A GIS based approach for predicting pavement deterioration on the UK road network.

J. Corker, R. Remenyte-Prescott*, M. Eskandari Torbaghan** and J. Ninic**

* *Resilience Engineering Research Group, Faculty of Engineering, University of Nottingham, University Park, Nottingham NG7 2RD: jonathan.corker@nottingham.ac.uk, r.remenyte-prescott@nottingham.ac.uk*

** *Department of Civil Engineering, School of Engineering, University of Birmingham B15 2TT: m.eskandaritorbaghan@bham.ac.uk, j.ninic@bham.ac.uk*

Abstract. The U.K. Road network is a complex and dynamic system, managed by a federated network of organizations, often with stringent constraints on resourcing. Remediating deterioration of the network is both costly and politically sensitive; resilience, sustainability, and utilisation impacts must be constantly balanced by asset managing organizations. An approach for localized prediction of deterioration is developed in this paper, combining remote sensing and automated survey data into a GIS-based model to support decision making. In addition, predictive models are proposed in this paper, based on random forest regression and classification approaches, which use the proposed data model to create localized deterioration profiles. The machine learning model has been trained with and validated against data from a 21km length of the UK Major Roads network, using ~10 years of condition data and authoritative traffic, cartographic, and environmental data from several UK government agencies. Both the regression model and the classifier can accurately predict condition metrics, described in industry standards. This approach allows for early detection and mitigation of pavement failure and support maintenance operations on the network whilst minimizing disruption and maximizing return on investment.

1. Introduction

In England, for areas outside of London, management of road infrastructure is carried out at two tiers: the Strategic Road Network (SRN) managed by National Highways, and the Local Road Network (LRN) managed by local authorities. Whilst the SRN mostly comprises roads of relatively modern construction, acting as arterial trunk roads between towns and cities, the LRN represents a more varied asset and comprises anything from minor residential roads constructed over 100 years ago to major connections within a region. To allow more effective funding for these intra-region networks, a middle tier of local authority-maintained roads, the Major Road Network (MRN), was established in 2017 by the U.K. government with and additional £3.5Bn. of funding allocated for improvements and capacity increases for these roads (Department for Transport, 2018).

Following the work of (Gong *et al.* 2018) and (Chen *et al* 2022) this paper looks to investigate the use of Random Forest approaches for prediction of pavement deterioration on the UK MRN and identify key datasets that can be used to support the prediction process within the context of the UK LRN. Novel use of authoritative spatial and contextual data alongside flexible predictors such as the Random Forest algorithm could support the predictive process in lieu of hard-to obtain information such as construction, materials, and historic maintenance. Part one of this paper acts as an introduction to UK Highways management and a primer on Random Forest algorithms. Part two provides details on data selection, preparation, and model configuration; the paper concludes in parts three and four, where results are discussed, and the conclusion placed in context.

1.1 *Condition and serviceability metrics for pavements*

Pavement roughness, measured as a function of vertical change relative to longitudinal distance, is a widely used indicator of pavement serviceability. High or changing roughness values (commonly described using measures such as the International Roughness Index, or IRI) indicate a poor ride quality, particularly for high-speed roads, and changes in roughness may indicate structural defects within the pavement (Adelinge and Gupta, 2013). Automated road surveys, which commonly capture roughness alongside several other condition metrics are commonly used by highways management organisations worldwide to assess and monitor the condition of their networks (Attoh-Okine and Adarkwa, 2013). These traffic-speed surveys may be undertaken alongside other survey techniques (such as visual inspections) to provide a more complete image of the road network over time and provide a quantitative measure of the pavement condition.

Since 2006, U.K., Local Authorities (LAs) have had a statutory obligation to undertake automated surveys of their network under the Surface Condition Assessment for the National Network of Roads (SCANNER) specification (Department for Transport, 2021). SCANNER surveys capture a range of information, including road geometry, transverse road profile, pavement texture depth, and roughness. Roughness is reported as Longitudinal Profile Variance (LPV) as 3m and 10m averages, not IRI, which complicates direct comparison of UK pavement data with pavement performance datasets, such as the Long-Term Pavement Performance (LTPP) database. Finally, for each surveyed 10m section of pavement an aggregated and weighted condition indicator – the SCANNER RCI - is calculated and reported by the local authority as a key performance indicator for the maintenance operations of the LA.

