Predicting Aquaplaning Performance from Tyre Profile Images with Machine Learning

Tillman Weyde[1], Gregory Slabaugh[1],
Gauthier Fontaine[2], and Christoph Bederna[3]

[1] City University London
[2] ENSTA Paristech
[3] Continental AG
{t.e.weyde,gregory.slabaugh.1}@city.ac.uk,
gauthier.fontaine@ensta-paristech.fr,
christoph.bederna@conti.de

Abstract. The tread of a tyre consists of a profile (pattern of grooves, sipes, and blocks) mainly designed to improve wet performance and inhibit aquaplaning by providing a conduit for water to be expelled underneath the tyre as it makes contact with the road surface. Testing different tread profile designs is time consuming, as it requires fabrication and physical measurement of tyres. We propose a supervised machine learning method to predict tyres' aquaplaning performance based on the tread profile described in geometry and rubber stiffness. Our method provides a regressor from the space of profile geometry, reduced to images, to aquaplaning performance. Experimental results demonstrate that image analysis and machine learning combined with other methods can yield improved prediction of aquaplaning performance, even using non-normalised data. Therefore this method has can potentially save substantial cost and time in tyre development. This investigation is based on data provided by Continental Reifen Deutschland GmbH.

Keywords: Image analysis, supervised machine learning, non-normalised data, tyre profile, regression.

1 Introduction

Aquaplaning occurs when a layer of water exists between a tyre and road surface, causing the tyre to lose traction and becoming unresponsive to user control. While the onset of aquaplaning depends on several factors such as the amount of water on the road surface and the road texture, the tyre contour and inflation pressure, as well as vehicle speed and weight, it depends critically on the tread pattern, which is the set of grooves, blocks and sipes, as shown in Figure 1 (a). Grooves provide larger channels through which water passes, and are often arranged both circumferentially and laterally. Sipes are narrow voids, located on the blocks, that allow the blocks to deform. Despite great progress in tyre development, tread pattern design in respect to aquaplaning performance still poses a difficult challenge.

M. Kamel and A. Campilho (Eds.): ICIAR 2013, LNCS 7950, pp. 133–142, 2013.

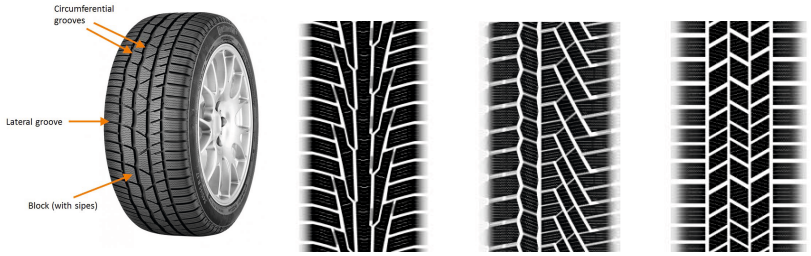

Fig. 1. A tyre tread consisting of grooves, blocks, and sipes is shown in (a). In (b), (c), and (d), we show a portion of example profile images for different tyres.

Geometries of tread patterns are typically described digitally using computer-aided design (CAD) software; examples are shown in Figure 1 (b), (c), and (d). Testing a new tread pattern requires fabrication of a tyre set using a custom mould made from the contour and pattern data, followed by experimental testing in a controlled setting using an aquaplaning rig. Unfortunately, this process is expensive and can take up to several months to perform, greatly limiting the number of tread patterns that can be evaluated.

To address this issue, approaches to predict tyre aquaplaning performance have been developed using finite-element methods (FEMs) [1,2], based on Lagrangian and/or Eulerian formulations along with fluid-structure interaction behaviours. These models have shown considerable promise but they face the problem of high effort in model generation and long simulation time. Therefore quick, analytical methods with focus on the pattern modification only are required especially in pattern predevelopment.

1.1 Our Contribution

This paper is the first, to our knowledge, to address the problem of prediction of tyre aquaplaning performance from pattern geometry reduced to pattern images only. We make several contributions:

- We extract a set of image features correlated with tyre aquaplaning performance, including a new frequency-domain feature called RAVLOMS that captures the radial variance of the log magnitude spectra.
- We develop a regressor using neural networks to map from the space of profile image features to aquaplaning performance.
- Our method can utilise and improve results, if available, predicted from other approaches like heuristic formulae or FEM.
- We address the problem of neural network learning on non-normalised data groups with variations in group size.

The rest of the paper is organised as follows. In Section 3 we describe the set of image features extracted from the images. In Section 4 we describe our

machine learning approach, based on regression using a neural network, and the challenge faced using non-normalised data. Experimental results are presented and discussed in Section 5.

