ID: 0004

Probabilistic Road Weather Forecasting

L. Chapman [1]

[1] The University of Birmingham,
Birmingham, UK

Corresponding author's E-mail: l.chapman@bham.ac.uk

ABSTRACT

The last 30-40 years has seen a step-change in the provision of forecasts to the winter road maintenance sector. The days of simple 'road danger warnings' are now a distant memory and highway engineers have had access to detailed site specific forecasts for over 30 years. More recently, the advent of route based forecasting has added increased detail and sophistication to the approach and further helps to strike the balance between securing winter roads and optimising salt usage. Any improvement to road weather forecasting has ultimately been driven by computational resources. As a result, road weather forecasts have historically been deterministic in nature. However, the potential now exists, as with other sectors, to potentially embrace probabilistic forecasting. All forecasts needs to be accompanied by credibility intervals of some sort to indicate the level of uncertainty and road weather forecasts are no different. The use of a Bayesian framework incorporating an emulator of an existing numerical model is one such means by which this can be achieved and could potentially bring greater confidence to decision making – *providing the sector requires and will use such information*. Furthermore, the superior computational speed of an emulator will allow models to be run for a greater time period into the future, possibly years.

Keywords: probabilistic forecasting, route based forecasting, decision support system

1 DETERMINISTIC ROAD WEATHER FORECASTING

Road weather forecasts are an integral component of road weather information system (RWIS). However, with a small number of notable exceptions, forecasts have historically been deterministic in nature. Such a forecast was a giant step forwards in the 1980's as previously most of the decision making was made based on ice detection sensors embedded into the road [1]. The move to ice prediction was prompted by ever-improving computational resources which enabled simple 0D energy balance models to downscale the regional weather forecast to predict road surface temperatures for the following 24 hour period [2]. Typically this was done for ice detection sites, which then became forecast sites, and subsequently allowed for some verification of the forecast. When the system was first introduced, two forecasts were issued (Figure 1). These were based on worst case and best case scenarios for a particular night and were termed pessimistic and optimistic respectively [3]. Two scenarios were not thought to be an appropriate level of detail, but ultimately was all that could be realistically achieved as at this stage as it typically took 15 minutes of runtime to produce each forecast!

However, due to a lack of confidence in the new technology, engineers tended to just base their decisions on the pessimistic forecast and were significantly over-salting the road network. This was clearly unacceptable and so a 'best stab' approach (producing a single forecast curve for decision making) was utilised which is still commonly used even to this day.

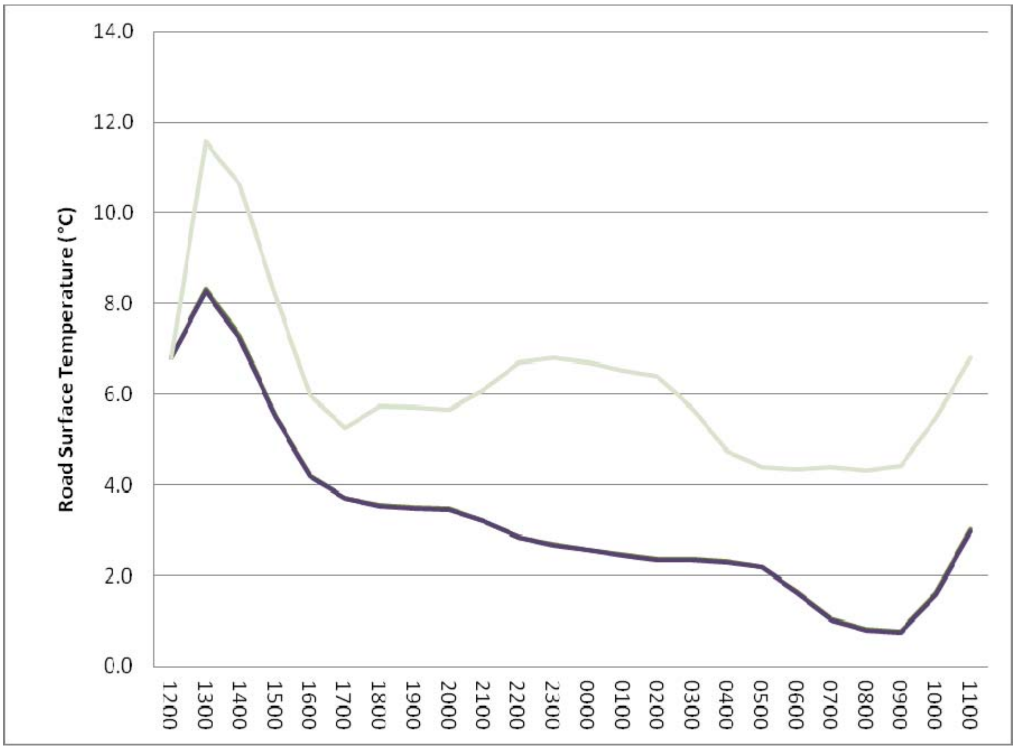

Figure 1. The original probabilistic road weather forecast?
In the early 1980s, two deterministic forecasts were issued, one pessimistic and one optimistic.

