



This is a repository copy of *A machine learning approach for real-time wheel-rail interface friction estimation*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/202330/>

Version: Accepted Version

---

**Article:**

Folorunso, M.O., Watson, M., Martin, A. et al. (3 more authors) (2023) A machine learning approach for real-time wheel-rail interface friction estimation. *Journal of Tribology*, 145 (9). 091102. ISSN 0742-4787

<https://doi.org/10.1115/1.4062373>

---

© 2023 The Authors. Except as otherwise noted, this author-accepted version of an article published in *Journal of Tribology* is made available via the University of Sheffield Research Publications and Copyright Policy under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>

**Reuse**

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# A Machine Learning Approach for Real Time Wheel/Rail Interface Friction Estimation

Morinoye O. FOLORUNSO<sup>1</sup>, Michael WATSON<sup>1</sup>, Alan MARTIN<sup>1</sup>, Jacob W. WHITTLE<sup>1</sup>, Graham SUTHERLAND<sup>2</sup>, Roger LEWIS<sup>1</sup>

<sup>1</sup>The University of Sheffield, Sheffield, United Kingdom

<sup>2</sup>Consulting Canetia Analytics Inc., United States of America

Corresponding Author's Email address: mofolorunso1@sheffield.ac.uk

## Abstract

Predicting friction at the wheel rail interface is a key problem in the rail industry. Current forecasts give regional level predictions, however, it is well known that friction conditions can change dramatically over a few hundred metres. In this study we aimed to produce a proof-of-concept friction prediction tool which could be used on trains to give an indication of the limiting friction present at a precise location. To this end field data including temperature, humidity, friction and images were collected. These were used to fit a statistical model including effects of local environmental conditions, surroundings and railhead state. The model predicted the friction well with an  $R^2$  of 0.97, falling to 0.96 for naive models in cross validation. With images and environmental data collected on a train a real time friction measurement would be possible.

Keywords: low-adhesion, wheel-rail interface, friction prediction, machine learning

## 1. Introduction

Rail travel is the most fuel and carbon efficient mode of transport and is a critical part of the strategy to achieve net zero carbon emissions in Europe [1]. A large barrier to its wider adoption is delays and seasonal restrictions caused by low adhesion conditions. It is estimated that these problems cost the rail industry £350 million each year [2]. In addition, delays, largely caused by low adhesion, are the highest cause of customer dissatisfaction with rail services [3].

Understanding the wheel rail contact is central to fixing this and many other problems, as such it has been modelled by many techniques. Models of tractive forces largely fall into three categories. Models which are used as inputs to multi body dynamics systems, such as FASTSIM [4,5,6] or Kalker's seminal CONTACT model for dry contact [7,8], typically assume the contact has a known friction coefficient or friction coefficient function. These models link the overall traction forces to the creep state of the wheel/rail interface, while resolving the contact patch into zones of stick and slip. The Extended Creep Force (ECF) model [9] represents an extension of these by assuming all slip is accommodated by a third body layer, however the properties of this layer cannot be directly measured and must be fitted from experimental data.

Other models, such as WILAC and LILAC [10,11], look at the effect of single contaminants or pairs of contaminants on the friction in the contact. These models are typically fitted to or presented with experimental data from test machines, which limit the creep state of the contact to one dimensional creep in the rolling direction and the contact patch to a single size and shape. These models may be analytical/numerical, treating the contamination in question as a lubricant with measurable properties, or statistical where the effect of the contamination or treatment is inferred from the model parameters. As the number of experiments needed to validate these models increases exponentially with the number of parameters investigated, these models are typically restricted to one or two contaminants and a range of creep values or speeds.

Lastly, forecasting models have been made which aim to predict the likelihood of real world events from weather data. Again these can be either purely statistical by fitting a model to past events and past weather data or analytical by, for example, predicting leaf fall and wind throw as a proxy measure for leaf contamination which is typically not directly measured or recorded in the real world. In the UK these models are produced by the Met Office [12] and others and they are used both to inform timetabling decisions and rail cleaning.

Though many models have been created, there remains an important gap in this spectrum. The “real world” traction value experienced by a train varies strongly with location, weather, contamination, train weight, and creep state. Clearly this value cannot be forecast effectively at a regional level, and it is not possible to measure the railhead state with the precision required for an analytical model, however all the necessary information is, in theory, present on the track or in the surroundings.

Machine learning models have gained enormous popularity in recent years for regression or classification problems on complex data. Convolutional neural networks (CNNs) have been used successfully in many image processing tasks and commonly outperform traditional image analysis and humans. Companies like Keen AI currently use sideways facing cameras mounted on rolling stock to collect and analyse images of the vegetation around the track, this images help identify intrusive plant species and vegetation, ultimately improving vegetation management for Network Rail [13]. One Big Circle have developed a system (Automated Intelligent Video Review, AIVR) to assist Network Rail in vegetation management and remote rail condition monitoring using videos collected on board a train analysed by AI packages [14]. In tribology, machine learning models have been used to predict friction in a variety of situations including effects from environmental and contact parameters [15].

