ignore non-relevant changes (Section 5.2)

Training the thematic classifier added in Step 1 to automatically identify textures anticipated to be common sources of non-relevant changes.

7.3 Technology Assessment

Within the context of the rail sector, the change detection software evaluated in this study has a technology readiness level of five, indicating that the technology has been demonstrated in a relevant environment. It is estimated that a change-based track inspection process, such as that described in Section 6, could be implemented on a pilot basis within 18 to 36 months, depending on the level of emphasis placed on achieving the goal. The first two areas for additional development listed in Section 7.3 are deemed to be the minimum set of improvements needed to deploy a commercial image-based change detection capability in the rail sector.

7.4 Recommendations

Based on findings from this study, ENSCO recommends continuing efforts directed toward ultimately deploying image-based change detection in the rail sector. Results achieved during this study are deemed strong enough to warrant consideration of deploying a pilot, change-based track inspection program in cooperation with a partner railroad. As a recommended minimum, any such pilot program should include applying change based processing to the following three, continuous, line scan image formats: 1) full-width track images, 2) rail web images, and 3) rail surface images. Additionally, ENSCO recommends considering the use of three-dimensional point cloud sensors and aerial imagery as potential components of any pilot, change-based track inspection program.

Recommended near-term steps related to moving toward a pilot program include:

1) Establishing cost estimates for the improvements listed in Section 7.3
2) Assessing potential benefits of combining two-dimensional and three-dimensional data during change-based track inspection
3) Assessing potential benefits of using aerial imagery during change-based track inspection
4) Establishing a plan for deploying a pilot, change-based track inspection program
5) Implementing the first two improvements listed in Section 7.3

Appendix A. Change Detection Software Assessment

A.1 Harris Report: Introduction

This report is the final deliverable of the Proof of Concept (POC) for Demonstration of Change Detection on Railroad Images project performed by Exelis VIS for ENSCO, Inc. ENSCO, a company that has been providing advanced technology solutions for government and commercial customers, is interested in demonstrating to the Federal Railroad Administration the potential benefits of applying off-the-shelf change detection to track-based images.

The goal of the POC was to assess potential benefits of applying off-the-shelf change detection to track-based images. Exelis VIS was tasked with validating that existing capabilities in Exelis VIS's commercial-off-the-shelf (COTS) software packages can successfully apply change detection to images of railway track to find areas of change. This investigation made use of Exelis VIS COTS products with optimizations for conducting image analysis on track-based imagery with the intent that all processing can be automated in the future.

A.1.1 Change Detection

Change detection is the process of identifying areas within an image that have undergone changes. The ability to identify these changed areas provides analysts valuable information of the possible areas of interest within imagery of the same scene over time. Exelis VIS has experimented with two different types of change detection for this project:

1. Pixel Intensity Based Change Detection
2. Thematic Classification Based Change Detection

A.1.2 Pixel Intensity-Based Change Detection

Pixel intensity based change detection is the process of comparing two images, taken at different times of the same geographic extent, and comparing the pixel intensity values to identify major changes. Thresholding is used to define the degree of the change in pixel values to limit the results to high levels of change.

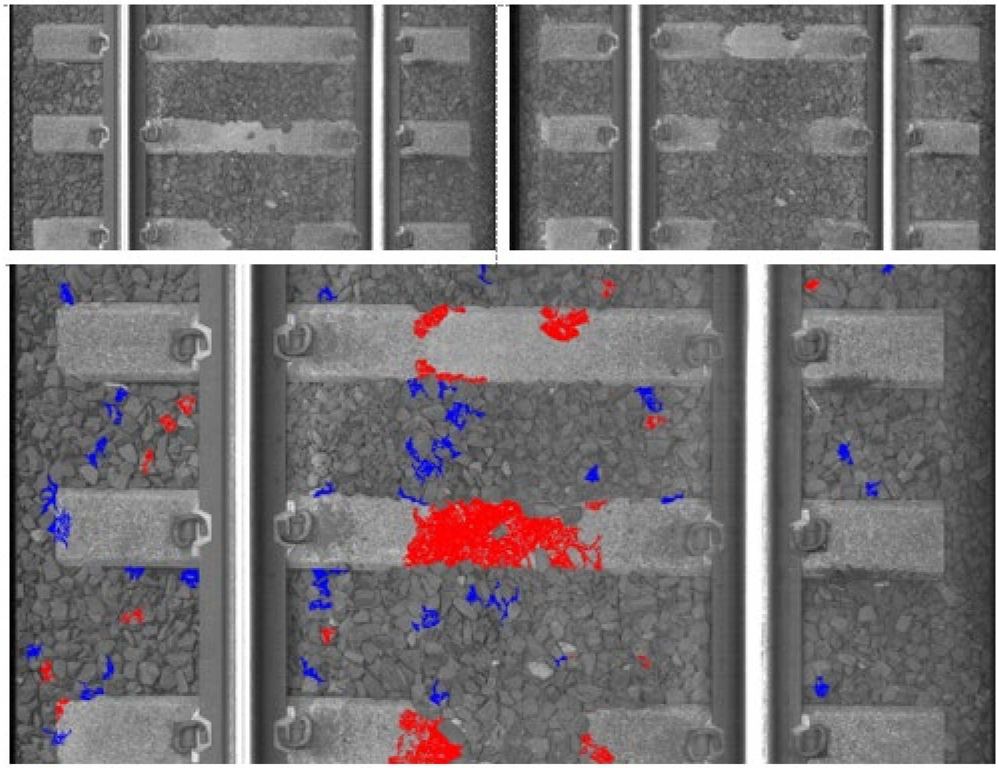

Figure A-1. Example Intensity-Based Change—TIME 1 & TIME 2 (top), Results (bottom)

Notes on processing:

1) Pre-processing included co-registration of the imagery as well as subtracting the mean along tracks for each pixel.
2) The intensity image from TIME 1 is subtracted from TIME 2 following co-registration. A positive difference means TIME 2 is brighter than TIME 1. A positive threshold is applied (+15), and pixels that meet this threshold are analyzed using cluster statistics. Connected groups of pixels are filtered based on the size of the group (number of pixels within the group). Groups that are less than a second threshold (10,000 pixels), are discarded and marked as no-change. This process is repeated for negative changes using a negative threshold (-15). Finally, the cluster statistics (i.e., region sizes) for each frame consisting of 8,192 x 2,048 pixels are recorded. This allows for displaying the largest cluster size for each frame, the average cluster size, the number of clusters, as well as the ratio computed dividing the largest by the average cluster size.

A.1.3 Thematic-Based Change Detection

Thematic-based change detection compares two images to find any differences in material type (or texture) between a "Before" and "After" image.

The first step in this process it is to generate the thematic map (classified images). Exelis VIS used a Machine Learning based technology, available in the ENVI suite of products, to accomplish this task. Machine learning algorithms are designed to simulate human learning processes by associating input and output through a training process. The training process consists of obtaining representative imagery of each target (objects similar in size, shape, and textures) to generate a model of what each target object looks like. For this project, Exelis VIS selected five different materials of interest in the track-based imagery for training:

Table A-1: Summary of Track-related Material Classifications for Training

Class #	Class Description	# of Training Images	Color
0	Rail clip	377	Red
1	Rail	1,064	Green
2	Cement (ties and slab)	958	Blue
3	Crushed stone (track ballast)	242	Yellow
4	Wood (ties)	177	Cyan

The second step in the process is to use the model based on training data to generate a classified image that classifies each pixel based on the closest match to one of the target objects.

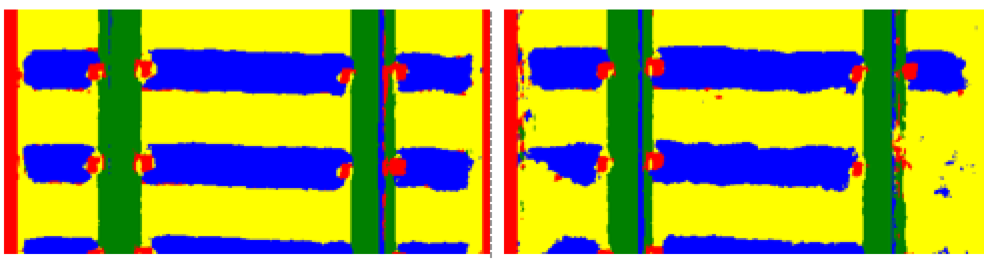

Figure A-2. Example of Thematic Map (Classified Image)—TIME 1 Left, TIME 2 Right

The third step in the process is to compare the two classified images to determine the change between the two.

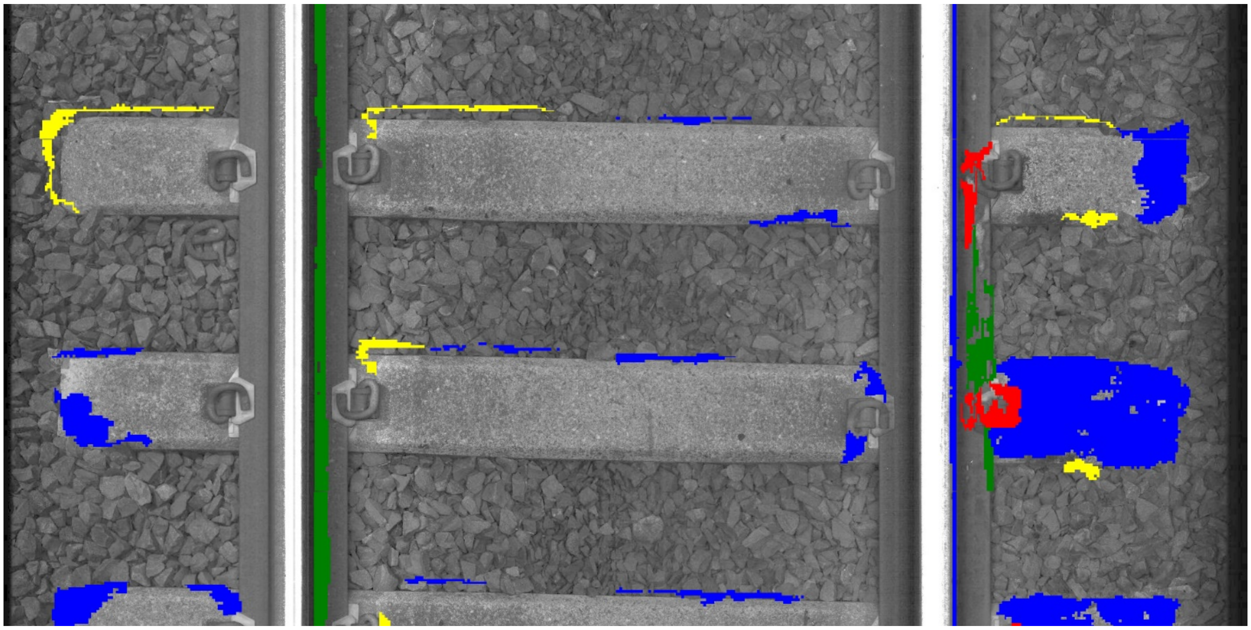

Figure A-3. Example of Change Detection on Thematic Maps

Notes on processing:

1) Unlike pixel intensity based change detection, instead of using thresholds to determine positive and negative changed pixels, any pixel that changed from one classification to another is grouped with other adjacent pixels making the same transition in classification. This means that with five possible classifications for each pixel, the possible transitions can be mapped in a table where the diagonal represents no change.

Table A-2: Classification Transition Matrix Example

From / To	Class 0	Class 1	Class 2	Class 3	Class 4
Class 0	No Change	-	-	-	-
Class 1	-	No Change	-	-	-
Class 2	-	-	No Change	-	-
Class 3	-	-	-	No Change	-
Class 4	-	-	-	-	No Change

1) Each of the 20 empty cells represents a transition that the cluster analysis is applied to.
2) The number of training images varied by class due to needed improvements. A small number of images were evaluated first. Additional images were added to improve the accuracies for each class until the desired results were met.
3) Note that the groupings of red pixels on the extreme right and left of the images that appear as straight lines are an artifact of processing near the edge of the image. These pixels are ignored in the change detection process since they do not change from image to

image and should be ignored when doing visual inspection of the imagery. This edge effect can be removed but was deemed low priority for this proof of concept project.

4) Finally, the cluster statistics are combined for all 20 possible transitions to allow for reporting the largest cluster, the average cluster, and the number of clusters for each frame that was analyzed.

A.2 Harris Report: Data Description

For this project, Exelis VIS used four data sets, two sets of data captures from two different times. Only the imagery was processed and analyzed, metadata (geolocation information) was not utilized.

A.2.1 Data Set 1 (2016/2017)

	2016	2017
Filename	2016062902_TCIS.stream.jpg	2017030801_TCIS.stream.jpg
Date	2016-06-29	2017-03-08
Length (pixels)	32,069,632	13,486,080
Length (meters)	13,819.8	5,810.7
Overlap (pixels)	11,802,624	11,802,624
Overlap (meters)	5,086.1	5,086.1

A.2.2 Data Set 2 (2012/2013)

	2012	2013
Filename	2012080715D_TRK01_DN0208.PGM	2013092409D_TRK01_DN0208.PGM
Date	2012-08-07	2013-09-24
Length (pixels)	5,296,128	3,928,064
Length (meters)	2,282.3	1,692.7
Overlap (pixels)	3,895,296	3,895,296
Overlap (meters)	1,678.6	1,678.6

A.2.3 General Notes

- Note that all distances in meters are estimated based on a square pixel assumption and rail-rail gauge of 1,435 mm.
- For the 2012–2013 pair, an additional pre-processing step was performed to compress the PGM files into the JPG format used in the 2016–2017 data sets. Non-uniformity correction was applied along with the compression. The compression ratio was about 4:1 comparing the PGM file to the stream.jpg.
- For both pairs the camera setup was similar in that there were four cameras each capturing 2,048 pixels in the X-direction (across track). Lengths along the track (Y-direction) were obtained using the motion of the vehicle while continuously capturing a narrow strip of 8,192 pixels across (2,048 x 4). Since the JPEG file format does not allow

for more than 65,535 pixels in any dimension, the long captures were divided into frames of 8,192 x 2,048 each. The “stream.jpg” formatted files consist of thousands of individual JPEG frames concatenated into a single file.

A.3 Harris Report: Pre-Processing Techniques

To obtain the best results, several pre-processing techniques were used on the data. These techniques are described in the following sections.

A.3.1 Frame-Based Processing

To facilitate and automate the pre-processing steps as well as the change detection calculations, the TIME 1 data set is divided into frame sizes of 8,192 x 2,048. The following processing steps are applied to each frame individually, with the exception that the co-registration uses an additional 400 pixels of overlap from one frame to the next frame to ensure continuity in the co-registration information. Overlap is defined when some of the same pixels are used for processing in two adjacent image frames. Because processing is done individually on frames, it means that the computations can be scaled by splitting workload between multiple CPU's to process data more quickly. It also makes processing more tolerant to errors or problems with an individual frame.

A.3.2 Co-Registration

To perform any type of change detection, an accurate co-registration is needed between the data sets being compared. A proven ENVI algorithm was used to compute co-registration information for roughly every 1,024 pixels along the driving direction (along tracks).

For the co-registration algorithm to run, 22 manual anchor points were selected for the 2016–2017 data sets, and a single anchor point was selected for the 2012–2013 data sets. The following graphs show the differences in positions found by the co-registration process. The difference is due to different image capture rates relative to the vehicle driving speed.

The following graph demonstrates the cumulative difference in 1,000's of pixels between 2012 and 2013, starting from a single manually selected anchor point (at 0,0).

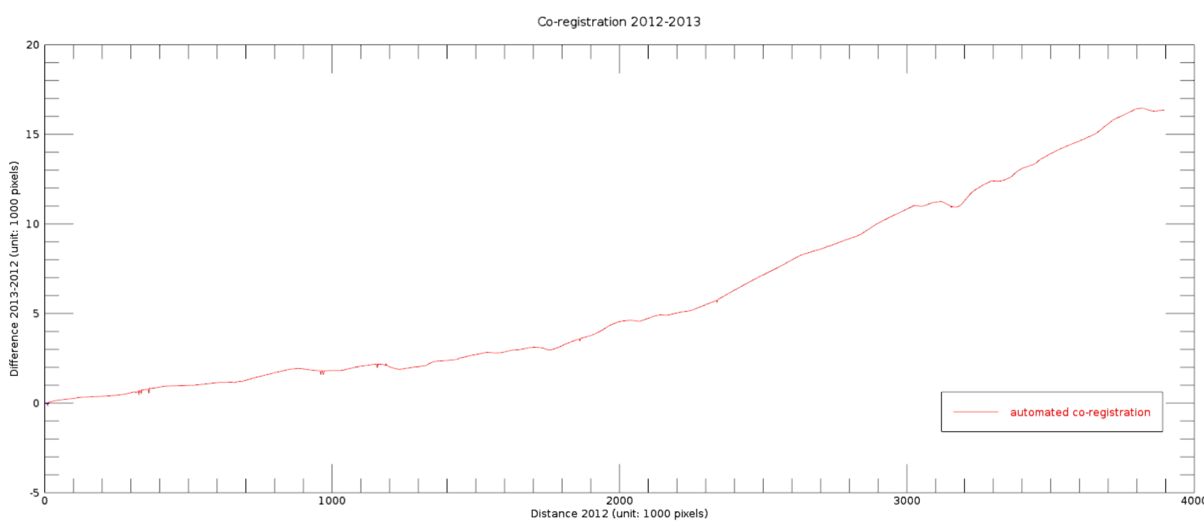

Figure A-4. Difference in Capture Rate between 2012 and 2013

Similarly, for the 2016 to 2017 image captures, the difference is shown in the graph below. The manually selected anchor points were needed in places where there were substantial changes in the imagery, such that the correlation metrics were low.

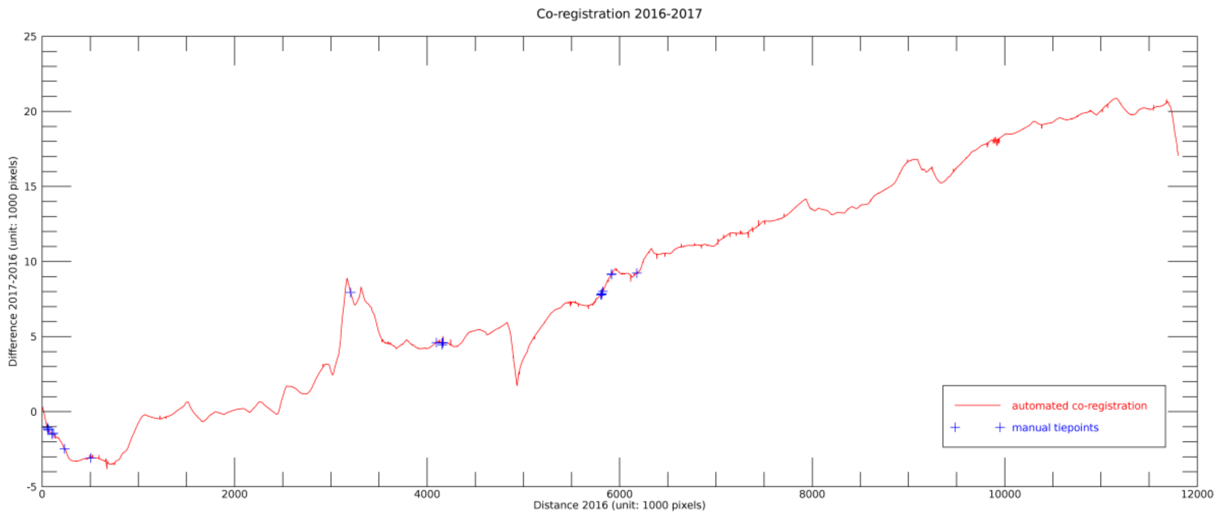

Figure A-5. Difference in Capture Rate between 2016 and 2017

In addition, a single shift in X (across the tracks) was calculated for each of the four cameras at intervals of 2,048 pixels along the tracks. The data sets are divided into frames of size 8,192 x 2,448 pixels, with a 400-pixel overlap along the tracks such that frames are spaced 2 048 pixels in the TIME 1 data set. The algorithm uses a windowing cross-correlation metric to obtain the best solution for each frame. The solution consists of four shift values (in X) for the cameras, and three Y-position pairs for each camera.

