

The research described was performed by Transportation Technology Center, Inc., a wholly owned subsidiary of the Association of American Railroads.

Key Findings:

- TTCI demonstrated the ability of a commercially available machine learning system to successfully analyze a library of railcar component images captured by wayside camera installations.
- A trained convolutional neural network application was 97 percent successful in classifying seven different truck types.
- Additional convolutional neural network applications were 100 percent and 91 percent accurate in detecting the spring group locations and load conditions of the cars, respectively.

Machine Learning Approach for Railway Machine Vision Systems

Abe Meddah, Matt Witte, and Anish Poudel

TTCI collaborated with a software provider specializing in computer vision to evaluate the merits of using machine learning technology to visually inspect railcar components and conditions. Specifically, TTCI researchers used the deep learning platform of the provider software to create three convolutional neural network-based applications that were capable of identifying a railcar's truck type, detecting the location of a spring group in freight vehicle trucks, and measuring the compression of spring coils to determine load conditions. To train these applications to accomplish these tasks, TTCI researchers input a total of 1,276 images obtained from the machine vision systems.

This *Technology Digest* discusses the design, performance, and results obtained from testing these three applications. This primary research effort focused on using existing provider platform features that enhance ease of operation by streamlining data input, such as automated model building, while artificial intelligence (AI) assisted in the labeling of images. TTCI researchers pruned, selected, and inputted all image data for application building.

BACKGROUND

The prevailing method for inspecting railcars in revenue service is periodic manual visual inspection while a train is stopped¹; however, this methodology is highly inefficient for several reasons. First, skilled inspectors may look at a component numerous times without referencing a history of inspection or documenting their own inspection for subsequent use. Consequently, a failing component only gets noticed upon complete failure. Second, manual inspection does not require recording and maintaining equipment health trends. By removing human subjectivity, near real-time automated inspection of in-service railcars (achieved through machine learning approaches of visual analysis) mitigates these deficiencies. Machine vision inspection systems are being installed in North American railroad revenue service. Through the Association of American Railroads (AAR) Strategic Research Initiatives Program, TTCI has supported the current needs of the railroad industry by assisting suppliers and other stakeholders in the development of machine vision technologies and related algorithms for evaluating railcar components and conditions.

DEEP LEARNING PLATFORM

The commercial provider platform is a cloud-based software designed to automate visual inspection tasks. Researchers can train convolutional neural network-based AI models to perform various visually oriented tasks such as classification, detection, detection with bounding boxes, point detection, and optical character recognition. The system platform is designed to apply deep learning techniques to automatically train and deploy convolutional neural network models, including all aspects of image parameter selection, pre- and post-processing, data augmentation (i.e., increase of image diversity in training data), and model selection. Thus, while the platform allows subject matter experts to manually label certain objects and/or patterns in images, the convolutional neural network models can use a lesser number of images to train themselves; thereby negating further human involvement.

CONVOLUTIONAL NEURAL NETWORK

Convolutional neural networks are a type of artificial neural network that use machine learning algorithms to analyze visual imagery. Unlike typical artificial neural networks, which transform an input by filtering it through hidden layers of mathematical manipulation, convolutional neural networks are convolutional and organized by height, width, and depth. Thus, the final output from a convolutional neural network will be to a single vector of probability scores, organized along the depth dimension. Moreover, with convolutional neural networks the neurons in one layer do not connect to all the neurons in the next layer, but instead connect only to a small region of it. These differences make convolutional neural networks more powerful and effective than traditional artificial neural networks at recognizing, interpreting and categorizing large, unstructured data sets— particularly those comprised of visual imagery. For instance, a convolutional neural network can be trained to analyze multiple images of railcar couplings and determine whether a coupler pin is missing (see Figure 1).




Figure 1. Simplified convolutional neural network scheme

Three different applications using the convolutional neural network technology were developed: 1) Truck Type Classification,

2) Spring Group Detection, and 3) Spring Group Load Classification.

CONVOLUTIONAL NEURAL NETWORK TRAINING

A convolutional neural network must be trained to recognize and categorize data. This training is accomplished “by example” and requires two components: 1) a sampling of images representative of the type expected to be processed, and 2) a specified desired output for when the convolutional neural network recognizes — or does not recognize — certain objects in each of these images. Grayscale JPEG images with 2048×2048 resolution were deemed adequate for developing and testing the applications. These images were obtained from TTCI’s manual viewer database system, which contains machine vision data from vision systems on site at the Transportation Technology Center, Pueblo, CO, and from actual revenue service.²

Users must initially label several images as they are uploaded to the system platform. However, when enough images are labeled, system automated processes can take over and finish the task. Figure 2 shows an example of the system’s labeling interface.