In this proof-of-concept study, we aimed to produce a statistical model of the friction at this interface based on real world data. We have simplified the system by leaving out the creep state, wheel/rail relative positioning, and train specific parameters, instead using a tribometer to gather data from the running band. We included railhead and surroundings state information through easily collected images. This is combined with local meteorological sensor data (temperature and humidity). The model is validated by cross validation on the initial data set and through collection of a further validation set. A flow chart of the processes involved in building the prediction is shown in Figure 1.

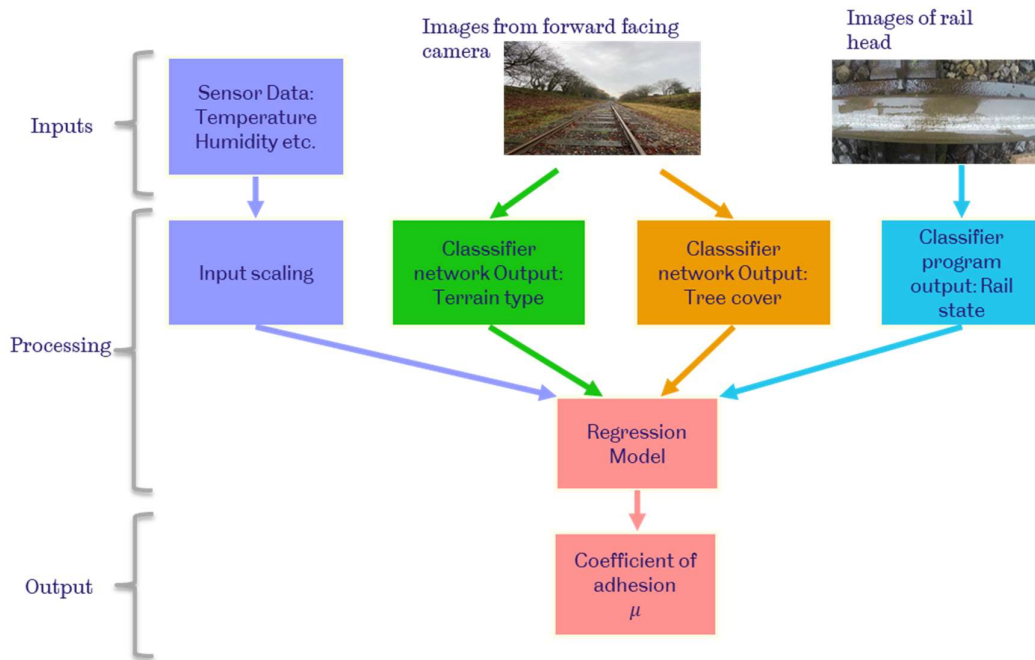


Figure 1: Machine learning friction prediction tool process

## 2. Background

### 2.1 Convoluted Neural Networks (CNNs)

In order to establish a relationship between non-linearly related parameters such as friction, relative humidity and temperature with non-mathematical parameters in this case images, an appropriate statistical model has to be chosen. Statistical models relate a set of independent parameters to one or several dependent variables. For any statistical model the complexity of the model and dimensionality of the input space need to be balanced by the amount of data available for training [16]. More complex models, or models of high dimensional data require more training data. The most appropriate type of model will depend on the complexity of the relations present, how much data is available and what other information is required with the prediction.

For image processing tasks, the dimensionality of the input is extremely large, for example, a one mega-pixel image with three colour channels requires a vector with three million dimensions to uniquely represent an input. In addition, the relations between individual pixel values and dependent variables of interest are often extremely complex, meaning that models need to be flexible, with many trainable parameters, to achieve good accuracy.

For these tasks CNNs are typically used, these leverage the structure of image data by filtering the input image, first producing maps of primitive features such as edges, then progressively higher order features such as simple shapes or objects. The result of this filtering is then fed into a further 'fully connected' network which is used to classify the input image. The network is trained by automatically adjusting the filter values and the weights in the fully connected network based on their derivative of the error, which can be found using the chain rule of calculus [17].

As the number of parameters in the filter kernels is high and many are needed to represent complex shapes, this process requires an enormous amount of data to avoid spurious correlations. This data requirement is a direct result of the high input dimensionality and the complexity of model, not a specific feature of neural networks. The resulting model consists of two distinct parts: a filtering network which takes an input image and produces "ratings" for each of a set of high level features and a fully connected network which links these features to image classes or other variables of interest.

## 2.2 Friction measuring approaches

Out of the various methods for taking friction measurements on the railhead, 3 methods were compared looking at their disadvantages and advantages based on their suitability for being used in the field. The outcomes are shown in Table 1. Folorunso et al. have extensively discussed different methods of friction measurements and describe the methods described in Table 1 in more detail and more analysis was done using the data collected (discuss further in this paper) to understand the environmental influences (such as relative humidity and railhead temperature) on railhead friction [18].