The output of the SCANNER RCI calculation is a number score between 0 and 315; due to the nature of the weighted calculation, a pavement with only slight defects may score 0, or several minor defects may compound to indicate relatively severe overall deterioration. For operational convenience, the SCANNER RCI score is categorised to indicate the urgency with which maintenance is required; scores below 40 are considered in good repair (‘green’), scores between 40 and 100 should be considered immediate maintenance candidates (‘amber’), and scores over 100 are considered in poor condition and likely to need immediate maintenance (‘red’).

1.2 GIS data and deterioration factors.

Flexible pavements (e.g., pavements constructed of asphalt and granular base materials) have several modes of failure. Once the waterproof surface layer of a pavement is breached, water ingress can cause rapid structural failure and disintegration such as potholes (Adelinge and Gupta, 2013); age can cause materials to degrade and become porous and simple wear can cause fretting and cracking. Climate effects such as freeze-thaw can cause delamination and ravelling, and deterioration can quickly spread points of weakness such as edges and joints (Al-Omari and Darter, 1995). These kinds of defects can be quick forming, costly to repair, and represent a safety risk to users of the road network. Datasets such as weather, geographic context, construction and maintenance history, and traffic profiles may help to target proactive interventions.

1.3 *Random Forest Classification and Regression*

Among common machine learning and inductive inference approaches, decision tree (and random forest of decision tree) algorithms are amongst the most widely used due to their relative ease of interpretation, their ease of configuration, and their ability to handle complex relationships between inputs and outputs. Whilst an individual decision tree is prone to overfitting and has a poor tolerance of outliers, using multiple decision trees acting in parallel, working on partitioned training data, and using an aggregated output, can improve overall accuracy and improve the tendency to overfit at the cost of interpretability. Random forest approaches can be used for either regression or classification without fundamentally changing the underlying algorithm; the key difference is at the final aggregation stage, where for example either majority vote (for classification) or mean value (for regression) could be used to predict a target variable (Liaw and Weiner, 2002). (Gong *et al.*, 2018) successfully used a random forest approach to predict IRI from the construction and condition data, and in their systematic review of ML approaches in civil engineering Chen (Chen *et al.*,2022) notes the potential for random forests for pavement prediction, and suggests key inputs to the predictive model.

Figure 1. Location plan indicating the A6097 area of study.

2. Proposed Methodology

2.1 Pavement and traffic data selection

Chen (2022) suggests that the key inputs for predictive models for pavement are existing condition, materials and construction, climate, and traffic characteristics. Definitive construction details for existing pavements can be difficult to determine, particularly on the LRN where roads may have been constructed using locally available materials and historic design specifications. Authoritative condition data for UK local roads is the purview of UK local authorities, and Nottinghamshire County Council provided SCANNER condition data for their network covering 2007 – 2021 in native UK Pavement Management System Interchange format. To allow for detailed analysis, a subset of this data was used covering one part of the Nottinghamshire Major Road Network.

The A6097 comprises an approximately 22 km long section of single- and dual- carriageway road, approximately 9km East of Nottingham, crossing the Trent Valley, and acting as a key internal artery for the county of Nottinghamshire. Much of the road was constructed in 1932, with successive modernization work undertaken to-date, most recently with major junction improvements in 2009 and with further upgrades to the road proposed as part of the MRN inaugural phases of work. Detailed construction and material records are unavailable, but general changes in construction methodology between sections of road have been captured by the local authority in their network hierarchy. In addition, traffic volumes (taken from U.K. Office of National statistics data) are consistent along the road's length. A location plan is shown in figure 1. Speed limit data was provided by Nottinghamshire as part of their network model and has been taken validated against speed limit signage.