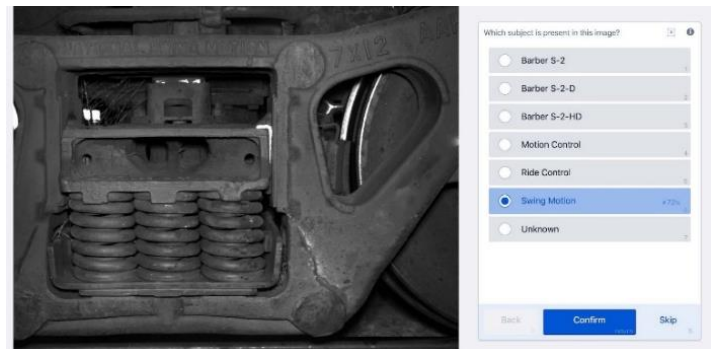


Figure 2. Labeling interface for Truck Type Application

For each application developed, the automated process randomly selected 75 percent of the images used as the training set. The remaining 25 percent of the images were withheld as a validation set and used later to evaluate for accuracy. Usually, successful convolutional neural network training necessitates the correct selection of many variables, generically referred to as “hyperparameters.” With the automated system, however, the hyperparameter selection processes are fully automated and do not require any specialized knowledge on the user’s part.

TRUCK TYPE CLASSIFICATION

Using 827 images from the database, the Truck Type Classification application was developed to identify the manufacturer and type of railcar truck present. To train the

convolutional neural network to sort out images and manually classify truck types, TTCI researchers relied solely upon the “Type of Truck” markings on truck sideframes. Although the “AAR Design Feature Code” and “Manufacturer’s Pattern Number Code” could have been used to determine truck types lacking the “Type of Truck” markings, doing so would have required extensive effort to create lookup tables and would not have been a good use of effort for this research and demonstration exercise. Seven classes of trucks were defined: Barber® S-2; Barber® S-2-D; Barber® S-2-HD; Motion Control®; Swing Motion®; Ride Control®; and “unknown.”

Figure 3a shows an example of a typical image of a truck where the class is clearly marked as a Barber S-2-HD (circled in yellow). Figure 3b shows an example of an “unknown” truck, where the class mark is either indiscernible or non-existent.




Figure 3. Example truck types: (a) Barber S-2-HD; (b) Unknown

With the Truck Type Classification application, the deep learning system was trained to analyse the content of the input image and correlate it to the single most representative truck class. Thus, for seven possible truck classes, there were seven potential output; each image in the training set had to be outputted as one, and only one, of these seven classes. For this reason, the “unknown” class became a catch-all; a sort of “none of the above” option relative to the other six.

Out of 827 images, the Truck Type Classification application achieved 97 percent accuracy. For the randomly selected 207 images in the validation set, 200 were correctly identified and 7 were misidentified; demonstrating 96.8 percent accuracy. Figure 4 shows this application’s accuracy for classifying each truck type.

To further validate accuracy, 265 additional, previously unseen images from TTCI’s machine vision database were inputted. This new group included a nearly even distribution of images containing the seven different truck classes. With this new group of images, only 11 of the 265 were misclassified, indicating roughly 96 percent accuracy, and making it consistent

with the 97 percent accuracy rate reported during model validation.

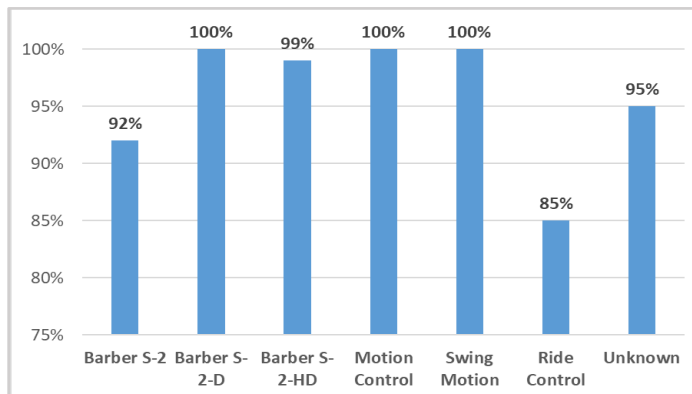


Figure 4. Application accuracy per truck class on validation set

SPRING GROUP DETECTION

TTCI developed a Spring Group Detection application using the “Box Detection” feature available in the platform. Using 28 randomly selected, grayscale JPEG images with 2048×2048 resolution, the convolutional neural network was trained to detect all the locations of a spring group. The application did so by providing an output bounding box with the pixel coordinates of each detected object. Training examples similarly annotated with bounding boxes to detect and locate the spring group in images were initially created and used by the application. Figure 5 shows an example of the feedback labeling interface.