Table 1: Methods of measuring friction in the field [18]

S/N	METHOD	DISADVANTAGES	ADAVANTAGES
1	Tribometer Train (Vehicle-borne Tribometer)	It has been decommissioned. It is difficult to transport on-time due to its size and it requires a train path.	It gives a real representation of the wheel/rail interface as it is a full-scale rig using the same wheel load experienced in the wheel/contact hence scaling effects are not present. It gives actual representation of environmental and weather data.
2	Pendulum Tribometer	The contact pad material is made of rubber which is different from the materials of the wheel/rail interface. Scaling effects are present here because of the small contact area. The skid resistance is measured by the pendulum instead of the rolling/sliding peak of adhesion levels.	It is a portable, relatively easy to operate and is accessible. It functions as an on-track and in-lab measuring method. For a range of contaminants, it provides a notable resolution between the values for skid resistance.
3	Hand Pushed Tribometer	Scaling effects are present due to the size of the contact area. Its operation is at walking pace, hence low speed. It only measures the friction at slip of the worst point, therefore cannot be used for measuring all points on the track.	It is portable, easy to use and accessible. It can measure different positions across the railhead because the measurement wheel can be adjusted. The measurement process per spot is fast, because it takes approximately 15 to 25 seconds to record one adhesion level measurement.

The pendulum tribometer was selected as the friction measurement method for this project because it is easily accessible, portable and simple to use. The numerical values obtained from the pendulum's have been correlated to the British Railway Research tribometer train [19]. The linear correlation has a  $R^2$  value of approximately 0.93 shown in Figure 2, meaning there was a good linear relationship between the friction values of the pendulum tribometer and the British Railway Research tribometer train at the conditions tested. Although the apparatus has

a smaller contact area compared to the actual wheel/rail contact area, the correlation discussed above implies the pendulum tribometer produce similar relative results to established friction measurement methods

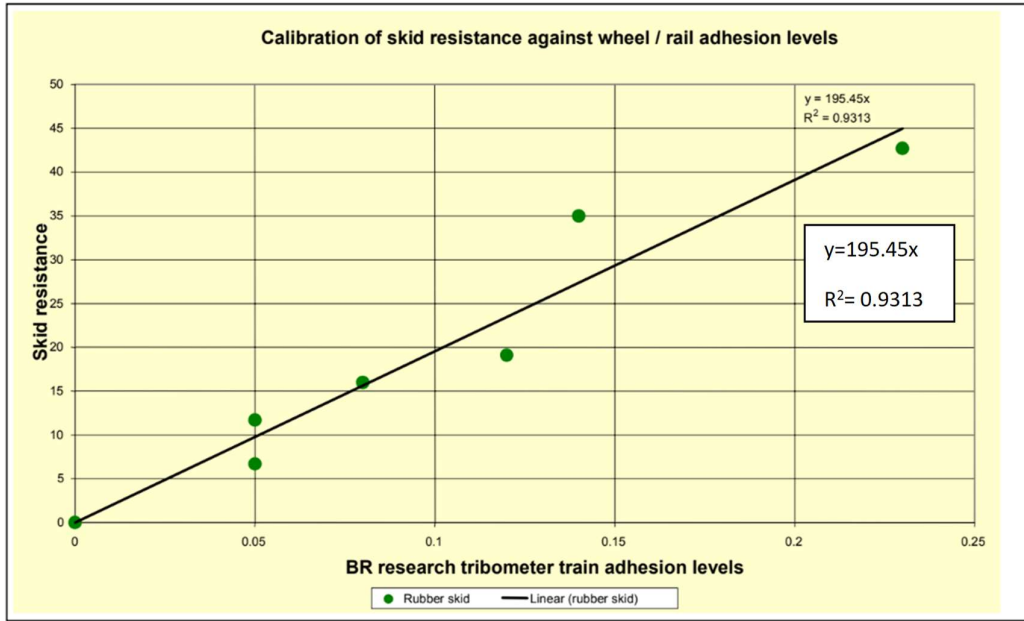


Figure 2: Pendulum to BR Research Tribometer train friction value conversion graph [19]

### 3. Data collection methodology

Data were collected throughout the Autumn period of 2019 and 2020 in a variety of locations and conditions. It was important to collect a complete dataset at each measurement location to build the model with all inputs. Data collected included photos, air temperature, humidity, and railhead temperature. In addition, a measurement of friction was collected, to train the model.

Data were collected at sites in the UK. Three heritage lines were used due to their ease of access, these were Ecclesbourne Valley Railway (Wirksworth, Shottle and Idridgehay stations), Midland Railway (Butterley and Swanwick stations), and Peak Rail (Matlock and Darley Dale stations) all located in Derbyshire, UK. As these were Heritage lines which are mainly run as tourist attractions mostly during holidays. Trains are driven over the tracks sparingly.

Tests which were done during Autumn 2019 were on run-in and ‘shiny’ railheads (which are representative of clean rail), as would be typical. Whilst those performed in Autumn 2020 were on a more oxidised railhead as fewer trains were running due to the COVID-19 pandemic.

#### 3.1 Friction measurements and methodology

The methodology established by Folorunso et al. was used in this work, and described briefly here. For a complete description refer to [18]. Friction measurements were taken by a pendulum tribometer. This apparatus, shown in Figure 3, swings a rubber pad against the rail. The distance which the pendulum travels after the contact is used to gauge the ‘skid resistance’ of the contact [20,21]. The skid resistance value was converted using the BR train conversion factor to give a coefficient of friction value. The apparatus must be set-up in the same position relative to the rail for every measurement and have the same pad/rail contact length.

