



IDEA

**Innovations Deserving
Exploratory Analysis Programs**

High-Speed Rail IDEA Program

Machine Vision for Railroad Equipment Undercarriage Inspection Using Multi-Spectral Imaging

Final Report for High-Speed Rail IDEA Project 49

Prepared by:
Narendra Ahuja and Christopher Barkan
Co-Principal Investigators
University of Illinois at Urbana-Champaign

December 2007

TRANSPORTATION RESEARCH BOARD
OF THE NATIONAL ACADEMIES

INNOVATIONS DESERVING EXPLORATORY ANALYSIS (IDEA) PROGRAMS MANAGED BY THE TRANSPORTATION RESEARCH BOARD

This investigation was performed as part of the High-Speed Rail IDEA program supports innovative methods and technology in support of the Federal Railroad Administration's (FRA) next-generation high-speed rail technology development program.

The High-Speed Rail IDEA program is one of four IDEA programs managed by TRB. The other IDEA programs are listed below.

- NCHRP Highway IDEA focuses on advances in the design, construction, safety, and maintenance of highway systems, is part of the National Cooperative Highway Research Program.
- Transit IDEA focuses on development and testing of innovative concepts and methods for improving transit practice. The Transit IDEA Program is part of the Transit Cooperative Research Program, a cooperative effort of the Federal Transit Administration (FTA), the Transportation Research Board (TRB) and the Transit Development Corporation, a nonprofit educational and research organization of the American Public Transportation Association. The program is funded by the FTA and is managed by TRB.
- Safety IDEA focuses on innovative approaches to improving motor carrier, railroad, and highway safety. The program is supported by the Federal Motor Carrier Safety Administration and the FRA.

Management of the four IDEA programs is integrated to promote the development and testing of nontraditional and innovative concepts, methods, and technologies for surface transportation.

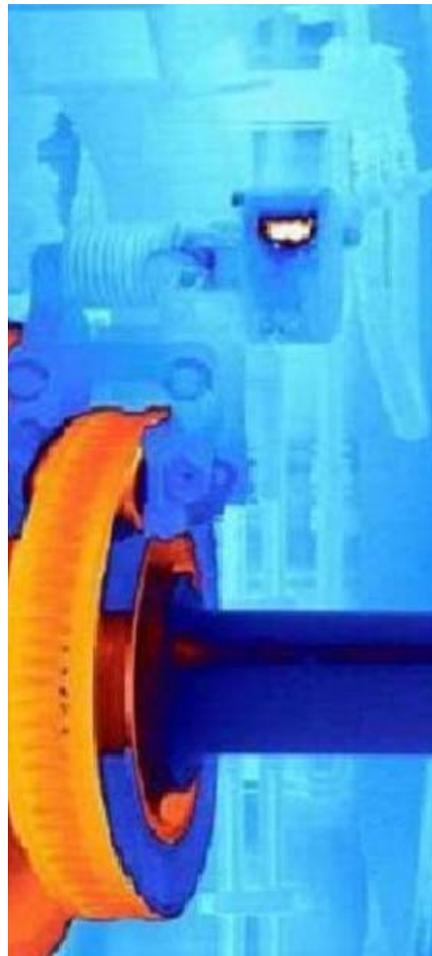
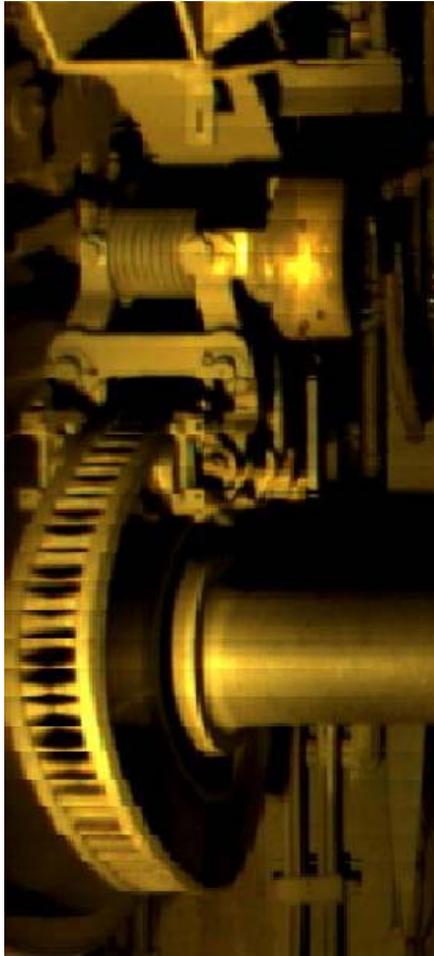
For information on the IDEA programs, contact the IDEA programs office by telephone (202-334-3310); by fax (202-334-3471); or on the Internet at <http://www.trb.org/idea>

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

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

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

Machine Vision for Railroad Equipment Undercarriage Inspection Using Multi-Spectral Imaging



**IDEA Program Final Report
September 13th, 2005 through August 15th, 2007
HSR-49**

**Prepared for
The High-Speed Rail IDEA Program
Transportation Research Board
National Research Council**

**Narendra Ahuja and Christopher Barkan
Co-Principal Investigators
University of Illinois at Urbana-Champaign**

December, 2007

ACKNOWLEDGMENTS

The research team would like to thank Amtrak and the Monticello Railway Museum for their willingness to let us use their facilities for testing of our image acquisition system. We are especially grateful for the help, time, and knowledge we received from Amtrak; specifically, Paul Steets from Wilmington, John Raila and Sarabpreet Bumra from Chicago, and Dale Kay from Beech Grove have greatly helped in this research. Gavin Horn was very helpful in arranging for the use of and providing training for the IR camera used in this project. We would also like to thank Derek Hammer of Hammer Motion Pictures for his consultations on lighting equipment and methods for even illumination.

Funding for this work has been provided by the TRB High-Speed Rail IDEA Program Project HSR-49, the University of Illinois Railroad Engineering Program, and the Beckman Institute Computer Vision and Robotics Laboratory.

ABSTRACT AND KEYWORDS

Current practices for inspection of railcars and locomotives include both manual and automated systems. However, inspection of railroad equipment undercarriages is almost entirely a manual process. Visual inspections by humans are performed either in a pit or trackside. The equipment is usually stopped over the pit or run slowly past the trackside inspector. In the latter case, it is not possible for a human to have an unobstructed view of the undercarriage as a train rolls by. Automated inspection by electronic systems has the potential to overcome certain limitations of human inspection.

The report describes the research conducted to develop a new approach to undercarriage inspection by means of machine vision analysis. This approach uses multispectral imaging from cameras viewing the undercarriage from a below-the-track perspective. Imaging using both visible and infrared spectra provides a means by which incipient failure detection can be addressed. Detection of missing, damaged, and foreign objects can also be identified using this approach. By extracting frames from video recordings in both spectra, panoramic images of the entire train can be created and analyzed. These images are further subdivided into individual railcar panoramas that can be matched to templates of railcars in known good condition to detect missing and foreign objects. More detailed diagnosis can be provided by using specific component-level templates allowing identification of damaged and overheated sub-components. In addition, comparisons can be made of duplicate component systems during operation, such as disk brakes, to discover thermal outliers indicating improper function. A prototype of this machine vision inspection system has been developed and tested at a passenger car service and inspection facility.

This investigation demonstrates the feasibility of a machine vision system to provide undercarriage inspection capabilities, as the train passes over the pit, aiding inspection crews and repair personnel. The system provides a clear and unobstructed visible spectrum assessment of the undercarriage in addition to an assessment from the thermal spectrum as well. The joint analysis of these undercarriage views can provide automatic detection of components in need of repair and also those that may be overworked or near failure. This allows the inspector to be aware of indications indicative of component problems that are developing, which may fail in the future. Therefore the system has potential for providing advanced warning, allowing additional time for repair personnel to plan repairs prior to possible in-service failures.

Keywords: Machine vision, railcar inspection, railroad car component failure, incipient failures, automated equipment inspection, anomaly detection, infrared spectrum, panoramic imaging, passenger trains.

TABLE OF CONTENTS

ACKNOWLEDGMENTS	4
ABSTRACT AND KEYWORDS	5
EXECUTIVE SUMMARY	7
BACKGROUND AND OBJECTIVES	9
OBJECTIVES OF THE PROJECT STAGES	9
BACKGROUND OF THE PROPOSED SYSTEM	10
IDEA PRODUCT	11
CONCEPT AND INNOVATION	12
INVESTIGATION	13
IMAGE ACQUISITION PROTOTYPE DEVELOPMENT	13
Undercarriage Lighting for Visible Spectrum Recordings	14
Equipment Setup	15
MODULE 1: IMAGE ACQUISITION AND PREPROCESSING	16
MODULE 2: PANORAMIC IMAGE PROCESSING	18
MODULE 3: MACHINE VISION INSPECTION	20
Module 3A: Global Anomaly Detection	20
Module 3B: Component-level Anomaly Detection	20
Module 3C: Relative Anomaly Detection of Like Components	28
TEST PLAN DEVELOPMENT AND TESTING RESULTS	30
PROJECT PANEL	31
CONCLUSIONS	32
FUTURE WORK	33
POTENTIAL IMPACT AND PAYOFF FOR PRACTICE	34
PLANS FOR IMPLEMENTATION	34
INVESTIGATOR PROFILES	35
PROJECT TEAM	35
GLOSSARY	37

EXECUTIVE SUMMARY

The objective of this project was to investigate the feasibility of a multispectral imaging system for automated inspection of passenger train undercarriages. Both visible and thermal spectra were used in our experiments on diagnosis of existing and incipient problems with undercarriage components. A panel of experts assisted in the selection of rail car types and components of interest. The car types included Amtrak Horizon and Amfleet as well as Genesis locomotives. The components were traction motors, disc brakes, air conditioning units, and bearings. In addition, detection of foreign objects was also of interest.

Early in the project, experiments were conducted using a camera with wide-angle lenses located in the inspection pit to obtain images of the entire undercarriage. However, we found that suitable wide angle lenses for infrared (IR) cameras were not available. The solution was to use the IR camera to record specific strips that showed the components of interest. Initially we also had problems with ineffective lighting and IR reflections back into the IR camera that caused difficulties in acquiring clear images. These problems were addressed by consulting motion picture lighting and IR camera experts. We chose to conduct our tests in a railroad inspection pit at the nearby Monticello Railway Museum. The close proximity allowed frequent visits. In addition, it eliminated the interference with normal railway operations that would result from working at a real railway site, e.g., at Amtrak's Chicago facilities.

After we had developed a system for successful acquisition of digital video of the undercarriage of trains, the next steps involved correcting certain distortions in the images extracted from the visible spectrum videos. These were caused by the wide-angle lenses and were removed using well established lens distortion correction methods. The undistorted images were then combined to produce a single panoramic image of the entire train. Panoramas using the IR images were produced in a similar manner. These train panoramas were then split into single car panoramas by using machine vision algorithms to first detect the location of the axles and then the couplers themselves. From the individual railcar panoramas, templates of entire railcars were generated to represent a specific railcar in good condition and used to match against incoming railcars of the same type. The number of different car types was limited to that of available in-service trains that passed over the pit during our acquisition sessions. In conjunction with railcar templates, more detailed templates of specific components of interest were also produced. In some instances, conditions that we needed to test were not observed so we artificially created them on the digital images to test the machine vision algorithms' performance.

Undercarriage inspection was achieved by using machine vision techniques to create image-based templates to represent an entire car in "good condition". These car-level templates were then compared to incoming cars captured by the system. This enables detection of overheated, missing or foreign object anomalies in both the visible and IR spectra. Once a car-level anomaly is detected, component-level templates can then be used to further isolate which particular component may be the cause of the car-level anomaly. Learning algorithms were employed to automatically classify the type of anomaly that was found at the component level into specific defect and no defect categories.

Another method of anomaly detection involved comparison of the temperature consistency among components that are repeated on a single car or entire train such as certain elements of the braking system. This enabled us to distinguish abnormal from normal temperatures for components whose temperature changes during normal operation. Examples of such components include brakes, bearings and traction motors. An example is provided that shows the detection and comparison of brake disc thermal signatures. From this procedure, hot and cold outliers can be found that can correspond to problems such as brakes stuck in the on or off position, or thin or missing brake pads. In this example, the brake calipers are also detected and their contraction is measured to determine the brake application status.

Although all image capture opportunities were conducted in good weather, an effort was made to simulate or consult experts on the effects of adverse weather. In addition, an algorithmic technique was developed to distinguish between the reflection or shine effects due to wet surfaces, and actual component anomalies.

Criteria for performance testing were established and documented in a field test plan and used during the final video acquisition that was completed under in-service conditions at Amtrak's passenger car maintenance and inspection facility in Chicago.

Overall, this study found that multi-spectral imaging (in particular visible and infrared) for inspection of passenger car and locomotive undercarriages is feasible and offers potential benefits in terms of both enhanced efficiency and effectiveness. Furthermore, some enhancements to the inspection process are possible even before fully-developed machine vision techniques are completed and refined, thereby making a phased approach to implementation possible while still accruing benefit. In addition, several organizations representing both railroads and the supply industry have expressed interest in further development of this technology.

