Developing New Methods for the Automatic Measurement of Raveling at Traffic-Speed

S. McRobbie, A. Wright & J. Iaquinta
Technology Development & Quality Assurance Group, Infrastructure Division, TRL Ltd, Wokingham, UK

P. Scott, C. Christie & D. James
Highways Agency, UK

ABSTRACT: In the UK extensive use is made of the information provided by traffic-speed survey vehicles to assess deterioration, project condition and identify needs for maintenance intervention. Until recently these surveys have been unable to provide information about the presence of raveling (or surface disintegration), a defect related to aging and trafficking, and which is an important driver of maintenance intervention. However, significant progress has been made in automating the detection of this defect. This began with the development of algorithms to identify raveling on Hot Rolled Asphalt (HRA) using laser texture measurement collected in the nearside wheelpath, but the limited coverage provided by the single laser restricted the accuracy and robustness of the measure. Therefore the approach was expanded to use multiple profile laser data to identify raveling on HRA over the width of the traffic lane. A robust measure of raveling should ideally identify the defect on all types of surfaces. This paper describes work undertaken to develop and implement a surface-type independent measure to identify raveling at traffic-speed, using multiple line profile data. The research has also investigated the use of images collected at traffic-speed to provide information about the presence of raveling. Testing has been carried out to compare these measures with the results of manual surveys, and it is anticipated that this new approach will be implemented within the regular network-wide traffic-speed condition surveys of the trunk road network.

KEY WORDS: Raveling, fretting, automatic, monitoring, texture, image.

1 INTRODUCTION

Raveling (as shown in Figure 1) is the removal of aggregate and binder from the pavement surface due to age, weathering and traffic. This loss of surface material can occur at any time. However, on a well built pavement, the defect is more likely to occur late in the pavement’s life as a result of fatigue or weathering. Raveling primarily affects pavement longevity, by permitting water ingress, typically, but can also affect road users, their safety, comfort, noise perception, and is therefore a defect of particular concern to the road owner. On the motorway and trunk road network raveling is often considered as one of the main defects indicative of a deterioration of the road surface.

In the UK, the loss of surface material from HRA surfaces (in which chippings are applied to the hot-laid material and rolled in to form a skid resistant wearing course) is known as fretting, and starts to occur as the binder hardens and loses its cohesion or adhesion. On Thin

Surfacing (TS) the surface disintegration begins with a localized loss of aggregates, which accelerates as adjacent aggregates are dislodged from the surface as a result of trafficking. This deterioration is often referred to as raveling, and can result in the total loss of the surface layer. In the remainder of this paper we will refer to this loss of surface material as raveling, regardless of the surface type.

Figure 1: Example of raveling on a HRA surface.

To obtain best value in the operation and maintenance of a road network it is imperative that road renewal schemes are accurately prioritized to ensure that the most beneficial schemes go ahead at the right time. For this reason a reliable way of detecting raveling is highly desirable, to efficiently plan maintenance programs in a proactive manner.

The surface condition of the UK motorway and trunk road network is assessed by vehicles operated under the TRAffic-speed Condition Survey (TRACS) contract commissioned by the Highways Agency. These surveys report every 10m a range of condition parameters including longitudinal profile variance at different wavelengths, rut depth, intensity of cracking and texture depth, as well as information about road geometry. A simple assessment of raveling derived from single laser texture measurements in the nearside wheelpath is also reported in TRACS, but this value is not representative of the whole carriageway.

Consequently, it is still necessary today to undertake visual inspections in order to measure the extent and location of raveling. These inspections can either be carried out as standard Coarse or Detailed Visual Inspections (CVI or DVI), or as bespoke visual surveys looking in detail at a particular site. However, the visual methods provide highly variable results. This is due to a number of factors, including the inherently subjective nature of the inspections, the effects of lighting, the inspectors’ levels of training or fatigue, the environmental conditions, etc. Besides, these visual surveys are costly, time-consuming, and require traffic management. Hence it would be desirable to develop a more objective and less variable way to detect and report the level of raveling, using data collected during traffic-speed surveys such as TRACS.