Despite ever-increasing computational power (Moore's Law: Figure 2), very little changed in the deterministic modelling of road surface temperature. Whilst there was considerable ongoing research and technological development in the sector, the focus of research was in improving model performance, particularly with respect to existing parameterisations (e.g. traffic, cloud, road construction etc). Indeed, the biggest improvements in the modelling did not occur in the downscaling process, but instead in the forecast models used to provide the input data for the energy balance model. There have been enormous advances in the science of numerical weather prediction over the last few decades. For example, the UK Met Office UKV model can now provide mesoscale output at a 1.5km resolution. This is pushing available computation to its limits and has significant benefits for road weather forecasting as this resolution now enables major local topographies to be resolved.

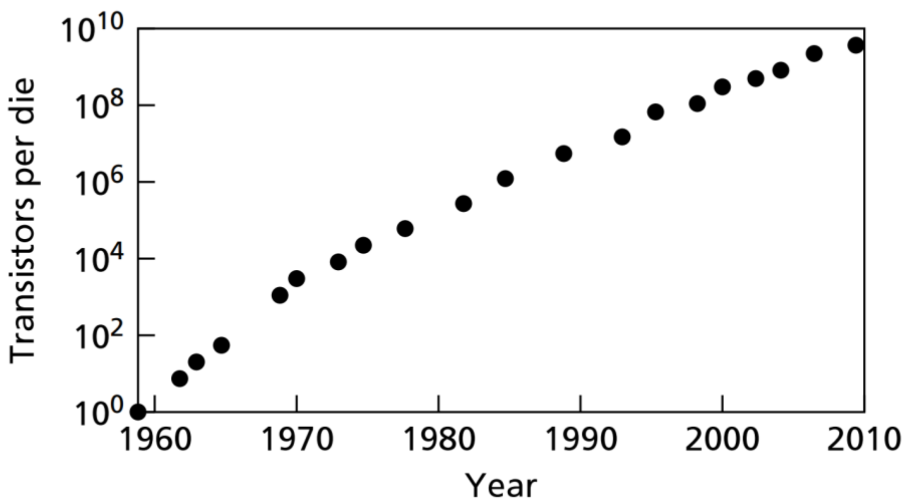

Figure 2. Moore's Law [4]

With respect to downscaling, the road weather community has only recently begun to take increased advantage of the available technology and computation. The best example of this is the shift by local councils and forecast providers across northern Europe to route-based forecasting techniques (e.g. UK Met Office: [5]; Meteogroup: [6]). The original route based forecasting model was the 'neXt generation' XRWIS [7] developed at the University of Birmingham. XRWIS takes into account how the local geography interacts with the regional climate to produce an accurate road surface temperature forecast for every segment (e.g. 50m section) of road. By knowing which sections of road are likely to fall below the 0°C threshold on a night-by-night basis, highway engineers can selectively treat just the affected routes and thus make significant savings in salt usage (Figure 3).