BACKGROUND AND OBJECTIVES

Current practices for inspection of railcars and locomotives include both manual and automated systems. However, inspection of railroad equipment undercarriages is an almost entirely manual process. Trained personnel perform a visual inspection of the equipment usually while it is stopped over an inspection pit, or from the wayside as it runs slowly past a trackside inspector. In the latter case, it is not possible for the inspector to have an unobstructed view of the undercarriage. Pit inspections do allow much better views of many undercarriage components. However, they require a specialized facility (the pit) and that the equipment be taken out of service for whatever amount of time is required to perform the inspection. Furthermore, these inspections are labor intensive, subject to the variability of inspection personnel capabilities, and limited to detection of defects that can be seen in the visible light range and in the direct view of the inspector. All of these factors limit the effectiveness and efficiency of the current inspection process.

Railroads are interested in automatic inspection technologies that have potential to overcome some of the limitations of the inspection process described above. The objective of this project was to investigate the feasibility of several possible elements of such a system by using multi-spectral imaging (visible and infrared range) combined with machine vision algorithms to integrate and interpret the information in the recorded images. These technologies offer the potential to perform inspections of passenger train undercarriages that are quicker and more effective than is currently possible. Also, the nature of such technologies is that they can systematically record and organize information for later comparison, referencing, and trend analysis, thereby further enhancing the utility of the information obtained. Both visible and thermal spectra were used in the experiments to identify existing and incipient problems with undercarriage components.

Although this project focused on the specific tasks and technologies described above, they are part of a larger effort in the railroad industry to use technology as extensively as practicable to perform a variety of inspection tasks. The University of Illinois has been conducting research and development on a variety of different inspection technologies, as have the AAR, FRA and various railway suppliers. Depending on the particular technologies and the inspection needs of a railroad, it may often make sense to integrate these technologies at a single site. This will enable railroads to take greater advantage of communications technology and wayside Automatic Equipment Identification (AEI) readers that will be needed at sites intended to identify specific pieces of equipment in order to link the information to them. A single AEI reader can be integrated with each inspection system to enable them to report defects on specific railcars, this information can also be sent to railroad databases (such as InterRISS) using the same wireless communication system hardware. The systems could also take advantage of the same train detection sensors and power installation, significantly reducing the cost of installing multiple independent sites. In addition, the cost of periodic maintenance would be reduced.

OBJECTIVES OF THE PROJECT STAGES

In Stage 1, a panel of experts was to be convened to provide guidance and support for the project. The purpose of the first panel meeting was to review the project objectives and investigative approach and deliverables, and to consider the most useful applications. Other Stage 1 objectives to be included were the acquisition of both visual and infrared images from cameras located such that they viewed the underbody of railcars. The systems intended for inspection were to include, but not to be limited to, critical disc brake, electrical, and bearing components. An agreement with the IDEA program regarding the range of equipment, car designs, and components was to be established. An investigation was to be conducted to see how images were likely to change as a result of components being damaged or missing, or how the images degraded as a result of weather conditions, e.g., accumulations of snow and ice. The main outcome of this first stage was to produce a system design, taking into account functionality under adverse weather conditions.

The primary objective of Stage 2 was to design and test a prototype system while addressing the challenges of variable car types and component conditions. In addition, the development of a series of laboratory tests using previously captured or simulated images to assess and refine the algorithms was required. These tests needed to include images of both normal and damaged or missing components and simulated or actual images degraded by weather conditions. A Stage 2 report was required that included a comprehensive field test plan. The main result of this stage was to provide a sense of the performance that could be expected from the proposed system in real life.

In Stage 3, feasibility tests were to be conducted in accordance with the Stage 2 Test Plan on images acquired at a passenger rail facility. This final report was prepared to include test results with real data and provide a sufficient basis to determine whether the proposed system will function acceptably for real-world trains. The panel will review and discuss project results, findings and strategies to facilitate implementation of the developed system in practice.

BACKGROUND OF THE PROPOSED SYSTEM

The proposed machine vision system is comprised of image acquisition hardware and computer algorithms for interpretation of each car and locomotive, and their individual components. The image acquisition system consists of visible and infrared (IR) cameras. The cameras are located between and below the rails, and are oriented upwards toward the undercarriage of the train. As the train rolls over, the cameras simultaneously record in the visible and infrared spectra to digital media. In our tests, information about the identity of each individual piece of equipment was recorded manually but this can be automated using Automatic Equipment Identification (AEI) or a similar technology. The visible and infrared videos were then processed frame by frame to correct for lens distortion and camera skew. From each frame in the visible and IR sequences, a middle strip was extracted, whose width is dynamically determined, so that each strip shows a distinct part of the passing undercarriage. The strips are algorithmically “stitched” together to create panoramic images of the entire train in both visible and IR. After identifying the car or locomotive ends in the two panoramas, they are split up into images of each individual car and locomotive. From the recorded information, each piece of equipment would be referenced to determine if it has been recorded before. If it has, the last known panoramic template of the car would be matched to the current one to inspect for any changes that might indicate incipient failures, or damaged, missing, or foreign components. If no previous recording of the subject piece of equipment was available, a template of an identical piece in known good condition would be compared. The inspection algorithms will also compare recurring components on the car for any relative anomalies between them. Those areas or components that do not match well with either the template or the other similar components would be marked as anomalies requiring further inspection.

In summary, the machine vision software system is organized into three principal operational modules (Figure 1). In Module 1 (Image Acquisition and Preprocessing), visible and IR videos are captured, the videos are converted into individual images, and lens distortion is corrected. In Module 2 (Panoramic Image Processing), images are combined to create both visible and IR panoramas of the entire train, and then the train panoramas are separated into individual equipment panoramas. In Module 3 (Machine Vision Inspection), anomalies are detected and identified at both the equipment and individual component level.

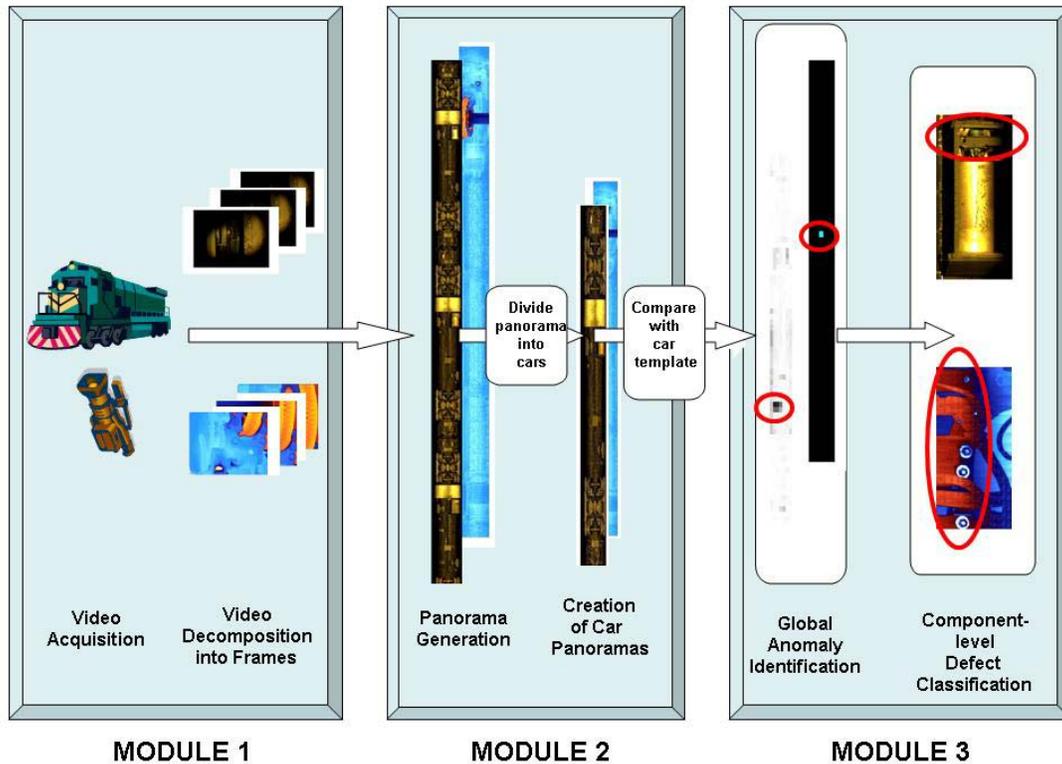


Figure 1. Flowchart for undercarriage inspection.

IDEA PRODUCT

The project objective is to develop a multispectral machine vision system for the inspection of rail equipment undercarriages to assess condition and detect incipient failures, as well as damaged, missing, or foreign components. The system will supplement and render more efficient standard undercarriage inspections by detecting and correlating thermal and physical anomalies. The likely locations where such a system could be installed are entrances to terminals, stations, or rail yards.

The system would automatically record the undercarriage of trains passing at slow speed, and then produce an inspection report. The results would be reported with textual and graphic information for further action. Anomalous items would be tagged for receiving a more thorough manual inspection and possible repair. The system offers several advantages over traditional inspection techniques. Inspection can be faster because it can be done while the train is moving in contrast with conventional inspection in which a train must be stopped on the pit track. It is also more efficient because only those parts of the equipment identified as needing more detailed inspection will be referred to inspectors. When automatic inspection can also identify the required tools or replacement items, these can be obtained and deployed in advance of the train's arrival at the repair facility, thus increasing the utilization of the facility for service-related activities rather than inspection. Furthermore, safety and reliability may be enhanced if incipient failures are detected and repaired before they manifest as physical damage or a service failure. Finally, the system can objectively inspect every piece of equipment in view of its health history stored in a central database. The history can also be used to detect trends in, and forecast, repair or replacement needs, and to detect unexpectedly high wear and tear and alert personnel about it.

CONCEPT AND INNOVATION

Several fundamental machine vision techniques were used and combined to develop the multi-tiered approach to anomaly detection and identification used in this project. An important factor contributing to progress was the iterative process whereby we would develop concepts in the lab and then take them into the field and determine how well they worked and what needed to be changed. The ability to repeat this process with relatively low cost in time and money enabled more rapid progress in developing increasingly robust, practical solutions that functioned under a variety of conditions. These innovations include panoramic image generation, robust identification of individual car boundaries, and inspection under a range of environmental conditions.

Module 1, in Figure 1, provides a reliable way to produce both visible and IR video frames for use in panorama generation in the presence of environmental and lens distortions. This involved calculating an appropriate focal length to use for visual cameras in order to capture the full width of the undercarriage of the train within the constraints of the pit, as well as determination of the proper shutter speed to capture satisfactory images of the moving train without blurring. Once a proper shutter speed was determined, we had to estimate the amount of light necessary, which was improved through experimentation. We also conducted experiments on the robustness of the visible-range video quality in varied weather conditions such as rain, and on the requirements for protective enclosures to shelter the cameras from the elements. Based on these tests, a prototype image acquisition system was developed and used to record visible and IR video. Frames were extracted from these videos to provide the input needed for panoramic generation of Module 2.

In Module 2, the lighting challenges, wide-angle lens distortion, and other real-world non-ideal conditions are compensated for by composing the panoramic image from only the central (minimally distorted) areas of each video frame. We used two well-developed algorithms in machine vision, the sum of absolute differences (SAD) technique and correlation. We experimentally found that SAD was more efficient than correlation for creating panoramic images in this application because there were minimal inter-frame illumination changes and it also performed better. We used Canny edge detection to form an axle edge template, and then used the distance transform (DT) to match the edges from the axle template to the axle edges in the undercarriage. Once the axles were located, the location of couplers was determined using spatial correlation. Matching couplers was only robust when done in a limited search space; therefore the initial detection of the axles through edge matching was critical.

In Module 3, the block-wise identification of global anomalies allowed a relatively rapid means to identify missing and foreign components by taking advantage of a stored template associated with each car. Canny edge detection with a DT was also used here to align components and investigate smaller, component-level defects. Edges were used in the location of many components (such as axles, water containers, and return spring of the braking system) since color and shine of many components varied across the images. Module 3 also uses a novel region-based scheme that uses Gaussian Mixture Models (GMM) to identify anomalies in regions, therefore producing a flexible way to classify even previously unseen defects.

In Stage 2, a method to locate component-level anomalies was also developed for Module 3, that involved spatially-based template matching using correlation or differences. However, during Stage 3, we realized that these methods caused too many false positives because mismatches of both physical and thermal images occurred frequently due to changes in external conditions, such as levels of cleanliness of the component that manifested itself as shine. Therefore, a GMM was implemented to classify anomalous regions so that anomalies that were not indications of defects could be distinguished. The GMM used a new feature space that allowed a dimensionality reduction, which had not been addressed in previous stages, but was found to be necessary in processing the data. It also allowed us to generalize classification of previously unseen defects based on defects that had been observed thus far.

INVESTIGATION

The primary objective of this section is to discuss the feasibility and effectiveness of the three modules of the proposed inspection system. These are Image Acquisition and Data Preprocessing, Panoramic Image Processing, and Machine Vision Inspection Algorithms. The final system software, when fully integrated, would be designed to perform in real-time immediately after a train passes.