2 DEVELOPMENT HISTORY

2.1 Stoneway algorithm

Several previous attempts have been made to automate the identification of raveling. In particular, van Ooijen et al., 2004, developed the Stoneway algorithm to detect raveling on Porous Asphalt (PA). In this algorithm the texture profile is inspected to identify deviations from the local mean texture depth. Where the deviations are deep enough and sustained over a long enough length, it is assumed that this indicates a loss of aggregate. In the algorithm the decision on whether aggregate has been lost is made by comparison with bespoke depth and length parameters. These were developed for PA which has relatively small aggregate size, and hence the depth and length parameters had relatively low threshold values. The algorithm was designed to use surface texture profile measured at a longitudinal spacing of 1mm collected with a single high resolution laser profiler.

2.2 Adaptation for HRA

In 2004 work was undertaken by TRL for the Highways Agency to identify raveling on HRA (which represented the largest proportion of road surfaces on the network at that time). The Stoneway approach was adapted for use on HRA surfaces by Wright (2004). The resulting HRA raveling algorithm again used texture data collected with a single laser profiler mounted in the wheelpath of the survey vehicle, and the parameters were tuned for HRA. The aggregates used in HRA being typically bigger than those used in PA, these values were selected to respond to larger deviations sustained for longer lengths.

2.3 Full lane width measure

The methods described above make use of information from a single laser located in the nearside wheelpath of the survey vehicle, hence raveling located elsewhere on the traffic lane was not detected. Work was therefore undertaken to adapt the method to cover the full width of the carriageway Although the lasers used to measure transverse profile on traffic-speed survey vehicles operate at a lower sampling rate than texture lasers, it was determined that these lower frequency lasers could provide a pseudo-texture measurement (at a longitudinal spacing of approximately 6mm) over several longitudinal measurement lines. Changes were made to the Highways Agency's HARRIS1 survey vehicle to enable the system to provide 25 lines of pseudo-texture measurement spread across the lane width. The algorithms developed for the measurement of raveling on HRA were then applied to this data to compute raveling values which were more representative of the entire width (McRobbie and Furness, 2008).

2.4 Limitations of current approach

The above sections detail the evolution of the automatic detection of raveling from a single line measure on PA, to a full lane width (multiple line) measurement for HRA. However, in each case the algorithms have been developed to target a particular surface type, and rely on parameters carefully chosen to ensure optimum performance on that surface. Therefore if the algorithms contained within the full lane width method are applied to non-HRA surfaces the results will be incorrect.

Over the last 10 years the proportion of the UK motorway and trunk road network surfaced with thin surfacing systems has significantly increased, and it has become necessary to have a method able to detect the presence of raveling on TS, in addition to HRA. For a practical measurement it is preferred that any automated measure of raveling for use on the network must be able to cope with both of these surfaces without producing unreliable or misleading results, or needing *a priori* knowledge of the surface type.

Research is being undertaken by TRL and the Highways Agency to develop an approach that would be independent of the surface type and able to identify the presence of raveling.

3 ENHANCED DATA SOURCES

The research has been carried out using data collected by HARRIS, the Highways Agency Road Research Information System. HARRIS has a downward-facing imaging system (at the rear) and a transverse profile measurement system (at the front), as shown in Figure 2.

Figure 2: HARRIS (the Highways Agency Road Research Information System)

3.1 LASER PROFILE DATA

The HARRIS1 transverse profile system comprises 25 profile lasers, one every 150mm across the vehicle, operated at 16kHz. Raw laser data collected at 100mm spacing longitudinally is used to characterize pavement shape and to assess rutting. However, a higher level of detail is required for the measurement of raveling. Therefore the system has been modified to provide height measurements from each of the 25 lasers at 6mm longitudinal spacing.

3.2 IMAGE DATA

The HARRIS image collection system utilizes three linescan cameras, mounted side by side in an enclosure, to capture images of the nearside, middle and offside of the pavement over a survey width of 2.9m, with a resolution of 2mm/pixel. The images are illuminated using an array of halogen lamps. Although every effort is made to remove variability when collecting images, they do suffer from some effects of uneven lighting, and the images from the three camera exhibit different grey level values. However, the pattern of variation is maintained longitudinally through successive images and can be characterized so that it is possible to overcome uneven lighting issues. By calculating the average pixel intensity at each location in a large set of images, changes in lighting can be quantified and the corresponding variability curves used to correct the image intensity and remove the effect of uneven lighting. Example HARRIS images are shown in Figure 3.