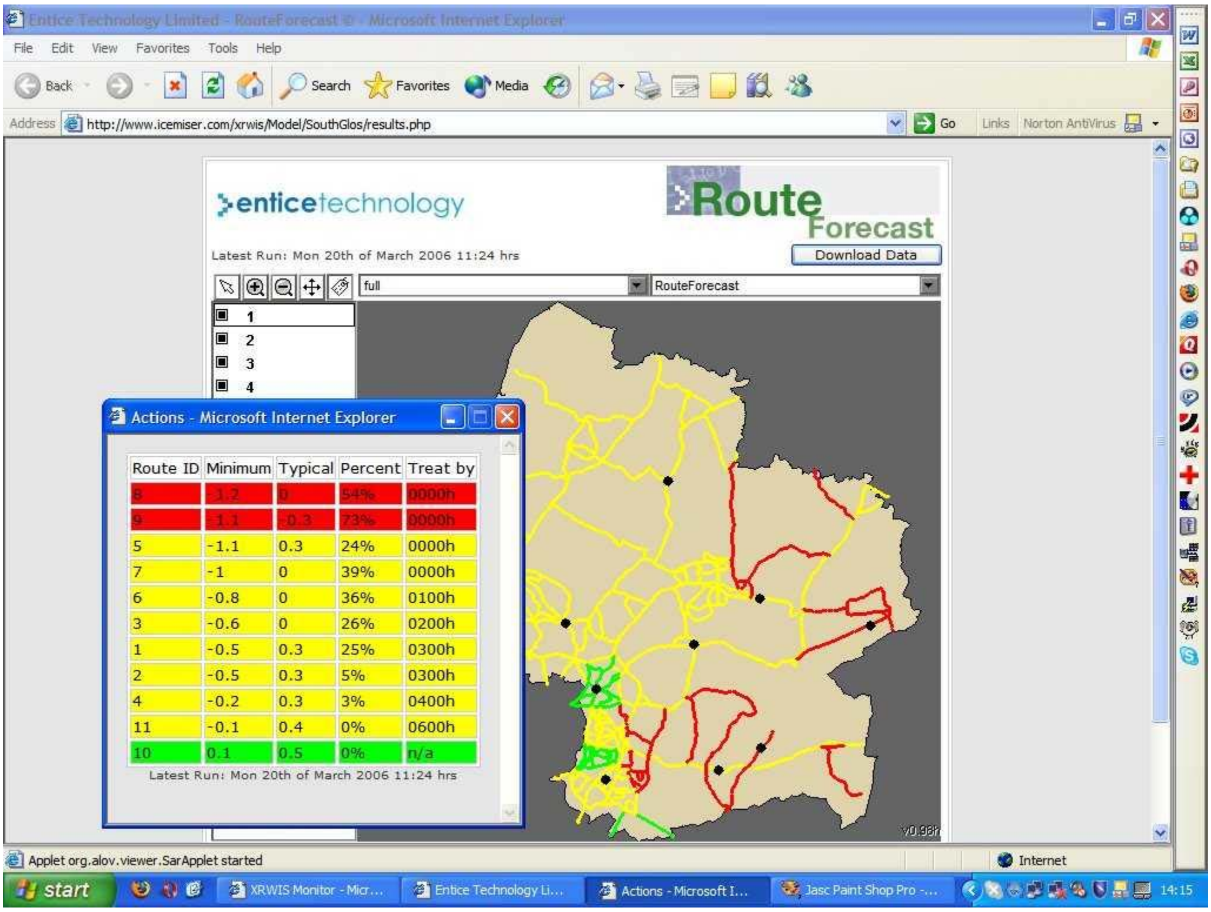

Figure 3. XRWIS screenshot. Some measure of the uncertainty of the forecast is provided to the user in the form of 'typical' (i.e. median) temperatures and the percentage of the route forecast to freeze

However, despite the increased information, route-based forecasts are still deterministic. A forecast curve (e.g. Figure 1) is essentially produced for the thousands of road segments in the forecast. This is a legacy of both the historical limitations on computational resources, as well as the early reluctance of decision makers to embrace more than one forecast. However, researchers are now starting to revisit the idea of producing probabilistic forecasts with early results indicating considerable savings [8].

2 A FRAMEWORK FOR PROBABILISTIC FORECASTING

As demand grows, there is a tendency to move from deterministic forecasts to probabilistic forecasts across a range of sectors. Probabilistic forecasts incorporate the physical and mechanistic knowledge of the systems under study. By quantifying the uncertainties in the forecast, the user should have improved confidence in that forecast. This is achieved by using credibility intervals of some sort to indicate the uncertainty. In many cases this can be expressed as a simple percentage, for example – tomorrow, there will be a 50% chance of snow.
The probabilistic component is the result of a model ensemble. This is essentially a Monte-Carlo analysis where the original conditions are perturbed across a range of plausible values given recent weather conditions, resulting in a range of output values (Figure 4)