In the U.K., traffic volume and character from automated traffic counts are reported by the Office of National Statistics (ONS). As part of the proposed modernization work, survey data from 2012 has been projected forwards to provide current and predicted traffic volumes for the northern and southern halves of the road and traffic volumes, reported as Average Annual Daily Traffic (AADT) and Average Annual Daily Flow (AADF), are relatively consistent along the road's length.

Table 1. List of datasets and inputs to the predictive model.

Dataset	Source	Metrics
Network Model	Nottinghamshire County Council	Speed limit, lane count
Traffic Volume	U.K., Office of National Statistics	AADT, AADF, and %HGV
Major maintenance works	Nottinghamshire CC Major schemes, 2007 - 2021	Resurfacing and realignment of carriageways.
Road condition data	Nottinghamshire CC SCANNER data, 2007-2021	Pavement condition and geometry
Topography	U.K., Environment Agency Digital Terrain Model (DTM)	Elevation (m AoD)
Topographic Position Index	Derived from DTM data (Wilson *et al.*, 2007)	Descriptive statistics (mean, median, maximum, minimum)
Terrain Ruggedness index	Derived from DTM data (Riley *et al*, 1999)	Descriptive statistics (mean, median, maximum, minimum)
Slope angle (degrees)	Derived from DTM (Zevenbergen and Thorne, 1987)	Descriptive statistics (mean, median, maximum, minimum)
Rainfall catchment	Derived from DTM (O'Callaghan and Mark, 1984)	Descriptive statistics (mean, median, maximum, minimum)
Geomorphic context	Derived from DTM (Jasiewicz and Stepinski, 2013)	Total coverage (m)

2.2 Geography and geomorphology data selection

The Ordnance Survey (OS) provide an authoritative source for U.K. mapping data via their statutory relationship with U.K., Local Authorities. The OS provides detailed topologic and network maps for the whole of Great Britain on a commercial basis, with licenses available for research purposes on a non-commercial basis. Due to the rural nature of the A6097, area coverage of permissively licensed alternatives such as OpenStreetMap was found to have comparatively lower levels of detail than OS alternatives, missing information such as road widths and minor junctions. OS Highways data also correlated more closely and completely to the surveyed data than either openstreetmap or the LA's own network hierarchy. Distance to the nearest junction on the carriageway was calculated from the OS Highways dataset.

A comparison of available Digital Terrain Models (DTMs) for the area of concern revealed that U.K., Environment Agency (EA) aerial survey data was both the highest available resolution (1m) and cheapest available. As part of their flood monitoring program, the EA undertakes ongoing aerial surveys of the route of the river Trent and its floodplain which the A6097 crosses in its southern extent. Permissively licensed DTM data was available from the OS, but at much lower resolution (10m), and higher resolution survey data was prohibitively costly or license-constrained along the full length of the route. A full list of datasets and metrics is included in table 1.

Underlying geology along the road length was considered as an input, and geological data from the British Geological survey was checked. In its southern and central areas, the road is mostly underlain by the Mercia Mudstone group, comprising weak interbedded sandstones and mudstones. In the northern extent, the road sits on the Sherwood sandstone group, a strong sequence of red sandstones. In both cases, the competent underlying bedrock is unlikely to be a source of failure that would be registered by the SCANNER survey. In the central floodplain area around the river Trent, the road is built onto an engineered embankment that rests on fluvial sands and gravels (the Trent Valley Formation). As with the northern extents, whilst mass movement and flooding may represent a risk to the road, these issues are unlikely to present in an immediately apparent way in SCANNER survey data. As such, the underlying geology was not included in the predictive model.

2.3 Data Preparation

SCANNER condition data was provided in U.K. Pavement Management System (PMS) HMDIF format. A loader was written using the Python programming language to parse the files, extract records relevant to the road of concern, and load the result into a Postgres database. Start and end points for each survey section and for each for each survey undertaken was generated from the HMDIF data, and a survey path generated by connecting the start and end points for each survey chainage.