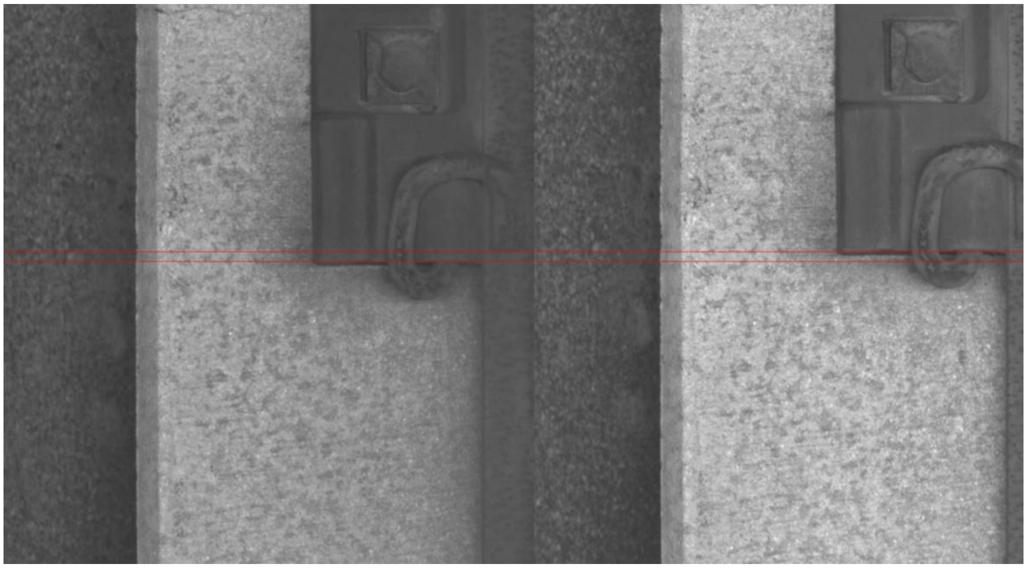

Figure A-6. Pixel Misalignment in the Y-direction (Along Track)

This shows a portion of camera 1 from both 2016 (left) and 2017 (right). In this case, there is an 18-pixel difference in the Y position between the two times. The co-registration algorithm is identifying this mismatch and storing the positional information in a table. We found that there is a need to identify registration for each of the four cameras, since they are not staying perfectly synchronized throughout the data collection. After applying an 18-pixel correction, the alignment is correct.

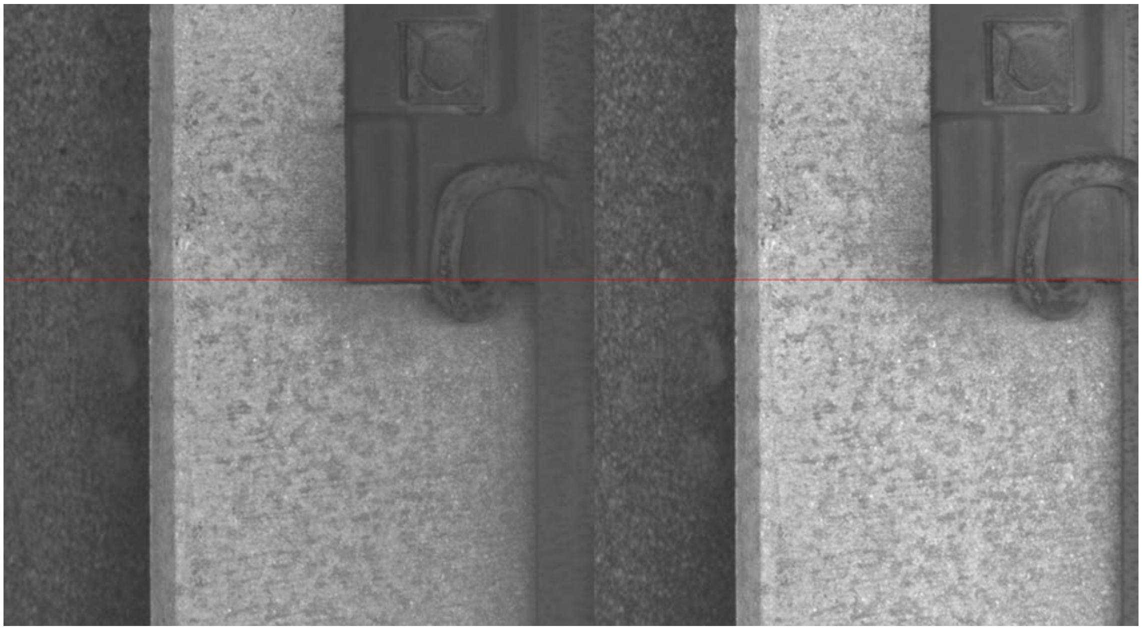

Figure A-7. Alignment Obtained Using ENVI Cross-Correlation Method

Similarly, correction is computed for X (across tracks) and the correction information is stored in a text file. The viewer software is reading the correction file and applying the correction when using the curtain function. The same correction is also applied as a pre-processing step to the change detection algorithm.

When co-registration is less accurate, the change detection will be affected in a way that results in a halo-effect around objects in the imagery. An example of a slight co-registration mismatch is shown in the image below. Notice the blue and red lines at the top and bottom of the ties. These lines indicate that the software found a change in that area, but it is most likely caused by slight errors in the co-registration rather than a real shifting of the ties.

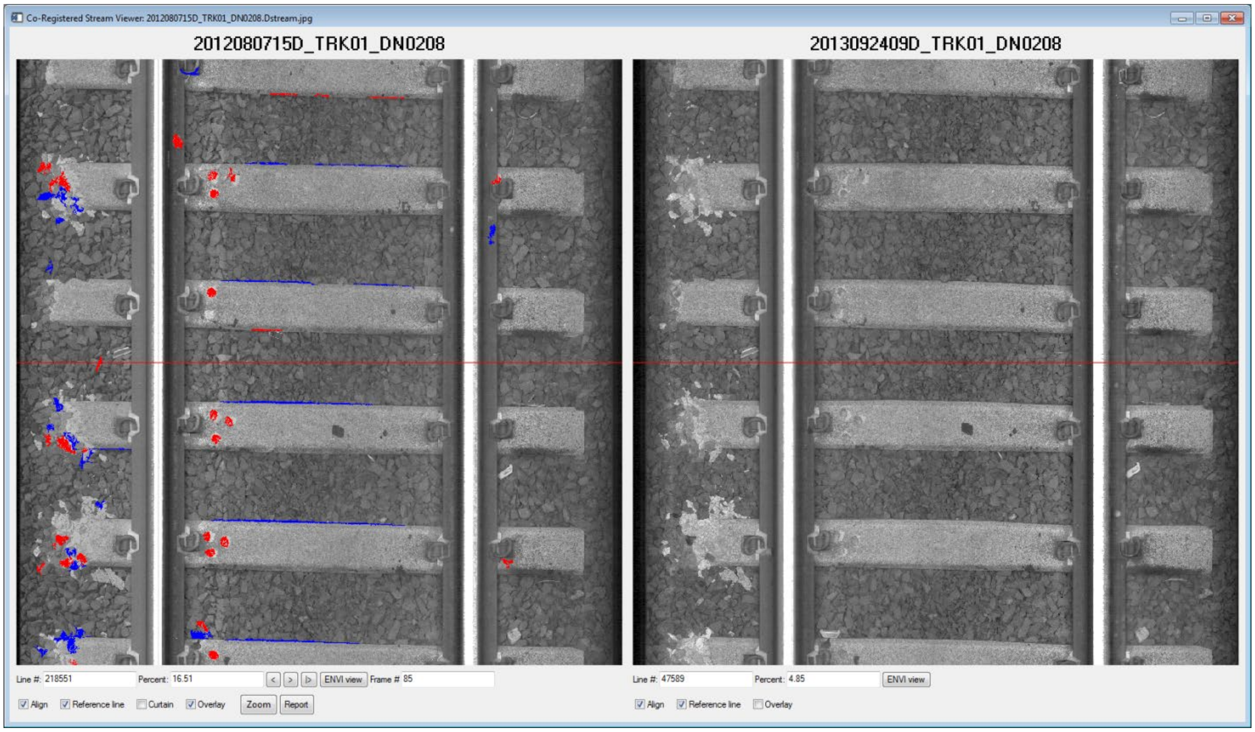

Figure A-8. Example of Small Co-Registration Mismatch

A.3.3 Intensity Normalization

The cameras used for this data collection have a detector array of either 2,048 x 1 or 2,048 x 2 pixels, and the data capture rely on movement along the tracks (Y) to obtain a continuous image stream. If the camera response is not uniform across the detector array, then that will result in a striping or banding effect in the imagery since the same detector pixels always are shown in the same X (across track) position in the image. The non-uniformity can be caused by either the lens (external) or the electronics in the detector array. In the example data that we received, there is significant non-uniformity visible in the 2012–2013 data set, but no visible effects are seen in the 2016–2017 data set. Presumably this is because better initial camera calibration was performed with the 2016–2017 data set.

Below is an example showing the 2012 data with a red line plot of the average pixel intensity for the whole file (about 5,000,000 lines of pixels). The plot curve is calculated individually for each of the 8,192 pixels across, by taking an average of all 5,000,000 pixels along at each X position.

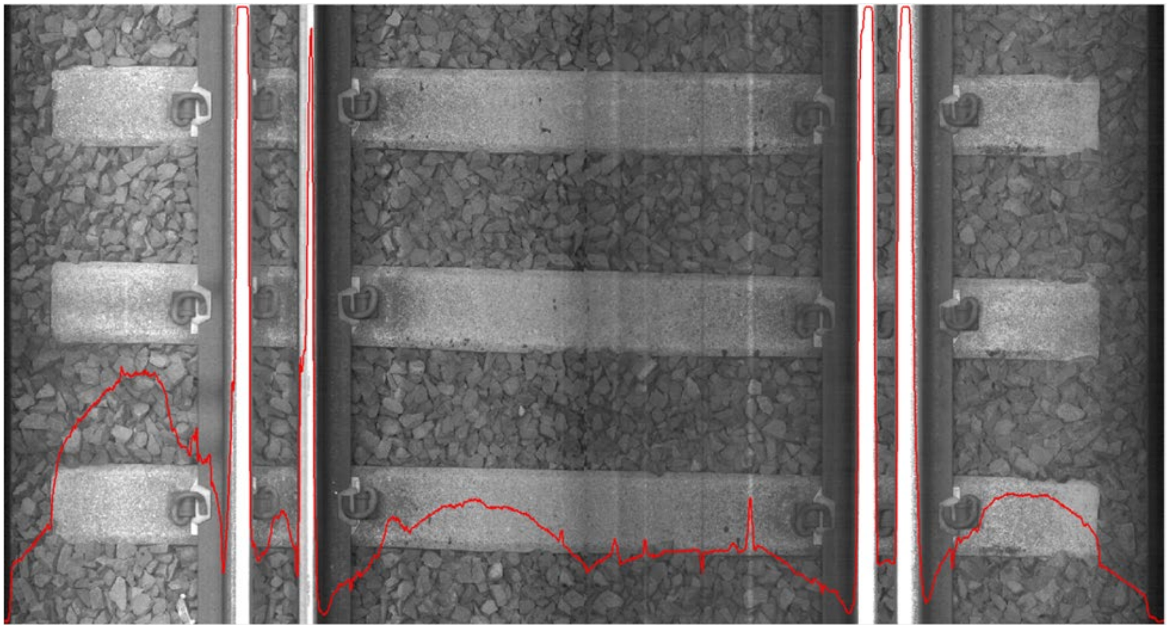

Figure A-9. 2012 Average Intensity Plot along 5,000,000 Pixels

Notice the bright vertical stripe to the right of the center. Also of interest is the relatively high brightness in the leftmost camera compared with the rightmost camera. For comparison, the 2013 data set is shown below.

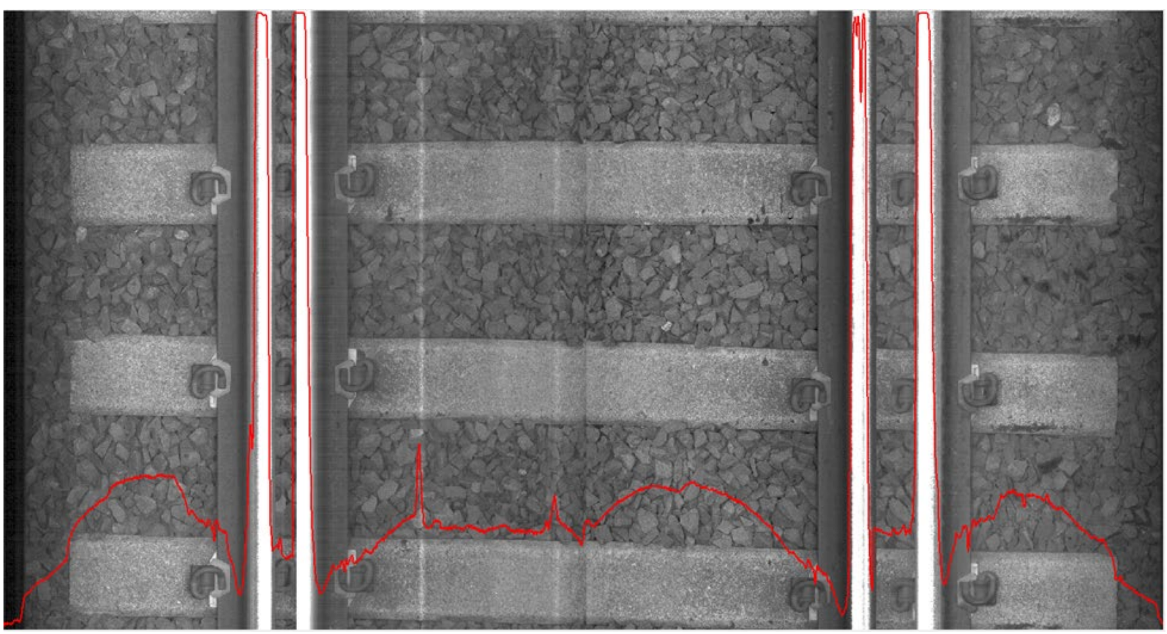

Figure A-10. 2013 Average Intensity Plot along 5,000,000 Pixels

Notice that the location of the brightest vertical stripe has been flipped and appears the same distance from the left edge as it was from the right edge in the previous year. Also, in the 2013 data set, the illumination is much more even between the leftmost and rightmost camera.

To reverse the effects of this non-uniformity, an inverse is computed for the areas away from the rails and this inverse is then applied to make the illumination uniform. The rails must be excluded, because they greatly affect the average intensity, but do not represent a sensor artifact.

The following figure shows the correction factor applied to the 2012 data across the 8,192 pixels. The dip around 5,300 causes the bright vertical stripe to be suppressed.

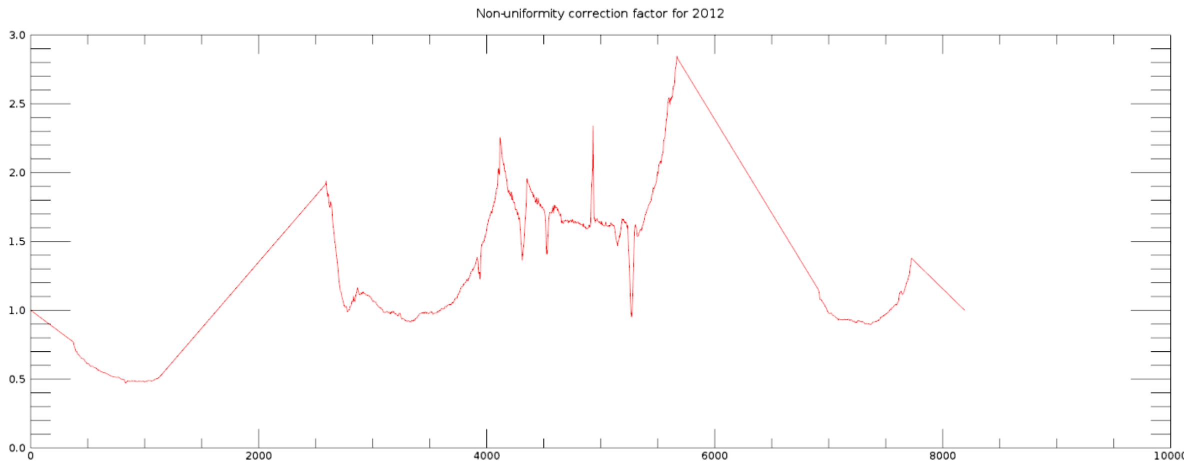

Figure A-11. Non-Uniformity Correction Factor for 2012 Data Set

To make the change detection ignore these types of systemic differences in illumination, a pre-processing step was performed to subtract the mean along Y for each frame being processed.

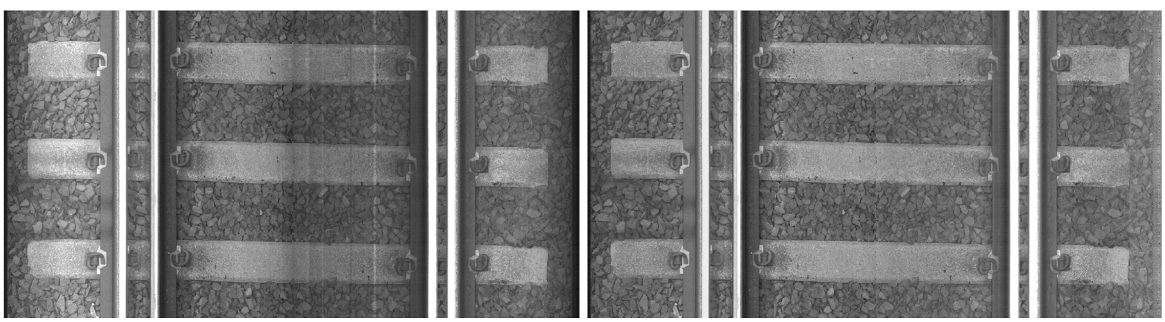

Figure A-12. Example of Removal of Illumination Differences—Before (left), After (right)

A.3.4 Cropping of Overlap Areas

While co-registration and change detection processing is performed on the full width of the four cameras, it is useful to crop out overlapping areas for visualization purposes. Cropping in the X direction (across tracks) is applied equally to the whole data set. A different correction is needed for each data set. The process for this cropping consists of manually selecting the following reference points:

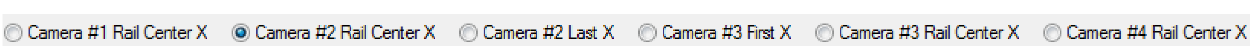

Figure A-13. The Reference Points Needed for Cropping

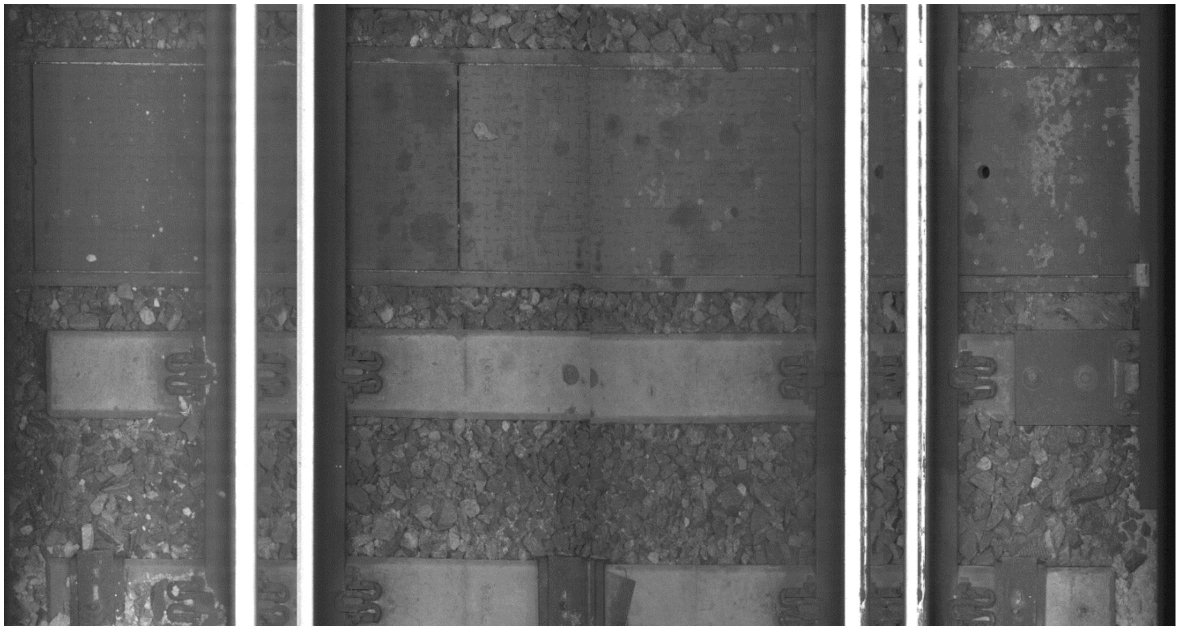

Figure A-14. Example of Four Images before Cropping

The following image shows selection (red lines) of rail center positions to be used for cropping overlap area between camera 1 and camera 2.