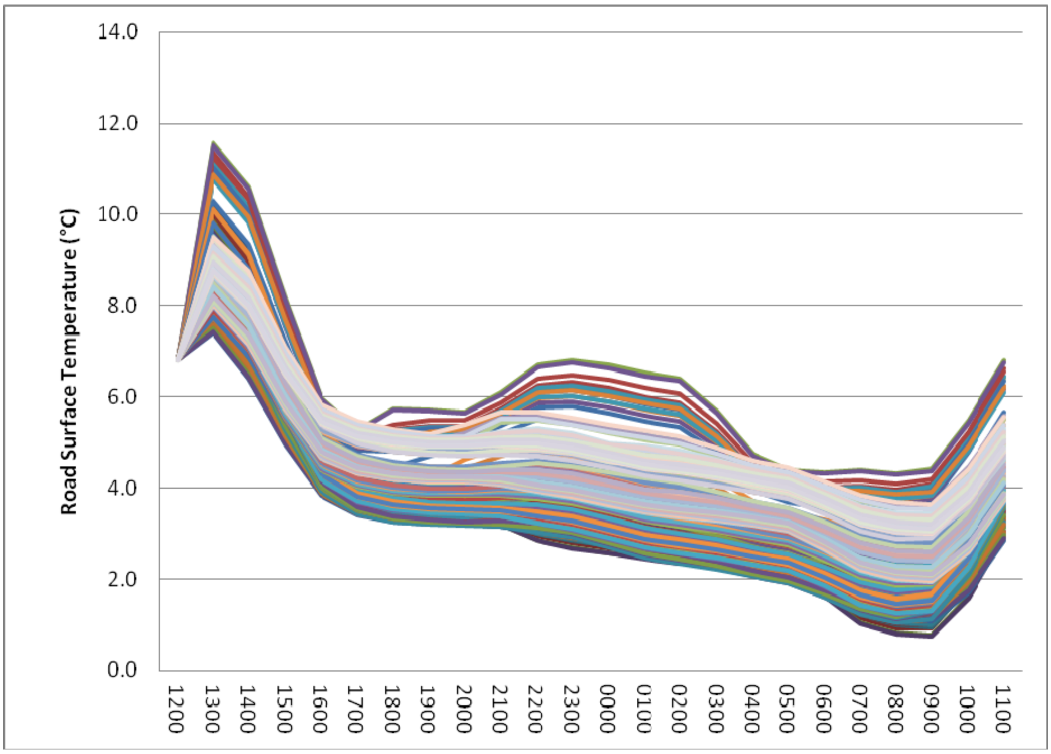

Figure 4. Example of a probabilistic road weather forecast using a 250 member ensemble.

Clearly this approach echoes back to the days of pessimistic and optimistic forecasts which essentially marked the extremes in the forecast distribution. The key difference here is that computational speed is now more than capable of dealing with more than two forecasts alone and can now be easily run for as many members of the ensemble as required. This approach can be greatly improved by using Bayesian probability theory to 'sharpen' the forecast. This will provide a full probability distribution and therefore a realistic estimate of uncertainty [9].

As a result, such forecasts can now easily be provided for the road weather sector. Indeed, there are several options available in the nature of the probabilistic forecast produced. A simple approach, which is perfectly suited to route based forecasting, is to provide a forecast graph such as Figure 4 which just demonstrates the variability around a salting route. In this example, the model perturbations are spatial / geographical. Alternatively, but perhaps more useful, the probabilistic forecast can be provided on a site by site basis where the perturbations are meteorological. Of course, there is no reason why a combination of both cannot be offered.

Overall, such an approach can be scaled up to multiple sites with relative ease, but may start to be increasingly challenging when approaching the number of sites used in a typical route-based forecasting system. For this reason, there may be some merit in adopting an emulator approach.

2.1 Emulators

Presently, a route based forecasting model typically takes under five minutes to produce a forecast for an average sized local authority (10000 sites / road segments). For the number of forecast runs required to produce probabilistic forecasts of such a non-linear system, it is advantageous to develop a statistical emulator of the model. An emulator, based on a Gaussian process approximation of the computer model, would be very fast to evaluate and thus could be run several million times per second even on relatively small computers [10].

However, building an emulator of route based forecasting models raises several technical challenges. In particular there is a need to consider emulation of strongly non-linear responses and switches typical of the effect of cloudiness on the model, the effect of discrete and categorical inputs (e.g. landuse or road classifications) and the joint emulation of temperature and moisture (both essential for the formation of ice). A novel hybrid emulator / ensemble approach to forecasting, which uses a specifically selected ensemble of model runs to help improve the emulator accuracy for a given forecast day / time could also have potential.