IMAGE ACQUISITION PROTOTYPE DEVELOPMENT

In machine vision installations for rail equipment inspection, the camera is generally stationary while the vehicle rolls past. Obtaining video images of the undercarriage of rolling stock imposes additional restrictions on the recording equipment not encountered in wayside machine vision systems. Initial recordings were taken with the camera mounted in inspection pits, below rail level and looking upward at the bottom of the equipment (Figure 2). The factor limiting the camera's view of the entire undercarriage is the gauge face of the rail. This creates an inverted triangle in which two corners are formed by the bottom-outside edges of the railcar or locomotive, and the third by the camera. Based on previous experience, we decided that the best orientation of the camera was perpendicular to the equipment undercarriage.

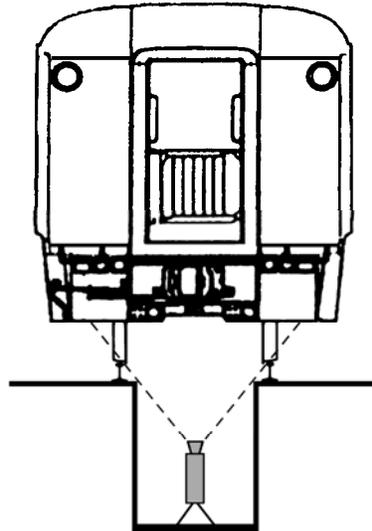


Figure 2: Camera Perspective from Beneath Railcar

Given the constraints on camera location, we needed to determine the focal properties needed to provide the necessary field of view in terms of image width and height. From the inverted triangle described above, we can determine a maximum depth of the camera lens below the railhead for a given piece of rolling stock. For a railcar with a maximum width of 10.5 feet and a floor 50 inches above the railhead, the maximum depth (l) is calculated using Equation 1.

	$\frac{50 + l}{126} = \frac{l}{56.5}$ $l = \frac{2825}{69.5} = 40.65'' = 3.39'$
--	---

Equation 1: Calculation of Maximum Depth

The video camera has two characteristics that determine the field of view: the camera's charge coupled device (CCD) size and the focal length of the lens. The camera we used has a ½ inch CCD. Equation 2 is a general equation that relates focal length (f) with distance from scene to lens (D) and horizontal width (w). Given that the horizontal width is 4 feet 8.5 inches at railhead height, with a maximum distance of 40.65 inches, and the constant (k) of 6.4 for a ½ inch CCD, the maximum focal length is 4.6mm, as seen in Equation 2.

$$f = k \frac{D}{w} = 6.4 \frac{40.65}{56.5} = 4.6\text{mm}$$

Equation 2: Calculation of Maximum Focal Length

We determined that a wide-angle lens with a 3.6-mm focal length yielded images that were satisfactory for our testing. This is one millimeter shorter than the maximum focal length in order to accommodate a camera enclosure over 12 inches long mounted on a tripod to be used in pits as shallow as 4 feet.

Changing the focal length of the camera is a balance between several factors. Increasing the focal length to its calculated maximum reduces the fish-eye effects of a wide-angle lens and therefore produces better quality at the edges of the image. However, a larger focal length also increases the distance from scene to lens necessary to obtain the requisite width and therefore requires positioning the camera farther below the railhead. Conversely, shorter focal lengths reduce the depth of the camera installation, but cause greater warping of the image.

Similar considerations are needed for infrared imaging. Current infrared cameras do not offer focal lengths as short as visible range cameras. The thermographic infrared camera we used for testing had a 25mm lens, much longer than the 3.6 mm lens used by the visible spectrum camera to capture the entire undercarriage. There are limited suppliers of wide-angle lenses for infrared applications, and therefore; these lenses are expensive. Since the camera cannot capture the entire width of the undercarriage, it can be positioned to capture a selected section of the undercarriage. Multiple cameras can then be employed to capture the remainder of the undercarriage as necessary or a wide-angle IR lens can be used.

Undercarriage Lighting for Visible Spectrum Recordings

After a number of initial test runs at Amtrak's Service and Inspection (S&I) pit facility in Chicago, we found that the initial lighting setup used produced motion blur in the videos. Moreover, the picture was much darker than expected. The lights did not adequately illuminate the outer edges of the car, whereas the low-hung items in the center of the car, such as water tanks, waste tanks, and truck components were over-saturated. There were also problems with shadows from car components. The solution to these problems was to use a faster shutter speed and stronger, more even lighting. After consultation with a motion picture expert in lighting methods, we upgraded the lights to studio-style lighting with 1,000 watts of incandescent light per fixture. These lights were also equipped with Fresnel lenses for more even light distribution.

With the help of the Monticello Railway Museum near the University of Illinois, we were able to conduct further testing of lighting using their inspection pit (Figure 3). Multiple passes were recorded of the same car to test the brighter lights in multiple orientations to evaluate their effect on video quality. We found that by using two lights and a combination of scrims (light inhibiting screens), we were able to provide fairly even illumination across most of the undercarriage. This reduced the problems of over-saturation and shadows.



Figure 3: Testing of lighting approaches at Monticello Railway Museum.

In order to provide strong, even illumination, we decided to use four studio lights surrounding the camera. However, because of their high power requirement each light needed a separate AC circuit. Four independent circuits were not available at either of our test locations; consequently, we were unable to use all four lights for several tests, and instead had to revert to two studio lights. Ultimately we solved this problem by supplying our own generators for the additional lighting fixtures and the improved lighting worked as planned.

Equipment Setup

The final lighting, video camera and IR camera setup can be seen in Figure 4 and Figure 5.



Figure 4: Final equipment setup at Amtrak

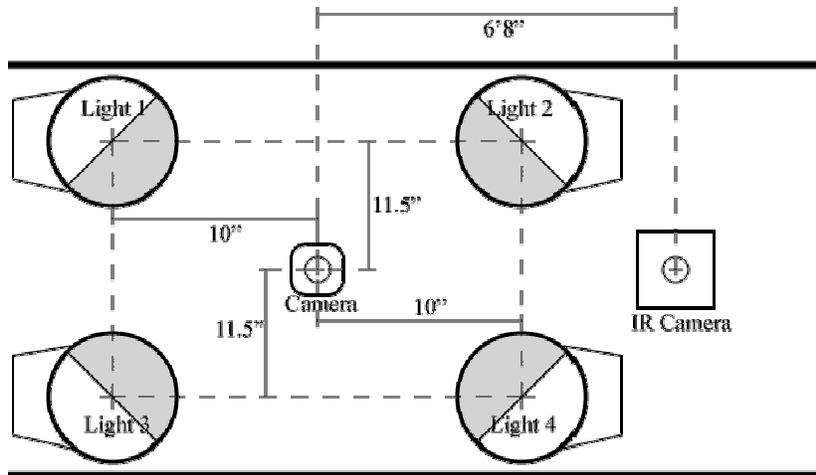


Figure 5: Final equipment setup diagram (not to scale)

In our final equipment setup, several scrims were used to further even out the illumination across the entire width of the undercarriage. All of the half scrims were placed at 45-degree angles, as seen by the gray shaded areas in Figure 5, and placed in the half of the light closer to the camera. This is necessary because the lights overlap most in the area above the camera which therefore needs to be dimmed. Lights 1 and 4 have one full single scrim and one half double scrim. Light 2 has one full single scrim and two half single scrims. Light 3 has one half double scrim. The evenness of the light was verified with a light meter on a flat target held 50 inches above the railhead. The IR camera has a tilt of 13 degrees toward the edge of the pit in order to cut down on IR backplane reflection and align the IR viewable area with the components of interest. These dimensions are not absolute, and adjustments can be made if necessary.

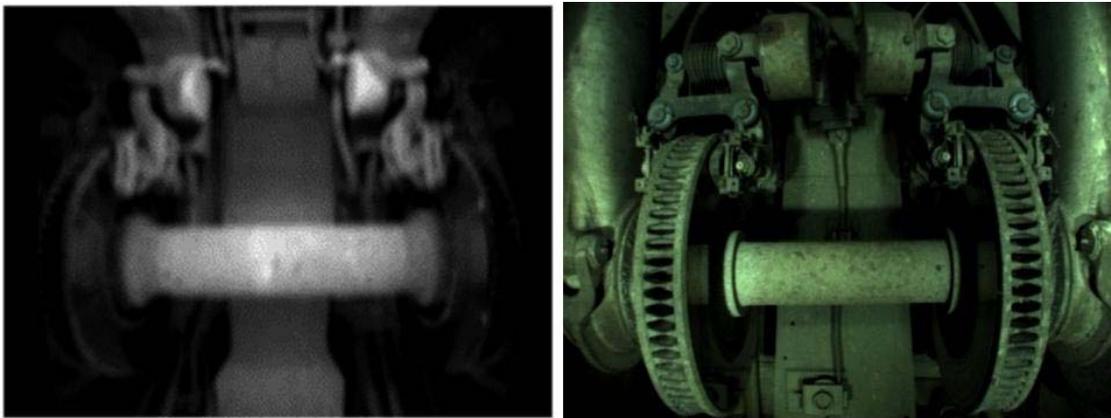


Figure 6: Comparison of initial image (left) quality to final image (right) quality

Image quality from the visible spectrum videos has improved significantly over the course of this project (Figure 6). The initial images were dark, blurry, and details were unclear, while the final ones show good detail, are well lit, and blur-free. These improvements are due to faster shutter speed and the improved lighting.

MODULE 1: IMAGE ACQUISITION AND PREPROCESSING

A digital video camera was selected that records images at a 640x480 pixel resolution. The camera is mounted inside a weatherproof enclosure to protect it from the effects of weather and other

environmental conditions such as dripping fluids, dust, dirt, and any other material from the undercarriage of the train.

A laptop computer controls the visible spectrum camera via a FireWire connection that is long enough to allow the computer to be located outside the pit. The computer contains software that allows the user to record and store digital video images as well as make adjustments to the shutter speed, white balance, frame rate, etc. of the camera. The video camera records at 30 frames per second. The lens aperture is fully opened to allow for maximum light to reach the charge coupled device (CCD). The protective case holding the IR camera has also been upgraded with an infrared-transparent protective cover placed over the view port to keep any dirt or dripping fluids from reaching the lens.

Image acquisition tests were conducted to verify that the videos collected were of suitable quality for both the preprocessing and inspection software modules. Adjusting the focus of both IR and visible cameras is accomplished by placing targets at a distance that represents the location of the undercarriage as the train passes. This procedure is required for the initial setup, and is not required once the equipment is installed. Although it is more difficult to see the same level of detail on the IR camera, this same method proved effective for focusing the IR camera. The IR camera can capture very good detail with proper adjustment of the camera parameters as seen in Figure 7.

Although the train speed was kept around 3mph (a normal pit speed), several tests were conducted with the train traveling twice as fast and there was minimal degradation in image quality. This indicates that the system could be modified to be used outside the yard, where trains travel faster. These modifications would include increasing the video frame rate and use of a shutter speed fast enough to prevent motion blur. More powerful lighting and/or a more sensitive CCD would also be required to accommodate these changes. In theory, doubling the frame rate should allow the speed of the train to be doubled as well. Since our off-the-shelf camera can only record at up to 60 frames per second (fps), the theoretical maximum for a system using this particular camera is 12mph. However, higher speed cameras are available that would enable higher train inspection speeds if that was desirable.

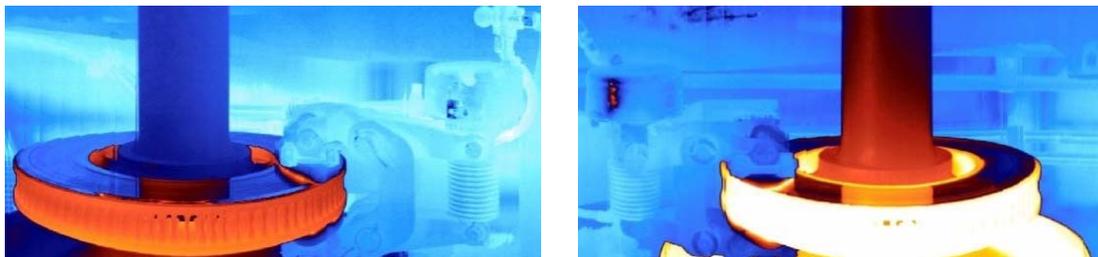


Figure 7: IR image quality at its best: two images from the same car type on the same train. The image on the right indicates a higher brake drum temperature which is thought to be due to a thin brake pad.

The optimal temperature range of the IR camera estimated in initial runs was used in our final testing at the Amtrak S&I facility. However, the ambient temperature during the testing sessions was hotter than during previous visits. This required readjustment of the IR camera to prevent over-saturation of the images. A production system for this technology would require development of an automated method of IR camera adjustment to keep the clarity consistent for different trains operating under different conditions.