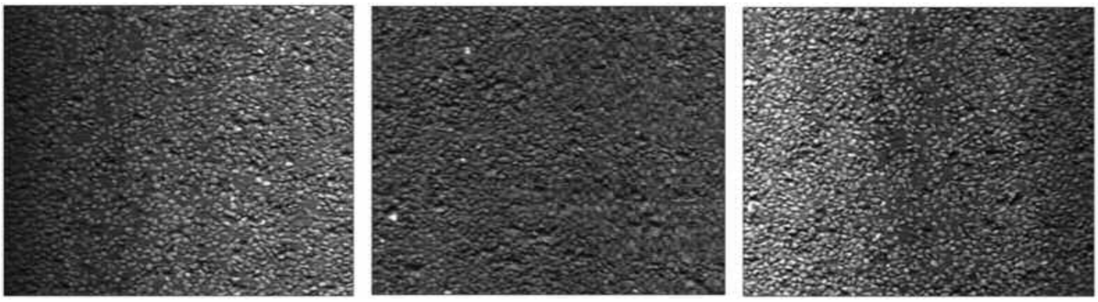

Figure 3: Example of nearside (left), middle (centre) and offside (right) HARRIS images.

These images illustrate the effects of uneven lighting and the levels of detail present. They also provide an indication of the ‘flattening’ effect which is a feature of the way in which the images are illuminated and collected. This flattening makes it hard to visually identify defects such as ravelling.

4 PROCESSING METHODS TO IDENTIFY RAVELING

In developing a new surface-independent measure of raveling we have revised the approach (in comparison with that described above) such that the algorithms do not attempt to identify specific lengths containing raveling. Instead the algorithms attempt to identify lengths that exhibit different texture characteristics than their surroundings, which should indicate the presence of a surface defect such as raveling. Hence it has been assumed that raveling is a localized defect, and it is further assumed that the presence of raveling affects the measured surface texture. By comparing the surface texture in a short length of pavement to that over a longer length of pavement we should be able to identify raveling.

In the research, methods have been developed to process both the pseudo-texture profile data and images to identify raveling. The processing methods developed for both data types have used similar approaches, in that the characteristic values calculated over a localized short length are compared to the characteristic values calculated over a longer length of the road (which surrounds the short length).

4.1 Pseudo-texture data

The pseudo-texture obtained from the profile lasers were filtered using a 2.5-100mm Finite Impulse Response band pass filter to remove short and long wavelengths. The Root Mean Square Texture (RMST) was then calculated from this filtered data over 100mm intervals. The distribution of these RMST values was calculated for each 10m length, and also for each 100m length. Comparison between the local (10m) values and the surrounding global (100m) values enables 10m lengths that have very different texture characteristics from the pavement around them to be identified. For example, we may expect a length of HRA containing raveling to exhibit a larger number of high RMST values than a length that does not contain raveling. Figure 4 shows typical histograms obtained from a short length of HRA containing raveling (blue), located in a longer length that is generally sound (green).

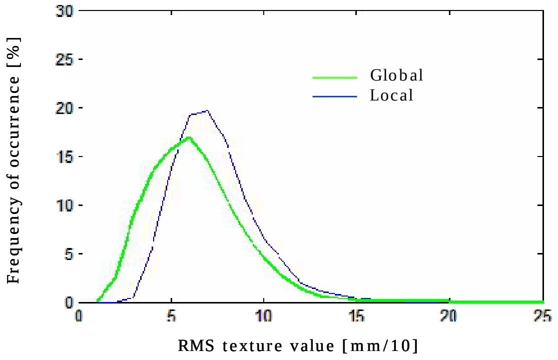

Figure 4: Histograms corresponding to a local length (blue) containing raveling, and a sound global length (green) of HRA.

By using the global texture as a baseline against which the local texture is judged, the method is self calibrating, and it is not necessary to define fixed parameters within the algorithms to detect raveling on different surface types. However, to develop a measure of raveling it was necessary to define a metric which would quantify the difference between the local and global texture histograms. The following three parameters were investigated.