2 Dataset

Our dataset consists of 21 tread profile images, arranged into three groups $G_{1,2,3}$ related to three test sets. G_1 consists of four images, G_2 consists of nine images, and G_3 consists of eight images. Example profile images are provided in Figure 1 (b), (c), and (d). Each image was made into a physical tyre; however, tyres were made with different tread compound. To account for that, rubber hardness measured in *Shore A* is a additional feature in our prediction system.

For each tyre, an experimentally measured aquaplaning performance value was provided, the so called AQF value, which is the target value for our prediction. The AQF is proportional to the speed at which aquaplaning starts to occur. In addition, a predicted aquaplaning value, the so called FT value, was made available by the tyre manufacturer using a heuristic formula based on tests and design experience. The FT value is the baseline over which we aim to improve.

The measurements for each group were conducted independently at different times and under different conditions and the available data are only relative to one tyre in each group that is taken as a reference. The tyre models tested in each programme were disjoint, therefore we had no common reference points for normalising the data.

3 Tread Profile Image Features

Based on input from experts in tyre development, we designed a set of seven features to capture physically meaningful profile design parameters that relate to aquaplaning performance. These features are described below.

3.1 Void Volume

Void volume refers to the amount of space forming the grooves and sipes. It is through these conduits, particularly the grooves, that the water will be channeled. We make use of two void volume features,

$$V = \sum_{x=1}^{N}\sum_{y=1}^{M} I[x,y]dxdy, \qquad Vc = \sum_{x=1}^{N}\sum_{y=1}^{M} I[x,y]\delta[x,y]dxdy \tag{1}$$

where $I[x,y]$ is the profile image, N and M are the image width and height respectively, dx and dy are the size of a pixel in mm. V measures the total void volume over the tyre. As is evident in Figure 1, the right and left sides of each profile image contain additional void volume, where the tyre is not making contact with the road due to its curvature. Vc provides a similar measurement to

V but only in the central portion of the image, represented with a mask $\delta[x, y]$. The mask is one in the central portion of the image and zero on the sides, and is automatically computed by thresholding (using a threshold T > 100) the image and setting the mask to zero for bright connected columns on the left and right side.

3.2 Angular Features

Empirical observations revealed that aquaplaning performance varied depending on the diversity of the angles appearing in the profile image. This is most evident when one looks at frequency spectrum of the profile image, as shown in Figure 3. In (a) we observe the image with the lowest aquaplaning performance, and in (b) the highest. We propose to capture this with a novel feature we call RAVLOMS (**RA**dial **V**ariance of the **LO**g **M**agnitude **S**pectra). To compute RAVLOMS, we sample the log magnitude spectrum along a ray starting at DC, producing a ray $r_\theta[t]$. A sample ray is shown in Figure 3 (a) for $\theta = 0$. RAVLOMS (denoted as R below) is then

$$R = \sum_{\theta} \sum_{t} [r_\theta[t] - \bar{r}(t)]^2, \tag{2}$$

where $\bar{r}(t)$ is the mean ray spectra averaged over all rays in 360 degrees, and is the sum of variance along the rays. RAVLOMS is zero for a spectrum that has no radial variance, and captures the degree to which rays vary along different directions in the spectrum.

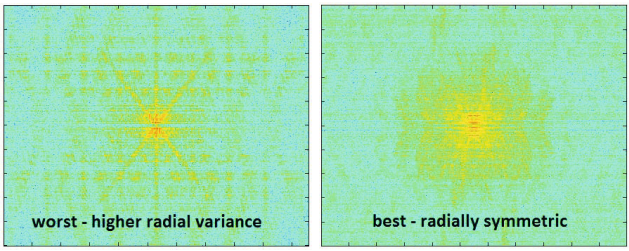

(a) (b)

Fig. 2. We show the log magnitude spectrum of two profile images with the worst (a) and best (b) aquaplaning performance. We note the the tyre with the best aquaplaning performance has a more radially uniform spectrum.

Two additional angular features, θ_l and θ_r measure the outgoing angles for the lateral grooves, as depicted in Figure 3 (b).