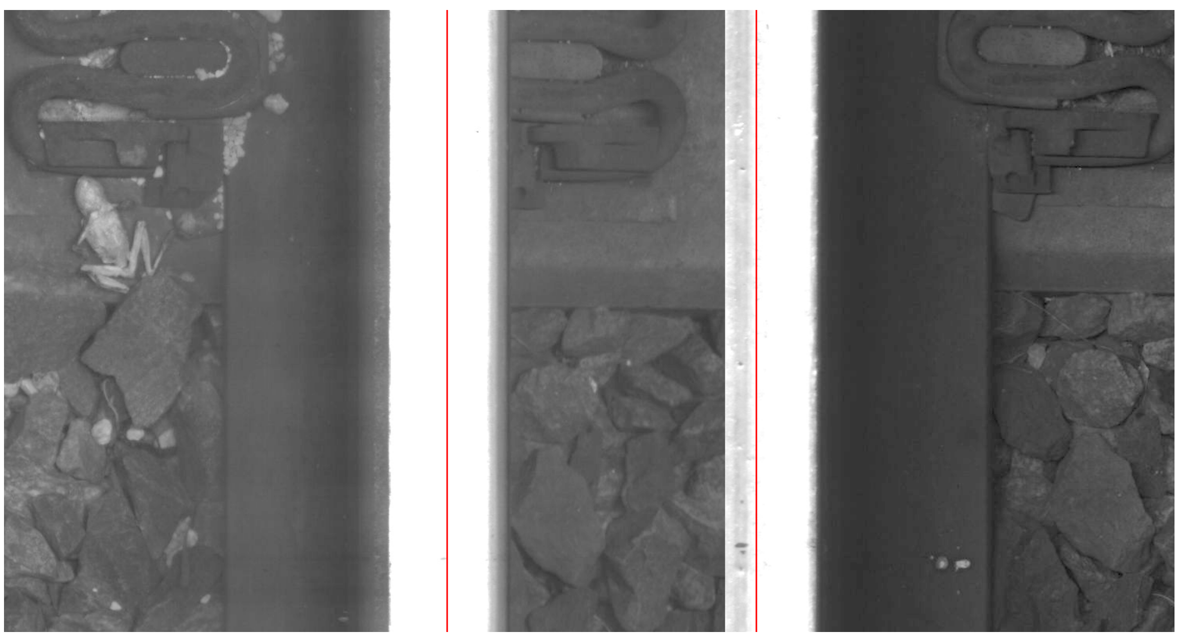

Figure A-15. Example of Cropping Lines to use Between Cameras 1 and 2

Similarly, between camera 2 and camera 3 in the center, there is an overlap area that should be removed for display purposes (red lines indicate selected reference points).

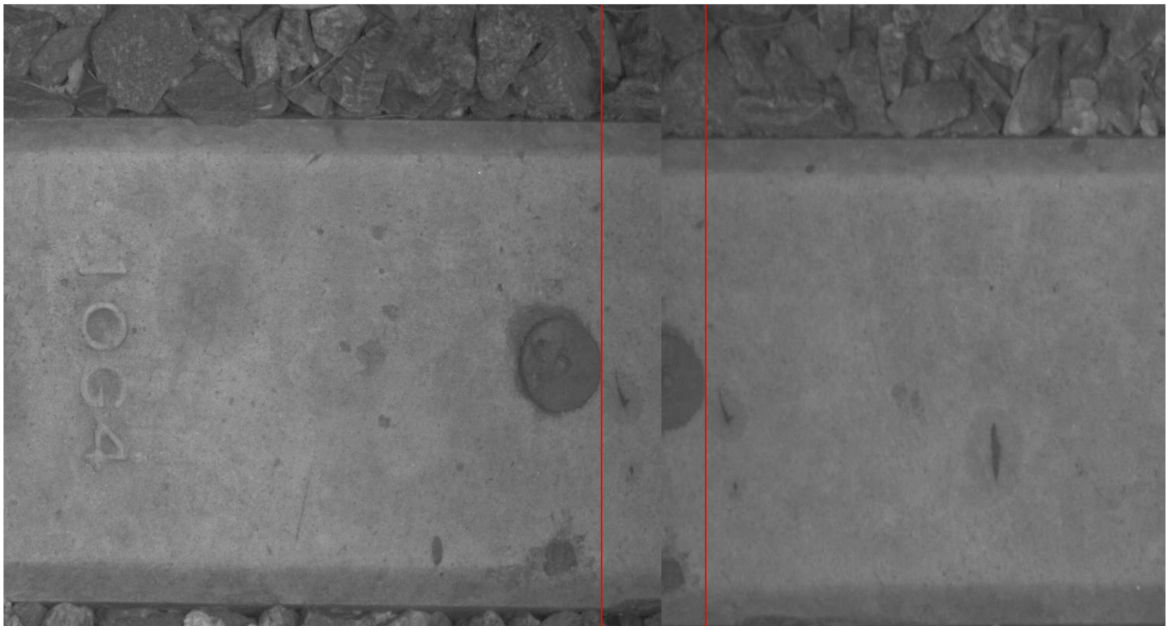

Figure A-16. Example of Cropping Lines to use Between Cameras 2 and 3

Applying this cropping between the cameras result in an image where the overlap is removed:

There are also shifts along track between the cameras. As noted in the co-registration section, these shifts vary along the capture, and are calculated using the automated co-registration algorithm. For an example, shifts in Y can be corrected in the example image above, and the resulting image is shown below.

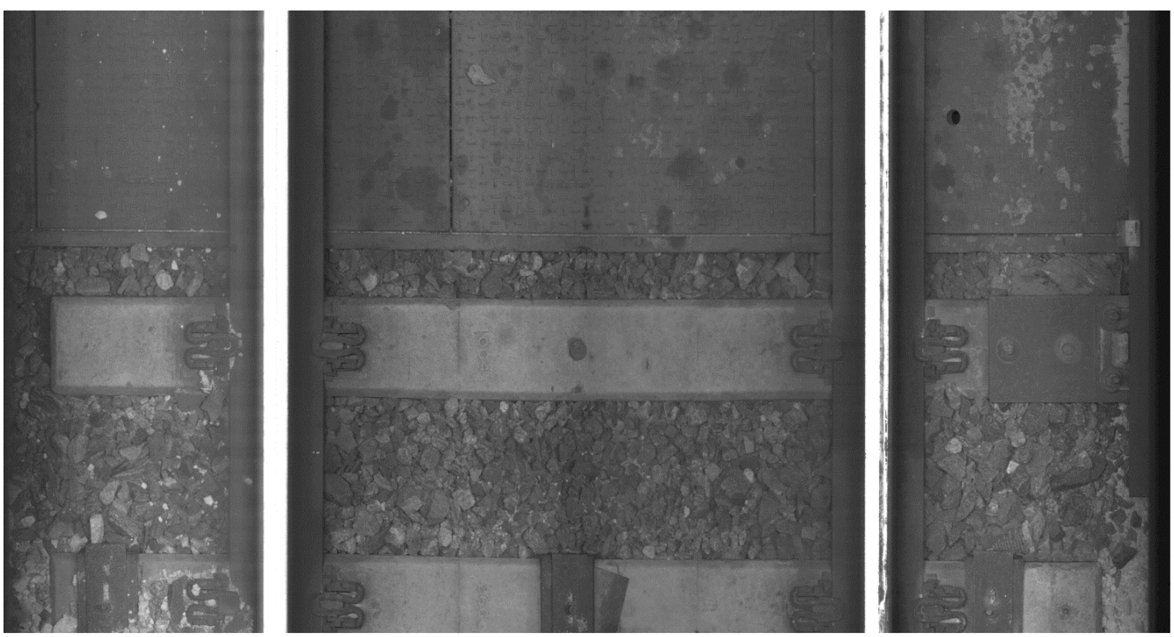

Figure A-17. Example of Cropped Image

The cropping line in the center, between camera 2 and camera 3 is only applicable to a specific ground height. This is because of the parallax between the two cameras. As a result, the cropping will look wrong in some areas where the ground was at a different height compared to the cameras. Specifically, when there are objects located on top of the ballast, they can appear cropped in the viewer.

A.4 Harris Report: Viewing Tool Interface Description

Visual review of the results of the change detection analysis is an important aspect of this project. Due to the nature of the captured imagery, the imagery being much longer than it is wide, an optimized viewer was configured for the project. The following is a description of the interface.

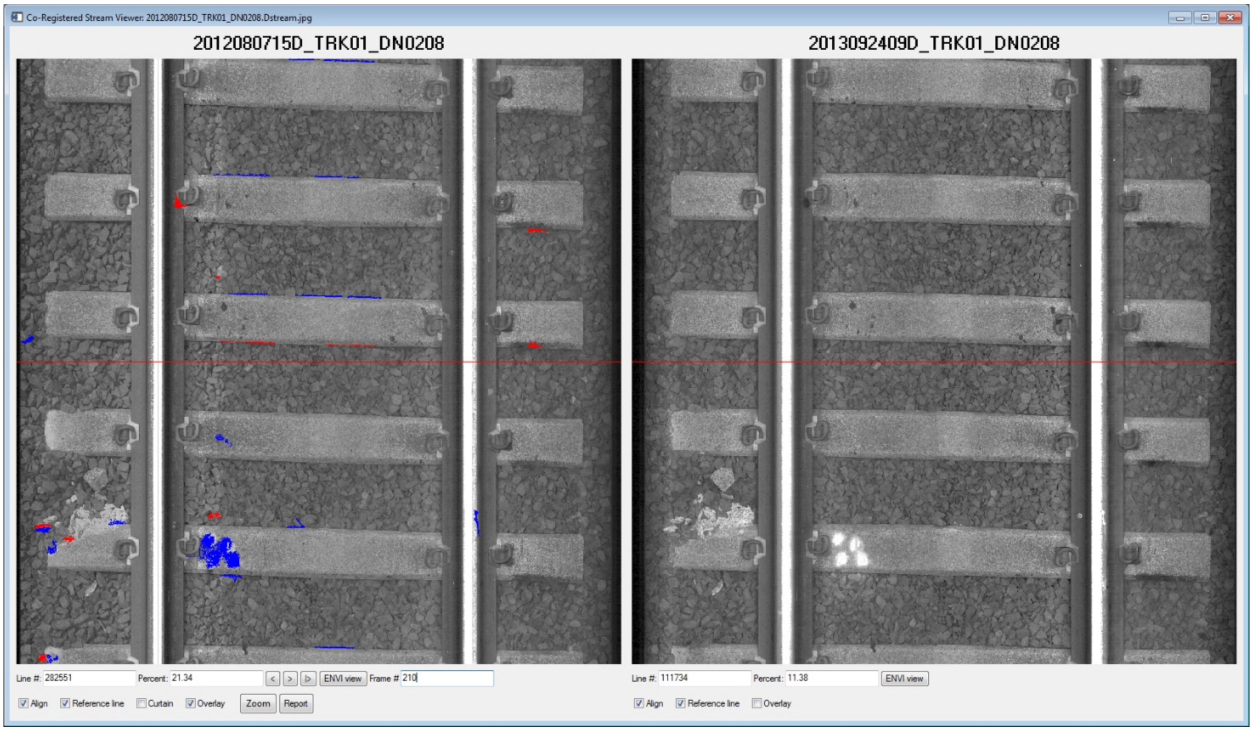

Figure A-18. Viewing Tool for Track-Based Imagery

The above display shows TIME 1 on the left and TIME 2 on the right. File names are shown at the top.

1) Image Display

    a. The display shows TIME 1 on the left and TIME 2 on the right. File names are shown at the top. The imagery is shown with a gamma correction applied to make the darker regions appear brighter. A gamma value of 1/(2.2) was empirically selected based on the appearance of the original imagery. Normally, JPEG image files are stored with gamma-encoded values and not linear intensities, but the example files we received appears to be stored as linear values. While processing was performed on the linear intensities, the following gamma correction was applied to enhance the image display:

    b. DV = floor((PV/255)^(1/2.2) * 255.99), where DV is the display intensity value resulting from the original (PV) pixel intensity value.

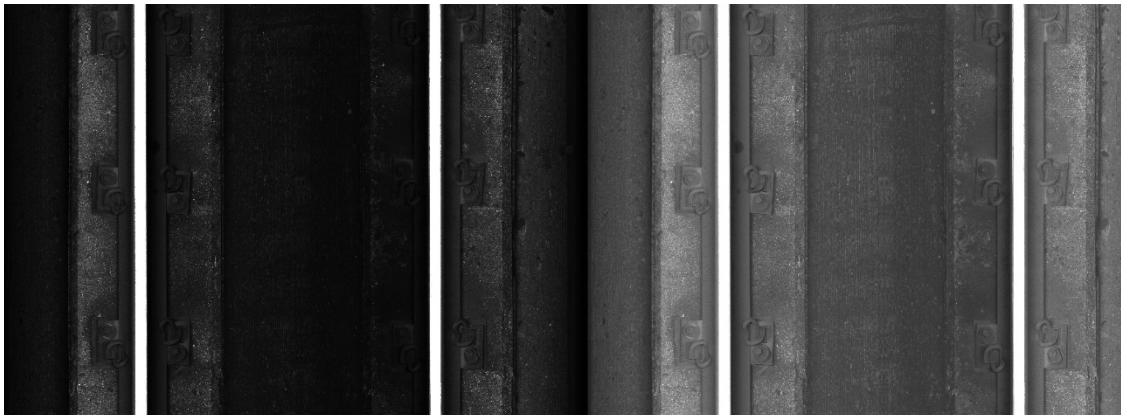

Figure A-19. Illustration of Gamma Correction (right) vs. Original Raw (left)

2) Navigation - There are several navigation controls in the main interface (in addition to navigation from the table described in the next section).
   a. Navigation by line number option, allows navigation to a specific line (pixel row) number in either the downsampled (Dstream) or full resolution (stream) image file. The user can type in a number and hit Enter to jump to a line.
   b. Navigation by percent option, allows navigating to a specific percentage of the full stream image file.
   c. The Arrow forward/backward, navigates forward/backward by one screen (four frames).
   d. The Play button starts automatic (animation) forward navigation at a constant speed. Pressing the Play button again, causes the animation to stop.
   e. Navigation by frame number, allows jumping to a specific frame number. Frames are relative to the first co-registration tie point (first common point). The frames are 2048 pixels in the TIME 1 reference image stream, and the TIME 2 right side image will follow to the same co-registered location. Change detection statistics are reported per frame.
3) Overlays – There are 3 types of overlays that can be shown. An overlay type is selected at the time the application starts, by selecting a corresponding .view file.
   a. Pixel intensity change can be selected by opening the “pixel-intensity-change.view” file when prompted. This overlay is only appropriate for TIME 1 because TIME 2 has been warped to match TIME 1 during the change calculation. As a result, this overlay will only be shown in the left view.
   b. Thematic change can be selected by opening the “thematic-change.view” file when first prompted. This overlay is also only appropriate for TIME 1 because TIME 2 has been warped to match TIME 1 during the change calculation. As a result, this overlay will only be shown in the left view.
   c. Classification overlays can be selected by opening the “thematic-classification.view” when prompted.

   d. Overlays (for both TIME 1 and TIME 2) can be turned on or off using the checkboxes on the interface.

4) Reference line – The reference line is used to verify proper co-registration between the 2 times. The reference line can be shown or hidden using the checkbox.

5) Curtain – The curtain checkbox allows the user to drag a vertical green line from side to side to gradually shift between TIME 1 and TIME 2 in the display. This functionality relies on the automated co-registration pre-processing to properly align TIME 1 and TIME 2 in order to show them in the same view. See curtain-step-1, -2, and -3.

6) Table Display - The table display allows for quick sorting of the frames (represented as one frame per row) based on any of the statistical metrics displayed in the columns. The columns are as follows:

   a. Frame number – This column shows the frame number. Each frame consists of 2048 pixels in the TIME 1 imagery, and the corresponding pixels in the TIME 2 imagery. Sorting by this column will simply sort the table spatially along the tracks.

   b. Correlation – This column shows a metric for how well the automated co-registration performed for each frame. A low number may indicate significant changes because the co-registration resulted in poor correlation between TIME 1 and TIME 2.

   c. # regions – This column shows the total number of connected regions. A region is defined as a group of pixels that met the threshold for change in intensity. Frames with a large number of regions represent areas of many changes.

   d. # pixels – This column shows the total number of pixels that met the change threshold criteria within each frame. If a large connected region changed, this number may be high while the previous column (# regions) may still be low.

   e. Largest (PI) – This column lists the size of the largest single region of change in pixel intensity within each frame. Sorting by this column will allow quick navigation to the largest connected regions of pixel intensity change, this feature can help reduce false positives in the change detection.

   f. Average (PI) – This column shows the average region size within each frame. The regions here refer to connected groups of pixels with a significant change in intensity.

   g. Ratio (PI) – This column represents the Largest (PI) divided by the Average (PI). This number will be low if all regions are of a similar size, and will be high if size varies greatly within the frame.

   h. Largest (Theme) – This column lists the size of the largest single region of change in thematic classification between TIME 1 and TIME 2. This allows quick navigation to the largest changed regions based on the thematic classification.

   i. Average (Theme) – This column shows the average region size for the change in the thematic classification. In this context, a region is defined as a connected group of pixels changing from a specific thematic class to another specific thematic class. Each pair of thematic classes can potentially have their own distinct regions that go into computing this statistic.

j. Ratio (Theme) - This column represents the Largest (Theme) divided by the Average (Theme). This number will be low if all regions are of a similar size, and will be high if size varies greatly within the frame.

k. Thematic type – This column shows the type change for the largest region in each frame. The naming format is class name for TIME 1 followed by class name for TIME 2.

7) Table Navigation – After selecting a sorting column, navigation to a specific frame (row) can be done by clicking on an individual table cell. Navigation is also possible using the keyboard arrow keys (up/down).