Figure 5. Example of box detection (green box) technology

A total of 28 images were labeled and inputted into the application. The model was 100 percent accurate in detecting the spring group on the validation set.

SPRING GROUP LOAD CLASSIFICATION

An application to identify and classify Spring Group Load conditions was also built. Using 156 randomly selected grayscale, JPEG images with 2048×2048 resolution, the

convolutional neural network was trained to use its ability to measure distance to visually evaluate the level of compression on the springs and determine load condition. The outputs identified for classification of load condition relative to the level of spring compression were “Loaded,” “Partially Loaded,” and “Unloaded” (see Figure 6).




Figure 6. Three loading condition classes: loaded (left), partially loaded (center), and unloaded (right)

The model achieved 91 percent accuracy on the randomly selected validation set. Figure 7 shows the application accuracy breakdown for each individual load condition.

The fact that the accuracy rate is not higher is attributable to the application’s difficulty with reliably distinguishing between “Loaded” and “Partially Loaded” conditions. It is expected that this issue could be mitigated, and accuracy rates would improve by using more data samples and consistent labeling.

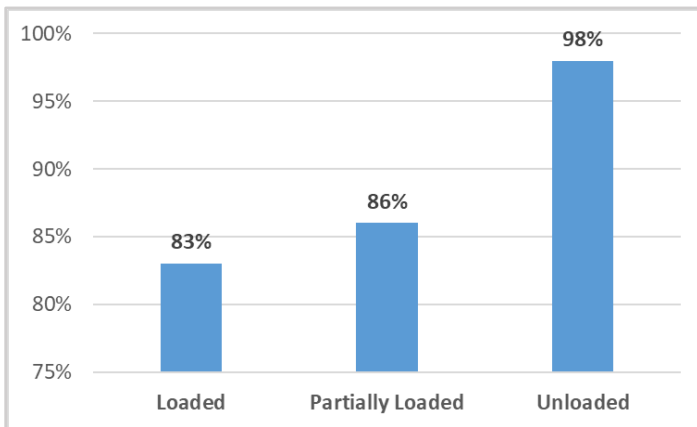


Figure 7. Application accuracy for each load condition

CONCLUSION

Using a software provider deep learning system platform, TTCI researchers were able to develop three different convolutional neural network-based applications that could identify railcar truck types, detect spring group locations, and measure coil spring compression from a visual analysis of an image to

determine load. A total of 1,276 images from TTCI’s machine vision database were used to test these applications and to evaluate the overall possibility of machine learning. Training these applications to identify and classify certain types of railcar hardware and conditions was achieved with relative ease, and the level of accuracy in the results was remarkable. The Truck Type Classification, Spring Group Detection, and Spring Group Load applications attained 97, 100, and 91 percent accuracy rates; respectively.

The testing and evaluation performed demonstrates the viability of using convolutional neural network-based, machine learning applications to detect, identify, and classify specific objects and patterns from images of railcar components. In addition, collaboratively creating and sharing uniform sets of railcar component/condition images by the industry stakeholders would help accelerate the development and deployment of more efficient AI-powered, convolutional neural network-based applications with higher rates of accuracy. This, in turn, would allow for reliable, real-time monitoring of in-service railcar components, as well as enhance the potential for maintenance planning. It should be noted that AAR subsidiary, Railinc, is leading industry efforts to build a machine vision image repository accessible to all stakeholders.

References

1. Witte, M. and Poudel, A. (2019). Review of wayside detection and monitoring technologies and their future for North American railroad applications. *Materials Evaluation*, Vol. 77(7), 885-896.
2. Witte, M., Meddah, A., and GeMeiner, B. (2017). “Accelerating machine vision inspection algorithm development.” *Technology Digest*, TD17-07. AAR/TTCI, Pueblo, CO.

For comments or questions about this publication, contact Abe_Meddah@aar.com

Disclaimer: Preliminary results in this document are disseminated by the AAR/TTCI for information purposes only and are given to, and are accepted by, the recipient at the recipient’s sole risk. The AAR/TTCI makes no representations or warranties, either expressed or implied, with respect to this document or its contents. The AAR/TTCI assumes no liability to anyone for special, collateral, exemplary, indirect, incidental, consequential or any other kind of damage resulting from the use or application of this document or its content. Any attempt to apply the information contained in this document is done at the recipient’s own risk. Unauthorized duplication or distribution is prohibited.