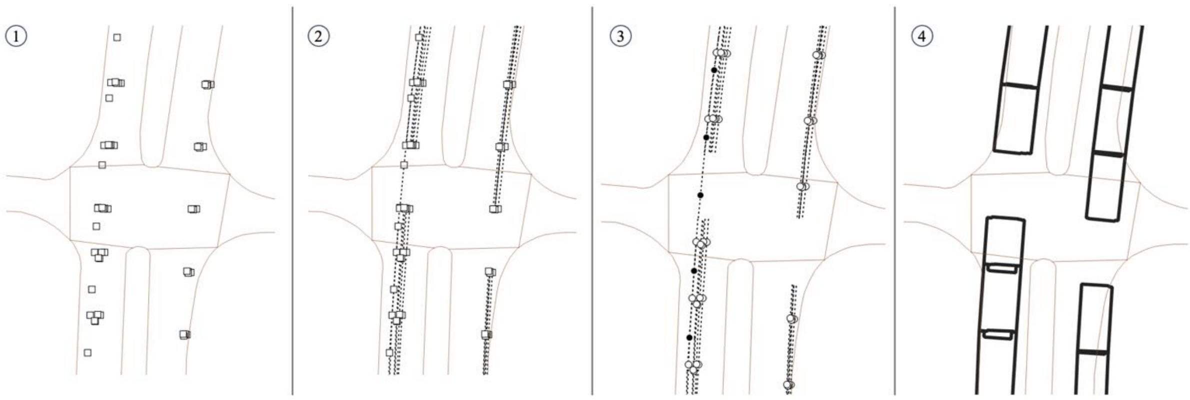

Figure 2. Extraction and positioning of survey locations from GPS tracks.

Centroids for each survey path were then calculated, and cluster analysis using the DBSCAN (Khan et al. 2014) algorithm (minimum cluster 3 points, minimum distance 2.5 meters) was used to filter outlier survey points that may have been the result of survey error or junction realignment. Survey path lines were then spatially buffered by 2m to create survey event polygons with a 10m length, resulting in 18,108 survey event polygons covering the full length of the road and containing approximately12 years of condition data. The stages of this process are illustrated in figure 2. SCANNER RCI values and contributing condition parameters were calculated following the SCANNER specification (UK Roads Board, 2011).

The DTM was used as source data to calculate geomorphological terrain characteristics and to calculate descriptive statistics for each survey event polygon – a full list of input parameters is given in table 1. Each quantitative value was then normalized with respect to other records in the dataset using the minimum and maximum values of that variable. The resulting dataset comprises 12348 records in 3242 clusters, recording time since previous survey, change in condition since previous survey, and ground context for a particular cluster of condition data.

2.4 Model development and training.

Decision tree regressors and classifiers from the python scikit-learn package (Pedregosa *et al.*, 2011) were used to build the Machine learning models. For each training cycle, 70% of the dataset was used for training the model, 20% for testing the model, and 10% was set aside for model validation.

Separate models using the same baseline data were trained to calculate regression models for each of roughness, cracking, transverse profile, texture depth, and Scanner RCI. In addition, simple linear regressions were calculated for each curve to allow for a baseline to assess relative performance.

To evaluate the predictive performance of the random forest classifier, the SCANNER RCI scores for each survey event were assigned their appropriate condition category (Red, Amber, Green). The time interval to the next survey within a cluster was then established for each record, and the condition category at the time of the next survey was added as a target variable for the classifier. This approach looks to establish whether, given a condition category at time ‘a’, can the classifier predict the condition category at a future time ‘b’.

3. Results

Table 2. Performance of random forest regression on the SCANNER deterioration dataset.