At track access points, measurements were spaced approximately 10m apart. In total data were collected from over 300 points.

The pendulum was used in accordance with the guidance provided by GM/GN2642 and is outlined below:

- The pendulum tribometer was visually inspected for damages and contaminants from previous tests on the contact pad before assembly.

- After inspection of the apparatus, it was set-up on the section of the rail chosen for measurement.
- An Infra-Red thermometer was used to record the temperatures and relative humidity.
- The pendulum tribometer was mounted on a wooden support (seen in Figure 3) to balance the tribometer on the railhead and it was set level using the three adjustable levelling screws (confirmed via a spirit level).
- The fittings of each part of the pendulum tribometer were checked by swinging the pendulum swing arm for any loose fits, or potential damages.
- The pendulum swing arm head was clamped securely in the spring-loaded release mechanism and the GoPro camera was clamped on the pendulum frame to take the railhead image.
- The GoPro camera was detached from the pendulum frame in preparation for the start of measurements, to prevent obstruction of the pendulum swing arm movement.
- The rubber slider was checked for damage, the energy loss scale was set zero and the contact length between the rubber pad and the railhead was set to approximately 127mm [19].




Figure 3: Pendulum tribometer measurement of friction [18]

### 3.2 Image and Environmental data collection

Other data collected at the same time as the friction values included environmental conditions (briefly highlighted in the pendulum usage procedures section 3.1): railhead temperature, air temperature, relative humidity and dewpoint temperature were measured with an infrared thermometer. Railhead contamination thickness was taken using a low-cost paint thickness gauge.

Railhead photos were taken (Figure 4), using a ‘Go-Pro Hero 4’ camera, from a consistent height and position by mounting the camera on the pendulum rig. Full images were 12 megapixels, however, only a 750 by 500 pixel region of interest was used for processing.



Figure 4: Samples of the Railhead images captured using a GoPro camera

Forward-facing images were also collected from the measurement sites. The forward-facing photos are similar to the image shown in Figure 5. A resolution of 343 by 609 or higher is required. At a higher rail industry readiness level these would be collected with a camera mounted on the front of the vehicle. As the forward-facing image is intended to encode only the location information, and not as an indication of weather, they were not taken for each individual measurement. This image will give information on the nature of the track surroundings and how close



vegetation is.


#### 4. Image processing

Image processing was carried out to reduce the images to a few features that are relevant for friction estimation. Different techniques were used for the forward facing images and the railhead images due to the different amounts of data available.

##### 4.1 Forward facing images

Forward facing images were collected from the test sites alongside the friction data. These images are complex, containing many different objects and situations. As such traditional image processing would be impossible and manual labelling would be too time consuming to scale to a full network system. Neural networks offer an attractive solution, but the small number of images makes directly training a network impossible. Instead, these images were augmented with a large set (~20,000) of visually similar forward facing images scraped from various UK sources on the internet. The resulting set provides a representative sample of images from UK networks and can be used to train a dimensionality reduction model which retains relevant distinguishing information from the images, while discarding information common to all images.

An example image from the augmented set is shown in Figure 5. Many parts of these images are common for every image, these areas are removed by cropping the image into two sub-images as shown. The sub images were then resized to the correct input size for a pre-trained CNN.



*Figure 5: An example forward facing image showing the sub images extracted for further analysis, image created by author*

These sub-images are passed through a CNN that has previously been trained on a large, labelled data set. In this study MobileNet V2 was used as it is suitable for high speed use on low cost devices. This network reduces the dimensionality of the data from millions of brightness values to a 2048 element feature vector. The feature vector consists of features which have been trained to be useful for common image classification tasks. These are high level features, many of which relate to familiar concepts (e.g. a human face).

Within this vector many features are irrelevant for our task or strongly correlated to each other. To further reduce the dimensionality, a Principal Components Analysis (PCA) was carried out on the feature vectors. This finds orthogonal, linear combinations of parameters which contain the most variation for the data presented [22]. Examples from the extremes of the first three principal components are shown in Figure 6. The values for the first 600 principal components were retained, these contained 90% of the total variation in the image data. These components are high level, abstract representations of the data and are unlikely to be summarised well by a description.

