Towards Cargo Wagons Brake Health Scoring through Image Processing

Andres Felipe Posada-Moreno[1][a], Thomas Otte[2][b], Damir Pehar[3], Marc Haßler[3][c], Holger Bartels[3], Anas Abdelrazeq[2] and Frank Hees[2]

[1]*Institute for Data Science in Mechanical Engineering (DSME), RWTH Aachen University, Germany*

[2]*Institute for Information Management in Mechanical Engineering (IMA), RWTH Aachen University, Germany*

[3]*Deutsche Bahn (DB) Cargo AG, Frankfurt am Main, Germany*

Keywords: Rail Transport, Health Estimation, Neural Networks.

Abstract: The increase of integrated logistics is generating the progressive integration of rail transport systems on a global scale. This raises the challenge of the safe and compliant operation of an increasing number of assets. Within this context, inspection of in-service cargo wagons becomes increasingly important. Among the wagon components, the brake pads are essential and must be constantly inspected and timely changed before any failure. This publication presents a novel system for the automated scoring of cargo wagon brakes through image processing and deep learning algorithms. The main goal of this system is to provide insightful information which can improve the observability of assets, as well as enable augmented decision-making in maintenance inspection processes. Through this work, a four-step novel approach is described. First, an image acquisition system was developed. Then, an object detection model is used to extract the important cargo wagon components. Next, images containing the extracted brakes are analyzed to extract the most relevant keypoints of the brakes. Finally, the ratio between the distances of multiple keypoints is used to score each brake and provide insightful information regarding their health. After implementation, the proposed approach is tested and the resulting scores are explored.

1 INTRODUCTION

With an increasing demand for integrated logistics in a globalized world, rail transport networks progressively converge. This extends to the physical connectivity of railway infrastructure, the coherence and interoperability of existing systems, as well as normative requirements and operational standards. As a consequence, previously isolated railway systems, are being integrated, multiplying the complexity of rail transport by the number of the integrated assets, their maintenance, and the information flows related to their operations.

A global example of these dynamics can be observed with the growth of the International Union of Railways (UIC), which, to date, is composed of 204 member organizations from all around the globe (Union Internationale des Chemins de fer, 2021). As part of their current technical priorities, the UIC has listed the usage of digital technologies and the leveraging of data to increase large systems' observability (Union Internationale des Chemins de fer, 2020).

At the same time, global biosecurity issues have highlighted the need for resilient operations and maintenance processes, which can adjust to limited personnel mobility situations. In this context, digitalization and the usage of artificial intelligence can contribute to more robust asset management in uncertain conditions. Similarly, the UIC has included the topics of "asset management and predictive maintenance" and "railway automation through artificial intelligence and robotics" as part of their agenda (Union Internationale des Chemins de fer, 2020).

Within this context, the German rail freight transport company DB Cargo AG, operates a fleet composed of over 2500 locomotives and 90000 cargo wagons (DB Cargo AG, 2021). Adding to asset diversity, this fleet contains more than 190 different cargo wagon types, which are operated, inspected, and maintained continuously. As a result, opportu-

[a] https://orcid.org/0000-0003-3751-0680
[b] https://orcid.org/0000-0002-4227-8938
[c] https://orcid.org/0000-0002-1545-1416

Posada-Moreno, A., Otte, T., Pehar, D., Haßler, M., Bartels, H., Abdelrazeq, A. and Hees, F.
Towards Cargo Wagons Brake Health Scoring through Image Processing.
DOI: 10.5220/0010824300003122
In *Proceedings of the 11th International Conference on Pattern Recognition Applications and Methods (ICPRAM 2022)*, pages 423-431
ISBN: 978-989-758-549-4; ISSN: 2184-4313

nities for digitalization not only contribute to the general objectives of the UIC, but are also directly translated into operational advantages of UIC members such as large rail transport companies.

Thus, the need arises for maintaining compliance with the General Contract of Use for Wagons from the UIC and the GCU bureau (GCU Bureau, 2021) at scale, and with more robust processes. To improve the inspection and maintenance processes required by the mentioned contract, railway companies have invested in digitalization technologies. More specifically, companies and researchers have been working on using state of the art artificial intelligence techniques designed to assess the individual integrity of the components of cargo wagons (Zhang et al., 2018; Rocha et al., 2018; Liu et al., 2016; Sun et al., 2017; Otte et al., 2020), railways (Zhang et al., 2021a; Rong et al., 2016; Shang et al., 2018) and other related infrastructure (Zhang et al., 2021b; Li et al., 2014; Peng et al., 2020). Through this goal, researchers and railway companies have joined hands in the development of multiple applications.