	Linear regression, mean $R^{2,}$	Random Forest Regression, mean $R^{2,}$ (RMSE)	
		Training	Test
Texture Depth	-19.34	0.66	0.64 (6.4)
Cracking	-20.49	0.45	0.34 (22.8)
Transverse Profile	-19.5	0.61	0.55 (7.3)
Longitudinal Profile	-19.27	0.68	0.64 (4.9)
Scanner RCI	-19.35	0.78	0.75 (11.2)

3.1 Regression results

The random forest regression outperformed the linear regression for all SCANNER parameters, although it should be noted that the linear regression is an exceptionally poor fit overall with negative average R^2 values for all condition metrics, indicating that a linear regression was a worse characterization than simply taking the average condition value over the survey period. Many sections (45%) exhibit no deterioration at all over the survey period or exhibit little change (+37%). Of key importance is the starting condition; if a section of pavement is in good repair, it tends to remain in good repair, and if a section has begun to deteriorate, the deterioration rate increases; this is consistent with our understanding of pavement behavior, where one defect represents a point of weakness from which further deterioration can spread. The plot of actual vs. predicted Scanner RCI values, based on results from the test dataset, is shown in Figure 3; The plot shows significant heteroskedasticity for low RCI values, which makes sense when considering that the underlying data comprises significantly nonlinear time series values, where each value is contingent on its predecessor.

In some cases, conditions on a section improved, which may be either due to spurious detections during survey, or spot-maintenance that was not captured in the major remediation schemes dataset. This is of particular concern for cracking, which has a poor regression result; indeed, the contribution of cracking to the overall SCANNER RCI has reduced weighting relative to other parameters to reflect this disparity (UK Roads Board, 2011).

Relative sparsity of data may also have been problematic for the regression model more generally (the average number of datapoints per section is 5), and aggregation of sections into comparably performing groups may significantly improve the quality of prediction by increasing available data. Longitudinal profile variance (LPV) and texture depth performed the best of the individual defects; as noted above, LPV is a compound metric that may reflect several underlying issues or construction defects and tends to remain static for adversely affected sections. Texture depth is perhaps the metric most reflective of general 'wear and tear', as vehicle traffic effectively polishes the road surface, and correlation with other condition metric such as surface coefficient of friction may be insightful. A summary of the results for all SCANNER parameters, including the aggregate RCI, is shown in table 2.

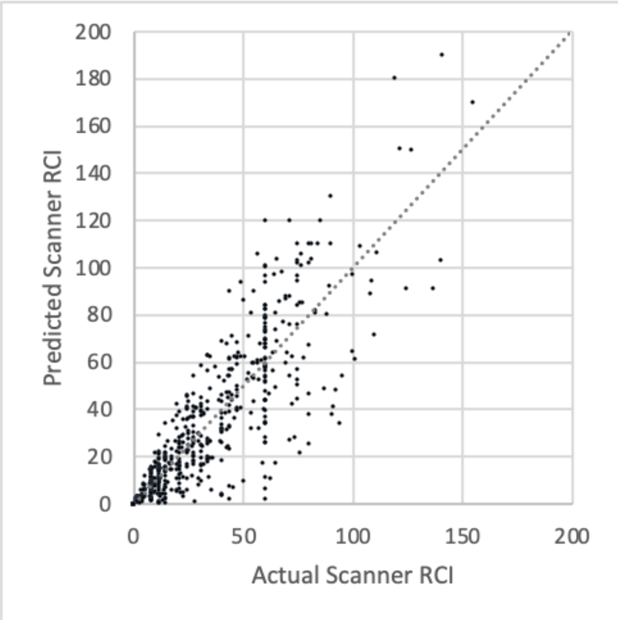

Figure 3. Actual condition vs. Predicted for the SCANNER RCI, from the test dataset.

3.2 Classifier results

Given a start condition and a duration of time to the next survey, the classifier correctly predicted the future condition category at a location in all except 4 cases from the total 2967 in the test dataset. In two of those instances the condition improved – indicating either a repair or an incorrect detection on previous readings – and in one instance the result was borderline (value of 98.4 vs. predicted condition 'Red', which has a 100 threshold). It should be noted that changes in condition are relatively uncommon (which is reflected in the results), and the average survey interval in the dataset is 2.7 years – meaning deteriorating pavements are likely to be in significant disrepair at the time of the second survey. As discussed in the methodology, records affected by major maintenance (e.g., resurfacing) were removed from the dataset, and it is logical that most significantly deteriorated areas are most likely to have major interventions.

Table 3. Confusion Matrix showing the performance of the random forest classifier on the on the SCANNER deterioration dataset.