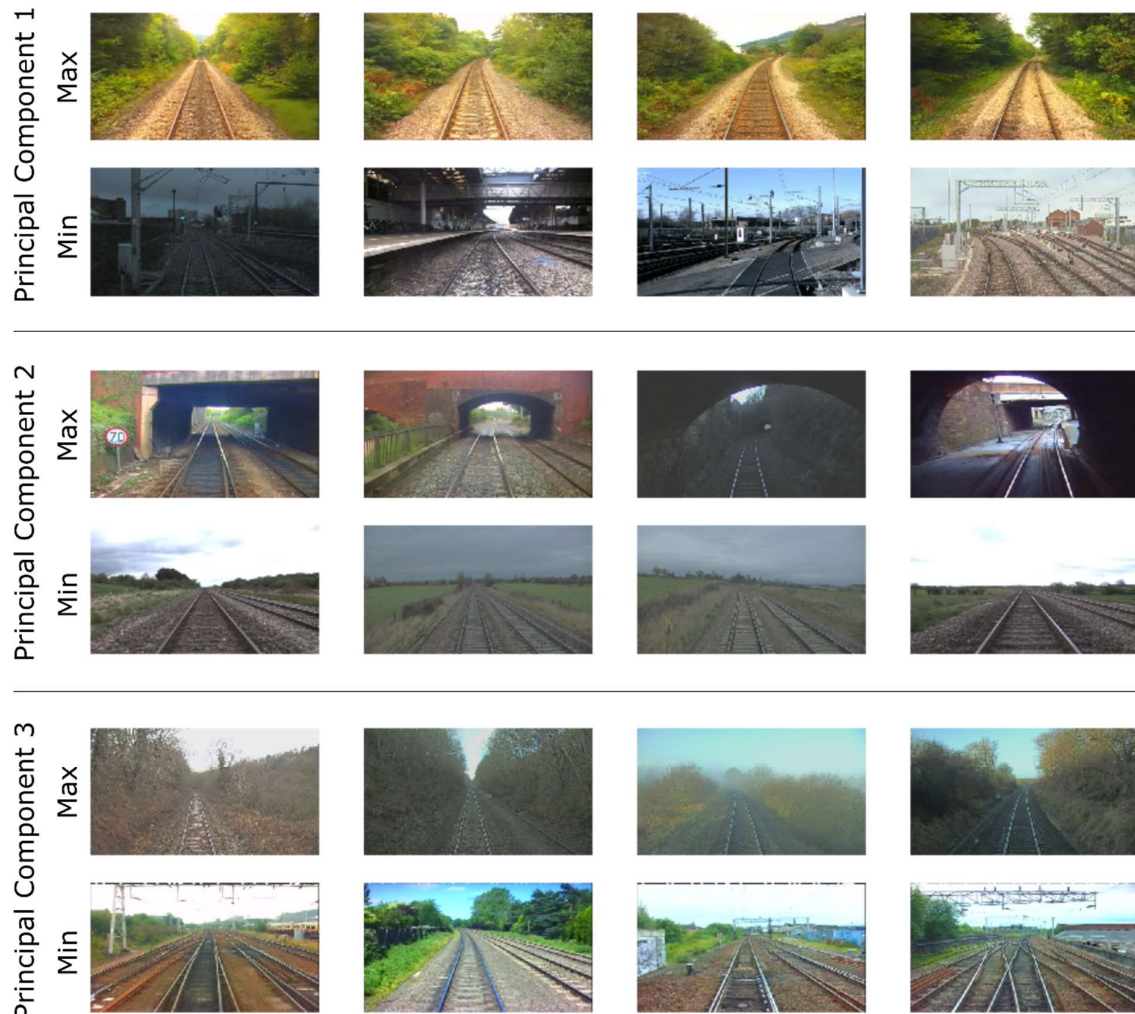


Figure 6: Images from the extremes of the first three principal components for the data set. Each pair of rows represents a principal component, image created by author

With this lower dimensionality and a large number of images, an unsupervised classification tool can be used. This splits the data into groups which are similar to each other. There are many methods of completing this task, here we have used a self-organising map (a type of neural network). The results of this classification are shown in Figure 7.





Figure 7: Examples of the different groups of images found by the unsupervised learning technique. Columns are separate groups: grassland, trees, rural, urban, image created by author

The result of this process is a pair of coordinates, which place a railhead image on the map. The groups presented in this map have been made from a representative sample of forward-facing images.

#### 4.2 Railhead images

Railhead images (see Table 2 showing samples of images collected to give an indication of the railhead images seen) were collected along with associated environmental and friction data. Unlike the forward-facing images there is no large source of representative railhead images. In addition, pre-trained networks are typically trained on images which are very different from the railhead images and are unlikely to be useful. However, the dimensionality of the images must still be reduced before the images can be used (see Figure 8). To achieve this, features were extracted from the images using traditional image processing techniques. Examples of the traditional image processing techniques and tool kit used include, but are not limited to an Edge detector for identifying the boundaries of the railhead in the image, Numpy for indexing operations to modify the pixel values of the image and SciKit-image that works alongside Numpy which was used for feature extractions [23].



Figure 8: Railhead Image before (left) and after (right) dimensionality reduction

Before feature extraction, the images were normalised and the rail was located in the image, the rail was then cropped out for further processing. The features were chosen as features likely to be correlated with the friction present. These were: the number of black pixels in the image, the number of orange pixels in the image, the average colour of the railhead and the sum of the first derivatives in the along-rail and across-rail directions. The

first derivatives are higher when many strong edges are present, such as when the rail is rough. These features could then be used directly in the prediction tool.

Table 2: Example of images collected under different conditions, classified under dry and wet

DRY CONDITIONS	WET CONDITIONS
	
	
	
	

### 5. Regression model

The forward-facing image map positions, railhead image features and sensor measurements were combined in a model to predict the railhead friction. A Gaussian process regression model was selected as it is flexible enough to accurately capture the relations which are likely to be present, and data efficient enough to be fitted well using a data set of this size. In addition, these models also provide an estimate of the error of the prediction given [24]. This mitigates the risk of incorrect estimation/prediction in new scenarios. Before fitting, all data have been linearly scaled to a unit scale, meaning that the highest value is scaled to 1 and the lowest value is scaled to 0.