Correlation parameter - The correlation coefficient between the two distributions is calculated. If the distributions are similar, as occurs when the local and global textures are alike, the correlation coefficient is close to one. If there are significant differences between the two distributions the correlation coefficient decreases. Obviously the correlation coefficient drops when the local texture is either rougher due to raveling or deterioration, or smoother because of patching or surface changes, than the global texture. Also, if there is extensive raveling in a length such that the global value has started to reflect the raveled nature of the pavement the local raveled texture values will look the same as the global raveled texture values and the correlation coefficient will be close to one. However, this is a practical indication and can be used as a first means of determining that the local texture has changed.

Texture ratio **-** The second comparison is based on the fact that ravelled areas tend to have longer tails in the histogram due to higher levels of RMST when aggregates are lost from the surface of the pavement. This method calculates the difference between the mean of the local texture and the 90th percentile value of the local texture, and the same difference for the global texture. The ratio of these two differences is used to determine whether or not the local texture exhibits more high values than its global counterpart**.**

Transverse variability - The third approach compares the RMST values reported by each laser with all the RMST values collected in the local area to find if any of the data collected by the individual lasers shows significantly higher texture. This should provide an indication about the transverse distribution of raveling across the carriageway. To achieve this, the histogram corresponding to all the data in the 10m length is subtracted from that of each laser within the same 10m length. Positive residual values relating to RMST texture greater than 1.5 times the mean of the entire 10m distribution are cumulated. This provides a dynamic way of overcoming some of the problems posed by the correlation parameter in presence of long lengths of raveling. Even though the global texture may reflect the presence of raveling there may be localized transverse differences in the RMST texture distributions.

Finally, a combined parameter was defined which considers all the above parameters to determine a single output parameter. To calculate the combined parameter all of the above parameters are firstly evaluated. If all report some degree of raveling then the value of the 'correlation parameter' is reported as the combined parameter. Otherwise the reported combined parameter is zero.

4.2 Image data

As with the profile pseudo-texture data the images were processed to look for discrepancies between the local and global characteristics. The goal was to develop one or more metrics that could be used to infer the presence of raveling. The following two parameters were found to be most promising.

Entropy - The entropy of an image is a statistical measure of the randomness of the pixel values within the image. An image which has a lot of 'texture' has a high entropy, while one with a low texture, containing few variations in pixel values, has a low entropy. The entropy is found by first calculating the histogram counts for the image. This counts how many of the pixels (*p*) in the image had each of the *n* possible pixel values. The entropy is then defined as:

$$E = -\sum_{0}^{n-1} p \log_2(p) \quad (1)$$

where *p* contains the histogram counts, and n is the number of grey levels in an image.
It was expected that images containing more raveling would have higher entropy values than those without raveling, and that by looking for regions where high entropy values were found the locations of the raveled parts of the network would be highlighted.

Edge density - In an image of a pavement an edge detector tends to detect the outlines of the aggregate within the binder. It was expected that an image of a raveled area of pavement would look different from an image of a pavement constructed of the same surface material which was in good condition, and that this could be indicated by the number of edges present within the images.

Among all the approaches available in image processing to detect an edge within an image, the Canny (1986) edge detector was chosen for this implementation because of its simplicity and its high level of performance.

5 PERFORMANCE ASSESSMENT

The texture-based and image-based parameters developed were considered either individually or combined together, and compared against reference data obtained on surfaces including HRA and TS and exhibiting different levels of raveling. It is very difficult to obtain absolute ‘truth’ data on the precise location and intensity of raveling. However, two methods were employed to provide reference data, the primary and secondary reference methods.

The primary reference data was collected via a walked manual inspection carried out on several test sites, where the inspector reported the location and severity of the raveling between 0 (no raveling) and 3 (a significant amount of raveling or severe raveling) at longitudinal intervals of 5m. However, it was not possible to collect this data on all of the test sites. The secondary reference data was obtained via manual analysis of the images collected by HARRIS, where an inspector assessed the raveling in $0.5\times0.5\text{m}^2$ squares overlaid on the images, using a scale from 0 (no raveling) to 3 (large amounts of raveling). Again, this data was not available on all of the test sites.

It has to be noted that the two types of reference data did not necessarily provide identical ratings, which introduces a degree of uncertainty in the results. In most cases it was felt that the site inspection provided a more robust reference measurement.