In recent years, multiple publications have introduced inspection applications for railway infrastructure (Zhang et al., 2021a; Rong et al., 2016; Shang et al., 2018), as well as in-service inspection systems for high-priority components in cargo wagons (Zhang et al., 2018; Rocha et al., 2018; Liu et al., 2016; Sun et al., 2017). These systems have explored classical techniques such as template matching (Ningning et al., 2016), and deep neural networks based techniques for recognition and inspection (Rocha et al., 2018) showing different levels of performance in controlled conditions and real-world applications.

The major challenges observed through the work of previous researchers can be summarized into two main topics. First, the usage of highly noisy data resulting from real-world conditions affects any trained algorithm. This concern arises from light conditions in varied environmental scenarios, asset degradation, and artifacts created through data acquisition processes. Second, the limitations of available sensors for post-hoc solutions make data acquisition cumbersome or impractical in real-world settings. This is the inherent practical restriction of having to implement a scalable inspection system, which adjusts and does not modify the normal operations of the cargo railway transport (post-hoc). This requirement is often the main concern on the implementation of new inspection systems, as any modification of existing infrastructure entails compliance risks as well as expensive downtime of operations.

In this context, the current work proposes an automatic system for cargo wagon brake scoring, which enables augmented decision-making processes regarding their health and the need for maintenance. To the best of the author's knowledge, the proposed system is a novel approach for the health scoring of cargo wagon brakes in the scientific and industrial community. The proposed system consists of four steps; (a) an image acquisition system; (b) an object detection model; (c) a keypoint extraction model; and (d) a brake scoring module.

The proposed approach and the necessary context will be presented as follows. First, other approaches related to cargo wagon and railway inspection will be introduced. Then, the proposed system and each one of its components will be described. Next, the results of the implementation of the system will be reported. Finally, the conclusions of the obtained results and the planned future work will be discussed.

2 RELATED WORK

Computer vision systems for process support in railway transport have been a focus of research for the last decade. From the various applications that have been previously explored, four main topics are highlighted. First, the inspection of critical infrastructure during their operations (Zhang et al., 2021a). Second, the identification of assets, based on the recognition of their wagon number (UIC ID) (Xiucai and Gongli, 2020). Third, the inspection of the cargo contained in cargo wagons through x-rays and deep learning (Rogers et al., 2017). Last, the inspection and identification of faulty components in cargo wagons as a mean to assist maintenance processes (Zhang et al., 2021b). Although this work leans heavily towards the last topic, other applications and approaches are highly related to sensors and models that can be used in the context of railway maintenance.

The inspection of critical infrastructure has focused on the detection of defects on rails. More specifically, rail surface inspection has been proposed through the usage of classical computer vision methods such as SVN-HOG (Gavai et al., 2019), as well as multiple types of neural networks for object detection and classification (Zhang et al., 2021a; Shang et al., 2018). Similarly, methods have been explored for the detection of related components such as tie plates (Li et al., 2014), hexagonal nuts in railway fasteners (Peng et al., 2020), and bondwires (Gavai et al., 2019). In these cases, a custom inspection vehicle has been used to acquire the rail images (Li et al., 2014).

On the subfield of cargo inspection, researchers have explored the usage of x-ray based image acquisition systems for inspection of illegal loads as well as cargo manifest verification. The first tackled issue is the acquisition of the x-ray images and the detection of the transported goods inside containers (Rogers et al., 2015). Afterward, multiple deep learning approaches have been developed for the detection of specific contraband objects (e.g. concealed cars) (Jaccard et al., 2016). One major issue related to these techniques is the lack of available data. In these cases, researchers have proposed the artificial creation of datasets based on mock-ups (Kolokytha et al., 2018) or synthetic generation of images based on CAD files (Visser et al., 2016). In consequence, it has also been proposed to verify transport manifests through the usage of the same detection techniques (Tuszynski et al., 2013).

Identification is one of the most important steps in the inspection process of an asset, either in-service or through the maintenance process. In this regard, researchers have used the required UIC ID of the cargo wagons as a direct identifier. Early techniques used image morphological operations and template matching to isolate and extract the wagon ID (Ningning et al., 2016). Afterward, hybrid approaches exploited the visual characteristics of white paint by extracting regions through MSER before using an OCR algorithm to interpret the painted characters (Xiang et al., 2016). The latest methods propose the usage of deep learning techniques to directly extract and interpret the UIC ID characters from side views of the cargo wagons (Xiucai and Gongli, 2020; Liu et al., 2019).