The Gaussian process is defined by a kernel function. This encodes the joint variability of the models parameters. This can be used to set prior information about how the data relate to each other, how much noise is present in the data and any underlying structure. The model consists of a summation of a constant kernel, a white noise kernel and a non-linear kernel. Multiple non-linear kernels were fitted and the one producing the highest marginal likelihood (rational quadratic) was chosen. The constant kernel is set to 0.5 while the further hyper-parameters of the model are set by optimisation during the fitting process. The optimisation aimed to achieve the maximum log marginal likelihood for the data given the model and a Gaussian error function.

### 6. Results

The fit achieved with the data is shown in Figure 9. The overall log likelihood of the model is 176.5 and the  $R^2$  value for the model, with this data is 0.97. In order to validate the system, data were left out of the fitting process and the prediction of the naïve model compared to the actual value at the left-out points. The first step in the validation process was to leave a single point of data out at a time. The prediction of the naïve model at the left-

out point can both be compared to the true value at that point as shown in Figure 10A and the prediction of the full model as shown on in Figure 10B.

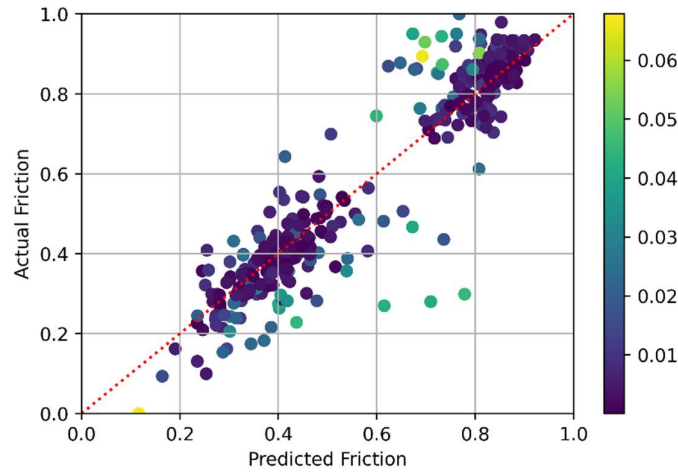


Figure 9: Model predictions compared to the actual value at each point, points are coloured by their leverage. Data on both axes are normalised to a 0-1 scale.

These results show naïve models are still able to explain 96% of the variation in the data (coefficient of determination = 0.96). Additionally, the average change in prediction between the full model and the naïve models is only 3% of the measurement range. This shows that, in general, the model is not over-fitted to the data, and that trends fitted by the model are likely to be real.

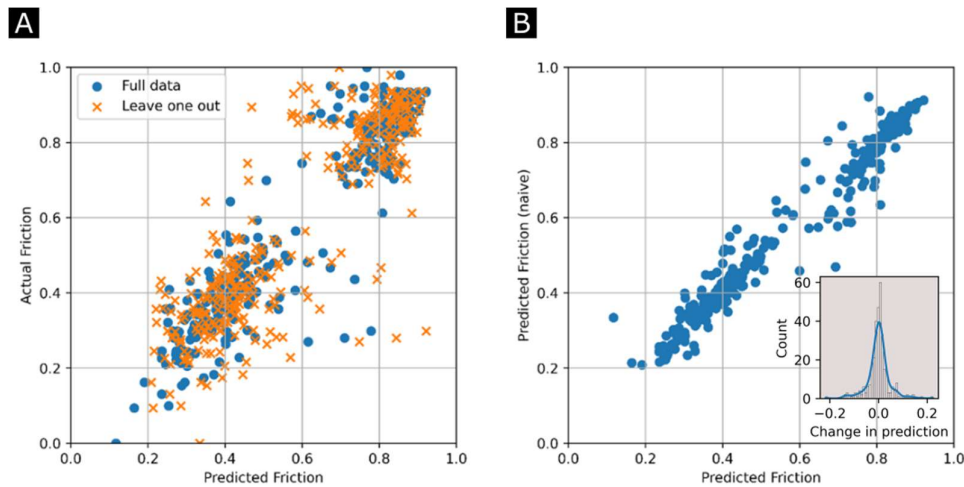


Figure 10: Results of the leave one out validation (A) and a comparison of the full model to the naïve model for each point (B), with a histogram of the change between the naïve and full models. Data on both axes are normalised to a 0-1 scale.

This process was extended to leaving groups of twelve points out. The models fitted leaving groups of twelve points out are compared to the true value at the left-out points in Figure 11A, this plot is for one set of groups which include the whole data set. These values are again compared to the result from the full model in Figure 11B. This process has been repeated for all possible groups of twelve points.