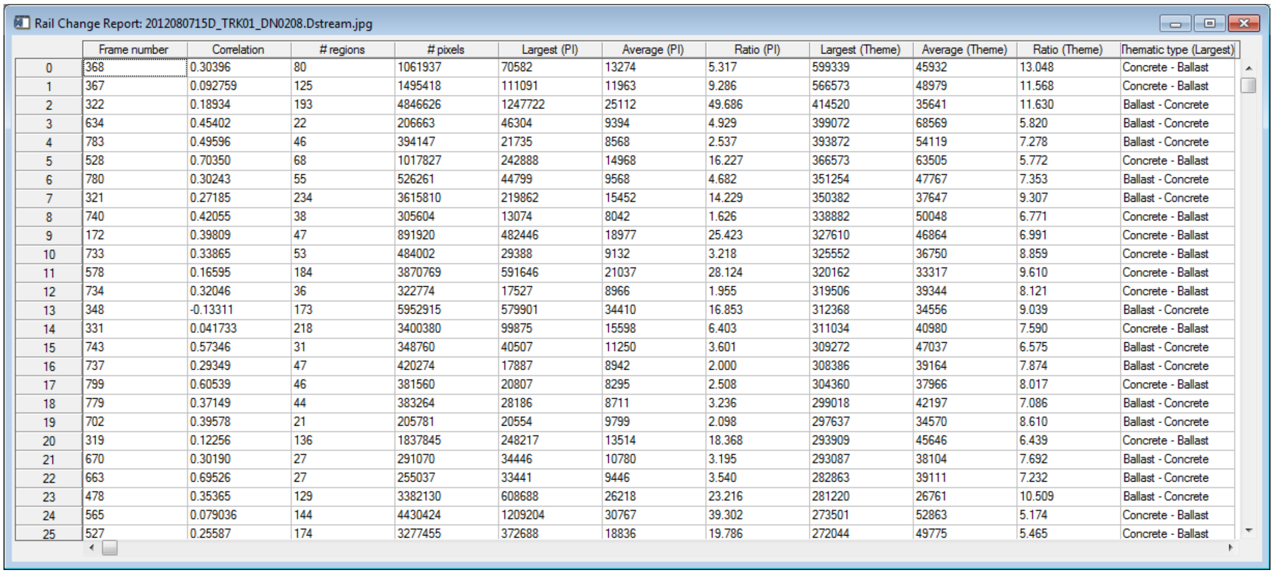

Rail Change Report: 2012080715D_TRK01_DN0208.Dstream.jpg

	Frame number	Correlation	# regions	# pixels	Largest (PI)	Average (PI)	Ratio (PI)	Largest (Theme)	Average (Theme)	Ratio (Theme)	Thematic type (Largest)
0	368	0.30396	80	1061937	70582	13274	5.317	599339	45932	13.048	Concrete - Ballast
1	367	0.092759	125	1495418	111091	11963	9.286	566573	48979	11.568	Concrete - Ballast
2	322	0.18934	193	4846626	1247722	25112	49.686	414520	35641	11.630	Ballast - Concrete
3	634	0.45402	22	206663	46304	9394	4.929	399072	68569	5.820	Ballast - Concrete
4	783	0.49596	46	394147	21735	8568	2.537	393872	54119	7.278	Ballast - Concrete
5	528	0.70350	68	1017827	242888	14968	16.227	366573	63505	5.772	Concrete - Ballast
6	780	0.30243	55	526261	44799	9568	4.682	351254	47767	7.353	Ballast - Concrete
7	321	0.27185	234	3615810	219862	15452	14.229	350382	37647	9.307	Ballast - Concrete
8	740	0.42055	38	305604	13074	8042	1.626	338882	50048	6.771	Concrete - Ballast
9	172	0.39809	47	891920	482446	18977	25.423	327610	46864	6.991	Concrete - Ballast
10	733	0.33865	53	484002	29388	9132	3.218	325552	36750	8.859	Concrete - Ballast
11	578	0.16595	184	3870769	591646	21037	28.124	320162	33317	9.610	Concrete - Ballast
12	734	0.32046	36	322774	17527	8966	1.955	319506	39344	8.121	Concrete - Ballast
13	348	-0.13311	173	5952915	579901	34410	16.853	312368	34556	9.039	Ballast - Concrete
14	331	0.041733	218	3400380	99875	15598	6.403	311034	40980	7.590	Concrete - Ballast
15	743	0.57346	31	348760	40507	11250	3.601	309272	47037	6.575	Ballast - Concrete
16	737	0.29349	47	420274	17887	8942	2.000	308386	39164	7.874	Ballast - Concrete
17	799	0.60539	46	381560	20807	8295	2.508	304360	37966	8.017	Concrete - Ballast
18	779	0.37149	44	383264	28186	8711	3.236	299018	42197	7.086	Ballast - Concrete
19	702	0.39578	21	205781	20554	9799	2.098	297637	34570	8.610	Ballast - Concrete
20	319	0.12256	136	1837845	248217	13514	18.368	293909	45646	6.439	Concrete - Ballast
21	670	0.30190	27	291070	34446	10780	3.195	293087	38104	7.692	Ballast - Concrete
22	663	0.69526	27	255037	33441	9446	3.540	282863	39111	7.232	Ballast - Concrete
23	478	0.35365	129	3382130	608688	26218	23.216	281220	26761	10.509	Ballast - Concrete
24	565	0.079036	144	4430424	1209204	30767	39.302	273501	52863	5.174	Concrete - Ballast
25	527	0.25587	174	3277455	372688	18836	19.786	272044	49775	5.465	Concrete - Ballast

Figure A-20. Table Display

A.5 Harris Report: Intensity-Based Change Detection Results

The following examples of change are provided as samples of overall results of the project for pixel intensity based change detection.

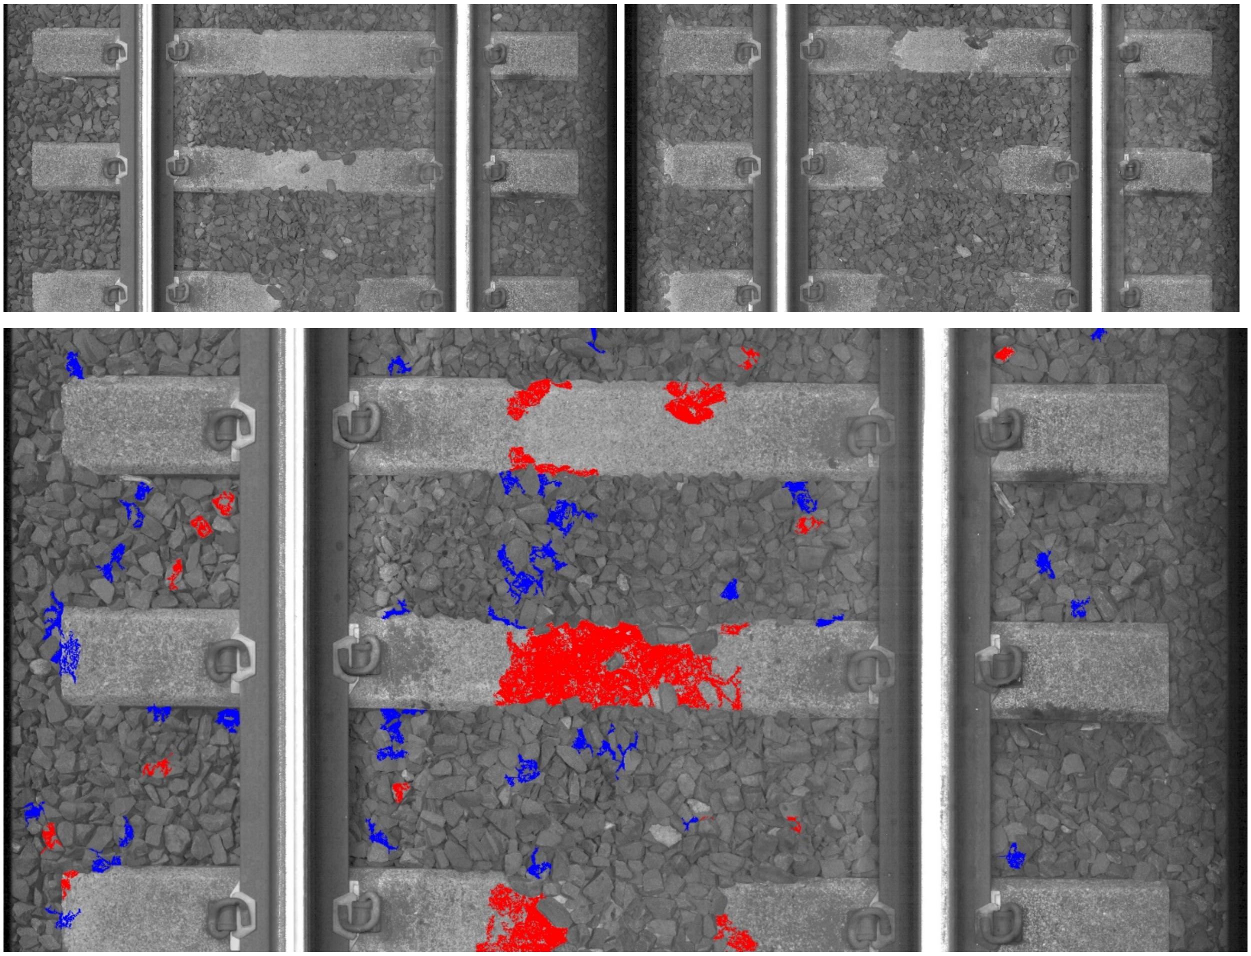

Figure A-21. Intensity-based Change Example 1: Changed Ballast (Set: 2012/2013, Frame: 0172)

Red showing change in ties and blue showing change in rock ballast.

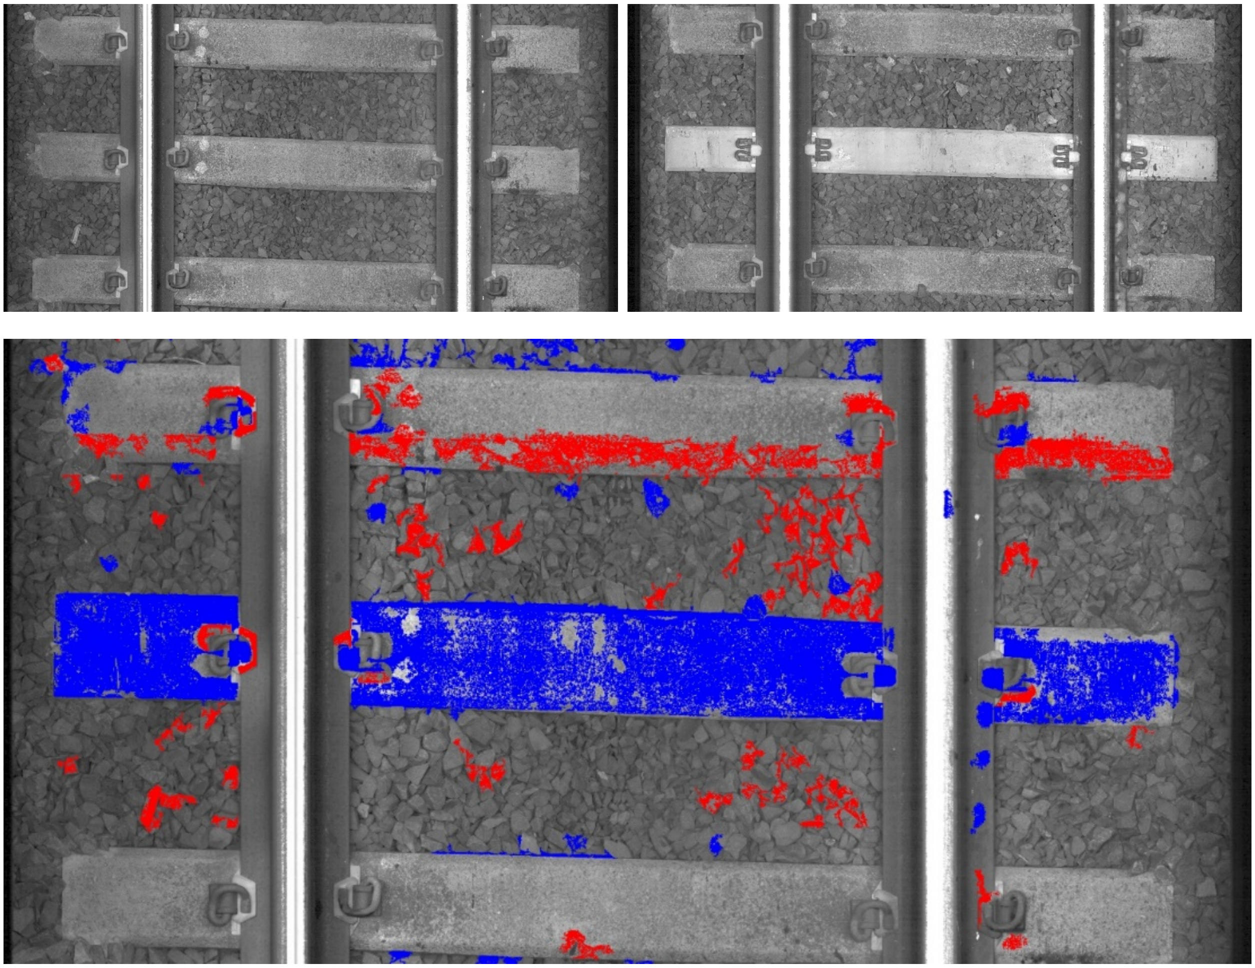

Figure A-22. Intensity-based Change Example 2: Replaced Tie
(Set: 2012/2013, Frame: 0468)

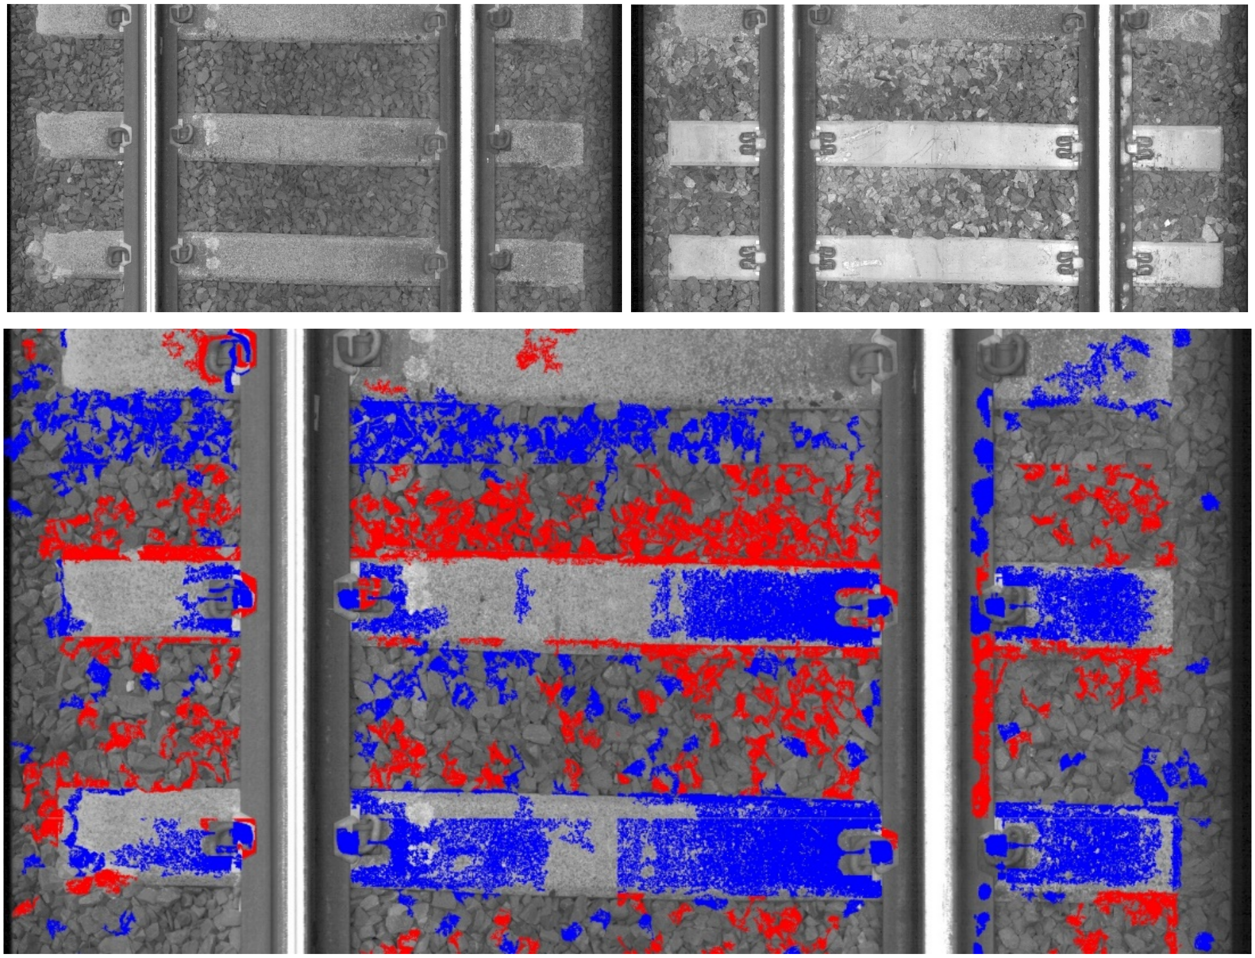

Figure A-23. Intensity-based Change Example 3: Replaced Tie (Set: 2012/2013, Frame: 0579)

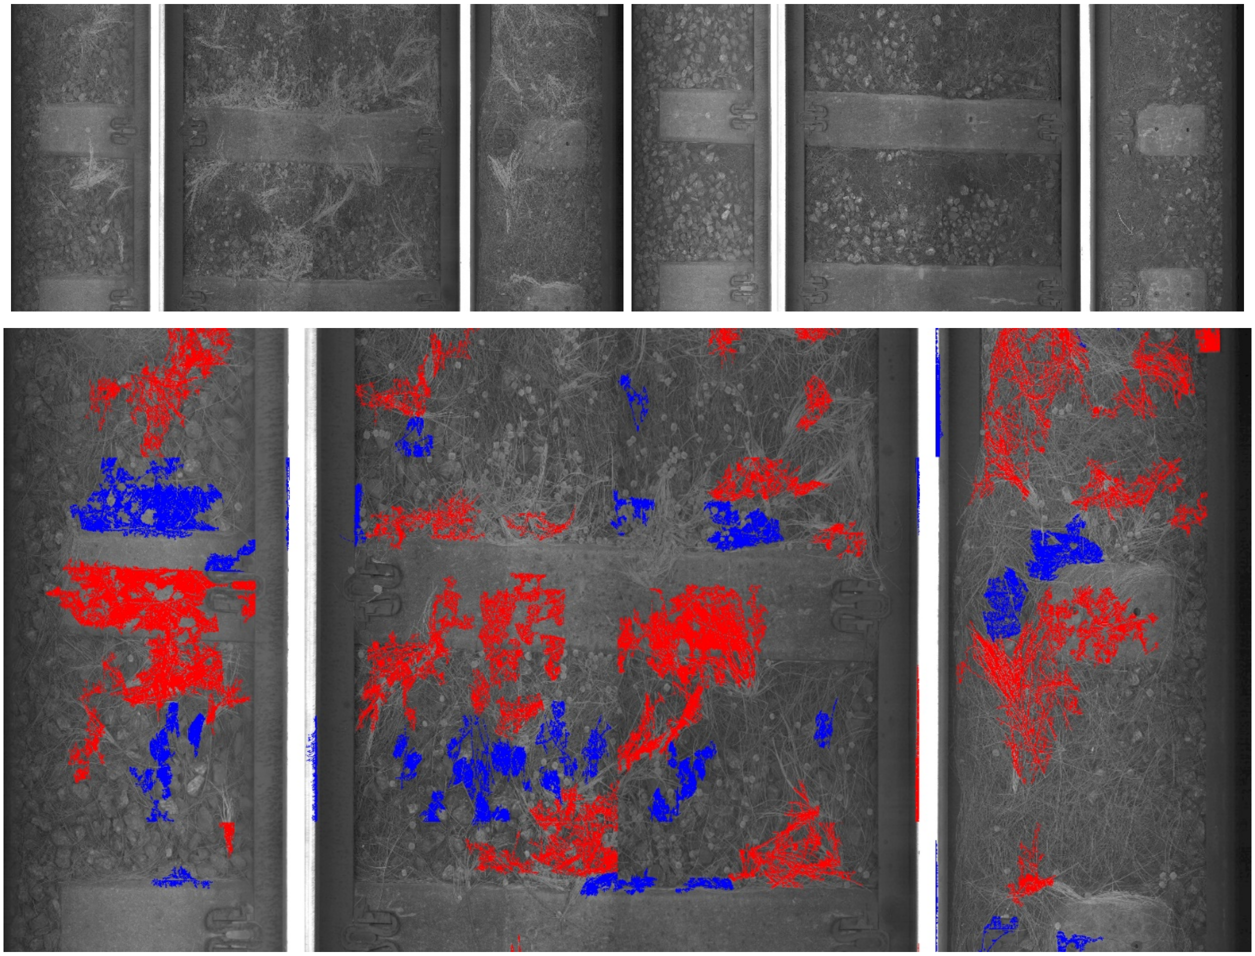

Figure A-24. Intensity-based Change Example 4: Vegetation Encroachment (Set: 2016/2017, Frame: 0326)