The inspection of cargo wagons has been explored through the analysis of each of their critical components. Initial approaches using template matching focused on the inspection of locomotive pantographs (Hamey et al., 2007). Similarly, approaches using other types of descriptors such as GLCM or HOG, combined with SVN for the classification of boogey block keys (Zhou et al., 2018) and locomotive speed sensors (Li et al., 2019) have been found. Aside from classical approaches, deep learning algorithms have proven to be the preferred approach for the inspection of single components such as bogie pads (Rocha et al., 2018) and valves (Pahwa et al., 2019). Similarly, frameworks for the detection and classification of multiple structural components have also been proposed as a unified approach for this challenge (Zhang et al., 2018; Posada Moreno et al., 2020).

Within the acquisition systems proposed in the literature, the most common ones are the stationary CCD cameras in railway stations (Hamey et al., 2007), the lateral high-speed cameras alongside railways (Zhang et al., 2020), and a line sensor camera bridges (Posada Moreno et al., 2020). Acquisition systems have shown illumination pollution and perspective variances. In perspective, the CCD approach has the limitations of working with slowly approaching locomotives. The high-speed cameras present a single lower perspective of the wagons, which limits the analysis to the inferior parts of the vehicles. Finally, the camera bridges captures a good perspective of the wagons, but also lead to a lateral elongation caused by their linear scanning cameras.

The proposed approach builds on previous paradigms to tackle the issue of brake inspection. Here, the brakes are not only detected for binary classification, but are scored concerning their deterioration state.

3 METHODS

The present work tackles the problem of how to enable, augmented decision-making on in-service cargo wagon brake inspection based on image data acquisition and artificial intelligence. Thus, the approach will cover a brief introduction to the hardware used in the data acquisition as well as the multiple steps of the data processing. In this regard, the proposed system can be described in four steps, (a) a multi-camera system that captures real-time images of cargo wagons, (b) an object detection system for extracting regions of interest and structural components from the acquired images, (c) a keypoint extraction model, for the localization of relevant component landmarks, (d) a scoring module, for the scoring of the health of brakes through a geometric relation.

3.1 Image Acquisition Pipeline

The image acquisition system is based on a bridge-like metal structure (Figure. 1) which serves as a support for a series of six line-scan cameras and six LEDs for the illumination of the arriving wagons. Furthermore, a laser scanner (Sick LMS111) allows the detection of arriving wagons as well as their incoming velocity. Two of the installed line-scan cameras (Basler racer raL2048-48gm) are responsible for acquiring images of both sides of the lower section of the cargo wagons.

Once a wagon approaches the camera bridge, the laser scanner detects their arrival and their velocity, automatically adjusting the acquisition rate of the line-scan cameras as well as the illumination intensity of the available LEDs. Using the velocity measurement and the individual scans, an image is composed by horizontally stacking the acquired one pixel wide scans. Thus, the obtained images have a fixed height of 2048 pixels and a variable length of between 10K

Figure 1: Image acquisition system.

and 50K pixels depending on the type of wagon that is being scanned and possible variations in speed. After the acquisition process, the images are temporally stored locally before uploading them to a cloud object storage. A partial example is shown in Figure. 2.

This type of setup allows the real-time acquisition of cargo wagon images, from a constant perspective, and without disturbing nor interrupting the daily operations of the railway system. Moreover, the implementation of this system does not require major structural changes in the existing infrastructure.

3.2 Object Detection Module

The object detection module receives images of the lower section of the wagons acquired by the camera bridge system and extracts regions of interest containing previously selected objects. The detected objects include brakes, axis, hooks, bumpers, wheels, and UIC-ID numbers. This step is shared between the cargo wagon identification and multiple health estimation analysis. Nonetheless, the object detection is performed at lower resolutions than each health analysis in order to reduce the computational requirements of the pipeline. The object detection of the current work focuses on the extraction of the four brakes present in each image.

Before performing the analysis, the images were normalized and resized to a constant width of 2400. Then a Faster RCNN (Ren et al., 2015) using as a backbone a ResNet-101 (He et al., 2016) is used to perform the object detection. This combination allows the extraction of smaller objects without compromising on performance. An example of the detected brakes can be seen in Figure. 2.

To compensate for the changing illumination conditions and the lack of faulty or degraded brakes, the process of training the Faster RCNN relied on data augmentations. Specifically, random scale changes, horizontal flips, and photometric distortions of brightness, contrast, saturation, and hue were added.