3.3 Cross-Sectional Features

The cross-sectional features characterise at the shape of the grooves. For a cross-section of a groove, we compute the area under the curve (AUC) by integrating

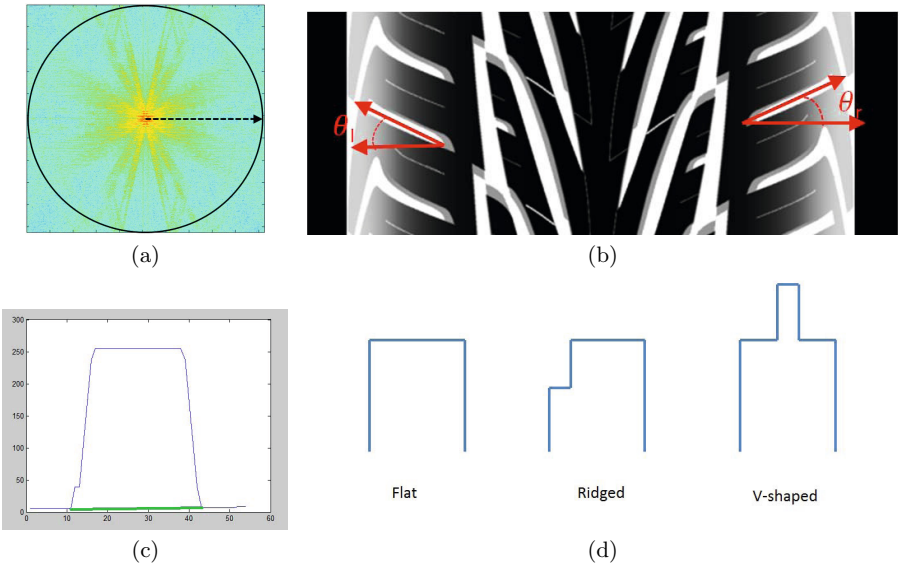

Fig. 3. Feature calculation. In (a), we show how the log magnitude spectrum is sampled along rays as part of the RAVLOMS feature. In (b), we show the two outgoing angles along the lateral grooves. In (c) we show a cross-sectional profile for a groove; the area under the curve is computed between the blue curve and green curve. In (d) we show a schematic the three cross-sectional shapes observed in the data.

the curve automatically by detecting the start and end of the groove based on peaks in the gradient; the region of support is shown in green in Figure 3 (c). Based on the observation that the cross-sectional profiles generally fall into one of three categories (flat, ridged, or v-shaped), as shown in Figure 3 (d), we also compute the cross-sectional shape (CSS). For this feature, a flat profile has a feature value of zero. Ridged profiles have a negative score that measures the void area that is missing compared to the flat case. V-shaped profiles have a positive score based on the area extending from a flat top.

4 Machine Learning of Aquaplaning Performance

Predicting the aquaplaning speed of the tyre based on a feature vector extracted from the tread profile is a regression problem, i.e. adapting a function to the data. The argument of our regression function is a vector that consists of the seven features introduced in the previous section plus the rubber hardness and the FT prediction. We use a standard feed-forward neural network with sigmoid activation functions in the input and hidden layers and back-propagation, i.e. gradient descent, for learning. This well known non-linear regression method,

developed independently by several researchers and made popular by [3]. Most non-linear regressors – like the neural networks used here – pose a non-convex optimisation problem. Thus a globally optimal solution cannot be guaranteed by gradient descent, as the method may converge to a local optimum. This is addressed by training the neural network several times starting from random weights.

We evaluated all models with "leave-one-out" cross-validation. That means we trained the model on n training sets consisting of the whole data set (of size n) without one item. We tested then each trained model on the one left-out item. For the neural networks, we use regularisation to avoid over-fitting (see [4]). The regularisation parameters as well as the number of hidden layers and neurons were determined by testing various values over validation sets, taken from the training sets. Cross-validation of model training and training parameterisation is especially important with the small data set at hand, as over-fitting can easily lead to error under-estimation.

We used several machine learning approaches, including a linear model as a baseline for learning. In terms of neural networks, we first used a single network for all data, ignoring the different origins. Secondly we used a separate network for every programme and a new method described in the next section.

4.1 Machine Learning from Non-normalised Data

For the initial approach we used a network with one hidden layer and one output neuron as shown in Figure 4 (a). Since the AQF values in the different groups are not normalised to a common reference, we would like to model a different regression function for each group but still exploit the information in the complete dataset. This is known as a multi-task scenario in machine learning, where one task means one group in our case.