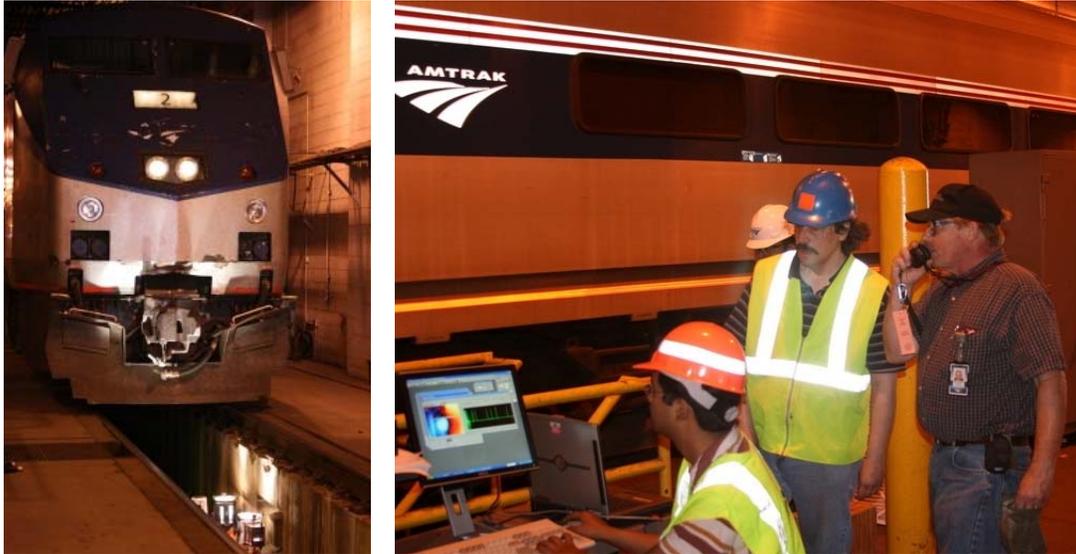


Figure 8. (a) Testing of system with train passing over inspection equipment, (b) Acquisition of both visible and infrared video to computers outside of pit as the train passes at the Amtrak S&I facility.

After the recording of in-service trains at the Amtrak S&I facility (Figure 8), the digital video data were taken back to UIUC for conversion into individual frames. The conversion of the IR videos into frames required the choice of a certain temperature range for IR display because there is a wide range of temperatures, but only a limited range of colors to display them within. As such, a well-selected range will enable high contrast between parts while minimizing unwanted artifacts. Once the individual frames were extracted, each frame was processed in a manner similar to the visual range images to remove the distortion created by the wider-angle lens used to maximize the field of view.

Image Alterations Induced by Weather

All of the video recordings completed have been under good weather conditions without effects of rain or snow. We simulated the effects of rain by spraying the underside of a car with water. As expected, the reflectivity of objects in the visible range was affected, but this did not substantially change the results. Both rain and snow present a drip problem for the camera equipment; however, adequate protection and cleaning can prevent this. A weatherproof enclosure already protects the camera, and a filter transparent to the infrared spectrum was used on the enclosure for the IR camera.

Concerning the effects of snow on inspections, we discussed this with an inspector with over 30 years of railroad experience. He told us that when a train comes in for inspection in winter, snow may be packed so hard on undercarriage components that it is impossible to inspect them. This is reported as “snow bound” by the inspector and only those parts that can be uncovered are inspected. Packed snow provides a challenge for both the visible and infrared range. As snow builds up, we expect that the ability to detect detail on objects on the undercarriage will diminish. Infrared cameras can only detect the temperature of the nearest physical object, so the snow buildup itself will register rather than the component behind it. Components that remain warm may not experience snow and ice accumulation, in which case the infrared camera should be able to record the component’s actual temperature.

MODULE 2: PANORAMIC IMAGE PROCESSING

Panoramic images are generated from video captured with train speeds typically observed (5 mph or less) result in an average inter-frame displacement of around 20 pixels per frame. Panoramic images are generated by first estimating the train’s displacement, in pixels, between two consecutive frames. Then, a center strip the size of the displacement is obtained from each frame, and all center strips from the video are stitched together to generate the panoramic images from both visible and IR videos. After this,

boundaries from the individual car panoramas are determined by detecting both the wheel axles and the couplers. These individual car panoramas formed in Module 2 will be used for MV inspection in Module 3.

The inter-frame train displacement was determined by matching each pair of consecutive frames, so that one frame is fixed and the other shifted by different pixel amounts until the two images are aligned (thereby undoing the change in motion-induced shift). The quality of alignment is measured by the sum of absolute difference (SAD). An alternative approach would be to use correlation between the images. SAD is acceptable in this case because inter-frame illumination changes are negligible.

Accurate pixel-level displacement is vital to creating a panorama without visible stitching artifacts (such as horizontal stripes). For certain types of passenger railcars, the varying distances of undercarriage components with respect to the camera lead to an appearance of the closer objects moving at a faster velocity. This is due to the effect of projecting objects in a three-dimensional world onto a single image plane. Closer components (i.e., axles) appear to travel faster than components further away (i.e., junction boxes). By computing the inter-frame displacement for each pair of consecutive frames, the length of the center strip changes dynamically to account for this effect. Undesired artifacts are only evident at boundaries surrounding low-hanging components. This should be explored in future work, but sufficient images were obtained for this study.

From the IR camera, panoramic images are generated at 30 frames per second. A similar methodology as previously described was used to stitch the strips of consecutive IR frames into the panorama. A common source of panorama error for the IR images is the presence of two dark semi-transparent layers, the presence of artificial horizontal edges produced by the IR camera, and the presence of an artificial bright light at the image's center that is an artifact of reflection back into the camera. These artifacts can be removed by preprocessing the IR video, so that the artificial bright light and horizontal edges are filtered out. Also, since the IR video from certain recording sessions has been error-free, use of a higher quality IR camera, or one meant for outside use, could eliminate these problems.

After the panoramic image of the train has been generated, it is parsed into panoramas of the individual pieces of equipment. To accomplish the detection of damaged, foreign, and missing components, a car-specific template is created and stored for each piece of equipment. This allows for each car to contain unique configurations within its railcar type, which will arise after repairs and other modifications. Storing a unique template for each car will lead to better sensitivity in the detection of missing and foreign components. For detecting defective components, a template of ideal operating conditions for each specific component of interest is also stored. Wheelset (wheel, axle and brake disk) detection is used to approximate the length and location of each car. This is followed by coupler detection to find the exact location of the ends of the cars. We use this two-tiered approach, because axle detection is more robust than coupler detection.

In the trains we recorded, there were two visually distinct types of wheelsets (one for the locomotives, and one for the railcars). Each wheelset type has the same shape. All axles are fully visible from the camera position used (i.e., there is no partial occlusion from the other components of the undercarriage). Couplers, however, have varied shapes and are often partially occluded by cables and air hoses. The axle detection routine is used to limit the search area in the panoramic image where the coupler is located. Once this search area has been estimated, it becomes easier to determine the most likely coupler type and location by matching the panoramic image with a variety of coupler templates showing different known coupler configurations.

Both axle and coupler detection require templates to be captured ahead of time. In our experiments, we used a manually edited axle edge image, and samples of several coupler configurations. In an actual system, the necessary templates could be generated by systematically recording each type of equipment in an operator's fleet and storing the data for subsequent retrieval as needed.

For axle detection, edges in the panoramic image are first detected by using the Canny edge detector. This edge panoramic image is then matched against the template representing the edge image of the axle. Much of the information necessary for object recognition is contained in the edges, or outline, of

the object. Edges are particularly important for this application because changes in colors and intensity, both in the visible and IR images, are expected due to external factors.

A distance transform (DT) is used to match each axle's edge image with the axle edge template images. Given two edge images of the same component, such as two wheel axles, it is intuitive that to match the components that are common in each image, one should find the location where the maximum overlap is achieved between the edges. However, such an exact match is not suited for real-world conditions due to variability in components, or even in the quality of the generated images. Matching using DT reduces the penalty of not finding one-to-one matches between edges by giving a low distance score to pixels that are close to edges so that a "close" match can also be considered. This allows the use of one axle edge template for a variety of images with varying quality.

We also use the axle edge template to divide the thermal panorama into individual cars. The thermal panorama appears at a different scale than the visible panorama, as the IR camera produces a more magnified view of the undercarriage due to the use of a different lens. Though the physical and thermal images contain different colored railcar parts at different scales, they share the shapes made by the edges. Therefore, to infer the scale change between the physical and the thermal images, the axle edge template is iteratively scaled to different sizes, and then matched with the thermal panorama. The best match indicates the scale and location of the axle in the thermal image. By using the detected axles as anchor points, the thermal panorama is registered and overlaid on top of the visible panorama. Further analysis of each component in the thermal image is done in Module 3.

MODULE 3: MACHINE VISION INSPECTION

Global and coarse-level anomalies are detected (Module 3A). This is followed by local inspection of individual components to detect physical and thermal anomalies (Module 3B). Finally, duplicate components within a single piece of equipment, or train, are compared and analyzed to identify operational outliers (e.g. non-uniform brake operation) (Module 3C).

Module 3A: Global Anomaly Detection

Foreign and missing components are best detected on a coarse level, where a global view of the railcar is compared to a car-level template. Foreign and missing components will result in a high level of mismatch between the equipment and the template in the area where they are (or should be) located. For both physical and thermal data, a global, car-level template is stored for each unique railcar.

To detect global physical anomalies, block-level correlation is performed between the undercarriage panorama and the railcar template. The size of each individual correlation block is approximately 60x60 pixels. Each panorama for a single piece of equipment was approximately 7040x640 pixels. Areas of mismatch have low correlation (shown dark), and a threshold over correlation values is set for declaring an anomaly. Once an anomaly is located, a subsequent step can be added where one identifies the nature of this anomaly. See Figures 9 (a-c) for an example of this process where the anomaly is marked with a red arrow.

To detect thermal anomalies, we compute the difference in color values between the thermal panorama and the equipment's thermal template. Ambient temperature is one important consideration. Colors associated with the thermal panorama of a currently captured railcar should be rescaled, so that the average (ambient) temperature matches with the average temperature of the railcar template. This ambient temperature is recorded, as it is used for calibration. Figures 9 (d-f) show the results of block-level global anomaly inspection for a thermal panorama.

Module 3B: Component-level Anomaly Detection

Individual components are compared against the corresponding templates that contain the ideal physical (or thermal) characteristics of each component. In general, there should be one general template

corresponding to each component type. However, sometimes components vary significantly across the railcar fleet. In this case, the car-specific templates should be captured in advance. Given the templates, the visible and thermal panoramas of a railcar, the components of the undercarriage are examined as follows:

1. Identify the scale and location of a given component in the video and thermal panoramas by aligning the template of that component using edge images.
2. Detect the areas where a component differs from its template.
3. Determine the categories of the anomalous regions using a statistical model of each anomaly category.

Step 1: Align Components with Template

In the first step, we align the components using edge matching with the DT applied to the Canny edge images. This is similar to the procedure of parsing the panoramic image of the entire train into individual panoramas of railcars, performed in Module 2. In our experiments, this alignment has proved sufficiently robust to partial occlusions.

While several alternative methods (e.g., corner matching or correlation) could be used with similar effectiveness to locate an undercarriage component in the visible image, edge matching is the most effective method for this purpose in the thermal images. When examining the thermal panorama, the components are best detected by their edges, because corner matching and correlation are susceptible to changing intensities occurring in the thermal image.

In Figure 10, there are two water containers. Figure 10 (a) is a template that was formed by merging all three water containers of this type that were obtained experimentally. We used the median-filtering technique to form a composite image from the median pixel values of three spatially-aligned water containers. Figure 10 (b) is a particular water container found experimentally, and it is evident that it is missing a faceplate. One important step in creating the template, and in comparing each component to the template, is to align them in a coordinate space. This is accomplished through the DT of the Canny edge images, as previously described.