3.3 Keypoint Detection Module

The keypoint detection module aims to extract landmark locations of the brakes, which are directly related to their usage and degradation. Having into consideration the horizontal stretching and compression generated by the data acquisition system, the extracted landmarks must suffice to indirectly compute the width of the brake pads, which diminishes with their remaining useful life (RUL). Thus, for each brake pad, six keypoints are detected, consisting of the upper, middle, and lower points of the inner and outer sides of the pad. Each brake contains two pads, an upper pad and a lower pad, which are denoted as A and B respectively. Similarly, the axis of the brake is also extracted as a keypoint.

The defined keypoints are structural points that exist in all the brakes of the considered wagon type. However, in different types of cargo wagons, pad A can be partially or completely occluded (Figure. 3) by other structural components such as the leaf springs. Thus, not all the brake images will contain a full set of keypoints for the brakes. The selected model for the keypoint detection is a variant of the Faster RCNN with a ResNet-101 mentioned before. While using the same model architecture as before, the bounding box regressor is modified to compute each keypoint. Similarly, each brake is extracted from the original (resolution) image and is normalized and the patch is resized to a height of 1333 and a width of 800 to fully utilize the available hardware. During training, the used data augmentations are the same as the ones mentioned in the previous step.

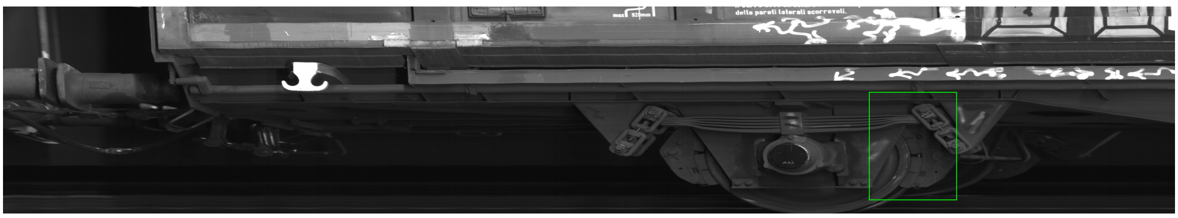

Figure 2: Section of cargo wagon image. The green square contains a single detected brake.

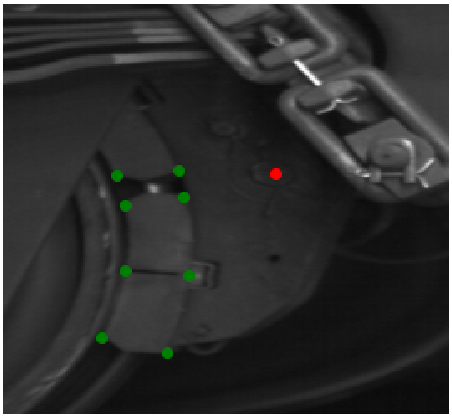

Figure 3: Keypoints of partially occluded brake. Brake pads border keypoints in green. Brake axis keypoint in red.

3.4 Brake Scoring Module

The brake score module has the aim of computing a score which is directly related to the RUL of the brake pads. The pads slowly get eroded while being used, reducing the effective width of the material which is in direct contact with the spinning wagon wheels.

While using scale invariant images, this score could be calculated as the distance in pixels between the outer and inner borders of the pads. Consequently, the scale could be used to transform this measure into metric units for better understanding. With the current setup, image stretching must be compensated. As the changes in velocity during the scan result in variations of scale along the horizontal axis, which are not constant at a global scale, but are stable locally. Thus, a reliable proxy measurement can be achieved, by computing the ratio between the width of the brake pad and the distance from the center of the pad to the axis of the brake. As the scale changes for different images, the distances between the points will change proportionally given Thales's theorem, ensuring that the ratio is maintained. Specifically, the ratio score $R_{B,mid}$ of the mid-section of the B pad is denoted in Equation 1. The numerator indicates the width of the brake B in the mid-section, and the denominator indicates the distance between the center of the brake B and the axis.

$$R_{B,mid} = \frac{||P_{B,mid,out} - P_{B,mid,in}||_2}{||\frac{P_{B,mid,out}+P_{B,mid,in}}{2} - P_{axis}||_2} \quad (1)$$

3.5 Implementation and Testing

The proposed system was implemented in two main components. The first one is the image acquisition pipeline, which consists of the described hardware and communication systems for the acquisition of the cargo wagon images and their transmission to an object storage.