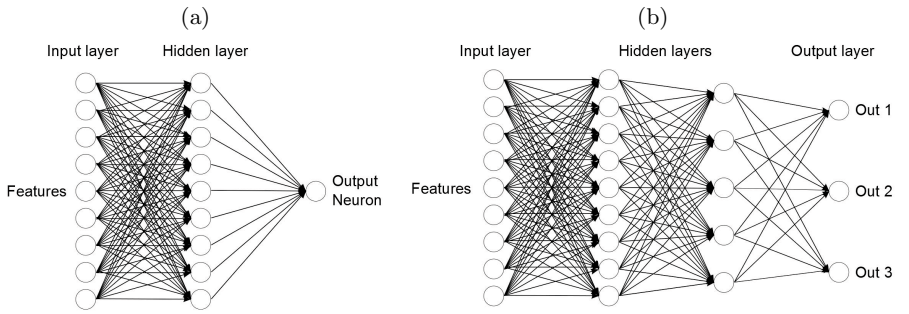

Fig. 4. Neural networks with a single hidden layer and a single output neuron (a) and with 2 hidden layers and 3 output layers as used for SOT (b)

To address multi-task learning, we used a network with 3 output neurons and 2 hidden layers, as shown in Figure 4 (b), where one output neuron corresponds to each group. We devised a learning method for this network, which we call Separate Output Training (SOT). SOT uses the learning method described by [5] but with support for groups of different sizes. We measured the error for any training sample from group i as the squared error with regards to the activation of output neuron n_i. In contrast to [5] we use the sum of squares error (SSE) without weighting and averaging as the overall loss function, to allow for different group sizes:

$$SSE(W) = \sum_k (n_{i,k} - y_k \text{ ,where } y_k \in G_i), \tag{3}$$

where W is the set of network parameters (weights), $n_{i,k}$ is the activation value of output neuron i for input from training sample k, y_k is the AQF value of sample k, and i is the index of the group that sample k belongs to. The error value for all other output neurons with $y_k \notin G_i$ is defined as 0. The hidden layers of the neural network are thus trained according to all the training samples, but the connections to an output neuron are only affected by training samples from the group that is relevant to that output. This approach has been shown theoretically and empirically to require less data for the same generalisation performance than learning for each group individually [5], which is particularly interesting on a small dataset like the one at hand. Multi-task learning has received some attention by researchers in recent years (e.g. [6,7]), who developed other learning models. These models were not applied here because they do not fit our task or they only address linear models.

5 Experimental Results

Linear Model. The baseline error from FT value is 3.4% average deviation of the prediction from the measurements. For comparison, we fit a standard linear model to the whole data set, minimising the sum of squared errors, with a resulting error value on the whole set of 2.8%. However, leave-one-out cross-validation yields an prediction error value of 6.41%. Therefore this method does not generalise well to new data. The generalisation could be improved by using a regularisation term, but reducing the performance on the training data, so a substantial improvement or predictions with a linear model is not achievable. We therefore tried using neural networks as non-linear models to reduce the prediction errors. The results in Table 1 show also that the FT value is the only significant component in the linear model. However, the newly developed AUC, CSS and RAVLOMS features show better contribution in terms of p-values than the volume and hardness based features.

Standard Neural Networks. With a standard neural network (NN), we trained a single output neuron on the whole dataset, assuming that the differences between the groups might be negligible. The results are presented in Table 2, with cross-validated error rates (NN CV) given as averages over 100 runs of the learning

Table 1. Linear model - mean error: 2.80%, cross-validated error: 6.41%. The p-values indicating significance at the .05 level are shown in bold print.

	Inter-cept	FT	Shore A	V	Vc	RAV-LOMS	θ_l	θ_r	AUC	CSS
Coeff.	-7.6909	0.1134	0.0202	-0.0103	-0.0186	-0.0198	0.0015	0.1887	0.0010	0.0002
p-Value	0.6486	**0.0010**	0.6711	0.6470	0.5591	0.4054	0.9459	0.6119	0.3657	0.3695

Table 2. Results of the single network approach

	Group 1	Group 2	Group 3	Entire set
FT error	0.0492	0.0361	0.0243	0.0341
NN CV error	0.0486	0.0771	0.0290	0.0533
% average error reduction	1.25	-113.5	-19.44	-56.49
% of trainings with NN < FT error	54	0	24	0

procedure. The average error is improved only for the first group by a small amount (1,24 %), otherwise is it worse than the FT error.

Removing Outliers. One peculiarity of the dataset was that there were two data items in the 2nd group, which had AQF values around 120, while all other values were between 90 and 103. Removing these outliers improved the learning success, as is shown in Table 3. The NN results are now improved by 28.5% over the FT error (which is also less than on the full set). However, this model is not robust against large variations between groups, and it can not predict well extreme AQF values. This is problematic, because the relevant information for recognising outliers can only be obtained experimentally at high cost, contrary to the objective of producing quick and reliable predictions.