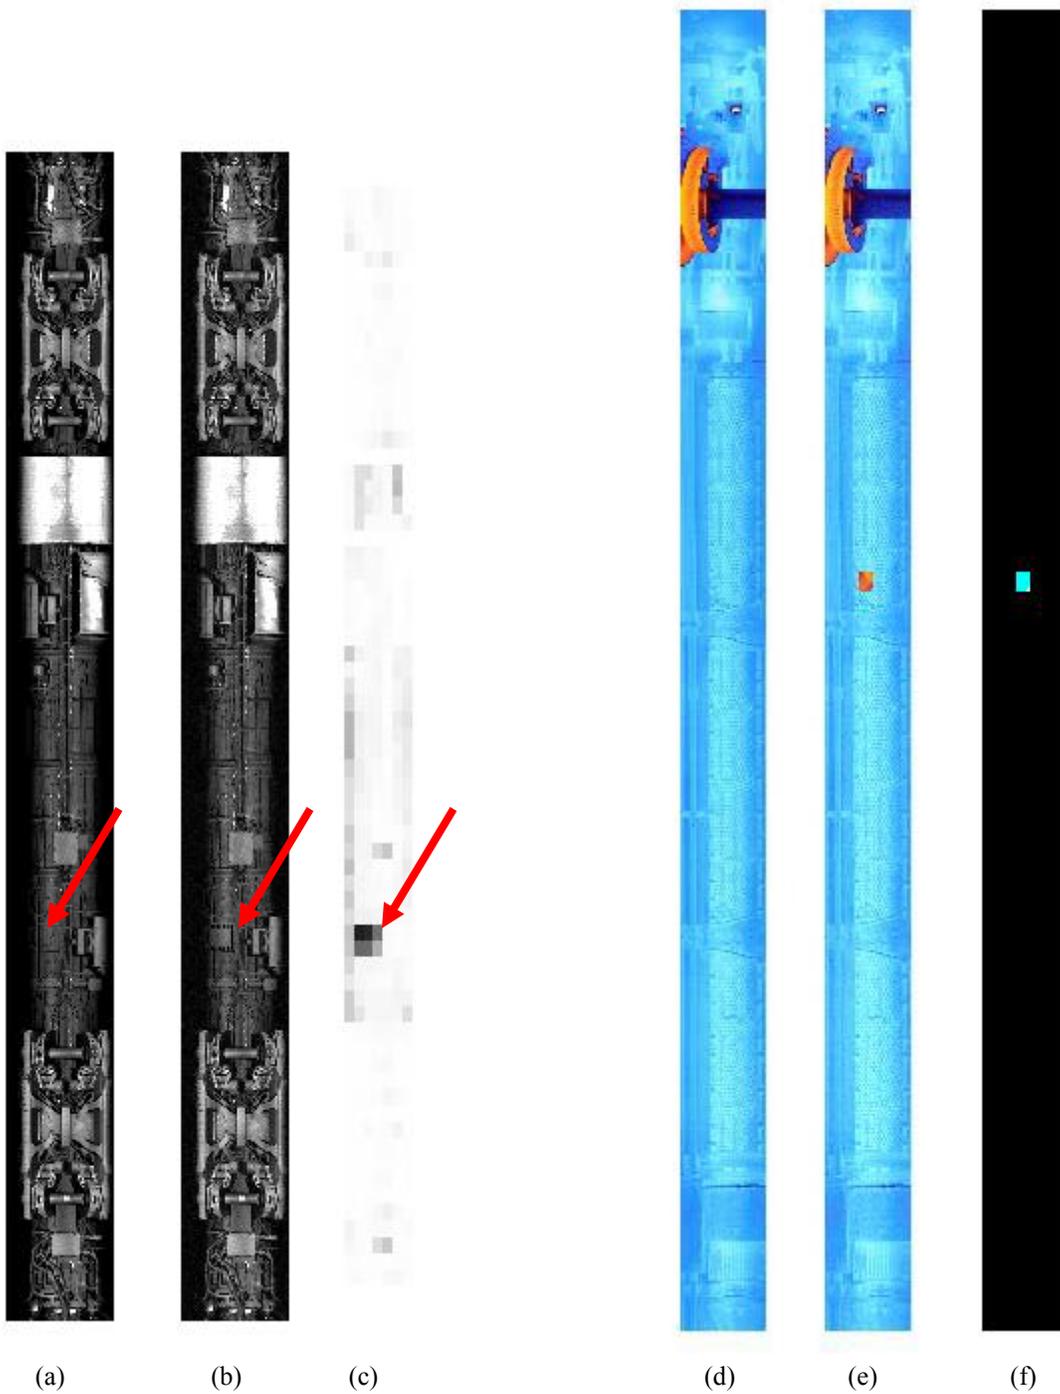
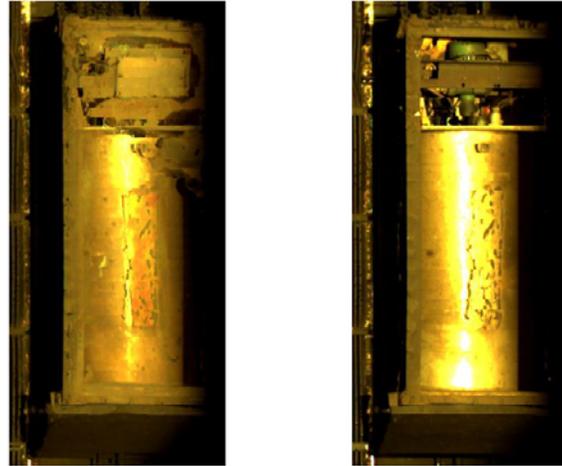


Figure 9. (a) Visible undercarriage template of an Amfleet-II Coach, (b) undercarriage with defect, and (c) the detected defect, shown as a dark block. (d) Thermal undercarriage template of a Superliner, (e) undercarriage with thermal defect, and (f) the detected defect, shown as a lighter block.

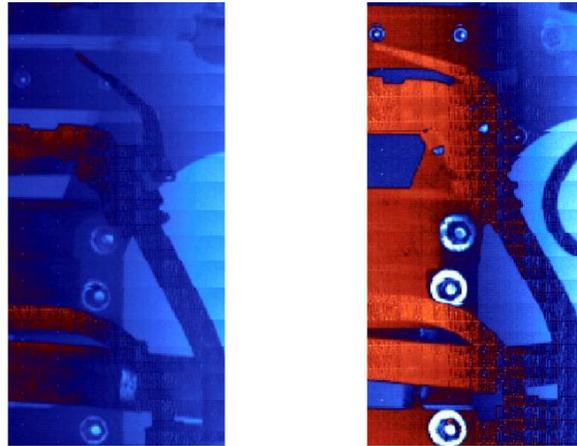


(a)

(b)

Figure 10. (a) Water container template and (b) defective water container.

Similarly, in Figure 11 there are two draft gear boxes in the thermal domain. The template shown in Figure 11 (a) was actually constructed from just one image due to insufficient data, and therefore there is an assumption that this is the “normal” thermal operating range of the component. The suspect draft gear box in Figure 11 (b) was found in a separate thermal image, and it is assumed that this is outside the normal operating range for the sake of experimentation.



(a)

(b)

Figure 11. (a) Thermal template of a draft gear box and (b) a suspect draft gear box.

Additionally, both the A.C. unit and traction motor were located in the thermal and visible domains, as shown in Figures 12 and 13. In Figure 12, the motor for the A.C. unit is hotter than its surroundings, and this provides a more complete picture of A.C. unit’s operating conditions. The portion of the motor that is occluded by the rectangular covering in the visible domain becomes apparent in the thermal domain. This demonstrates that the IR camera can capture the thermal properties of some occluded objects, and not merely the temperature of the closest object. It was previously believed that the thermal properties of the object closest to the IR camera would only be detected. The full extent to which the IR camera can overcome occlusion has yet to be tested.

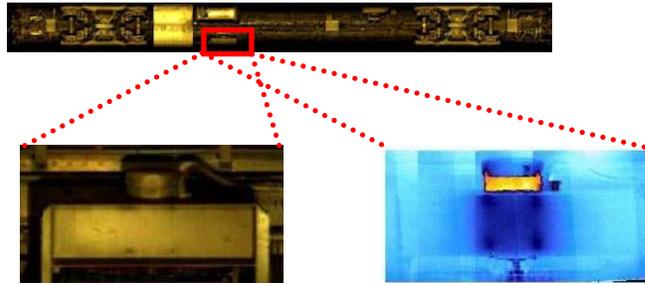


Figure 12. A.C. unit

In Figure 13, a partial view of the traction motor is shown. Only one side of the train could be captured due to the IR camera placement constraints. Not only is the traction motor heated, but the surrounding areas are heated as well. Uneven heating on the traction motor surface is also apparent.

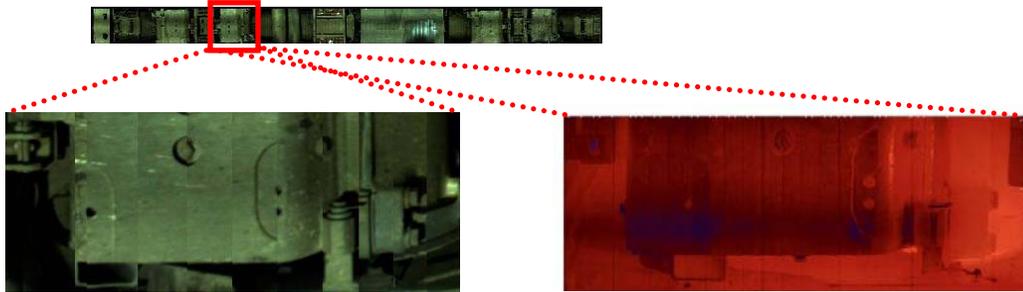


Figure 13. Traction motor

Due to limited data acquisition, we did not obtain multiple images of these A.C. units or traction motors. Module 3B will continue with the example from Figures 10 and 11. In future work, the method developed in the second and third step of Module 3B will also be used for classifying anomalous regions of the A.C. unit and traction motor.

Step 2: Identify Areas of Mismatch

In the second step, the areas of dissimilarity between the template and the detected component are identified. This is shown in Figure 14. Since the dissimilarity is measured by overlaying the template and aligned component, it justifies the necessity of alignment in the first step. By using only these areas of dissimilarity for the third step, the problem of anomaly classification is conveniently split into two parts: determining if a component-level anomaly has occurred (second step) and classifying this anomaly (third step). To identify the area of anomaly, i.e., the area of mismatch between the physical and thermal images of a given component and its template, we compute the correlation of the template and the visible image, as well as the difference in red, green and blue (RGB) color values between the template and the thermal image. In the visible domain, areas with small correlation values below a given threshold are anomalous, and in the thermal domain, areas with large RGB color differences above a certain threshold are classified as anomalous.

Figure 14 (a) shows the areas of low correlation between the water container template and the defective component that were shown in Figure 10. The parts of the defective component that are located in this low-correlation area are labeled as areas of high visual mismatch. Similarly, Figure 14 (b) shows the areas where the temperature difference between the template and the suspect draft gear box shown in Figure 11 are the greatest. The contents of the suspect draft gear box that are located in this area of high temperature difference are labeled as areas of high thermal mismatch.

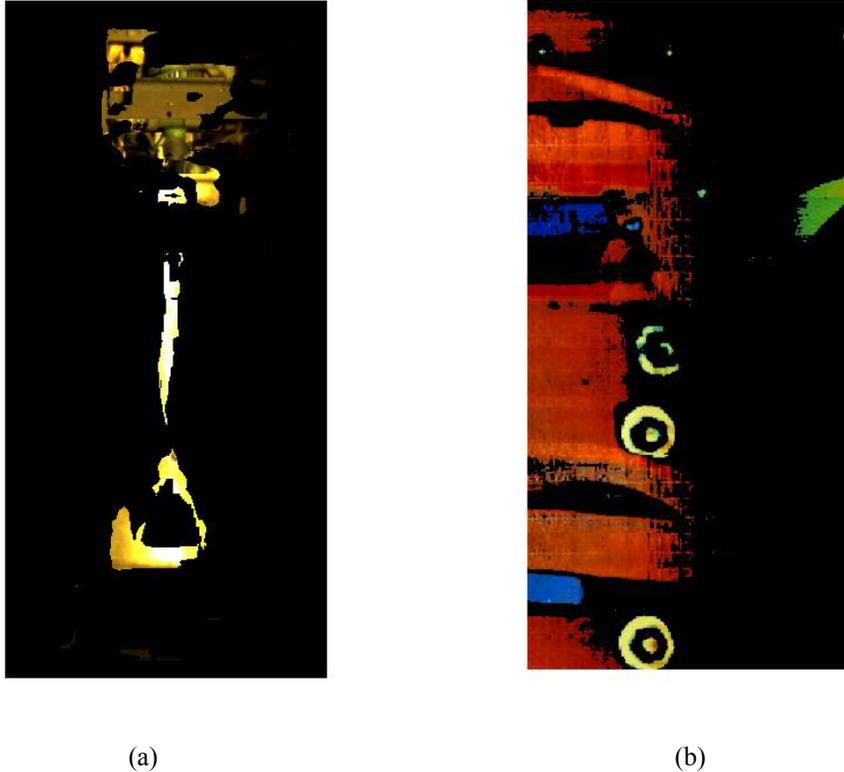


Figure 14. (a) Visual mismatch and (b) thermal mismatch.

In the third step, the areas of high mismatch are grouped into contiguous regions of similar pixels. In our experiments, we have used the well-known K-means clustering algorithm to form these regions, though more sophisticated region-forming techniques can also be employed (e.g., watersheds, region agglomeration, etc). These anomalous regions are then characterized by the following features: 1) x and y coordinates of the region's centroid; 2) the mean red, green, and blue color values; 3) region area in pixels; 4) a simple shape estimation of the region (eccentricity); and 5) a measure of how compact the region boundary is (compactness). Thus, each anomalous region is represented by a five-dimensional (5D) feature vector. This representation allows us to formulate the problem of identifying the nature of a detected anomaly as vector classification in the 5D feature space.

Step 3: Learn Defect Categories

Each defect category is statistically represented in the 5D feature space using the Gaussian mixture model (GMM). In the GMM we used, a single anomaly category is represented by a weighted combination of several Gaussian probability distributions; therefore, each category can have multiple domains scattered in the feature space. It follows that the weighting coefficients, as well as the mean and variance parameters of the individual Gaussian distributions, have to be learned in training on a sufficiently large number of examples of anomalous regions.