The second component is the joint pipeline of object detection, keypoint detection, and brake scoring. To obtain a scalable system, both deep learning models were implemented in Pytorch. The complete analysis was implemented as a five-step task, which (a) downloads an image; (b) detects the objects of interest and filters the brake images; (c) computes the brake pad scores for each brake instance; (d) sends the results of the analysis to a specified webhook. To enable parallel execution, the Python library Ray was used and exposed through a REST API.

Subsequently, the wagon type Hbbillns 305 was selected for the training and validation of the proposed approach. This reference is one of the most prevalent in the analyzed fleet, containing two axes per wagon and two brakes per axis, which only have minor structural variations. Through the operation of the camera bridge, a constant supply of cargo wagon images was achieved. Over 2000 of the acquired images were used for training, and testing of this approach. From these images, 425 were used for the training and testing of the object detection algorithm. Next, 480 of the labeled brakes were extracted and labeled for the training and testing of the keypoint detection algorithm. Once both algorithms were trained, 1642 images were used for validating the complete pipeline.

4 RESULTS

4.1 Image Acquisition Pipeline

During the implementation of the current concept, two camera bridge systems were built in Germany. Each camera bridge was installed at a marshalling yard, with the objective of scanning the cargo wagons during their reorganization process. On average, a single image of the lower sections of a cargo wagon has a volume of 5 megabytes. Usually, between 1500 and 2000 images of said type are acquired daily per camera bridge, depending on their location and traffic. The implementation of the camera bridge systems also contains an RFID reader, which aids in the identification of over half of the cargo wagons which contain RFID tags. After the acquisition, the images are composed locally after being transferred to a cloud object storage periodically.

4.2 Object Detection Module

To train the Faster RCNN object detection algorithm, a set of 425 lower images of cargo wagons type Hbbillns 305 were labeled. Subsequently, these images were split into 340 images used for training and 85 images used for testing. An available pre-trained model on ImageNet (Deng et al., 2009) (from pytorch zoom) was used, and all layers were trained for 10 epochs. The model was implemented using the Pytorch framework and trained in an Ubuntu server with an Intel Xeon Silver 4116 CPU and a Tesla V100 GPU. Through this process, a mean Average Precision mAP75 of 0.84 and a mAP50 of 0.973 were achieved.

4.3 Keypoint Detection Module

From the previous step, a total of 480 brakes were extracted. These images had varying dimensions of roughly 800x1200 pixels. To label the previously shown keypoints, the images had to be equalized, as the human eye has difficulties with low contrast images, which still contained brakes. Then, the images were split into 384 images used for training and 96 images used for testing. Similar hardware and a pretrained backbone were used during a 10 epochs training. As a result, an mAP75 of 0.917 and a mAP50 of 0.968 were achieved. Additionally, the detected keypoints had a vertical and horizontal Mean Average Error of 4.33 pixels in comparison with the annotations. An example of the detected keypoints can be seen in Figure. 4.

4.4 Brake Scoring Module

After the training of the previous models, a new set of 1642 lower images from cargo wagons type Hbbillns 305 were analyzed. In these images, a total of 5911 brakes were extracted. The keypoints of the images were detected and the mid-ratio of the brake pad B was computed. As seen in the histogram of Figure. 5, an average ratio of 0.292 was observed. It can be observed that the distribution of brakes follows a skewed normal distribution. Most of the brakes which are completely new, or recently installed, will have a ratio in-between 0.35 and 0.45. Similarly, brakes with a ratio between 0.15 and 0.35, still have a long RUL. Finally, the brakes with ratios lower than 0.15, are soon to require a change. These observations were performed with a domain expert from DB Cargo AG. It must be clarified, that no brake in the dataset approached critical levels, as all are changed with spare time. The direct detection of their health score will enable better planning and observability of the maintenance process, which will allow better scheduling of the required resources.

The distribution shown in Figure. 5, allows a more comprehensive understanding of the state of the brake pads, in contrast with the single measures of brake width which are currently distorted by the scale variance of the images. This scale variance can be observed in the scatter plot from Figure. 6, in which each analyzed brake is a point. The horizontal axis is the distance between the center of the brake to the axis, and the vertical axis is the width of the brake pad, also called delta.