5.2 Results of the texture-based approach

Figure 5 shows an example of the measurements obtained over a 100m length of one of the test sites. The primary reference data (onsite survey) obtained over the 100m is shown as the left of the three vertical displays, the secondary reference data (manual analysis of images) in the middle, and individual RMST values over the 100m length to the right. For these displays the color scale is from dark blue (low RMST or raveling) to red (high RMST or raveling). The plots to the right of strips compare the distribution of RMST values for the global (green) 100m lengths with local (blue) 10m lengths.

Figure 5 a-d corresponds to an unraveled location where there are no strong areas of colour in any of the three vertical plots. The two histograms are almost identical, indicating that the distributions of texture values are almost the same in the local and global areas. Figure 5 e-g show a length containing raveling where clear areas of raveling shown in red and yellow can be seen in the primary reference data. These areas match with the secondary reference data. The RMST values show similar patterns, with 3 distinct areas (circled in purple on the graph). Compared to the global histogram, the local histogram seems shifted to the right. This means that locally there are proportionally more high texture values than there are in the surrounding area and hence this 10m length is very to have a higher level of raveling.

The performance assessment was undertaken by calculating the severity of raveling every 10m for each parameter described in section 4.1, above. For comparison with the reference, the 99th, 95th and 70th percentile values of each parameter reported over the whole data set were calculated. These were used as thresholds to define category 1 (below 70th percentile), category 2 (between 70th and 95th) and category 3 (above 99th) deterioration. The reference data were compared against both the individual parameters, and the combined parameter. The combined parameter exhibited the most notable agreement, and is discussed below.

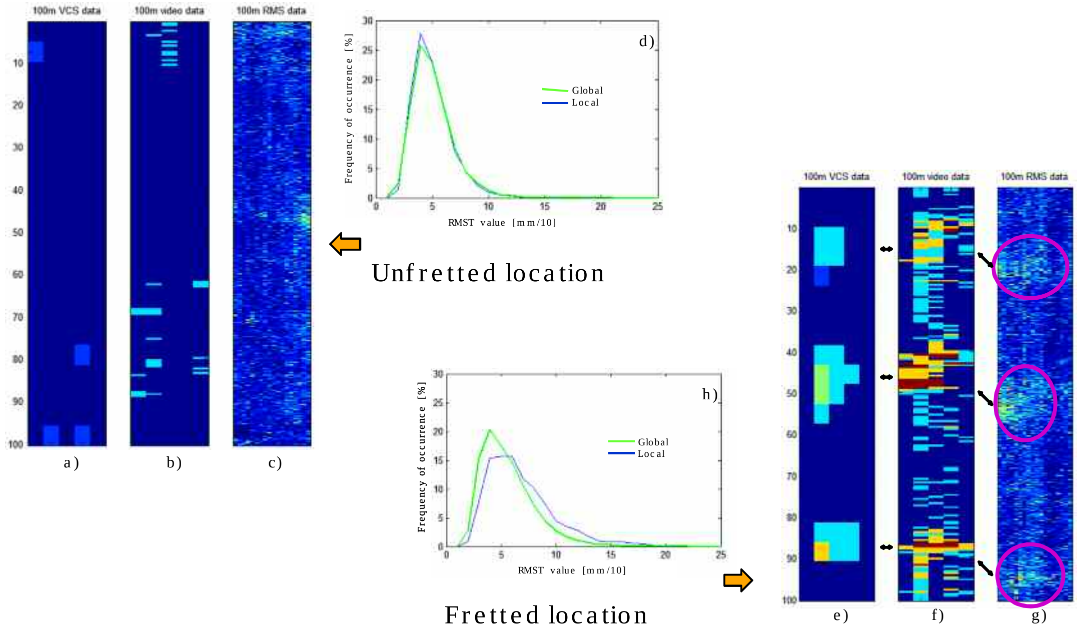

Figure 5: Example of data from 100m length of a test site: (a & e) output from onsite surveys; (b & f) manual analysis of images; (c & g) individual RMST values; (d & h) RMST distribution in global and local lengths for an unraveled site (on the left, a to d) and a raveled site (on the right, e to h).