Human intervention is required during the training phase to classify the available training examples. The class labels can be rough (defect/no defect) or more specific (defect-corrosion, no defect-shine, defect-missing faceplate). In case an external condition causes a mismatch with the template (e.g., snow buildup), it can be added as a class; therefore, whenever this condition is encountered again (e.g., snow is viewed as a white region in the visible image, and a cold region in the thermal image), it can be

correctly classified. Although the experiments in this report have examined physical and thermal defects separately, the results can be integrated in future work by using the (x,y) coordinates of regions' centroids to achieve a more robust anomaly classification. Figs 15 and 16 show an example of physical and thermal anomalies, where mismatched regions are detected and classified. The successful results we have obtained demonstrate the robustness of the approach to different methods used for forming anomalous region groupings.

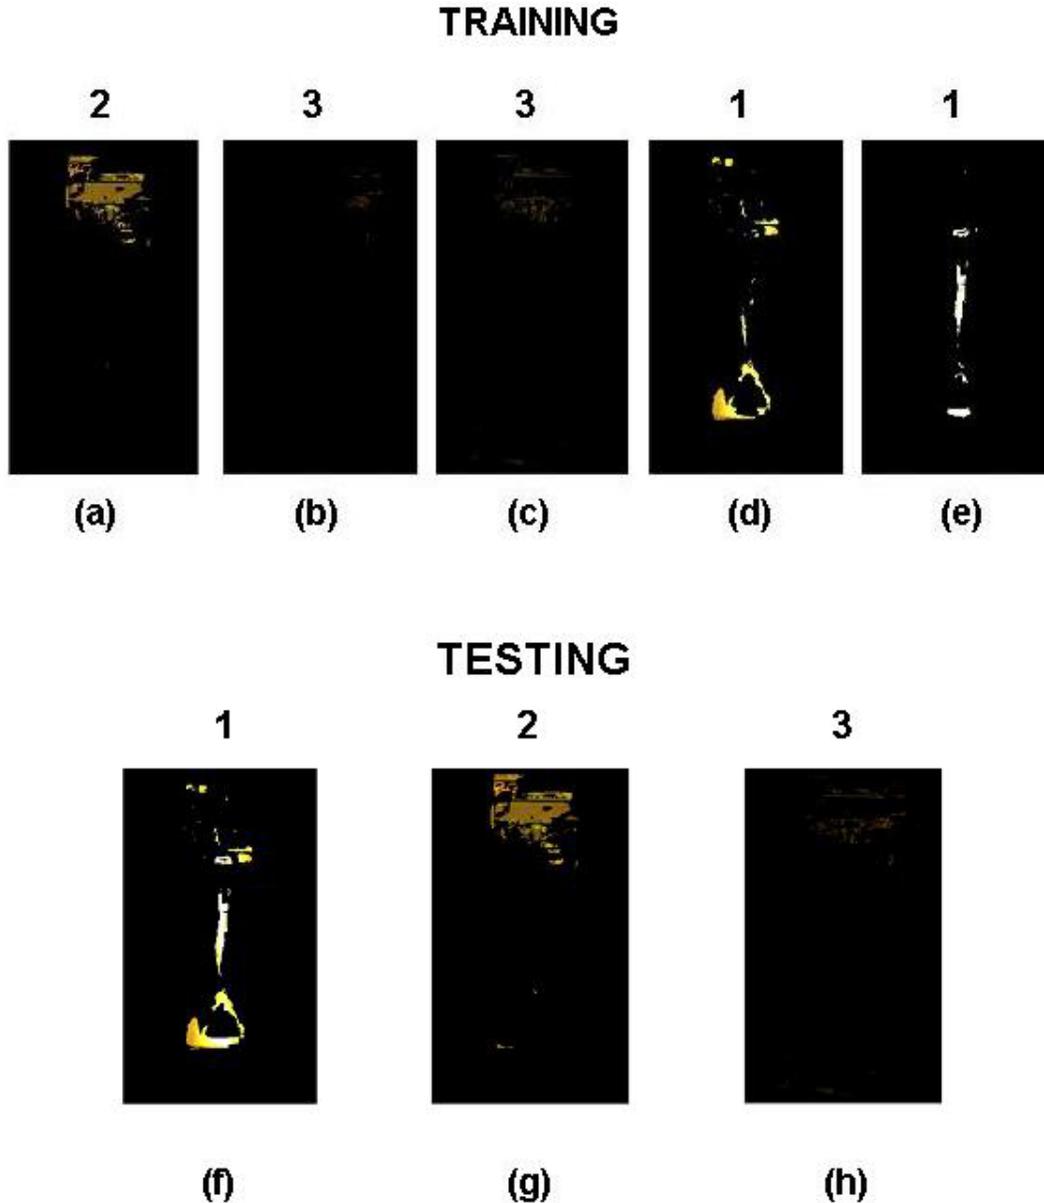


Figure 15. Results of physical component defect classification. Regions are labeled in training (a)-(e) and detected in testing (f)-(h) as: 1=shine (no defect), 2=missing faceplate (defect), and 3=no defect.

In Figure 15, during training (a)-(e), K-means produces five regions from the area shown in Figure 14 (a). As each region is produced, a user labels them 1, 2, or 3 based on the nature of the mismatch. Label 1 means a harmless mismatch (in this case, there is more shine than usual), label 2 is a missing faceplate, and label 3 means there was no source of mismatch worth noting. A Gaussian Mixture Model (GMM) is trained with these labels. A legitimate concern is the sensitivity of this scheme to the regions

that are formed, since various factors will cause future regions to be different in practice. For this reason, the experiment was rerun with K-means producing only three regions, as shown in Figure 15 (f)-(h). Note that the region in (f) is not in

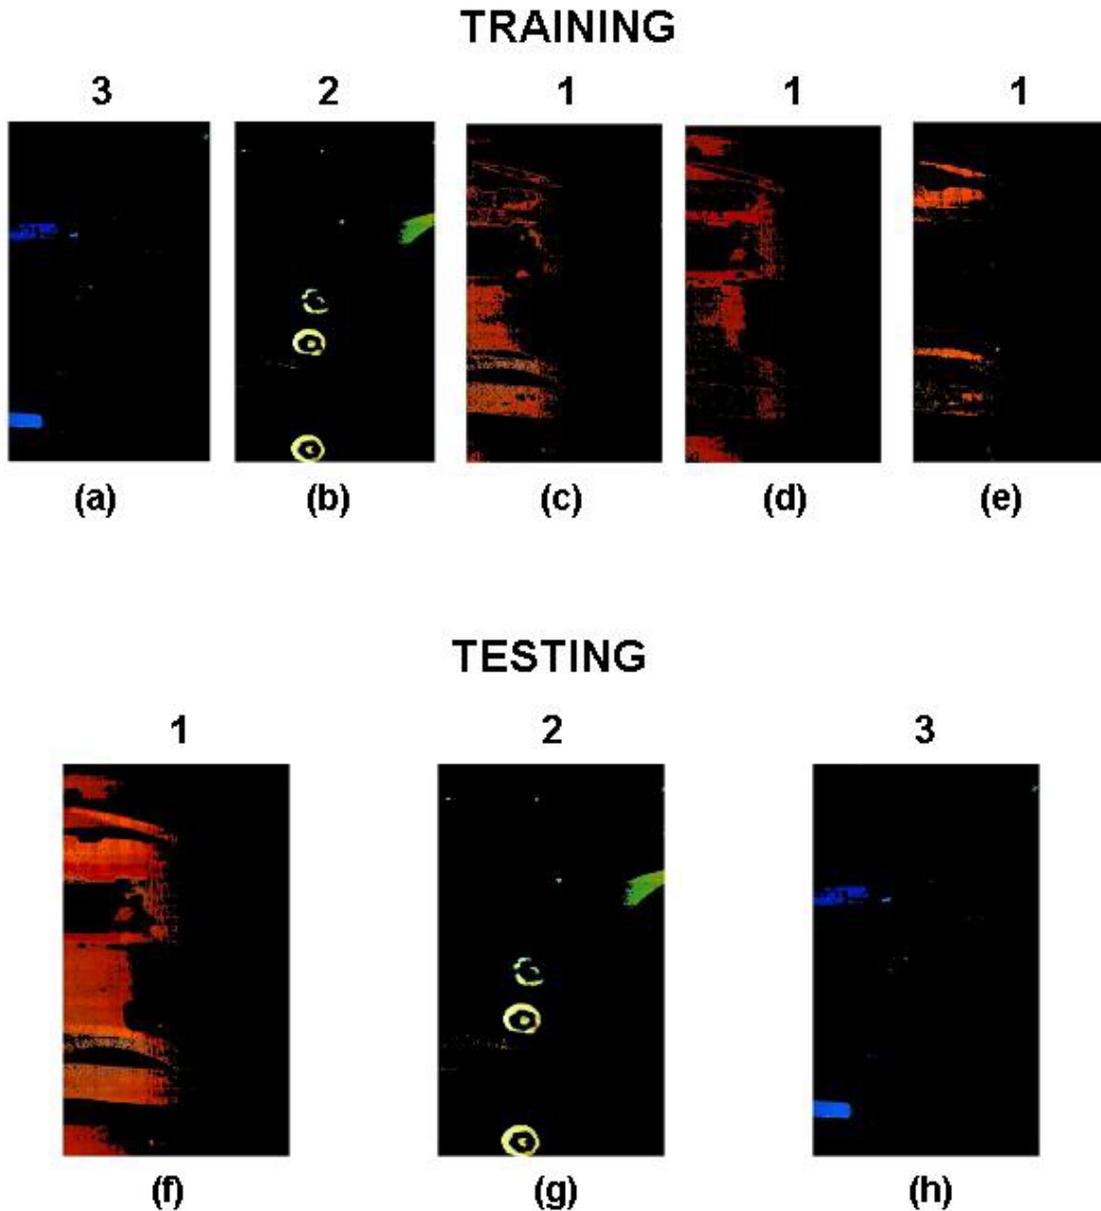


Figure 16. Results of thermal component defect classification. Regions are labeled in training (a)-(e) and detected in testing (f)-(h) as: 1=hot surface (defect), 2=shine from nuts and bolts (no defect), and 3=no defect.

the training set, but it is similar to both (d) and (e). The GMM model is run with these three as test data, and it successfully classifies all three regions (f)-(h). This demonstrates the ability of the GMM model to generalize to previously unseen regions.

Similarly, Figure 16 (a)-(e) shows the results when K-means produces five regions from the area shown in Figure 14 (b). The regions are manually labeled, as shown in (a)-(e) and used to train a GMM. K-means is rerun using three regions, as shown in (f)-(h). Again, the resulting labels that the GMM creates for (f)-(h) show the ability of the model to generalize to previously unseen regions.

Module 3C: Relative Anomaly Detection of Like Components

In Module 3, individual components can also be inspected for consistency of operation within a single railcar. For example, the brakes should be applied uniformly, and there should also be an even distribution of heat produced during the uniform operation of these brakes. To begin with, the brakes are located as shown in Figure 17.

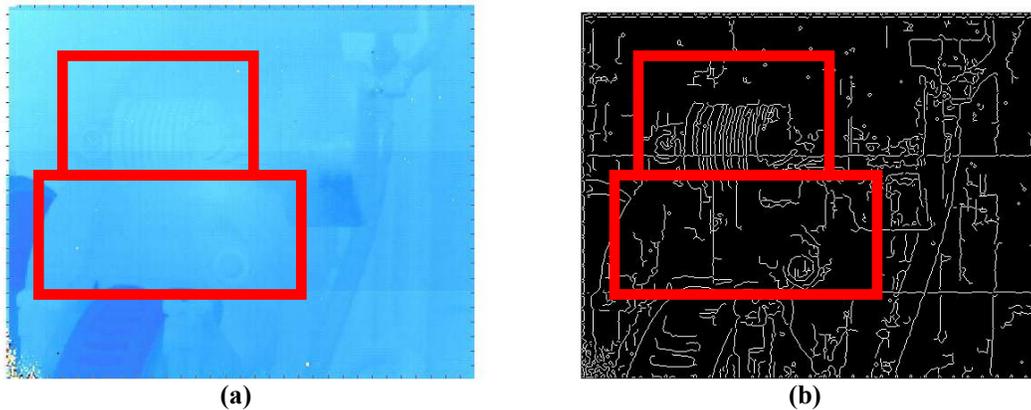


Figure 17 (a) IR image of one brake (b) the Canny edge transform of one brake.

Within the edge transform, measurements are made to determine proper application of each brake. The measurements shown in Figure 18 measure the operating size of the return spring (Figure 18 (a)) and the brake caliper (Figure 18 (b)). By measuring these distances, we measure the amount of contraction applied to the return spring device, and if the brake caliper portion is reacting within the range of acceptable motion to the heat being generated by it. Also, since these components are now located, measurements can be made of the surrounding thermal areas. For example, the thermal temperature of the disk brake, brake shoes, and other points of interest are made. All of these measurements can then be placed into a hybrid feature space that contains both distance and thermal measurements.