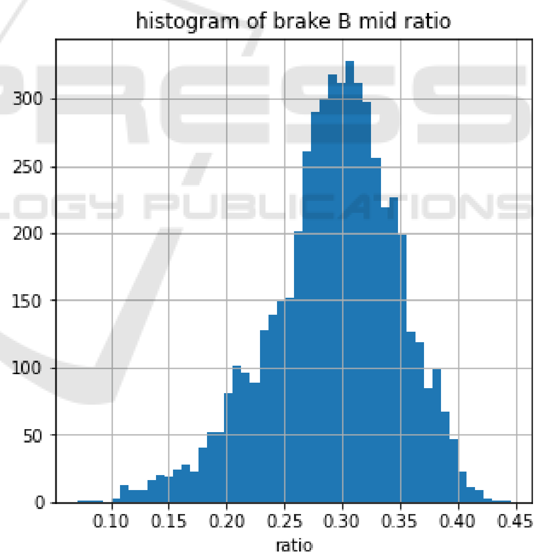

Figure 5: Histogram of (scores) ratios from the mid-section of the brake B.

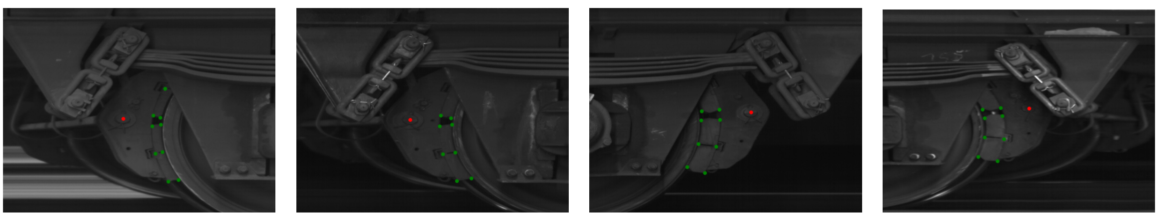

(a) Brake with score 0.15. (b) Brake with score 0.24. (c) Brake with score 0.33. (d) Brake with score 0.42.

Figure 4: Result of keypoint detection on cargo wagon brake and scored brake pads.

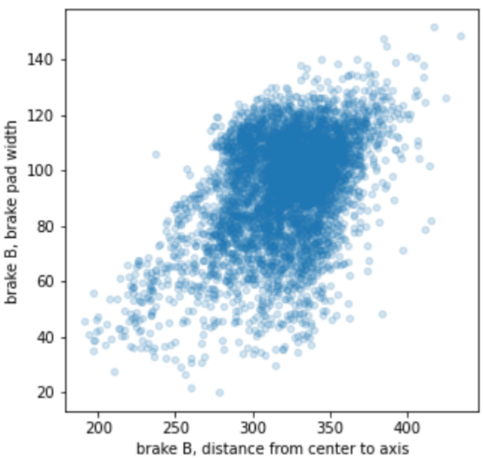

Figure 6: Scatter plot of distance from the center of a brake to the axis vs the width of a brake pad.

Interpreting the distance between the center of the brake B and the axis, it can be noticed that the scale difference between two images can be up to a factor of two. This representation shows how images, which have different scales, can be processed to obtain a stable ratio of the state of the brake pads. Example images of different brake scores (ratios) are shown in Figure. 4. In these results, brakes with a lower ratio have been worn out and should be scheduled for maintenance. Similarly, brakes with a higher ratio have a longer RUL.

5 CONCLUSION

In the current work, a novel system for cargo wagon brake scoring is presented. This system aims to automatically extract relevant information from in-service cargo wagons in order to enable better decision-making with respect to their maintenance requirements. It tackles the challenge of computing a score proportional to the RUL of a brake instead of implementing binary classification in search of defects. The novel system is robust to scale variant image acquisition techniques and allows the processing of cargo wagons under varied illumination conditions.

The proposed system is composed of four main steps, (a) a camera bridge for image acquisition, (b) an object detection algorithm for the extraction of regions of interest, (c) a keypoint detection algorithm for the extraction of landmarks in a brake, (d) a brake scoring module for the estimation of the degradation of brakes.

The proposed system was described, implemented, and tested. The object detection system achieved an mAP50 of 0.973 and the keypoint detection system obtained an mAP50 of 0.968. The scoring module was then tested with 1642 new lower images of cargo wagons. The exploration of the results shows that a mean ratio of 0.292 indicates healthy brake pads, and ratios close to 0.1 are related to worn-out brake pads. The results of the implementation show the viability of the proposed approach.

Future work will focus on five different topics. First, although the current selected models show an outstanding performance, migration to single shot detection model can reduce significantly the required computational resources. Second, explore the transferability of this approach to other wagon types within the cargo fleet. Third, developing methods for the compensation of the scale variance of the images, in order to achieve metric measures of the state of the brake pads. Fourth, exploring other approaches to merge multiple data sources in order to relate the obtained scores with the RUL of the brake (time remaining before requiring maintenance). Fifth, exploring other unsupervised or semi-supervised approaches for the discovery, extraction, and labeling of objects and landmarks in large cargo wagon databases. Finally, by taking into account the RUL and the expected deterioration of each trajectory, the obtained data can be used to better plan train schedules to include prescriptive maintenance.