Figure 6 compares the primary reference data with the 'combined parameter' (see §4.1) on a 20km site, primarily surfaced with HRA. It can be seen that there is reasonable agreement between the datasets in terms of lengths containing generally high and generally low levels of raveling, although the agreement is not always good. When the comparison is extended to multiple sites we found that there were lengths of good agreement between the automatic and reference data. However, there were still locations with unexplained differences and research into this method is therefore continuing.

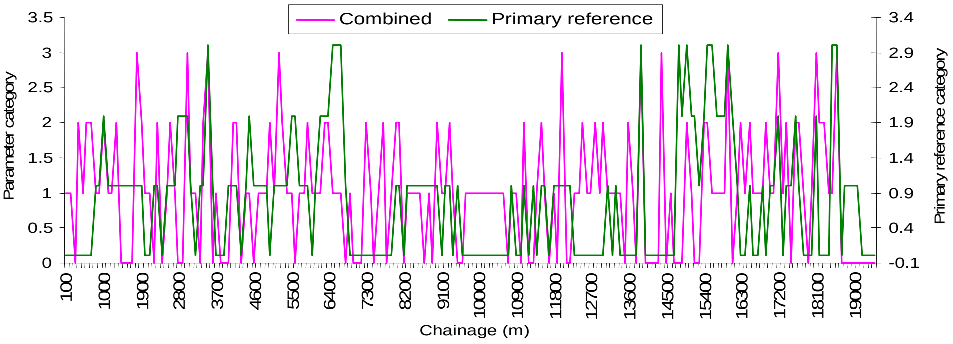

Figure 6: Comparison of the primary reference and the combined profile-based parameter.

5.3 Results of the image-based approach

Figure 7 shows examples of HARRIS images with (clockwise, starting top left) raveled HRA, raveled TS, sound TS and sound HRA. These images were split into an array of small grid cells and the entropy and edge pixel density were derived for each cell (these are displayed as smaller images to the right of the original image in Figure 7). On an image by image basis it is often possible to identify raveled images, and even in some cases the locations of the raveling within those images.

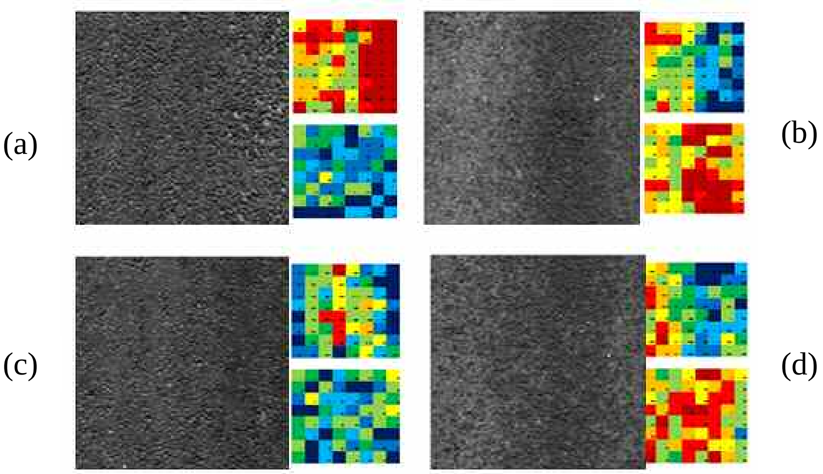

Figure 7: Example HARRIS images with (a) raveled HRA, (b) raveled (TS), (c) sound HRA, (d) sound TS. The smaller images alongside show how the entropy (top) and the edge pixel density (bottom) vary within each image.

Figure 8 compares the primary reference data (reported as sum of the intensities of raveling found in each length) with the automated image analysis (edge density and entropy, calculated by comparing the variations between the local edge pixel density and entropy parameters with the globally determined values, given a unitless measure scaled for display purposes). It can be seen from Figure 8 that the agreement between the image-based raveling parameters and the reference measurement is not consistent. The automated image analysis frequently fails to identify the raveling, although there are occasional lengths of agreement. There are a number of possible explanations for the poor performance of the image-based parameters. These include surface type changes and patches, road markings and debris or water on the road. However the single most important factor is that it is often difficult to determine the presence of raveling with any certainty even when inspecting the images manually, due to the single viewing angle and the lighting used in the collection of the images. These factors combine to make it very difficult to develop an automated method to detect a depth related defect such as raveling.