In this hybrid feature space, a GMM can be created to identify acceptable operating conditions of the brakes with respect to all brakes on the train (using all data) or it can also be used to identify single outliers with respect to the other four brakes on the railcar. This is done in a similar fashion to component-level defect detection previously described; only the measurements in Figure 18 are now in the feature space. Additionally, thresholds could be determined before inspection so if any brake is mismatched with the other brakes by more than a predetermined amount, it should be deemed defective. By using this hybrid feature space, the automated inspection system can verify proper brake operation with respect to the entire train, with respect to other brakes on the same car (to check for balanced application), and with respect to any preset threshold.



Figure 18 (a) The return spring portion detected, and the amount of compression measured with the red line, and (b) the brake caliper portion detected, and amount of brake application measured with the red line.

Once the return spring and brake caliper were located (shown in a box in Figure 18), the landmark components shown in Figure 19 were used to create the measurements (shown as lines with dotted endpoints in Figure 19).

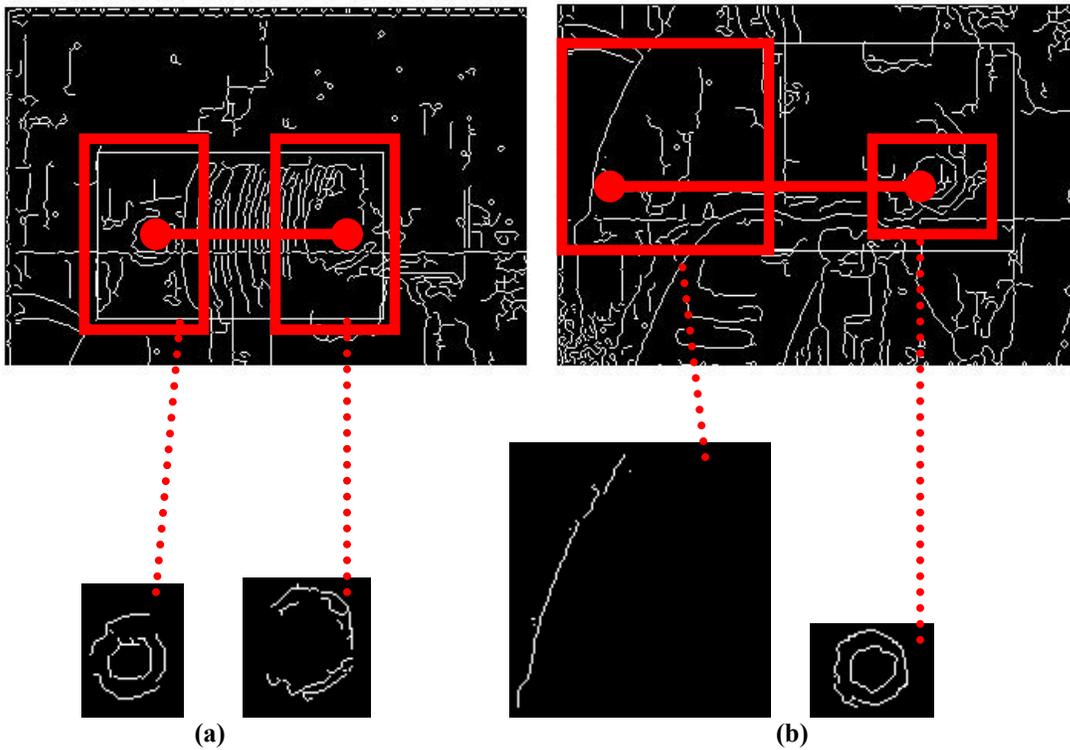


Figure 19. (a) The landmarks used for measuring return spring compression, and (b) the landmarks used for measuring brake caliper width.

The viewpoint of the IR camera strip limited the number of visible brakes to only two per car; therefore our data was limited in determining “balance”.

Table 1 gives a summary of two brakes (one front and one rear) found in four railcars. Pixel displacements were rounded to the nearest 5 pixels (values 1,6,11,..) to expedite testing. For the return spring, up to a five-pixel differential between two brakes in the same car was observed. For a car with two evenly applied spring components, the brake caliper could have up to a ten-pixel differential. Lower pixel values in Table 1 correspond to a more closed caliper.

Once all components in Figure 19 were located, a thermal snapshot of the brake disk area was taken. Since the IR camera calibration scale is stored, this can be converted to a temperature. For these settings, a brighter orange corresponds to a hotter temperature. If this chart is viewed in black and white, the median luminance value for the image is also given. A higher value approximately corresponds to a hotter temperature. In addition to the mean luminance, the distribution of the colors could be used to detect uneven distribution of colors, such as the dark spot on Brake 2 of Car 3 that appears in the IR image.

	Car 1		Car 2		Car 3		Car 4	
	Brake 1	Brake 2	Brake 1	Brake 2	Brake 1	Brake 2	Brake 1	Brake2
Return Spring compression measurement (pixels)	121	121	116	121	116	121	116	116
Brake Caliper opening measurement (pixels)	204	199	199	209	209	199	204	204
Snapshot of Disk Brake Area								
Luminance Median of Disk Brake Area (temperature)	0.37	0.39	0.58	0.49	0.47	0.45	0.32	0.37

Table 1. Brake operating properties

In Car 2, because of the smaller brake caliper opening and higher temperature of Brake 1, a thinner brake pad (with respect to Brake 2) is suspected. User defined thresholds will have to be set for declaring outliers in the braking system. As stated previously, all four brakes of a railcar should be examined (we only could acquire either the left or right side of a train with the current IR setup). Then, the numerical values in Table 1 would be used in a GMM, along with the data that represents any shape detected in the snapshot (such as eccentricity and compactness, as done for the case of component detection). Options should be available to set relative thresholds based on the resulting distributions, or hard thresholds based on railcar standards.

TEST PLAN DEVELOPMENT AND TESTING RESULTS

A Field Test Plan was developed in Stage 2 of the project and reviewed by the TRB IDEA Program. The primary objectives of these tests were to show the feasibility and effectiveness of the three main modules of our proposed inspection system. The modules were tested sequentially at different times during the field-testing. Module 1 was tested at the Amtrak site and Modules 2 and 3 were performed in the laboratory at UIUC. Since we were bound to testing only during the normal operating hours of Amtrak, our system evaluation was limited to the trains that were available at the time. Three trains in total were captured with multiple passes of the last train in an attempt to induce changes in temperature. Performance

specifications in each module were designed to verify that the functions performed satisfactorily in each category. The following paragraphs present an overview of the criteria for each module.

The main test criteria for the Image Acquisition and Preprocessing Module were developed to verify that the images captured were suitable for both the panorama generation and inspection modules. The criteria examined the focus, exposure, and field of view of the video systems from both spectra. The focus of the cameras must be set so that the images are sharp enough for proper panoramic image generation and edge detection. Proper exposure was tested to ensure the images were well lit so that components to be inspected are seen in high enough detail without being washed out by overexposure. Testing at different train speeds was also conducted to ensure that changes in pit speed would not affect the clarity of the images. Tests were conducted showing the effect of temperature changes in the field on the temperature range of the IR camera output images. Testing also ensured that the image extraction, reformatting, and dewarping were performed correctly to provide proper input to panorama module.

When generating images of entire trains, using the Panorama Generation Module, verification was made to ensure that the amount of missing or duplicated pixels was minimized during the stitching process; this was done by visual inspection and also by verifying proper railcar proportions. An evaluation was also made to verify that panorama generation did not affect the contents of the images, to make sure that they were a true depiction of the component condition present when the train passed over the camera system. Verification of proper separation of individual railcar panoramas from the original train panorama was also made and subsequently improved by adding a prior wheel/axle detection step to limit the area searched for the identifying the coupler.

In the Machine Vision Inspection Module, foreign and missing components were found to be best detected at a course-level using the global railcar template. Verification was made to ensure a high level of mismatch was produced when the equipment panorama and the railcar template were compared in areas where foreign or missing components were, or should be, located. This was also verified in the thermal anomaly detection where the difference in color value is computed between the railcar's thermal panorama and the thermal template. At the component-level, individual templates containing the physical (or thermal) component characteristics were tested to make sure a proper scale and location could be identified to match the component in the equipment's panorama. Testing was also performed to confirm that the system could learn to properly categorize anomalies into specific defect classifications. Difficulties in the identification of false positives were later addressed by utilizing the learning algorithm to train the system to properly classify these into categories of actual defects or into new categories that were based on the condition that caused the false identification.

PROJECT PANEL

The panel convened for the project includes Paul Steets of Amtrak, Jim Lundgren of TTCI and Gavin Horn of AMTEL and IFSI at UIUC. Paul Steets offered Amtrak's knowledge, facilities, and equipment to work on this project. Jim Lundgren has had principal responsibility for TTCI's development of railroad machine vision technology. He is knowledgeable about the variety of inspection technologies currently available and the various applications for which there is interest in both freight and passenger rail. Gavin Horn is the Illinois Fire Service Institute Research Program Manager where he works in the area of firefighter health and safety; he also holds a Research Scientist position with the Advanced Materials Testing & Evaluation Laboratory at UIUC. In both of these research areas, he works directly with IR cameras and analysis of the resulting data.

Gavin Horn arranged for the use of the IR camera for the work on this project. He also provided training for our group on the calibration and operational procedures for the IR camera. In addition, he conducted preliminary tests at both the Monticello Railway Museum and at the Amtrak Facility in Chicago.

With the understanding that the video recording is in the visible and infrared (thermal) range, the group from Amtrak decided the priorities for initial inspection from beneath the train were 1) traction

motors, 2) air-conditioning units, 3) brakes, and 4) axle bearings. Traction motors can be monitored for overheating of either the bearings by friction or electrical overheating of the motors themselves. For air-conditioning units, thermal imaging can be used to determine if the system, especially the compressor, is being overworked. The disc brakes common on Amtrak equipment are more easily visible from underneath than from a wayside location. The specific car types recommended were the Horizon, Amfleet, and Genesis locomotives.

This application could be used for freight trains as long as appropriate components are designated for inspection. The system could be adapted to inspect the traction motors on freight locomotives. Outside of that, further discussions with participants in the freight railroading industry could provide knowledge about which components would be visible from between and below the rails that would benefit from physical and thermal imaging.

CONCLUSIONS

Advice from our panel was instrumental in arriving at the initial definition of our system. Our design was based on the range of equipment, car designs, and components specified by the panel early on. This assistance also included the recommendation and authorization for testing at the Amtrak Service and Inspection (S&I) Facility in Chicago using in-service equipment.

Initial image acquisition at the local Monticello Railway Museum exposed difficulties such as image blur due to inadequate lighting. Further experimentation with additional lighting, recommended and provided by a motion picture consultant, allowed the exposure time of the camera to be reduced thus eliminating the motion blur from the moving train. Creating early panoramas led to other restrictions on the equipment setup. A sensitivity to camera rotation (skew) with respect to the track caused difficulties in matching consecutive frames of the video for panorama generation. Automatic image adjustments to slightly rotate the frames, prior to extraction of the center of the image, were added to the software to compensate for this. To provide an adequate amount of overlap in the center strips of consecutive frames, a faster IR camera frame rate was needed for the velocity of trains moving across the inspection pit. Similarly, to adjust the amount of overlap selected by the algorithm, the velocity of the train in the video was continuously monitored to make adjustments when the train speed increased or decreased.

During the development and testing of the prototype system, we addressed the challenges of car types and varying component conditions. Ideally, hundreds of car images are needed to analyze the large number of possible variations from not only one type to another, but also the many variations within a single car type. To deal with this during our investigation, templates were created based on fusion of data from several cars of the same type, thus encompassing any small variations; efficacy of such fusion was brought out in the successful detection of the absence of the cover plate in the example shown earlier. It is recommended that the final system should capture each car and create a unique template for that car instead of using a single car type template. This would be the case for the IR templates as well. Concerning variations in component conditions, a classification algorithm was developed to learn to decipher subtle differences (such as shine from new metal) and actual damaged or missing components. Our experiments showed that global template matching was most effective for finding missing and foreign object anomalies. However, the system needed more detailed templates and a more robust matching algorithm in order to effectively find damaged and thermal component details. These were later implemented and the results have been presented in this report.

The software completed in this study serves as a platform for further development. It has the flexibility to easily incorporate other data sources and features.