ACKNOWLEDGEMENTS

Research supported by the Federal Ministry of Transport and Digital Infrastructure of Germany (BMVI) – Project QUISS (19F2060).

REFERENCES

DB Cargo AG (2021). Your Logistics Partner.

Deng, J., Dong, W., Socher, R., Li, L., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, pages 248–255. IEEE Computer Society.

Gavai, G., Eldardiry, H., Wu, W., Xu, B., Komatsu, Y., and Makino, S. (2019). Hybrid image-based defect detection for railroad maintenance. *Electronic Imaging*, 2019(9):360–1.

GCU Bureau (2021). General contract of use for wagons general contract of use for wagons GCU. Technical report.

Hamey, L. G. C., Watkins, T., and Yen, S. W. T. (2007). Pancam: In-service inspection of locomotive pantographs. In *Proceedings of the International Con-*

ference on Digital Image Computing: Techniques and Applications, DICTA 2007, 3-5 December 2007, Adelaide, Australia, pages 493–499. IEEE Computer Society.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society.

Jaccard, N., Rogers, T. W., Morton, E. J., and Griffin, L. D. (2016). Detection of concealed cars in complex cargo x-ray imagery using deep learning. *CoRR*, abs/1606.08078.

Kolokytha, S., Flisch, A., Lüthi, T., Plamondon, M., Visser, W., Schwaninger, A., Hardmeier, D., Costin, M., Vienne, C., Sukowski, F., Hassler, U., Dorion, I., Gadi, N., Maitrejean, S., Marciano, A., Canonica, A., Rochat, E., Koomen, G., and Slegt, M. (2018). Creating a reference database of cargo inspection x-ray images using high energy radiographs of cargo mock-ups. *Multim. Tools Appl.*, 77(8):9379–9391.

Li, B., Zhou, S., Cheng, L., Zhu, R., Hu, T., Anjum, A., He, Z., and Zou, Y. (2019). A cascade learning approach for automated detection of locomotive speed sensor using imbalanced data in ITS. *IEEE Access*, 7:90851–90862.

Li, Y., Trinh, H., Haas, N., Otto, C., and Pankanti, S. (2014). Rail component detection, optimization, and assessment for automatic rail track inspection. *IEEE Trans. Intell. Transp. Syst.*, 15(2):760–770.

Liu, L., Zhou, F., and He, Y. (2016). Automated visual inspection system for bogie block key under complex freight train environment. *IEEE Trans. Instrum. Meas.*, 65(1):2–14.

Liu, Z., Wang, Z., and Xing, Y. (2019). Wagon number recognition based on the yolov3 detector. In *2019 IEEE 2nd International Conference on Computer and Communication Engineering Technology (CCET)*, pages 159–163. IEEE.

Ningning, L., Yuehui, M., Wang, Y., and Wang, N. (2016). An algorithm of freight train number locating based on template matching and morphological operations. *Metallurgical and mining industry*, (2):102–107.

Otte, T., Bartels, H., Posada Moreno, A. F., Wittenburg, G., and Haßler, M. (2020). Holistic data infrastructure and analytics system for rail freight transport. In *14th ITS European Congress*. 14th ITS European Congress via 1st Virtual ITS European Congress, online, 9 Nov 2020 - 10 Nov 2020.

Pahwa, R. S., Chandrasekhar, V. R., Chao, J., Paul, J., Li, Y., Nwe, M. T. L., Xie, S., James, A., Ambikapathi, A., and Zeng, Z. (2019). Faultnet: Faulty rail-valves detection using deep learning and computer vision. In *2019 IEEE Intelligent Transportation Systems Conference, ITSC 2019, Auckland, New Zealand, October 27-30, 2019*, pages 559–566. IEEE.

Peng, Z., Wang, C., Ma, Z., and Liu, H. (2020). A multifeature hierarchical locating algorithm for hexagon nut of railway fasteners. *IEEE Trans. Instrum. Meas.*, 69(3):693–699.

Posada Moreno, A., Klein, C., Haßler, M., Pehar, D., Solvay, A., and Kohlschein, C. (2020). Cargo wagon structural health estimation using computer vision. *8th Transport Research Arena, TRA2020*, pages 04–27.

Ren, S., He, K., Girshick, R. B., and Sun, J. (2015). Faster R-CNN: towards real-time object detection with region proposal networks. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R., editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 91–99.