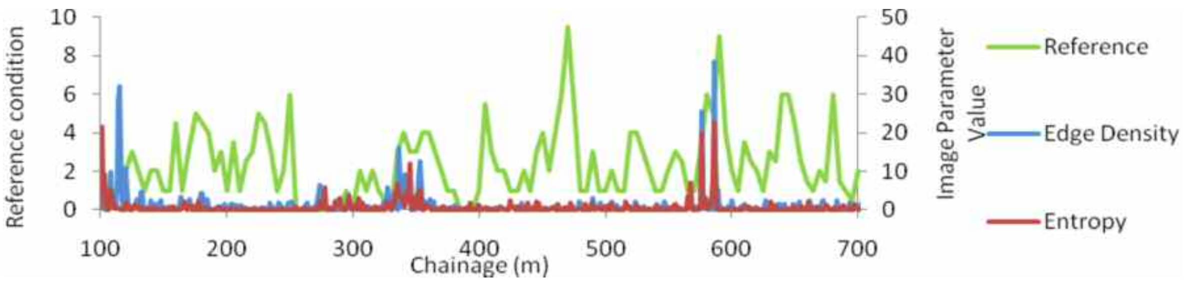

Figure 8: Comparison of primary reference data with image entropy and edge pixel densities over 600m of data.

6 CONCLUSIONS

Raveling is a significant defect affecting not only pavement longevity but also road safety, ride quality and noise. On the motorway and trunk road network in the UK raveling is becoming the main defect indicating deterioration of the road surface. It is therefore desirable to develop an objective method to detect and report raveling, using a technique that can be applied without the need for road closures. The information provided by such a method could be applied in the efficient planning of maintenance programs. The research described in this paper has built on previous work to develop methods that utilise texture profile data for the detection of raveling. The use of image collection and analysis was also investigated.

Texture-based parameters show considerable potential for the detection of raveling. However, a practical objective method to detect raveling should be able to be applied over the wide range of different surface types found on the road network. In this work we have developed an algorithm to detect raveling that utilises texture measurements obtained over the whole width of the traffic lane, in multiple measurement lines. To reduce the effect of surface type on the level of raveling reported, the method compares local texture characteristics with global texture characteristics, and hence does not require *a priori* knowledge of the surface type. Parameters have been developed to highlight ‘anomalous’, lengths likely to contain raveling, that could be investigated in more detail by highway engineers. Network level tests have shown that this is often successful, such that lengths reported to containing raveling in manual inspections are often reported to have a high value by the texture based method. However, the development is continuing, to further develop and refine the method. Work is also underway to turn the results of this research into practical tools, in particular by confirming the technical and practical requirements for routine collection of pseudo-texture data needed to implement the raveling detection algorithm. It is intended to introduce in the next Highways Agency’s TRACS contract a full lane width method to highlight lengths likely to be affected by raveling.

An investigation into the use of image-based methods to detect raveling has found that, although image-based parameters can be developed that identify raveling in specific images, these are less successful when applied over longer lengths of test data. It is suggested that imaging methods providing depth information (e.g. stereo vision) could be applied to resolve this problem, but these have not been investigated in this work.

REFERENCES

Canny, J., 1986. *A computational approach to edge detection*. IEEE Transactions on Pattern Analysis and Machine Intelligence. 8(6):679-698.

McRobbie, S., and Furness, G., 2008. *Automated detection of fretting on hot rolled asphalt surfaces*. TRL report PPR 299, TRL Limited, Crowthorne, UK.

Scott, P., Fairclough, R., Radband, K., Zohrabi, M., Sanders, P., McRobbie, S., and Wright, A., 2008. *Measuring surface disintegration (ravelling or fretting) using traffic speed condition surveys*. 7th International Conference on Managing Pavement Assets (ICMPA), Calgary, Canada.

van Ooijen, W., van den Bol, M., and Bouman, F., 2004. *High-speed measurement of ravelling on porous asphalt.* 5th Symposium on Pavement Surface Characteristics - Roads and Airports (SURF), Toronto, Canada.

Wright, A., 2004. *The use of surface texture measurements to identify fretting on hot rolled asphalt.* TRL report PR/INN/41/04, TRL Limited, Crowthorne, UK.