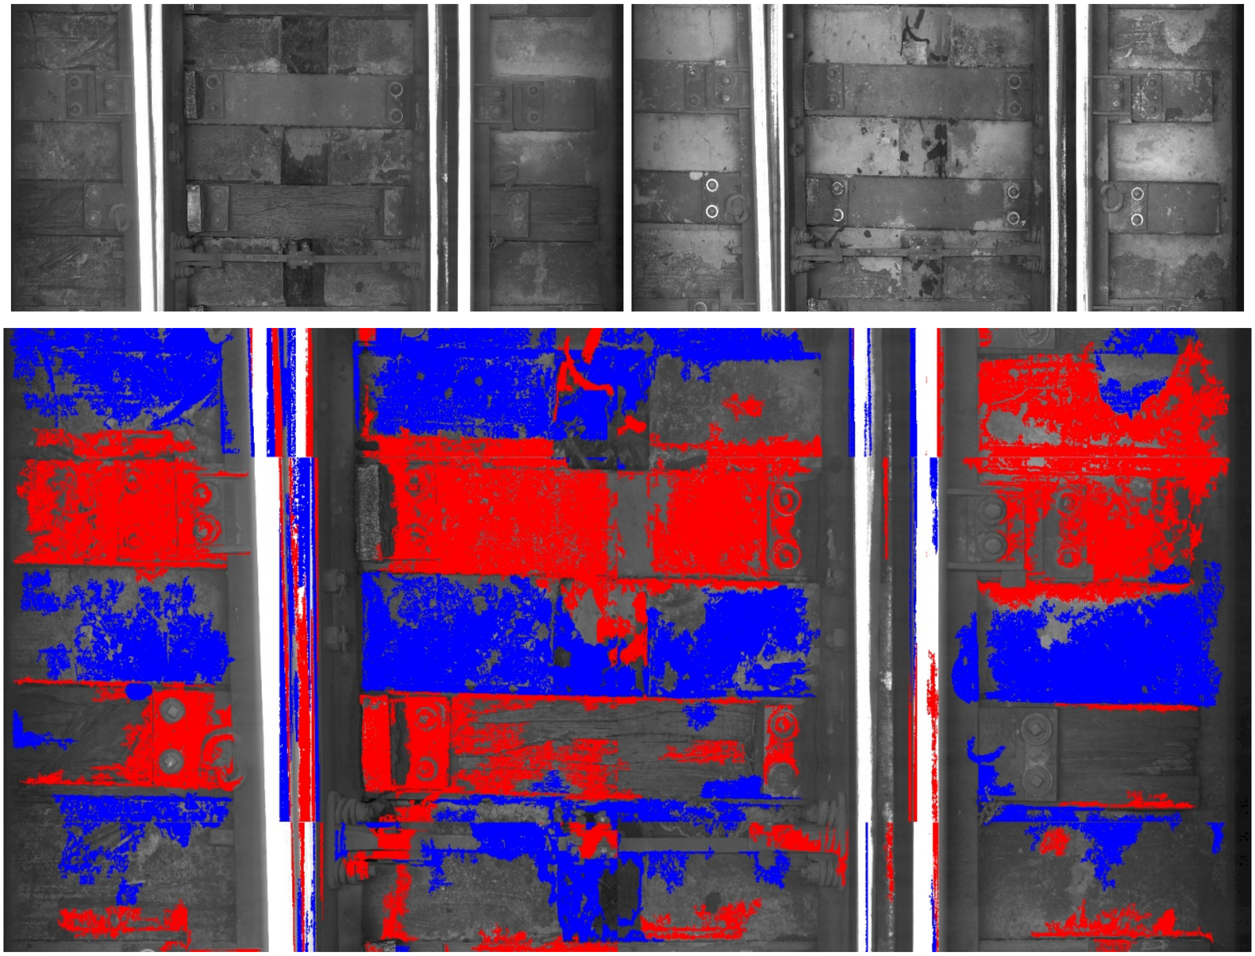

Figure A-25. Intensity-based Change Example 5: Tie Change (Wood to Concrete) (Set: 2016/2017, Frame: 2029)

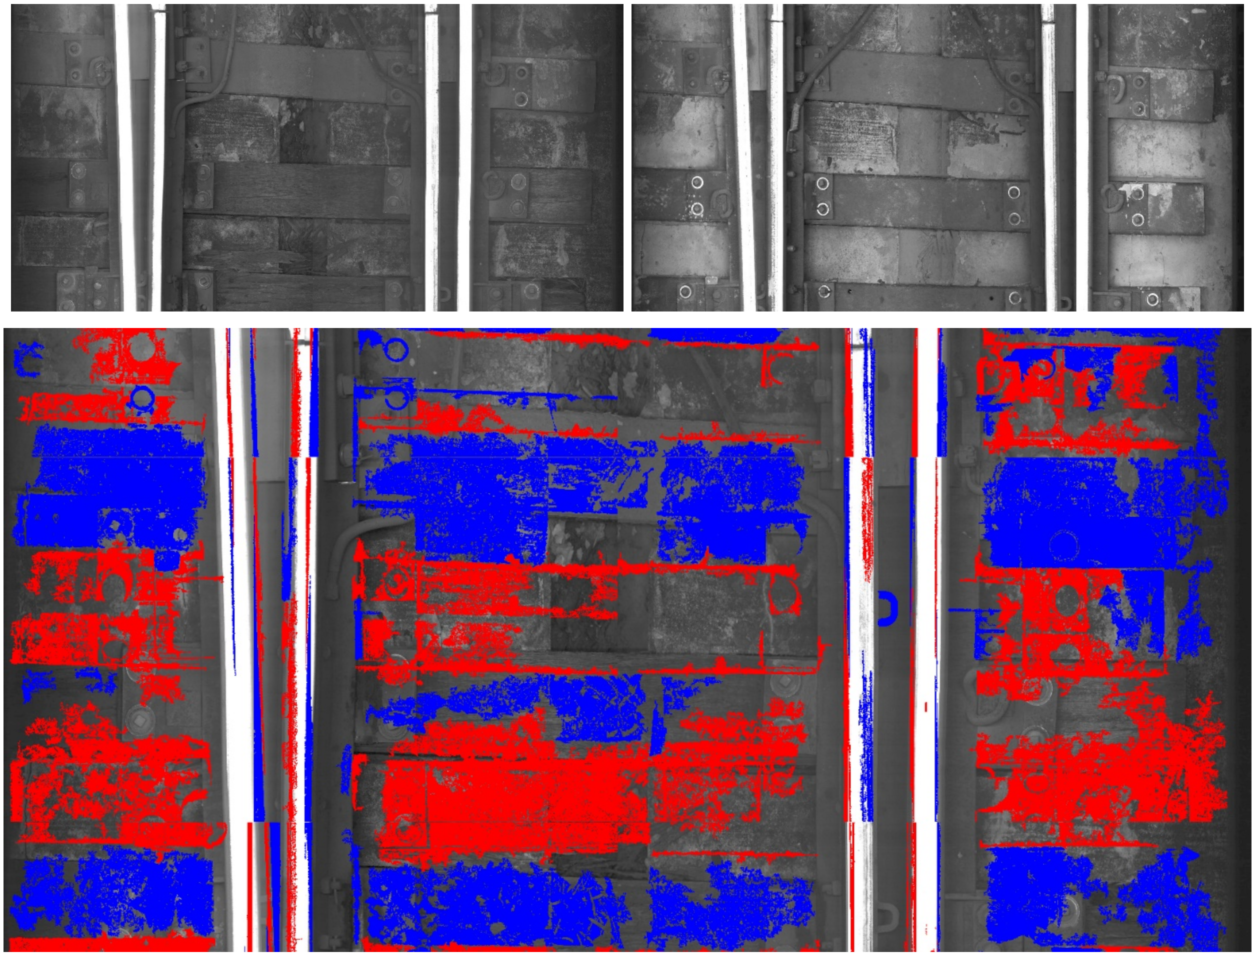

Figure A-26. Intensity-based Change Example 6: Tie Change (Wood to Concrete) (Set: 2016/2017, Frame: 2032)

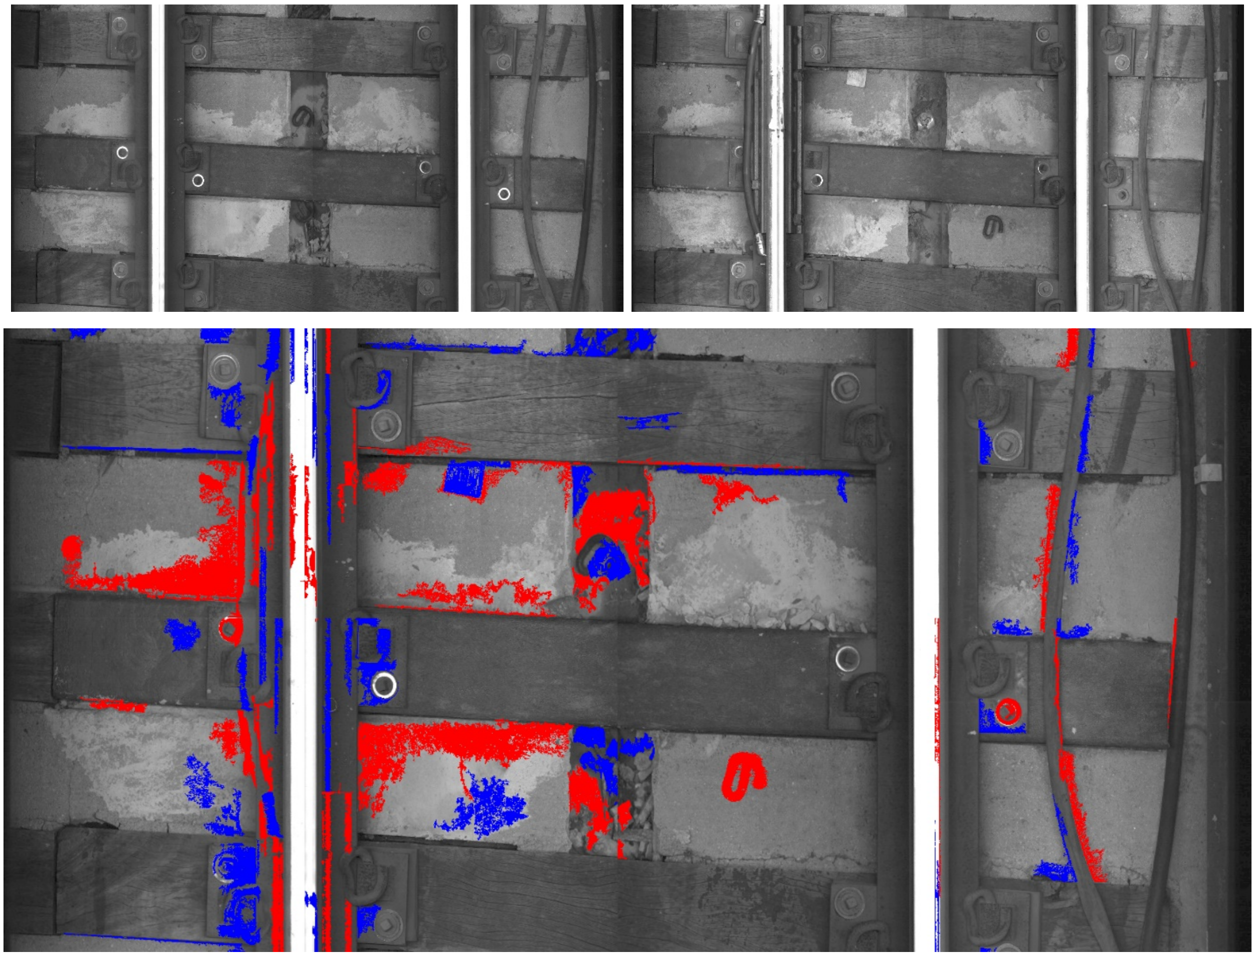

Figure A-27. Intensity-based Change Example 7: New Equipment/Missing Clip (Set: 2016/2017, Frame: 2788)

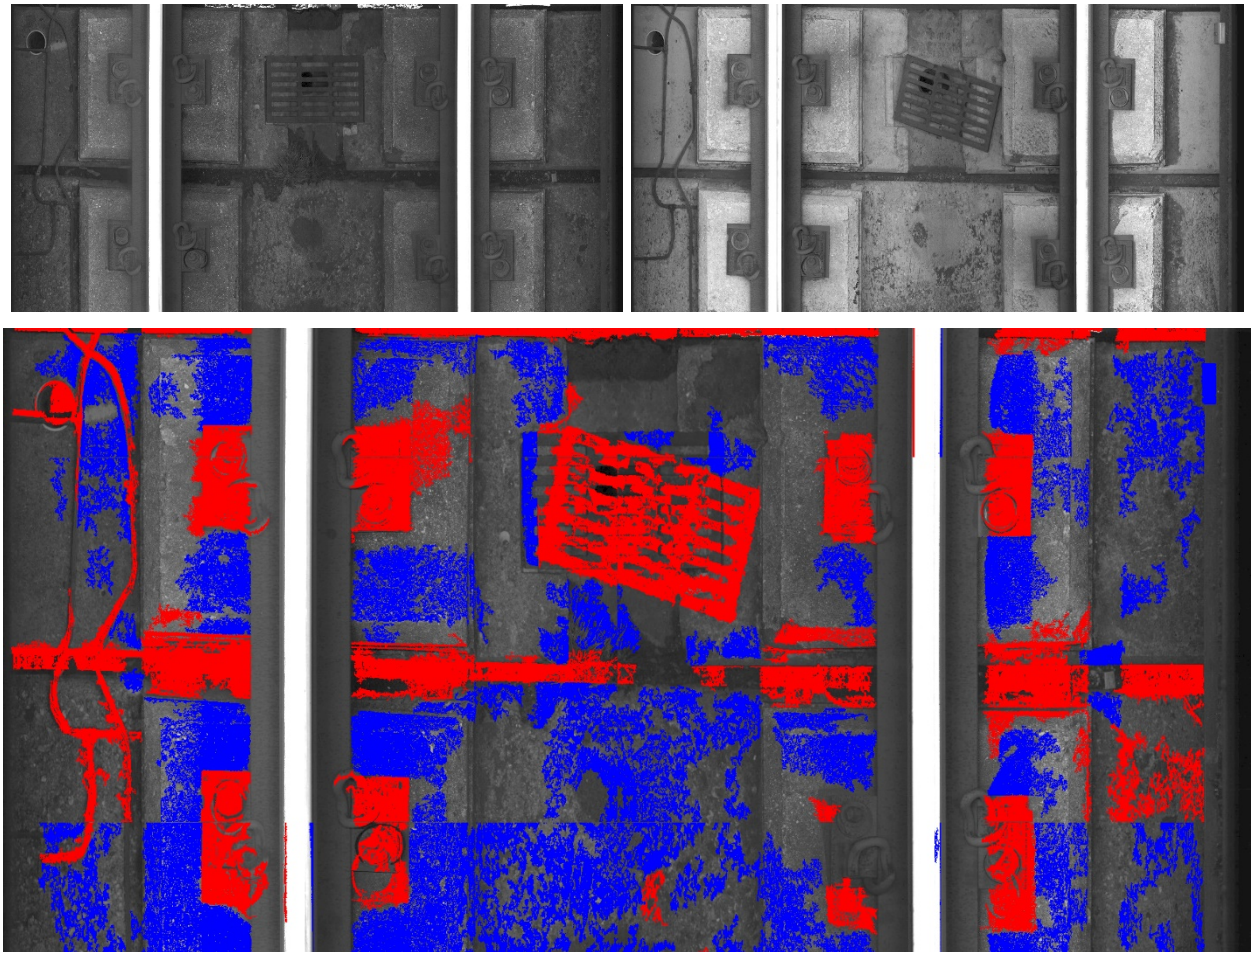

Figure A-28. Intensity-based Change Example 8: Misaligned Grate (Set: 2016/2017, Frame: 3046)

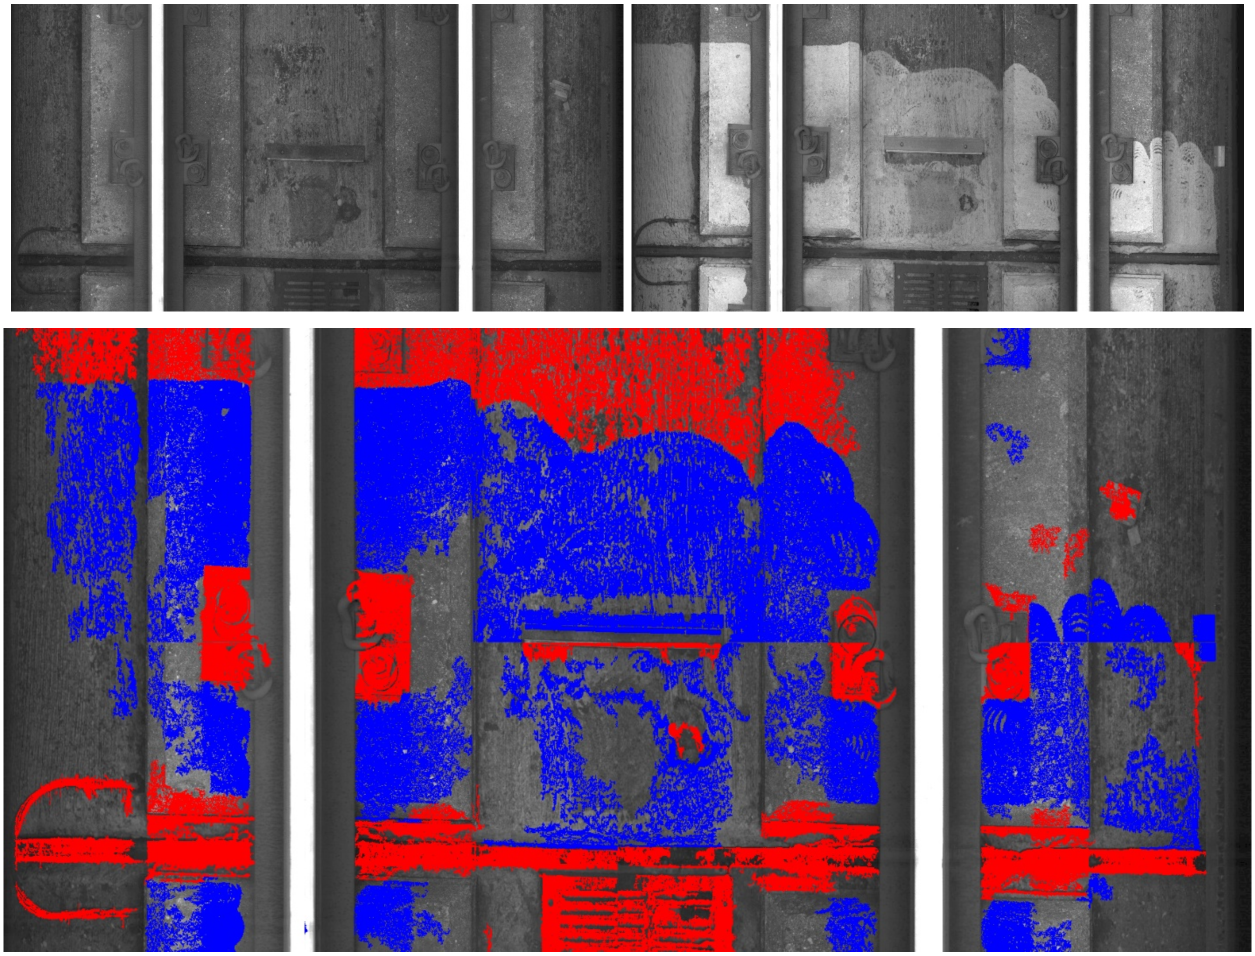

Figure A-29. Intensity-based Change Example 9: Substance on Concrete (Set: 2016/2017, Frame: 3118)

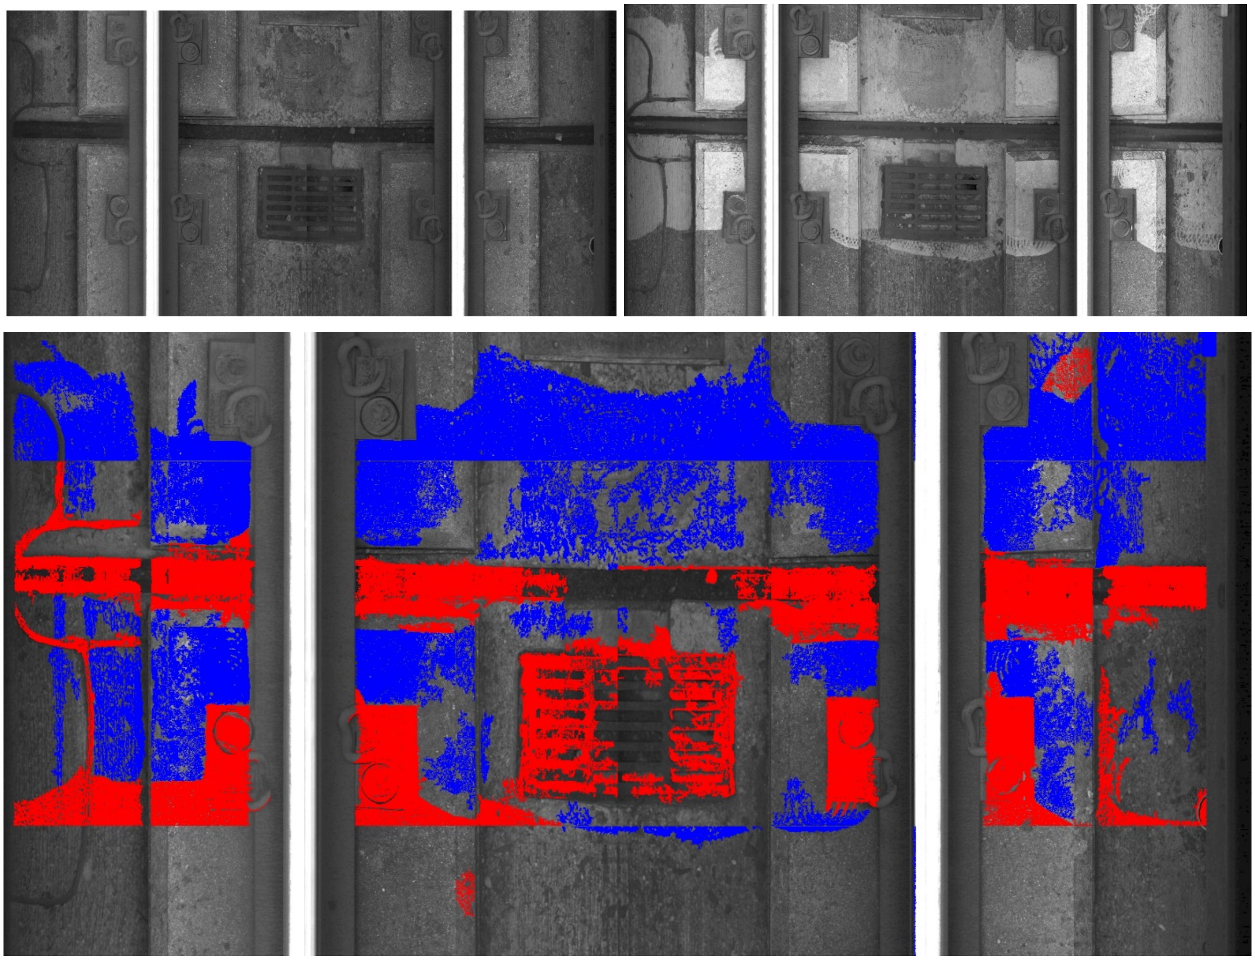

Figure A-30. Intensity-based Change Example 10: Substance on Concrete (Set: 2016/2017, Frame: 3189)