Rocha, R. L., Siravenha, A. C. Q., Gomes, A. C. S., Serejo, G. L., Silva, A. F. B., Rodrigues, L. M., Braga, J., Dias, G., Carvalho, S. R., and de Souza, C. R. B. (2018). A deep-learning-based approach for automated wagon component inspection. In Haddad, H. M., Wainwright, R. L., and Chbeir, R., editors, *Proceedings of the 33rd Annual ACM Symposium on Applied Computing, SAC 2018, Pau, France, April 09-13, 2018*, pages 276–283. ACM.

Rogers, T. W., Jaccard, N., and Griffin, L. D. (2017). A deep learning framework for the automated inspection of complex dual-energy x-ray cargo imagery. In *Anomaly Detection and Imaging with X-Rays (ADIX) II*, volume 10187, page 101870L. International Society for Optics and Photonics.

Rogers, T. W., Jaccard, N., Morton, E., and Griffin, L. (2015). Detection of cargo container loads from x-ray images. In *2nd IET International Conference on Intelligent Signal Processing 2015 (ISP)*, pages 1–6. IET.

Rong, J., Song, S., Dang, Z., Shi, H., and Cao, Y. (2016). Rail track irregularity detection method based on computer vision and gesture analysis. *Int. J. Online Eng.*, 12(12):55–59.

Shang, L., Yang, Q., Wang, J., Li, S., and Lei, W. (2018). Detection of rail surface defects based on cnn image recognition and classification. In *2018 20th International Conference on Advanced Communication Technology (ICACT)*, pages 45–51. IEEE.

Sun, X., Ding, J., Chiara, G. D., Cheah, L., and Cheung, N. (2017). A generic framework for monitoring local freight traffic movements using computer vision-based techniques. In *5th IEEE International Conference on Models and Technologies for Intelligent Transportation Systems, MT-ITS 2017, Naples, Italy, June 26-28, 2017*, pages 63–68. IEEE.

Tuszynski, J., Briggs, J. T., and Kaufhold, J. (2013). A method for automatic manifest verification of container cargo using radiography images. *Journal of Transportation Security*, 6(4):339–356.

Union Internationale des Chemins de fer (2020). uic work programme 2020-2022. Technical report.

Union Internationale des Chemins de fer (2021). Vademecum (List of UIC members).

Visser, W., Schwaninger, A., Hardmeier, D., Flisch, A., Costin, M., Vienne, C., Sukowski, F., Hassler, U., Dorion, I., Marciano, A., Koomen, G., Slegt, M., and Canonica, A. C. (2016). Automated comparison of x-ray images for cargo scanning. In *IEEE International*

Carnahan Conference on Security Technology, ICCST 2016, Orlando, FL, USA, October 24-27, 2016, pages 1–8. IEEE.

Xiang, X., Yang, F., Wang, M., Bao, W., and Sheng, Y. (2016). Id localization and recognition for railway oil tank wagon in the industrial scene. In *2016 12th World Congress on Intelligent Control and Automation (WCICA)*, pages 826–829. IEEE.

Xiucai, G. and Gongli, C. (2020). Railway freight train number detection method based on deep learning. In *2020 IEEE 9th Joint International Information Technology and Artificial Intelligence Conference (ITAIC)*, volume 9, pages 927–930. IEEE.

Zhang, D., Song, K., Wang, Q., He, Y., Wen, X., and Yan, Y. (2021a). Two deep learning networks for rail surface defect inspection of limited samples with line-level label. *IEEE Trans. Ind. Informatics*, 17(10):6731–6741.

Zhang, Y., Lin, K., Zhang, H., Guo, Y., and Sun, G. (2018). A unified framework for fault detection of freight train images under complex environment. In *2018 IEEE International Conference on Image Processing, ICIP 2018, Athens, Greece, October 7-10, 2018*, pages 1348–1352. IEEE.

Zhang, Y., Liu, M., Chen, Y., Zhang, H., and Guo, Y. (2020). Real-time vision-based system of fault detection for freight trains. *IEEE Trans. Instrum. Meas.*, 69(7):5274–5284.

Zhang, Y., Liu, M., Yang, Y., Guo, Y., and Zhang, H. (2021b). A unified light framework for real-time fault detection of freight train images. *IEEE Trans. Ind. Informatics*, 17(11):7423–7432.

Zhou, F., Song, Y., Liu, L., and Zheng, D. (2018). Automated visual inspection of target parts for train safety based on deep learning. *IET Intelligent Transport Systems*, 12(6):550–555.