Although the majority of our testing was completed under good weather conditions, we made an effort to discover how images would change under varying weather conditions. During initial testing, the undercarriage was sprayed down to simulate rain conditions to produce reflections off of wet components. The effects on the panoramic image generation were minimal but the resulting change in appearance due to the reflections could now be taught to the classification algorithm in a similar manner as was done for the

shine off of new or cleaned metal. An investigation was also conducted to determine the effects of snow conditions. In discussions with an inspector with 30 years of railroad experience, it was discovered that packed-on snow can result in an inability to manually inspect certain components. In these cases, components still visible are inspected but on cars where components cannot be uncovered from beneath the snow, the car is then reported as "snow bound". Packed-on snow will be problematic for the visible spectrum and will be detected as an anomaly. The IR camera would detect the temperature on the surface closest to the camera, i.e. the snow or ice, unless the temperature behind the snow is significantly hotter, in which case it may detect it but with inaccurate results. However, the system can be taught that a combination of a visible mismatch along with very low temperature readings from the thermal spectrum be classified as "snow bound". On the other hand, a cloud of water mist or blowing snow generated from fast moving trains would significantly degrade the image acquisition process and ultimately be problematic to comparison and detection algorithms. On the mechanical side, adequate protection was provided to the equipment under the track from weather related affects. These included a water-proof enclosure with a glass window for the visible camera and a custom enclosure for the IR camera with a plastic filter cover which allowed the infrared spectrum to pass through.

A comprehensive field test plan was written and approved by the IDEA program. This test plan was carried out during the last visit to Amtrak's S&I Facility in Chicago and the results were reported. During these tests, the task of producing even illumination across the undercarriage was time consuming even with the wide-beam lights. In the future, an algorithm could be developed to assist in automatically evening out the illumination by providing real-time feedback from the unevenness in the image brightness. This setup will need to be executed only once for each system installation. The acquisition of the IR data was complicated by several manual calibration steps. This resulted in uneven quality of the data from day to day. An IR camera, made for operating in outdoor conditions, for the actual distances to components, and temperature variations expected to be encountered should be acquired.

Obtaining data for cars having incipient failures, damaged, missing, or foreign components was difficult. This was in part because we only imaged in-service cars during normal operations which as one would expect would contain very few components with these characteristics. It is recommended that future data collection be done at a repair facility so bad components can be imaged prior to replacement, and components on cars can be replaced with damaged or foreign components in areas of interest on specific car types.

The performance of the developed system on real world data captured from in-service trains is very promising. We were able to identify wheels, couplers, brake components, water tanks, a/c units, and traction motors from both the visible and IR video data. Due to insufficient knowledge of normal operating temperatures or lack of the availability of similar components across the car/train (for temperature comparisons), we were not able to conduct the entire breadth of inspections of each component as suggested by the panel. Therefore, several inspection examples used actual data while others required adding simulated data to the actual images captured.

FUTURE WORK

The results to date indicate that the machine vision undercarriage inspection system that we developed and tested is feasible. The research team believes the following steps should be undertaken to further enhance its capabilities and effectiveness. A wide-angle infrared camera that has a lens with a focal length similar to that of the visible spectrum camera, which is available from a few manufacturers, should be used. This would allow for one camera to cover the entire undercarriage and thus enable a more efficient comparison and integration with visible-range video recordings. Further, the current system has been tested in the relatively controlled environment of an inspection pit. If the system is to be used in places other than repair pits, then it would need to be weather-hardened.

There is also a need for a current template to be kept of the geometric and photographic appearance of each individual car in the fleet as it undergoes changes due to repairs. This could be achieved if, after each repair or modification, the car is rescanned by the system and its template updated.

This will also increase the confidence level in the detection of foreign objects.

The most important enhancement to the machine vision inspection system would be further automation of various aspects. Many of the manual processes should be replaced by automated procedures. The goal should be an inspection system that could be setup and configured to automatically work through all of the processes described and send inspection reports to appropriate personnel. Further automation enhancements should also include availability of feedback from the inspector to the system, as a source of continuous training of the algorithms so their performance improves continuously. Moreover, consideration should be given to integrating this undercarriage inspection system with the car end and side inspection system currently being developed.

POTENTIAL IMPACT AND PAYOFF FOR PRACTICE

It is vital to rail safety to ensure that critical mechanical components are in good working order at all times; however, the present system is highly inefficient for several reasons. Visual inspections by humans are inherently inefficient, as humans are not well suited for inspection tasks, i.e., vigilance tasks for low probability events. Because of the large number of cars and components to be inspected, it is difficult to satisfactorily inspect them all. Under these conditions the potential exists for certain defects to be missed, particularly difficult-to-detect items. Another problem is that the present system has no “memory”. Coupled with inspection regulations this means that the same components are being inspected repeatedly, at relatively short intervals, even if they were judged more than satisfactory in the previous inspection. Components have service lives far in excess of the typical inspection interval. Consequently, much of inspectors’ time is expended inspecting items that do not need inspection. Automating as many of the tasks as possible would enhance both efficiency and safety. In the parlance of one railroader this would enable many more of the “finders to become fixers” which is what actually affects safety. The system developed here is a critical step in the development of technology to achieve this goal.

As the technology matures, algorithms for more and more tasks that lend themselves to automated visual inspection can be developed. Another key element of the system envisioned is to integrate the data gathered into a transportation company’s information technology (IT) system. This will enable tracking of component wear and performance so that programmed maintenance can be optimized and fewer service disruptions occur due to a car being unexpectedly bad-ordered. Integration with the IT system may also allow improved understanding of components’ quality, interactions with the operating environment and the effect on service life. This in turn can provide insights into improved component design and practice.

PLANS FOR IMPLEMENTATION

During the execution of this research, members of the team made a number of presentations on the work at industry meetings and conferences. Several people and organizations expressed interest in further development and application of the technology. Amtrak was closely involved assisting us in conducting this work. Our senior contact there was Paul Steets who was very supportive of the project and was interested in seeing the technology further developed as he believed it would be useful for Amtrak. He indicated that they would be interested in further development of this technology in collaboration with UIUC if we continue this work. Another organization, KLD Labs is a supplier of machine vision and other inspection technologies to the railroad industry. They too were impressed with the results and potential applications of the technology described in this report and expressed interest in collaborating with us on its further development. Mike Iden of the Union Pacific (UP) Railroad was also interested in the possible benefits of the use of multi-spectral machine vision technology as a means of monitoring thermal condition and detecting anomalies in freight locomotive traction motors. He was especially interested in its potential use for detection of overheated pinion gears, a problem that they have recently experienced at UP.

INVESTIGATOR PROFILES

NARENDRA AHUJA

**Donald Biggar Willet Professor of Engineering
Department of Electrical and Computer Engineering
Beckman Institute University of Illinois at Urbana-Champaign**

Narendra Ahuja received his B.E. degree with honors in electronics engineering from the Birla Institute of Technology and Science, Pilani, India, in 1972. He earned his M.E. degree with distinction in electrical communication engineering from the Indian Institute of Science, Bangalore, India, in 1974, and his Ph.D. degree in computer science from the University of Maryland, College Park, USA, in 1979. From 1974 to 1975 he served as the Scientific Officer in the Department of Electronics, Government of India, New Delhi. From 1975 to 1979 he was at the Computer Vision Laboratory, University of Maryland, College Park. Since 1979 he has been with the University of Illinois at Urbana-Champaign where he is currently Donald Biggar Willet Professor in the Department of Electrical and Computer Engineering, Beckman Institute, and the Coordinated Science Laboratory. His research interests are in computer vision, robotics, image processing, image synthesis, sensors, and parallel algorithms. His work emphasizes integrated use of multiple image sources of scene information to construct 3-D descriptions of scenes; the use of integrated image analysis for realistic image synthesis; parallel architectures and algorithms and special sensors for computer vision; extraction and representation of spatial structure, e.g., in images and video; and use of the results of image analysis for a variety of applications including visual communication, image manipulation, information retrieval, robotics, and scene navigation.

CHRISTOPHER P.L. BARKAN, PH.D.

**Associate Professor - Department of Civil & Environmental Engineering
Director - Railroad Engineering Program
University of Illinois at Urbana-Champaign**

Christopher Barkan is an Associate Professor in the Department of Civil & Environmental Engineering and Director of the Railroad Engineering Program at the University of Illinois at Urbana-Champaign. He received his Bachelor's degree from Goddard College in 1977 and his M.S (1984) and Ph.D. (1987) degrees from the State University of New York at Albany where he conducted research on environmental applications of stochastic optimization models. He held a postdoctoral fellowship at the Smithsonian Environmental Research Center before joining the Association of American Railroads (AAR) Research and Test Department in their Washington, DC office in 1988. At the AAR he had principal responsibility for the railroad industry research program in risk, environmental and hazardous materials transportation safety until moving to the University of Illinois in 1998. Dr. Barkan's current research is focused on safety, energy efficiency, and risk analysis of railroad transportation systems. Current projects include safety and optimality analyses of railroad tank car damage resistance in accidents, risk factors affecting the probability of major railroad derailments, and risk to human health and the environment from hazardous materials shipped in tank cars. He is also collaborating on several projects developing machine-vision technology for application in the railroad industry. Dr. Barkan serves as the director of the AAR Affiliated Laboratory at the University of Illinois and in this role maintains frequent contact, coordination and collaboration with the railroad research staff at the Transportation Technology Center, Inc. in Pueblo, CO, the Safety & Operations staff at the AAR, and with research and engineering staff among North American railroads. Barkan also serves as Deputy Director of the RSI-AAR Railroad Tank Car Safety Research and Test Project, a long-term, cooperative effort of the North American railroad and tank car industries to improve railroad tank car safety.

PROJECT TEAM

In addition to the co-principal investigators, the research team consisted of Benjamin Freid (Railroad Specialist), Esther Resendiz (Machine Vision Developer), Sinisa Todorvic (Machine Vision

Developer), Nicholas Kocher (Imaging Specialist), Steven Sawadisavi (Imaging Specialist), and John M. Hart (Project Leader).

JOHN M. HART, M.S.

Research Engineer

**Beckman Institute for Advanced Science and Technology
University of Illinois at Urbana-Champaign**

John M. Hart is a Research Engineer in the Beckman Institute for Advanced Science and Technology and the Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign (UIUC). He received his Bachelor's degree in Electrical Engineering Technology at DeVry Institute of Technology, Chicago in 1984 with honors. During graduate degree course work at UIUC from 1985-87, he served as Head Teaching Assistant in the Advanced Digital Systems Laboratory. He then worked in industry as an Engineer for Frontier Engineering Inc. from 1987-91. He completed his Master's degree in Electrical and Computer Engineering from UIUC in 1992 where his research involved biologically inspired control of walking robots. Since then he has been a Research Engineer in the Computer Vision and Robotics Lab at the Beckman Institute developing new camera technologies and leading projects involving the application of machine vision in field research. He is also the Manager of R&D at Vision Technology Inc. (VTI) where he is involved in the development of advanced camera products based on technologies transferred from the university. He also is Principal Investigator for VTI on Small Business Innovative Research grants and has successfully lead two through their Phase II completion. His research interests include advanced camera technologies, machine vision wayside detection systems, walking robotics and prosthetics, and cybernetics.

GLOSSARY

Canny Edge Detector - The Canny edge detector is one of the most popular methods for edge detection. The output of the Canny detector is a black-and-white image of the same size as the original one, where white pixels denote the detected edges, i.e., the change in the image's intensity. It works by first Gaussian smoothing an image to eliminate noise, and then computing the 2-dimensional gradient (rate of intensity change) in the grayscale image. Then, post-processing is performed to quantize the gradient of the image and produce thin, 1-pixel width white curves against a black background to represent the edges.

Correlation – a method for estimating the quality of match between two images at each pixel.

Distance transform – a measure of difference between two images at each pixel.

Dimensionality reduction – the performance of many algorithms substantially downgrades if the number of features characterizing each data sample is large. To alleviate this problem the features of samples are analyzed so that a few most relevant features are selected. This procedure is called dimensionality reduction, since by pruning out irrelevant features the feature selection algorithm reduces the dimensionality of the feature space.

Feature space – each data sample is characterized by a number of features. These features can be interpreted as coordinates of a space in which the data represent points. This space is called the feature space.

Gaussian Mixture Model - The GMM is a commonly used statistical model in computer vision because it requires only a small amount of training data to estimate its parameters, and it is powerful enough to capture the underlying distributions of a wide variety of data. The GMM represents a weighted sum of Gaussians. Thus, the model parameters are the mean and variance of the Gaussians and their corresponding weighting coefficients. The GMM has multiple modes where the dominant mode reflects the Gaussian with the largest weight.

Gaussian smoothing – represents a filtering method, where the input image is filtered with a Gaussian filter to eliminate noisy changes in intensity across the image.

K-Means Clustering – Clustering, in general, is used for grouping samples of data into a number of clusters, where samples that belong to a cluster are more similar to one another than samples belonging to different clusters, where the measure of similarity used may vary. In particular, in the K-means clustering, the algorithm iteratively assigns each sample to one of the K clusters, by measuring the similarity of a sample with the each cluster's centroid.

Partial occlusion – objects in the scene may be positioned at different depths from the camera. If the objects also lie along the direction in which the camera captures the scene, then the objects closer to the camera partially or completely occlude those farther away. This phenomenon is called partial occlusion.

SAD – is abbreviation for the sum of absolute difference. Specifically, given two sets of data, $\{x_i\}$ and $\{y_i\}$, SAD represents the sum of terms $|x_i - y_i|$

Scrims - small metal screens that can be used in front of lights to lower the intensity of the lighting. Half scrims only cover half the light, while full scrims cover the entirety of the light. Double scrims cut out twice as much intensity as single scrims.