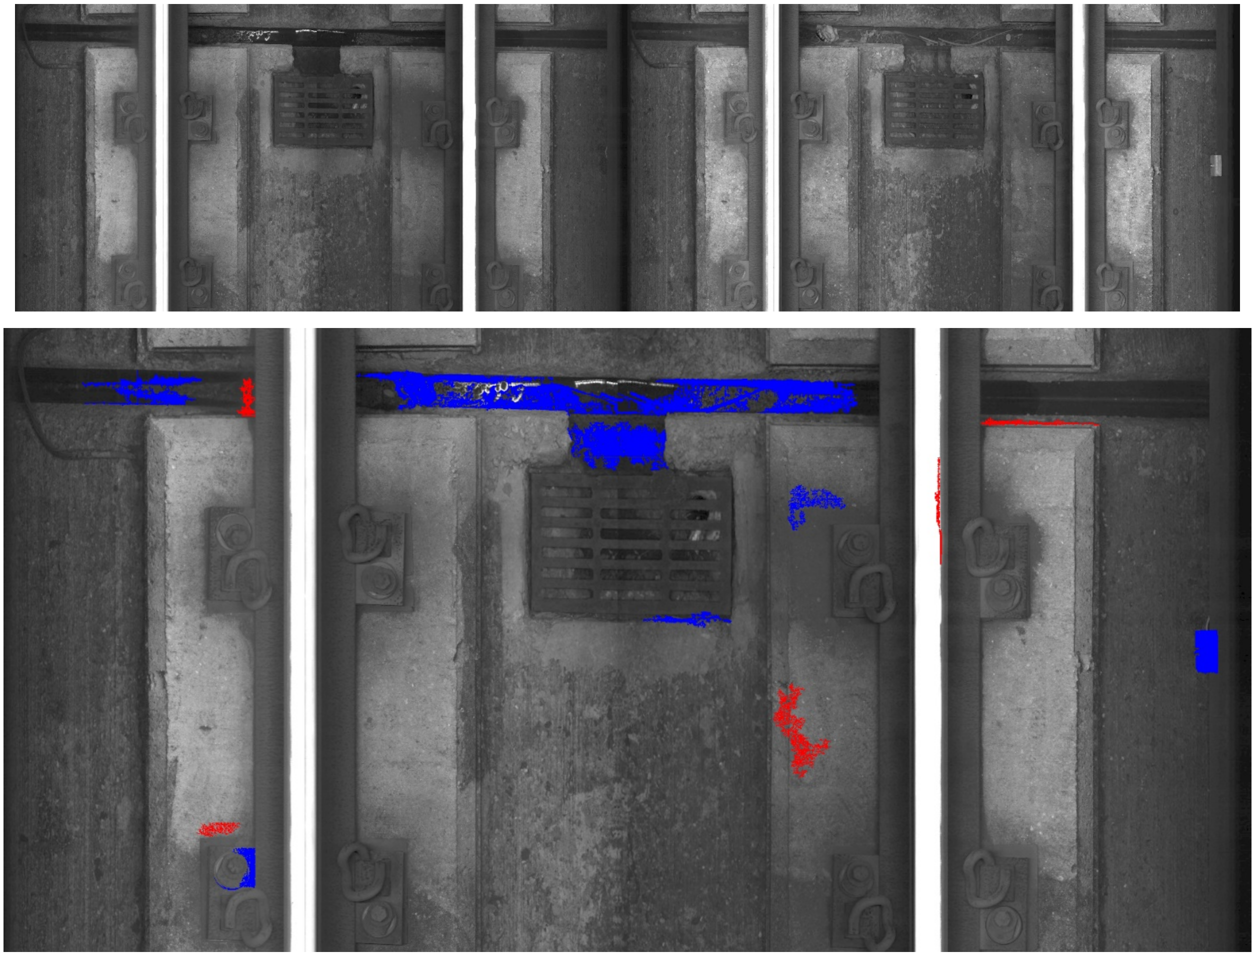

Figure A-31. Intensity-based Change Example 11: Standing Water (Set: 2016/2017, Frame: 5634)

A.6 Harris Report: Thematic-Based Change Detection Results

The following examples of change are provided as examples of overall results of the project for thematic classification based change detection.

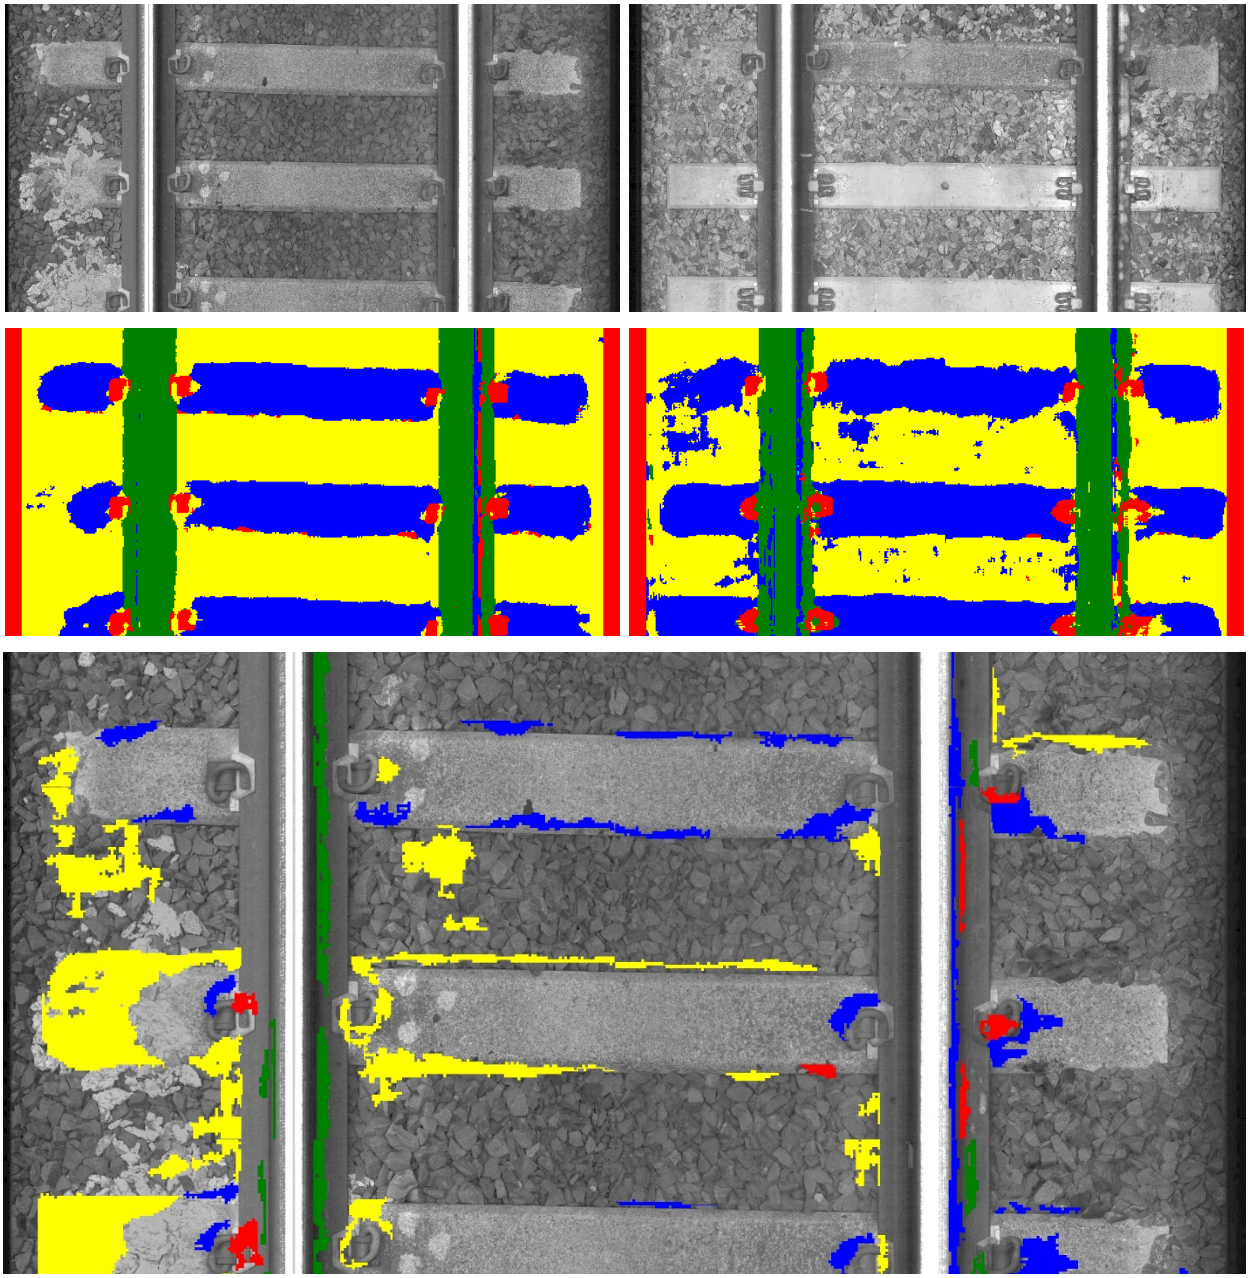

Figure A-32. Thematic-based Change Example 1: Damaged Tie (Set: 2012/2013, Frame: 0322)

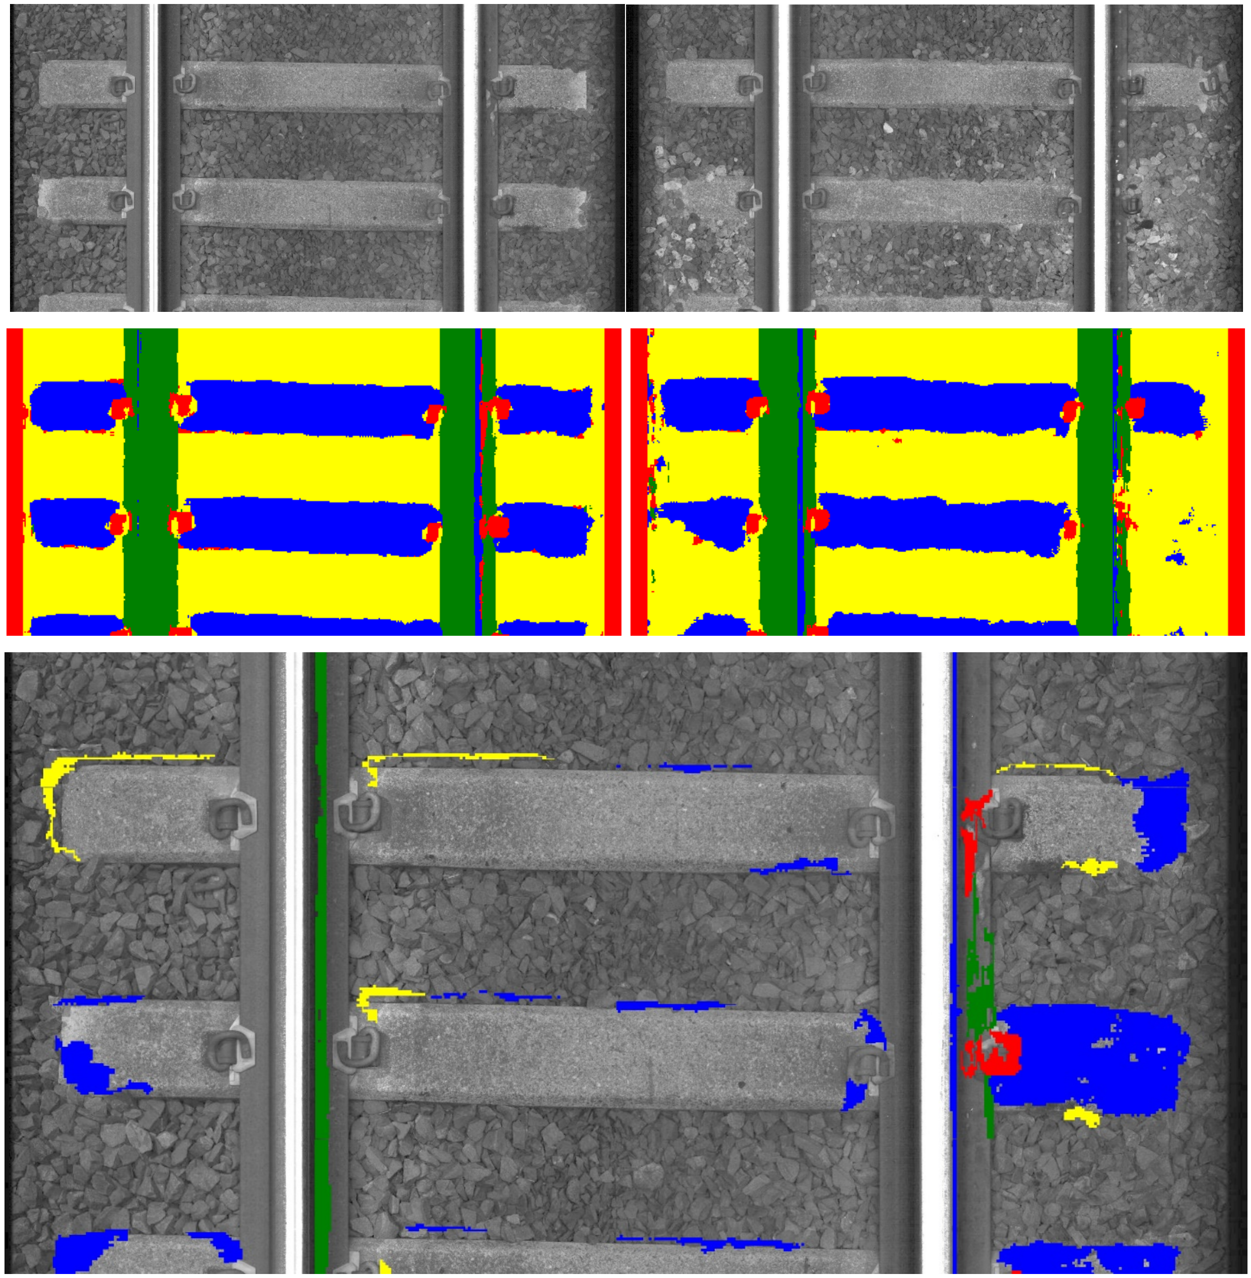

Figure A-33. Thematic-based Change Example 2: Damaged Tie (Set: 2012/2013, Frame: 0368)

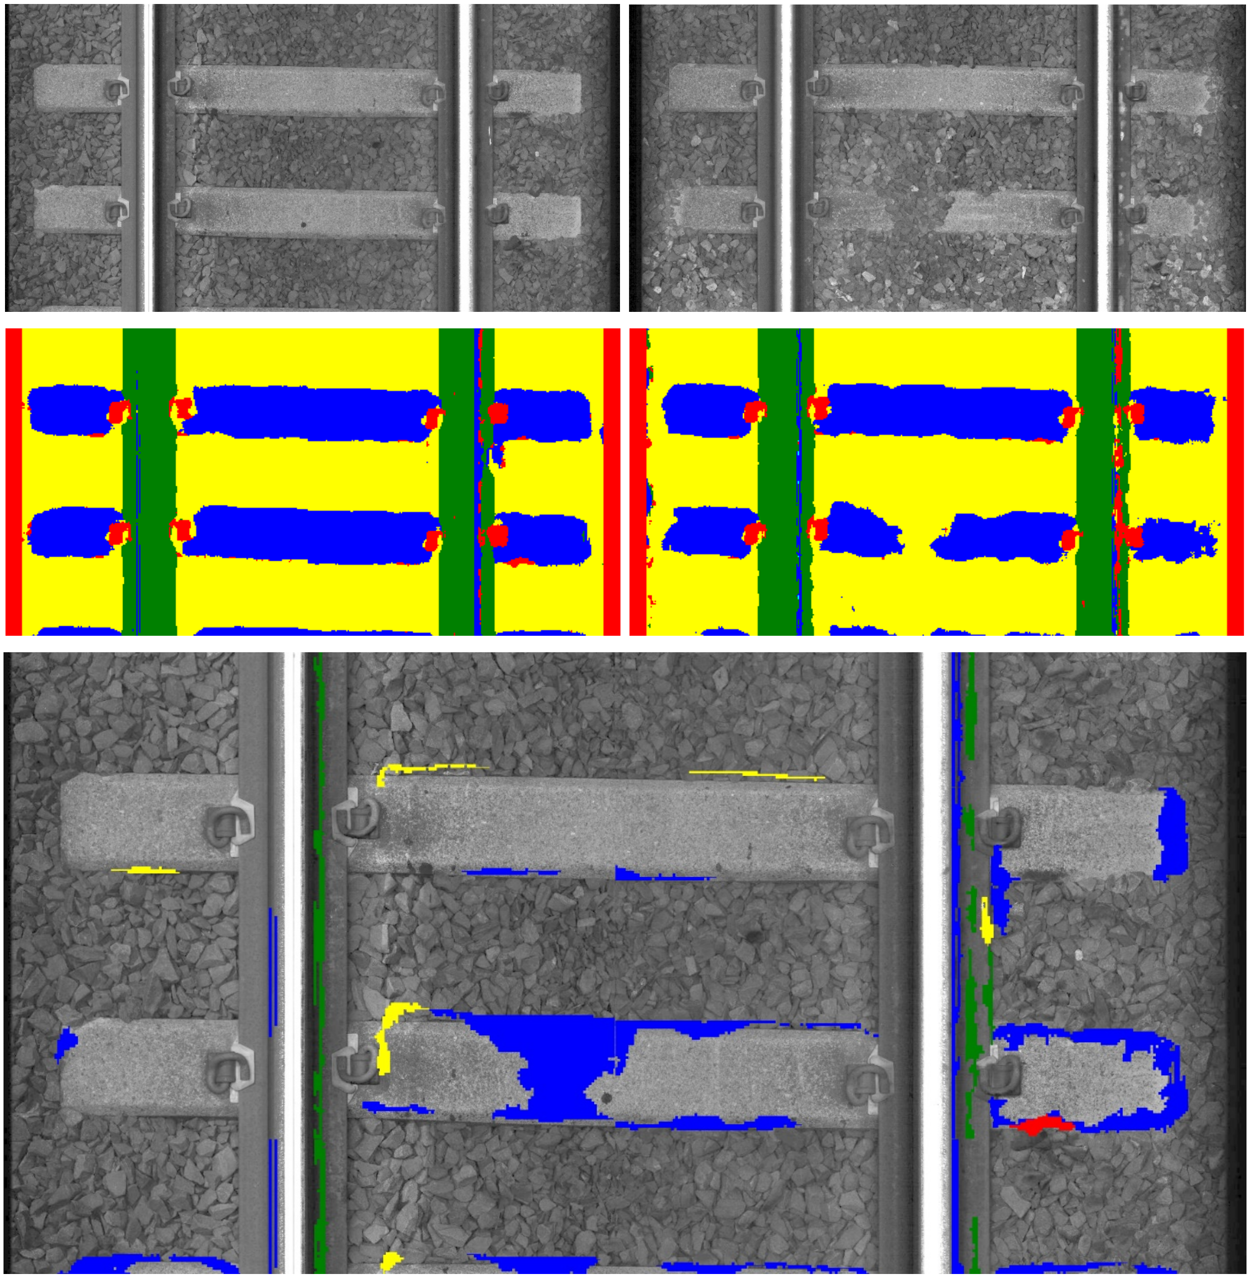

Figure A-34. Thematic-based Change Example 3: Ballast Change (Set: 2012/2013, Frame: 0528)