Table 3. Results of the single network with 2 outliers removed from Group 2

	Group 1	Group 2	Group 3	Entire set
FT error	0.0492	0.0238	0.0243	0.0293
NN CV error	0.0251	0.0309	0.0103	0.0210
% of average error reduction	48.98	-29.59	57.69	28.51
% of trainings with NN < FT error	99	9	100	100

Separate Networks per Group. For comparison, we also trained a separate NN per data group. In this constellation we conducted regression on the AQF and the residuum AQF relative to the FT prediction. Although the latter method performed better than the former, both produced greater errors than the FT prediction. We also tried to predict the AQF only from the features, not using

Table 4. Results achieved with three separate networks

	Group 1	Group 2	Group 2 w/o outliers	Group 3	Entire set
FT error	0.0492	0.0361	0.0238	0.0243	0.0341
a) AQF with FT prediction as an input					
NN CV error	0.0337	0.0863	0.0347	0.0145	0.0490
b) Residual AQF with FT prediction as an input					
NN CV error	0.0432	0.0490	0.0240	0.0166	0.0355
c) AQF without FT prediction as an input					
NN CV error	0.0349	0.01095	0.0415	0.0198	0.0611

the FT prediction, but the results were way worse than the FT predictions. The results are summarised in Table 4, showing that this approach performs worse than the FT alone, probably because of the small amounts of data for training the network.

Separate Output Training. The SOT training as described above was a way of including the outliers and making use of the complete information in the dataset. The results are presented in Table 5, showing that with the SOT learning on the whole dataset the predictions can be improved over the FT. This means that the SOT method provides a regression method that is robust against variations in the output mappings and that was reasonably successful in predicting even extreme AQF values.

Table 5. Results of the SOT approach

	Group 1	Group 2	Group 3	Entire set
FT error	0.0492	0.0361	0.0243	0.0341
NN CV error	0.0328	0.0446	0.0066	0.0279
% of average error reduction	33.21	-23.55	72.86	18.18
% of trainings with NN $<$ FT	100	0	100	70

Discussion. The results presented show that the NN learning can yield improved models compared to the FT baseline in Tables 3 and 5, but outliers and non-normalised data groups present challenges. Good results were achieved by removing data from group 2 identified as outliers and training a standard NN, achieving 28.5% average error reduction. However, to include all data and to address the different AQF levels in different test groups, we devised the SOT method for training multi-task mappings. This method allows to make full use of the available data, and requires no prior recognition and removal of outliers. It is robust to different reference levels per group, and showed still good results even for extreme data items with 18.18% average error reduction (including the outliers).

6 Conclusion

Conventional fast methods of predicting tyre performance, such as the FT baseline used here, are based on physical models and heuristic hand-crafted formulae based on engineering experience. Our results show that by using machine learning it is possible improve the prediction performance of the currently used heuristic formula in a cross-validated test.

To this end we developed a set of features, that extract relevant information from the profile image based on engineering experience and inspection of the available data. The features developed, in particular the newly developed AUC, CSS and RAVLOMS seem promising. The machine learning experiments showed that outliers and non-normalised data are challenging problems for predicting aquaplaning. The SOT learning method presented here has proven to deal with these problems. It allows successful modelling without the need to detect and remove outliers and thus proving a quick and robust prediction.

The results of this initial study show that tyre performance prediction from profile images with machine learning is a promising direction for further research. Machine learning can potentially increase the speed and reduce the cost of tyre design significantly, as more accurate predictions can lead to more focused development and a reduced number of necessary tests.

References

1. Fwa, T.F., Kumar, S.S., Anupam, K., Ong, G.P.: Effectiveness of Tire-Tread Patterns in Reducing the Risk of Hydroplaning. Journal of the Transportation Research Board 2094, 91–102 (2009)
2. Jenq, S.T., Chiu, Y.S., Wu, W.J.: Verification and Analysis of Transient Hydroplaning Performance for Inflated Radial Tire with V-shaped Prototype Grooved Tread Pattern Using LS-DYNA Explicit Interactive FSI Scheme. In: ISEM-ACEM-SEM (2012)
3. Rumelhart, D.E., McClelland, J.: Parallel Distributed Processing: Explorations in the Microstructure of Cognition, vol. 1. The MIT Press, Cambridge (1986)
4. Bishop, C.M.: Neural networks for pattern recognition. Clarendon Press, Oxford (1997)
5. Baxter, J.: Learning internal representations. In: Proceedings of the Eighth International Conference on Computational Learning Theory, pp. 311–320. ACM Press (1995)
6. Argyriou, A., Evgeniou, T., Pontil, M., Argyriou, A., Evgeniou, T., Pontil, M.: Convex multi-task feature learning. In: Machine Learning (2007)
7. Chen, J., Liu, J., Ye, J.: Learning incoherent sparse and low-rank patterns from multiple tasks (2010)