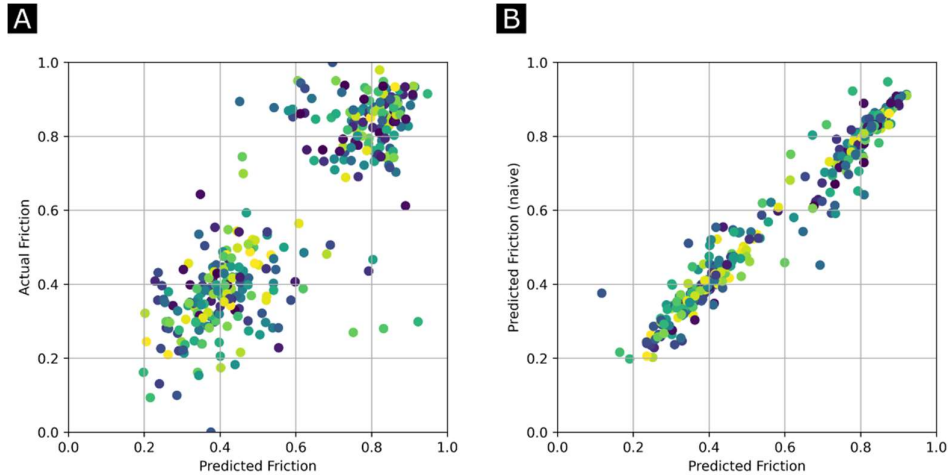


Figure 11: Results of the leave groups out validation for one set of groups (A) and a comparison of the full model to naïve models for each group (B). Data on both axes are normalised to a 0-1 scale.

As above for the leave one out validation, the mean change in prediction between the full model and the naïve models was only 3% of the measurement range. The mean coefficient of determination for the left-out points was 0.96.

While it is clear from Figure 9 and Figure 10A that several points are over leveraged, and not predicted well when left out, the majority of the points are not. In collecting these data, we have aimed to collect from locations and in conditions likely to cause low adhesion, as such much of the data are from low adhesion conditions and are well predicted when left out of the fitting process.

### 5. Discussion

This project has shown that it is possible to accurately predict the friction conditions on the rail using low-cost sensor data and images of the railhead and surroundings. The overall coefficient of determination of the model was 0.97. This is in line with or better than many machine learning models which have been applied to other tribology systems, despite the complexity of the wheel rail system [25,26].

At its current state, the tool has some limitations which restricts its use in its current form. Firstly, and most importantly, the friction value which is predicted is the friction measured by a pendulum tribometer. Previous field results have shown that the readings of this tribometer are correlated with the adhesion a wheel would experience on the same track, however, as with all friction measurement devices, they give slightly different results for tests carried out on the same section of rail. This happens due to the different operating mechanisms of each device, geometry and material of contacting parts (with the railhead), contact forces etc. therefore the pendulum results alone could be correlated directly with wheel/rail friction, but the BR Research conversion gives more confidence as this involved a full-scale and actual measurement of wheel/rail interface friction. With the use of the BR Research conversion the tool is expected to predict friction results representative of the wheel/rail friction values.

Extra care must be taken when using the measuring apparatus, because residue on the rubber contact pad, or improper placement of support platform, can cause false friction readings to occur. Ultimately leading to training the prediction tool on wrong information

Secondly, in its current state much of the image processing is not optimal, features extracted from CNNs are optimised for the tasks which the CNNs were originally trained for. Many of these will relate to faces or other features we do not expect to see in a lineside context. This is reflected in the PCA being able to effectively remove hundreds of these features with little loss in variation of the data set. Likewise, the groups from the self-organising map are selected only to keep variation in the images. With a large amount of data these processes could be optimised for friction estimation as could the image processing of the railhead images.

The issues discussed above highlight the need for large sets of friction data from on train data recorders. To be useful in a statistical model these data must additionally tie the friction values to other factors, raw friction values with only a time stamp are not usable except to be tied to historical weather forecasts, which would limit

predictions to a regional level. However, there are substantial difficulties in collecting such data, due to safety considerations when attaching electrical equipment to in-service trains. Additionally, the type of data which is needed is not simple to collect and not collected by default. These problems lead to a disconnect between the quality and content of data collected by train operators and the needs of modellers.

Work is underway, testing the image sensitivity of the model. This will inform us of the robustness of the prediction tool's image processing capability, highlighting areas of weaknesses such as image resolution, lighting conditions and focus distance that may need further attention. The model is also undergoing continuous retraining with new data sets to optimise the railhead image processing for the friction estimation.

Next steps in this work involves:

- Additional model-training from a diverse range of locations to increase the spectrum of data the tool is trained with, therefore increasing scenarios for accurate prediction
- Developing and testing an image capturing system with sensors, that can be placed on-board a train
- Collection of railhead and forward facing images (using the system developed) on board a moving train with corresponding friction measurements. This will be a step moving towards a fully implemented train-borne system.

## 6. Conclusions

In conclusion, we have shown that:

- Friction conditions on the rail can be estimated/predicted from images of the surroundings, images of the railhead and some easily obtained sensor data. However, progressing the system described above to a real-world context primarily requires further data gathering from the real-world.
- Retraining the tool with specific data to an environment is possible and straightforward with python and basic ML knowledge.
- The  $R^2$  value for the prediction model was 0.97 showing the precision at the moment is very good and the naïve model's  $R^2$  value was 0.96. Although the  $R^2$  value is likely to change with increased training, the precision of a neural network is not fully measured by it.
- The use of a Gaussian process (GP) is applicable in the analysis of non-mathematical relations such as images characteristics, conditional statements and established non-linear relationship used in friction, railhead and air temperature. The GP model has aided the establishment of a link between the properties and made prediction possible.