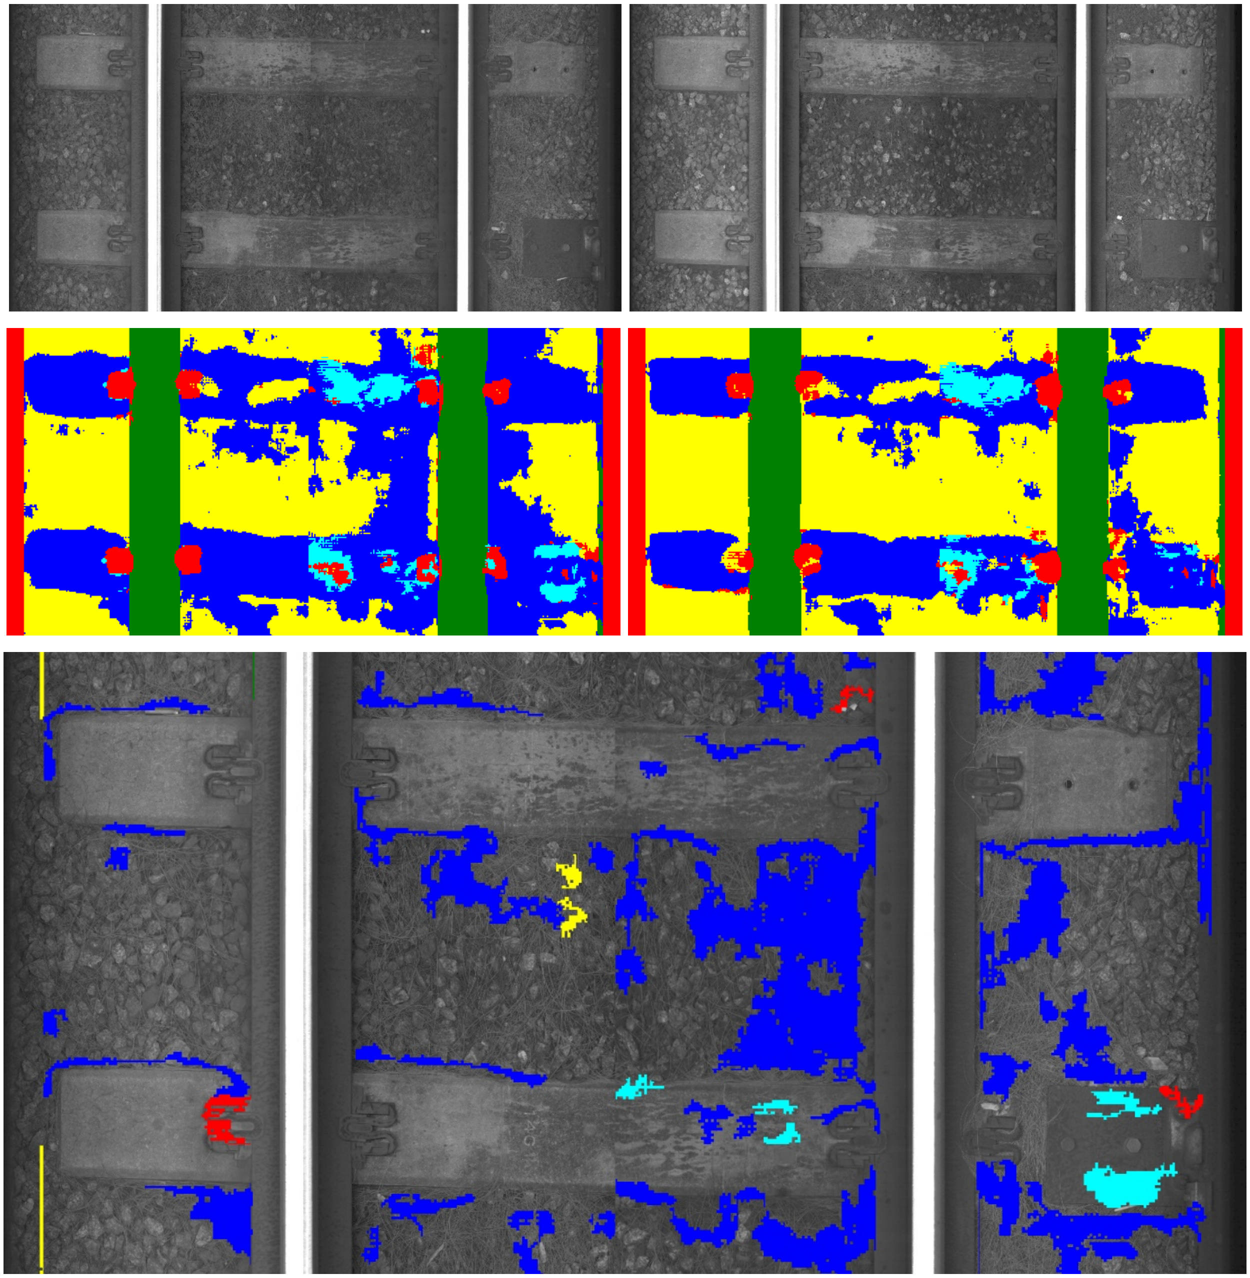

Figure A-35. Thematic-based Change Example 4: Vegetation Encroachment (Set: 2016/2017, Frame: 0343)

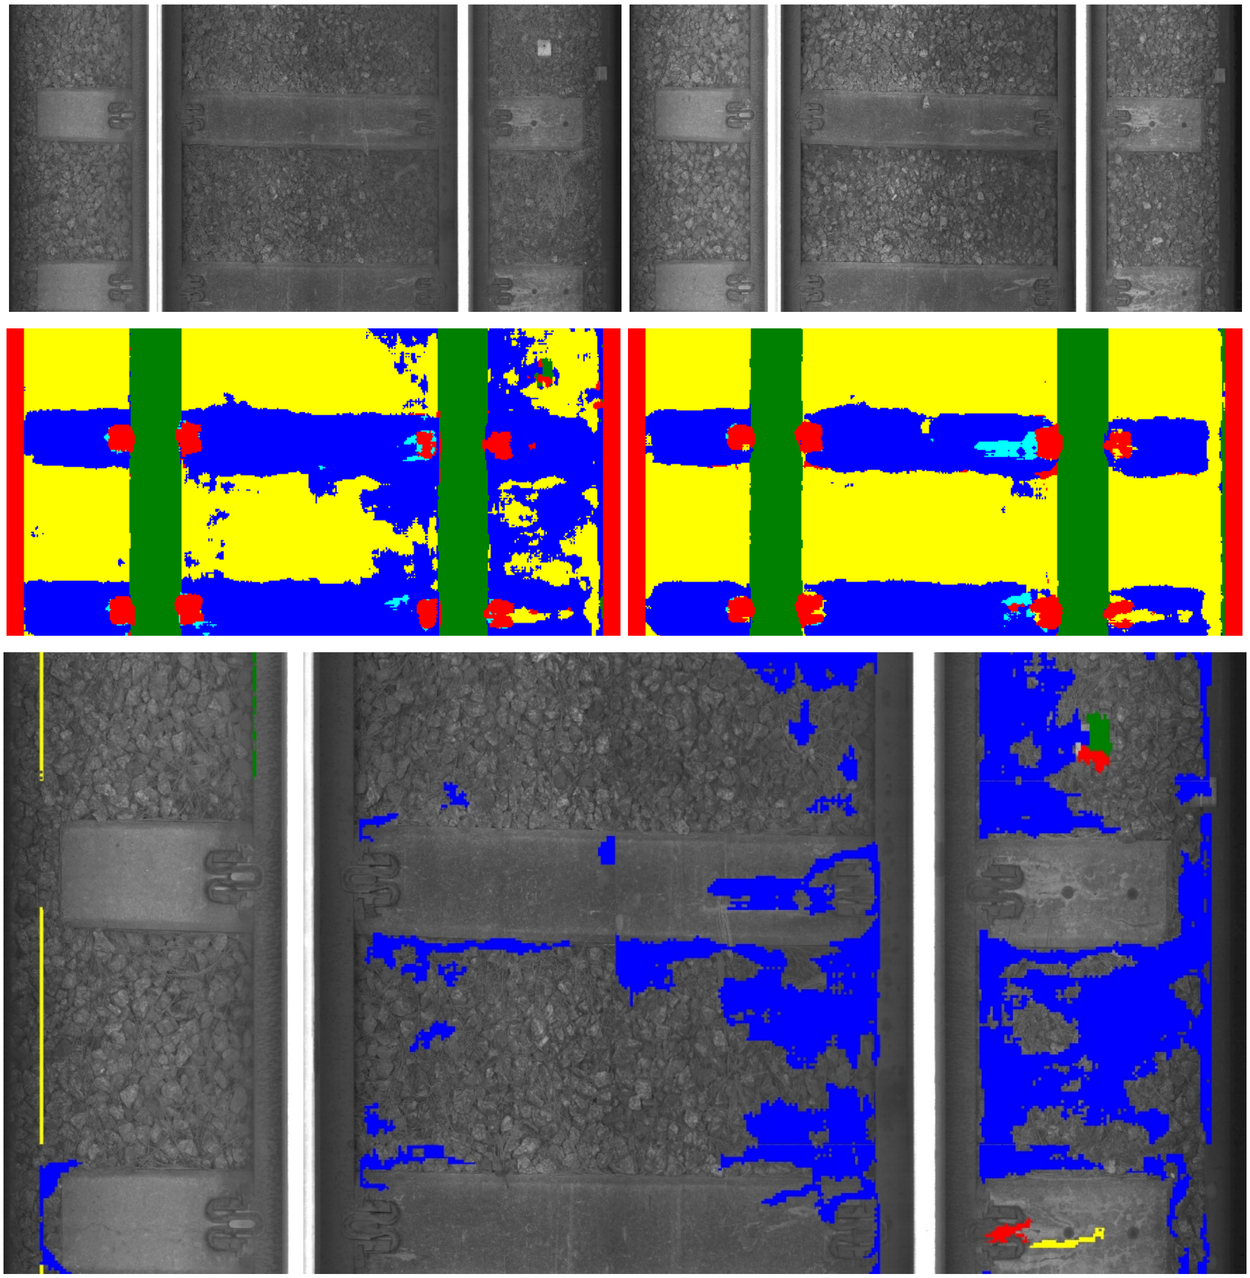

Figure A-36. Thematic-based Change Example 5: Vegetation Encroachment (Set: 2016/2017, Frame: 0353)

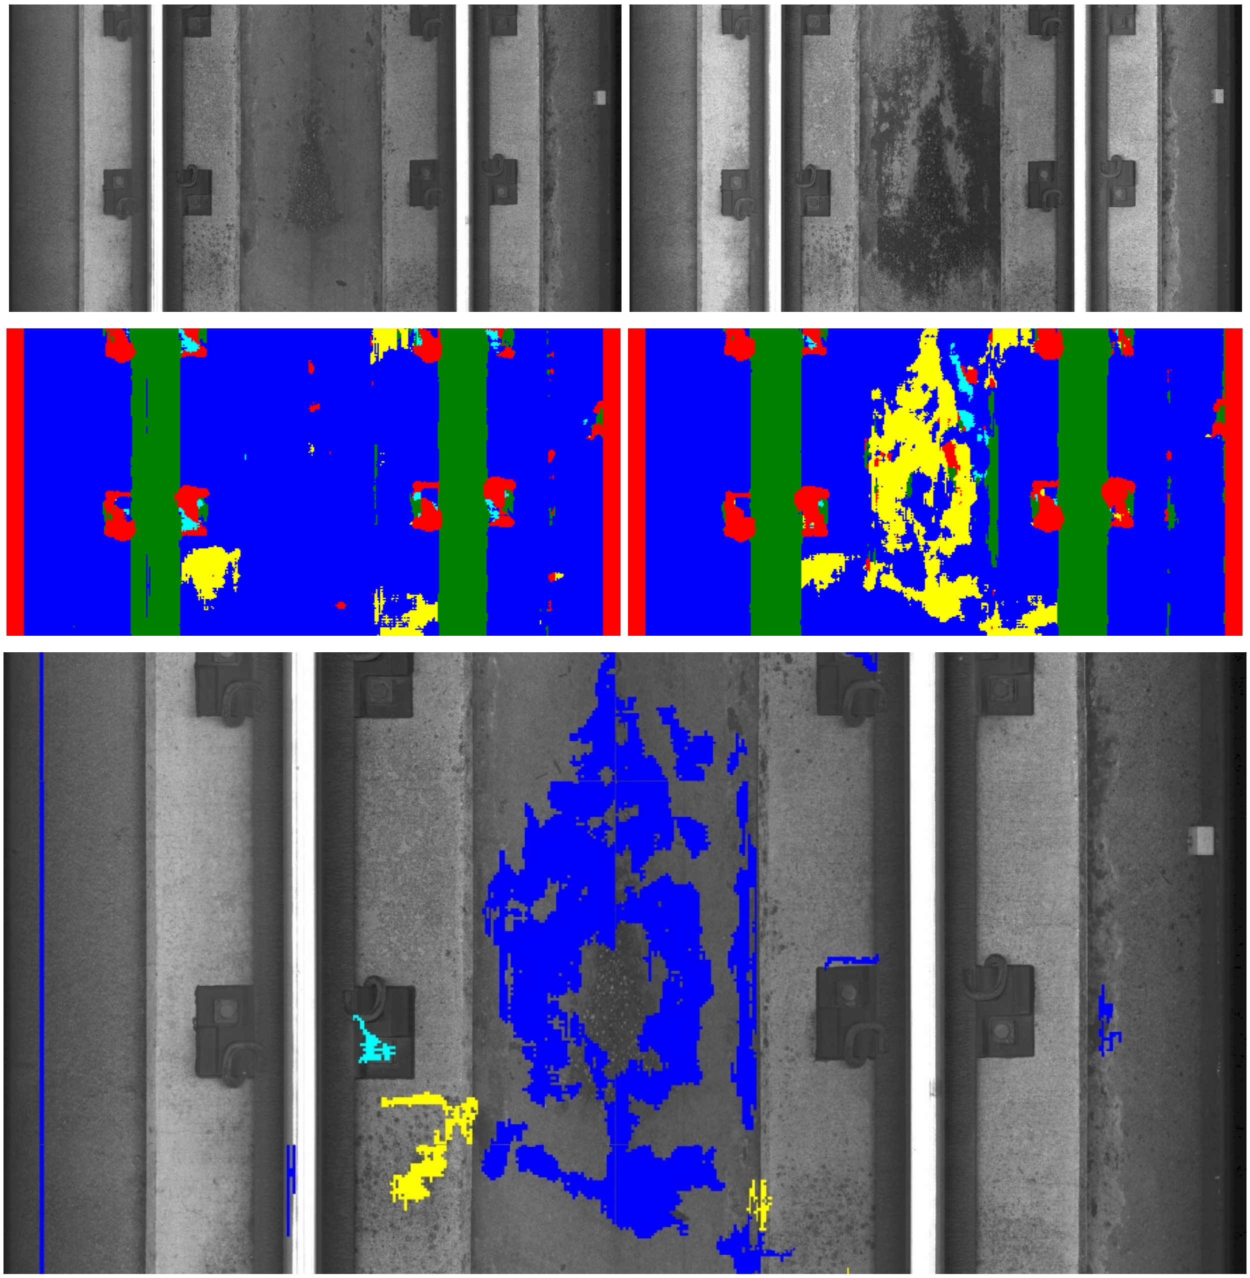

Figure A-37. Thematic-based Change Example 6: Water Damage (Set: 2016/2017, Frame: 1234)

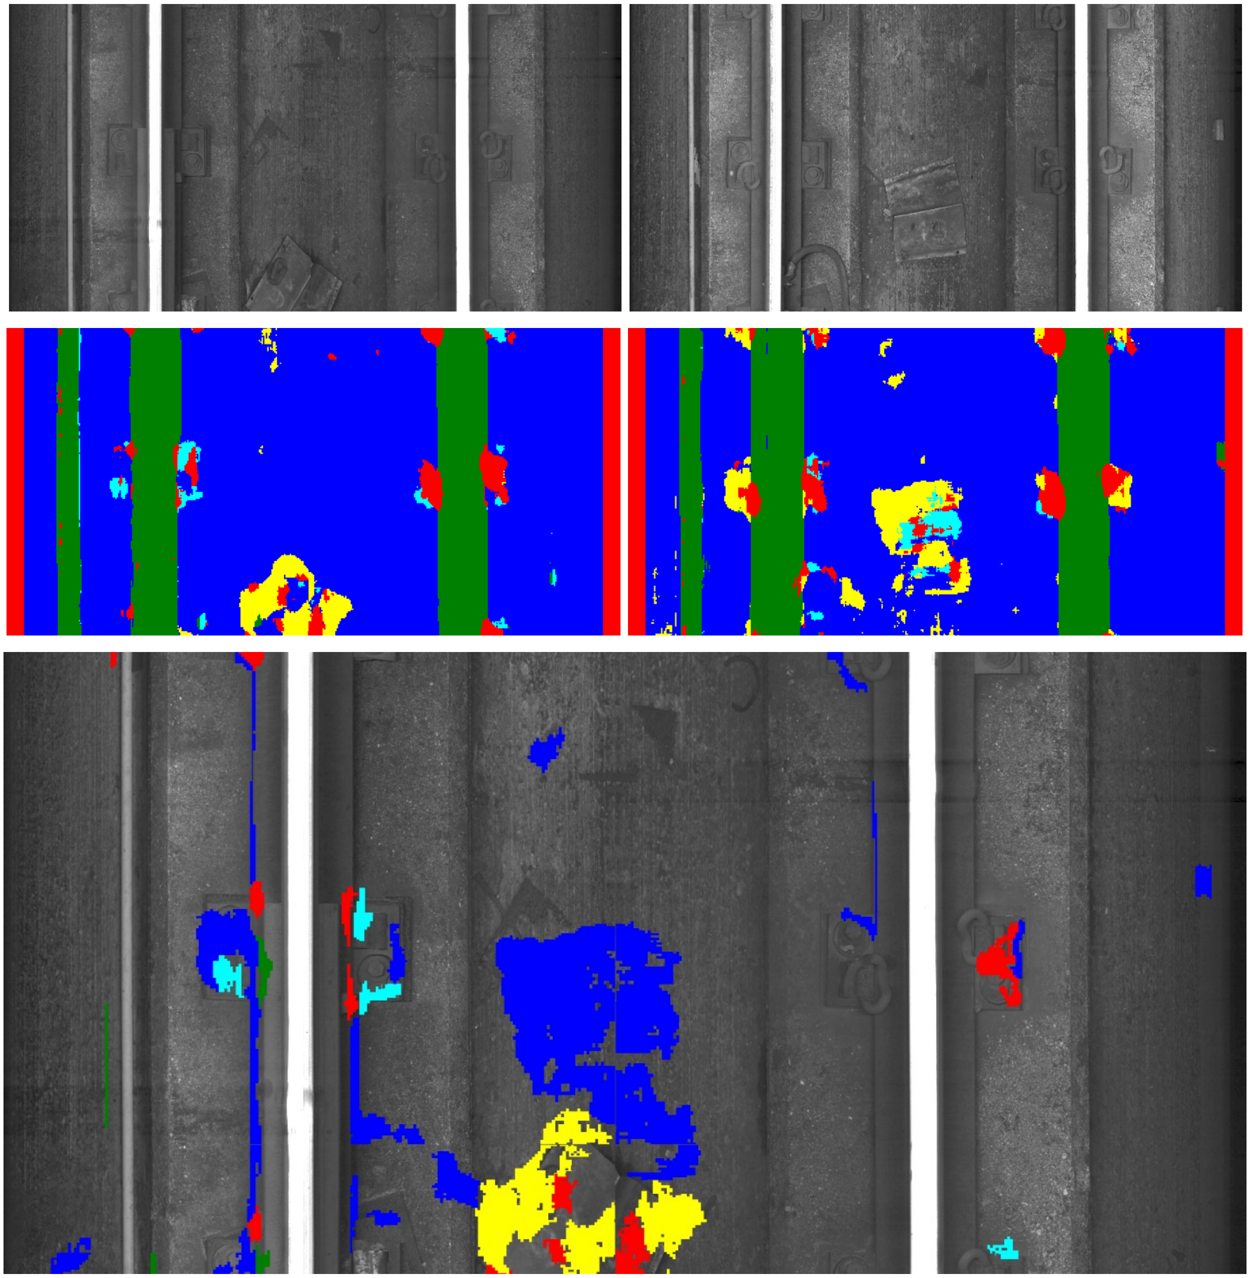

Figure A-38. Thematic-based Change Example 7: Missing Clip (Set: 2016/2017, Frame: 2015)