		Predicted section condition		
	as	Green	Amber	Red
Actual section condition	Green	2673	0	0
	Amber	0	283	1
	Red	0	3	7

4. Conclusion

Using the above methods, it has been shown that data-rich models incorporating GIS data can predict deterioration across several parameters, and a method for correlating condition metrics tied to the logical network back to the physical network has been demonstrated. The capability of Random Forest approaches to model both regressions and classifications has been demonstrated. The ease of deployment and ready availability of this model makes it potentially applicable to highways management organizations with comparable datasets, particularly UK local authorities with extensive major rural road networks as a low-cost supplement to their existing systems; however additional training will be required where there are significant variations in traffic volume.

Future work is focusing on models with greater input flexibility such as graph-based models, and on sensitivity analysis to establish which factors correlate most to with deterioration. Most sections do not significantly deteriorate, and future work will consider a shift to identifying deteriorating sections (i.e., classification and outlier extraction) on a per-defect basis, rather than regression for all sections. In addition, incorporation of climactic factors, drainage, buried assets, and geotechnical features are likely to improve predictive model accuracy and will be considered in future approaches with more diverse road.

References

UK Department for Transport, 2018. ‘Major Road Network and Large Local Majors Programmes: programme investment planning’. UK Department for Transport.

Attoh-Okine, N. and Adarkwa, O., 2013. Pavement condition surveys–overview of current practices. *Delaware Center for Transportation, University of Delaware*.

UK Roads Board, 2011. ‘SCANNER Surveys for Local Roads-User Guide and Specification’. *UK Roads Board*.

Adlinge, S.S. and Gupta, A.K., 2013. Pavement deterioration and its causes. *International Journal of Innovative Research and Development*, *2*(4), pp.437-450.

Wada, S.A., 2016. Bituminous pavement failures. *Journal of Engineering Research and Applications*, *6*(2), pp.94-100.

Al-Omari, B. and Darter, M., 1995. Effect of pavement deterioration types on IRI and rehabilitation. Transportation research record, 1505, p.57

Department for Transport. *Technical Note: Road Condition and Maintenance Data*. Network Condition and Geography Statistics Branch, Department for Transport, London, 2021

Chen, K., Eskandari Torbaghan, M., Chu, M., Zhang, L. and Garcia-Hernández, A., 2022. Identifying the most suitable machine learning approach for a road digital twin. *Proceedings of the Institution of Civil Engineers-Smart Infrastructure and Construction*, *174*(3), pp.88-101.

Gong, H., Sun, Y., Shu, X. and Huang, B., 2018. Use of random forests regression for predicting IRI of asphalt pavements. *Construction and Building Materials*, *189*, pp.890-897.

Jasiewicz, J. and Stepinski, T.F. (2013) ‘Geomorphons — a pattern recognition approach to classification and mapping of landforms’, *Geomorphology*, 182, pp. 147–156.

Khan, K., Rehman, S.U., Aziz, K., Fong, S. and Sarasvady, S., 2014, February. DBSCAN: Past, present and future. In *The fifth international conference on the applications of digital information and web technologies (ICADIWT 2014)* (pp. 232-238).

Liaw, A. and Wiener, M., 2002. Classification and regression by randomForest. *R news*, *2*(3), pp.18-22.

O’Callaghan, J.F. and Mark, D.M. (1984) ‘The extraction of drainage networks from digital elevation data’, *Computer Vision, Graphics, and Image Processing*, 28(3), pp. 323–344

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V. and Vanderplas, J., 2011. Scikit-learn: Machine Learning in Python, Journal of Machine Learning Research, 12.

Wilson, M.F., O’Connell, B., Brown, C., Guinan, J.C. and Grehan, A.J., 2007. Multiscale terrain analysis of multibeam bathymetry data for habitat mapping on the continental slope. *Marine Geodesy*, *30*(1-2), pp.3-35.

Zevenbergen, L.W. and Thorne, C.R. (1987) ‘Quantitative analysis of land surface topography’, *Earth Surface Processes and Landforms*, 12(1), pp. 47–56.