## Acknowledgments

The authors would wish to acknowledge RSSB for initial project funding through project CF-UOS-02. For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) license to any Author Accepted Manuscript version arising.

The work was also supported by the EPSRC Programme Grant "Friction: The Tribology Enigma" (EP/R001766/1).



## References

1. Agency EE. Rail and Waterborne—Best for Low-Carbon Motorised Transport. Publications Office Luxembourg, 2021.
2. Lightoller A and Purcell L. Low Adhesion and Sanders. *PWI* 2020; 138: 28–33.
3. Monsuur F, Enoch M, Quddus M, et al. Modelling the impact of rail delays on passenger satisfaction. *Transportation Research Part A: Policy and Practice* 2021; 152: 19-35.
4. Kalker J. A fast algorithm for the simplified theory of rolling contact. *Vehicle system dynamics* 1982; 11: 1-13.
5. Bernal E, Spiryagin M, Stichel S, et al. Friction-Slip Curves – The Pathway From Twin-Disc Tribo Measurements to Full-Scale Locomotive Multibody Simulations. In: 2022 Joint Rail Conference 2022, V001T07A005
6. Chevalier L, Eddhahak-Ouni A and Cloupet S. On a Simplified Model for Numerical Simulation of Wear During Dry Rolling Contacts. *Journal of Tribology* 2008; 131. DOI: 10.1115/1.3002322.
7. Kalker JJ. *Three-dimensional elastic bodies in rolling contact*. Springer Science & Business Media, 2013.
8. de Mul JM, Kalker JJ and Fredriksson B. The Contact Between Arbitrarily Curved Bodies of Finite Dimensions. *Journal of Tribology* 1986; 108: 140-148. DOI: 10.1115/1.3261134.
9. Meirhofer A. A new wheel-rail creep force model based on elasto-plastic third body layers. 2015.
10. Trummer G, Buckley-Johnstone L, Voltr P, et al. Wheel-rail creep force model for predicting water induced low adhesion phenomena. *Tribology International* 2017; 109: 409-415.
11. Buckley-Johnstone L, Lewis R, Six K, et al. *T1149: Extension of WILAC Low Adhesion Mode*. 2020. RSSB Project Report.
12. El Rashidy R, Stow J, Chapman V, et al. *Feasibility of Integrating Operational Data with Adhesion Forecasts*. Report no. COF-FCA-03, 29-07-2020 2020. Institute of Railway Research, University of Huddersfield, The Met Office.
13. Bauroth S. AI Holds Key to Improving Biodiversity by Britain’s Railway Tracks, <https://keen-ai.com/blog/ai-holds-key-to-improving-biodiversity-by-britains-railway-tracks/> (2022, accessed 15-11 2022).
14. OneBigCircle. Case Study: Network Rail, <https://onebigcircle.co.uk/case-study-networkrail/> (2022, accessed 15-11 2022).
15. Marian M and Tremmel S. Current trends and applications of machine learning in tribology—A review. *Lubricants* 2021; 9: 86.
16. Chambers JM and Hastie TJ. Statistical models. *Statistical models in S*. Routledge, 2017, pp.13-44.
17. Bishop CM and Nasrabadi NM. *Pattern recognition and machine learning*. Springer, 2006.
18. Folorusno MO, Lewis R and Lanigan JL. Effects of temperature and humidity on railhead friction levels. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*; 0: 09544097221148236. DOI: 10.1177/09544097221148236.
19. RSSB. *Guidance on Wheel / Rail Low Adhesion Measurement*. 2008. 160 Euston Road, London NW1 2DX.
20. Lewis R, Lewis SR, Zhu Y, et al. The modification of a slip resistance meter for measurement of railhead adhesion. *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of rail and rapid transit* 2013; 227: 196-200.
21. White B. *Using Tribo-Chemistry Analysis to Understand Low Adhesion in the Wheel-Rail Contact*. University of Sheffield, 2018.
22. George A and Vidyapeetham A. Anomaly detection based on machine learning: dimensionality reduction using PCA and classification using SVM. *International Journal of Computer Applications* 2012; 47: 5-8.
23. Van der Walt S, Schönberger JL, Nunez-Iglesias J, et al. scikit-image: image processing in Python. *PeerJ* 2014; 2: e453.
24. Lee J, Bahri Y, Novak R, et al. Deep neural networks as gaussian processes. *arXiv preprint arXiv:171100165* 2017.
25. Hasan MS, Kordijazi A, Rohatgi PK, et al. Triboinformatics Approach for Friction and Wear Prediction of Al-Graphite Composites Using Machine Learning Methods. *Journal of Tribology* 2021; 144. DOI: 10.1115/1.4050525.
26. Xie H, Wang Z, Qin N, et al. Prediction of Friction Coefficients During Scratch Based on an Integrated Finite Element and Artificial Neural Network Method. *Journal of Tribology* 2019; 142. DOI: 10.1115/1.4045013.