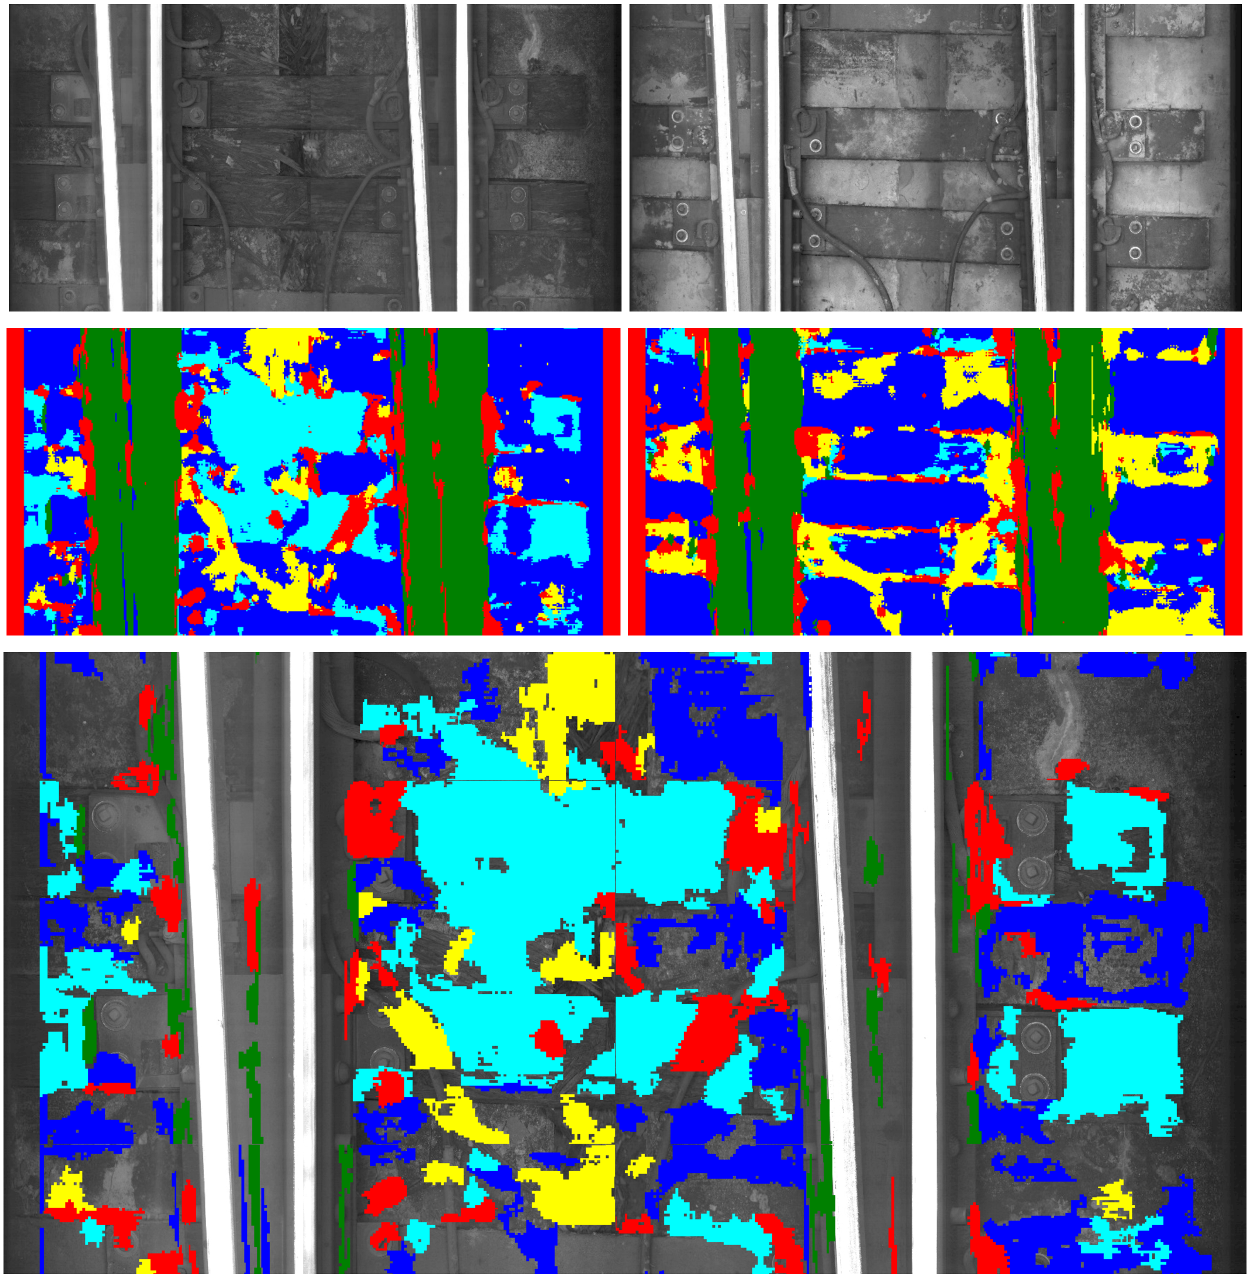

Figure A-39. Thematic-based Change Example 8: Tie Change (Wood to Concrete) (Set: 2016/2017, Frame: 2034)

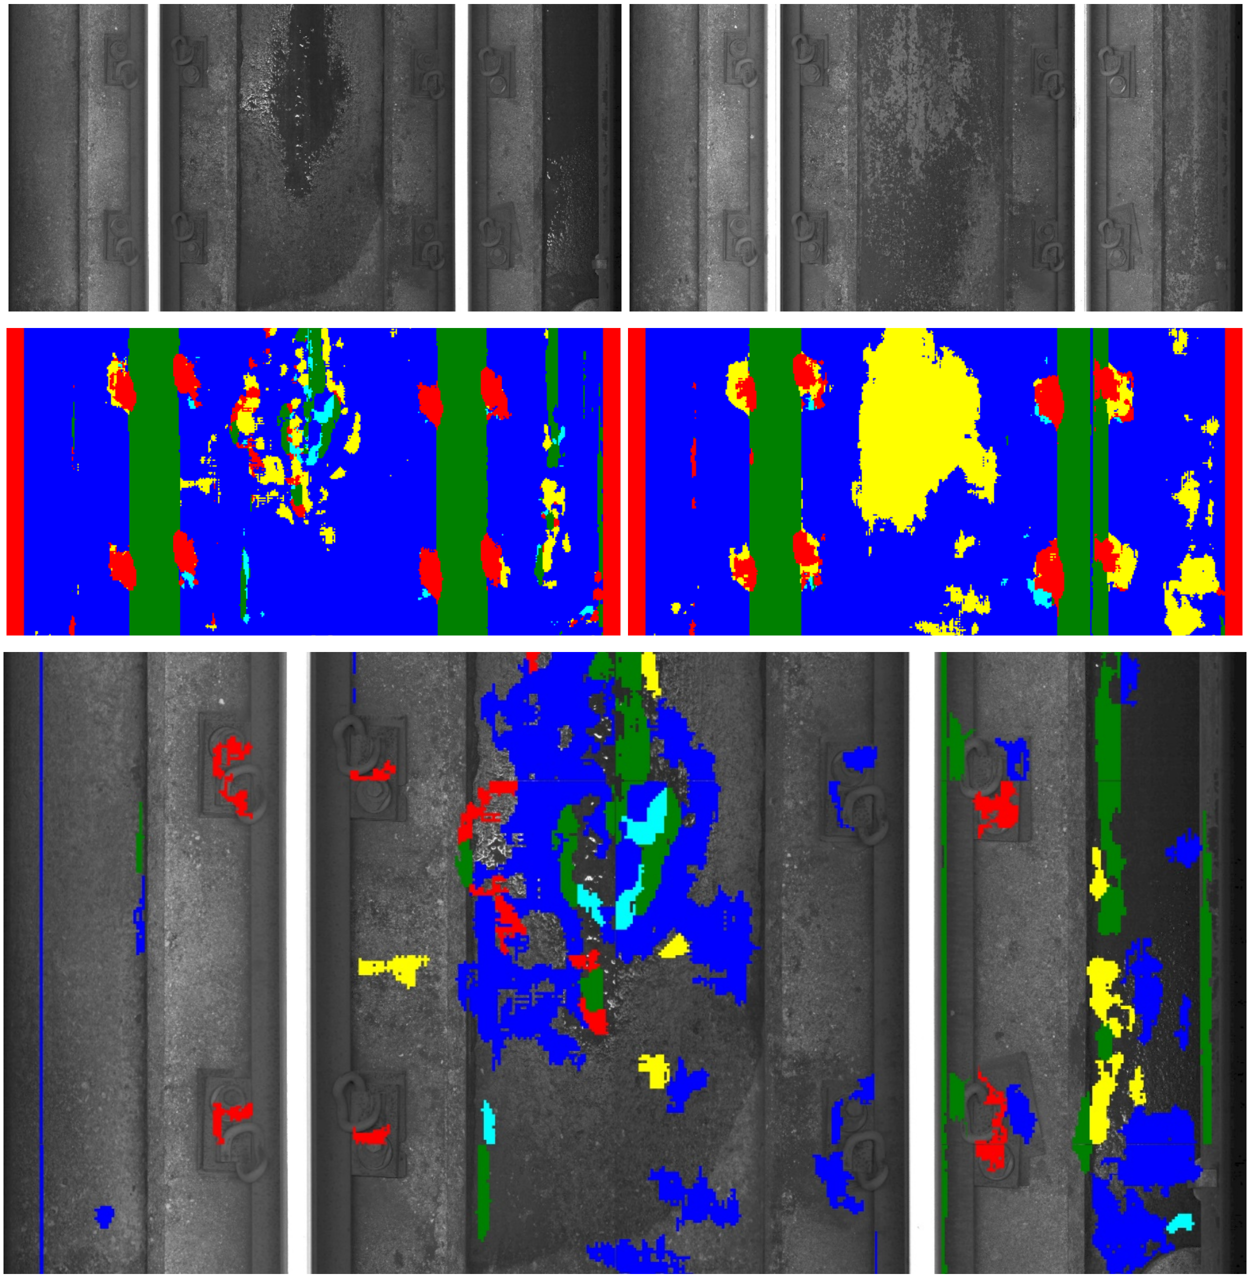

Figure A-40. Thematic-based Change Example 9: Standing Water (Set: 2016/2017, Frame: 4745)

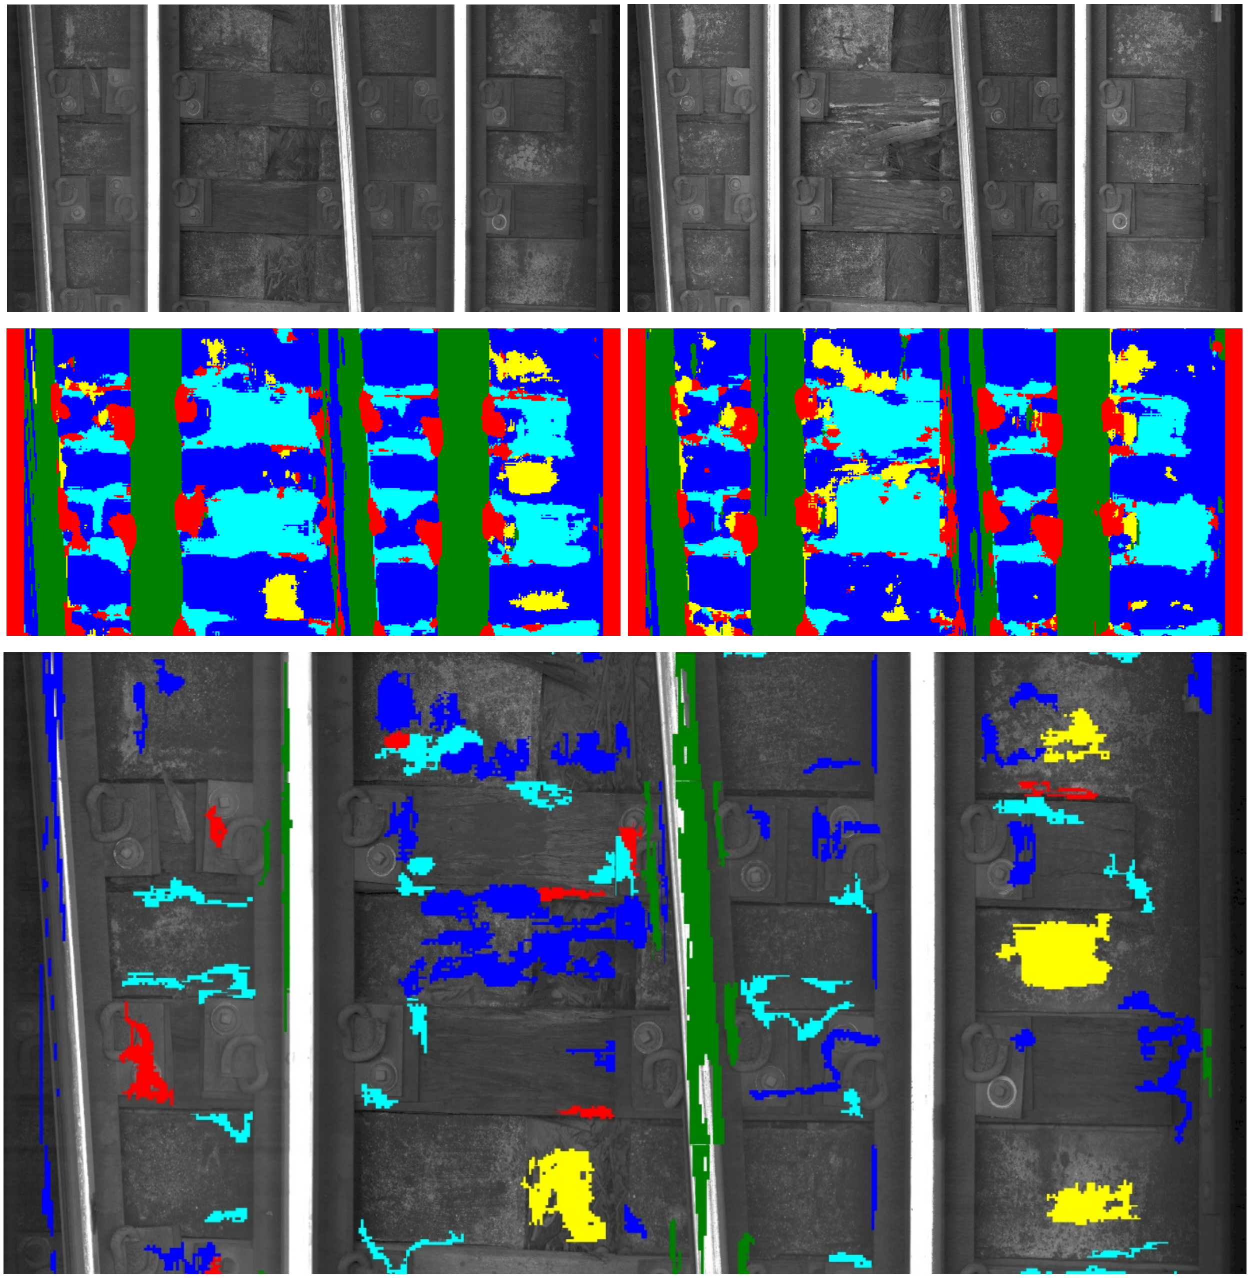

Figure A-41. Thematic-based Change Example 10: Damaged Tie
(Set: 2016/2017, Frame: 2041)

A.7 Harris Report: Conclusions

Exelis VIS has determined that change detection is a viable option for finding regions of interest on track-based imagery. The ability to find a variety of objects and/or conditions was proven to be possible given the collected data.

Exelis VIS offers the following general findings:

1) Pre-processing was instrumental in the preparation and the removal of artifacts from the data to reduce false positives in change detection. Exelis VIS used many features of its ENVI product line as well as its image science subject matter expertise to prepare the data for analysis. Exelis VIS is confident that these steps are critical and can be generalized and automated.
2) Exelis VIS employed several methods to sort the changes based on overall change in a frame, clustering of changed pixels, and changes in class (thematic only). While these methods were successful in identifying large areas of change, it wasn't optimal for small/relevant changes. Other techniques should be investigated to properly sort the changes based on the needs of the customer.
3) The image display interface configured for this project and specific to rail data was very useful in scanning the data to look for change. The ability to see the imagery side by side with the change detection overlay along with the curtaining feature was very useful. This tool could be employed to rapidly verify automated findings.
4) Due to the amount of processing required, automated processing is highly advised. Exelis VIS has extensive experience dealing with automatically processing very large data sets. In this case, the data was divided into many frames in order to utilize processing on many CPUs and GPUs in order to optimize processing times.

Exelis VIS offers the following findings on Pixel Intensity based change detection:

1) The algorithms were successful in finding areas of significant change. For example: vegetation encroachment, water damage or presence of standing water, and new/changed infrastructure.
2) The sorting techniques employed in this proof of concept were prone to highlight areas of larger change but were not adaptable to areas of small change.
3) Sorting cannot take any context within the scene into account, which limits the ability to find certain conditions. It can only detect areas of change.

Exelis VIS offers the following findings on Thematic Classification based change detection:

1) The algorithms were successful in finding areas of large and small change based on cluster statistics and a class change. For example: vegetation encroachment, damaged infrastructure, water damage or presence of standing water, and changed infrastructure.
2) The sorting techniques employed in this proof of concept were prone to highlight areas of larger change but are adaptable to areas of small change. The addition of a spatial component to sorting could assist in finding smaller features, for example missing clips.
3) The model could be improved with the addition of more examples and more classes for higher accuracy of the thematic classification.

4) The contextual information offered by the thematic change approach provides more detailed change while still employing an automated approach.

Exelis VIS recommends the following:

1) Further investigation into pre-processing automation. Being able to quickly process data to identify potentially hazardous situations as well as maintenance issues is critical.
2) Improvements to the thematic classification model. Exelis VIS believes a very accurate model could be developed that identified all targets of interest within the scene. This would provide a highly accurate thematic map from which many conditions could be detected.
3) Improvements to the sorting algorithm to be able to identify specific conditions that are important to the industry. For example, using the thematic images to identify changes in class that identify important conditions like a change from clip to tie that would identify a missing clip or a change from ballast to another class could identify eroding ballast.
4) Investigate other methods for utilizing the data captures such as identifying rail canting and skewed ties using distance and angle measurements. Due to the high resolution of the data and the ability of the thematic map to identify ties and rail, distance or angle measurements would be easy to extract from the imagery.
5) In addition to the two-dimensional processing demonstrated under this project, a spatial component or an additional three-dimensional data set would enhance the results:
    a. Ability to identify more accurately discriminate relevant from irrelevant changes by identifying high priority regions within the two-dimensional data to look for change. For example, areas directly over the junction of rail/tie to look for presence or absence of clips.
    b. Ability to identify more conditions by augmenting the analysis with three-dimensional data. This would allow for detecting issues that are elevation dependent instead of the current visual analysis. For example, detecting a change in elevation of a tie or ballast.
    c. Improve detection accuracies by leveraging both two-dimensional and three-dimensional data inputs. It is obvious from the results of the proof of concept that certain conditions are only detectable through imagery but that three-dimentional data could improve the accuracy of any spatial components needed to identify important regions of interest.

The additional spatial information could be used as yet another input into the thematic classification in order to assist in identifying classes or as yet another layer to augment the change detection, such as a change in height of a tie. Note that imagery and elevation data together would provide much more information than either in a standalone fashion. The ENVI suite of tools supports the latest two-dimensional and three-dimensional sensor data including panchromatic imagery, multi and hyperspectral imagery, Light Detection and Ranging (LiDAR), synthetic-aperture radar (SAR), and Full Motion Video (FMV) and can handle multiple inputs of the same scene for analysis purposes.

Going forward, based on the findings of this proof of concept, Exelis VIS believes that an automated workflow to process rail data (two-dimension and/or three-dimension) to identify areas of interest for maintaining the rail infrastructure is realistic. Using Harris expertise and its suite of ENVI tools for automated data analysis, storage, and visualization would highly benefit the rail industry.

Abbreviations and Acronyms

Abbreviations & Acronyms	Definition
COTS	Commercial-off-the-Shelf
DV	Display Intensity Value
FRA	Federal Railroad Administration
FMV	Full Motion Video
LiDAR	Light Detection and Ranging
PV	Pixel Intensity Value
POC	Proof of Concept
RFIs	Requests for Information
RFPs	Requests for Proposals
SAR	Synthetic-Aperture Radar
UAVs	Unmanned Aerial Vehicles