2.2 Uncertainty Analysis

An emulator will also allow for a detailed uncertainty analysis of models to be performed. This would assess the uncertainties in the model outputs given the uncertain inputs. Again, a simple Monte Carlo method would be the starting point, however this will not capture all the uncertainties in the system. No forecast model itself is perfect, and thus a model for the model error, often called a *discrepancy model*, also needs to be constructed. Bayesian methods could also be used to learn the discrepancy model based upon thermal mapping and road weather site observations. This approach could have enormous potential for diagnosing areas where the existing forecast model could be improved in future research.

3 CHALLENGES

3.1 Visualisation

A pertinent question is how to effectively visualise all this new potential model output? When faced with a forecast graph such as Figure 4, it is easy to see how a winter maintenance engineer may feel somewhat bewildered. In a typical route based forecasting system, there may be 10000 such graphs sitting behind the user interface each containing 250 (or more likely more) ensemble members. Clearly this is inappropriate, and the scientific community has an obligation to provide information, not just masses of data to users [11]. For this reason, constraint needs to be shown if there is to be any chance of adoption of the technology by the user community. There is no simple answer to this problem.

The first step is to ensure that the output is suitable for use in a simple decision support system. The challenge then it to reduce the data to a minimum so that it is meaningful at a glance by an engineer. One approach would be to present a forecast curve such as that shown in Figure 4 for each salting route. However, there is a need to take this a step further and to summarise the findings. Presently, XRWIS provide summary statistics for each route (see table in Figure 1). A sensible approach could be to just add another column which expresses the uncertainty in the forecast for that route – i.e. the percentage number of ensemble members which forecast the route to fall below freezing. This, when considered in conjunction with the percentage of the route forecast to fall below freezing on the mean (i.e. deterministic) forecast should be sufficient to help engineers reach a sensible decision without too much deliberation.

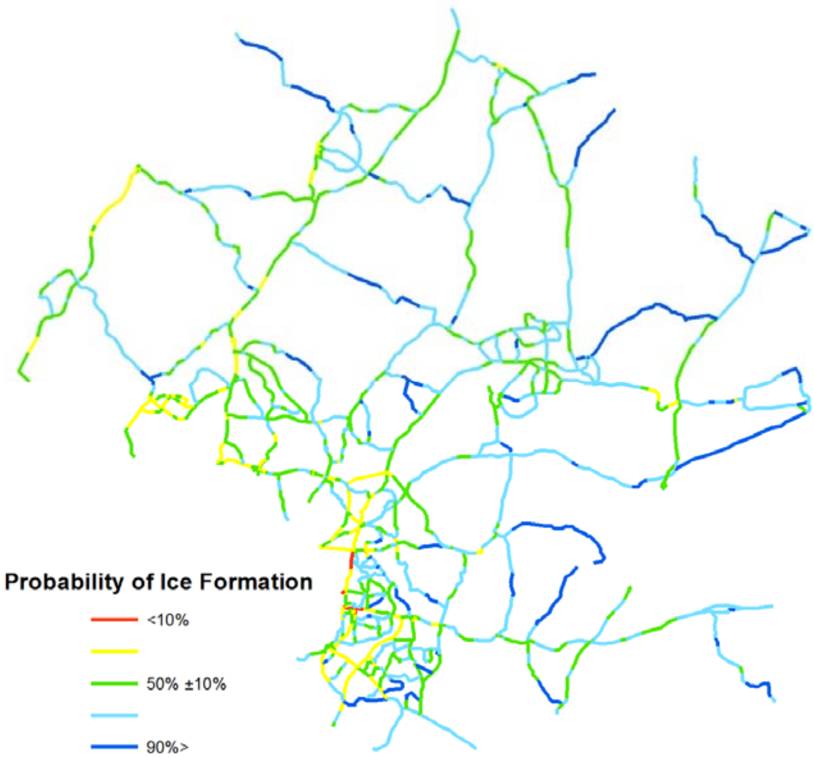

Figure 6. A thermal map for the 21st Century!?

However, this approach may be oversimplistic and engineers may instead prefer to visualise the spatial component (i.e. the percentage of ensemble members which forecast ice at each point). This is easily achieved and most route based forecasting and decision support systems would support this approach (Figure 5). As a result, it is also easy to visualise how such data could then be used in a bigger system based upon dynamic routing.

3.2 Operational Use

Regardless of how the information is displayed to the engineer, a significant level of training will be required by the end-users. The early experiments with optimistic and pessimistic forecast curves were not promising but that was 30 years ago and there is now a greater understanding and acceptance of uncertainty of forecasts, even by the general public [12]. Hence, a professional should now have no problems in dealing with the information, provided it is presented in an appropriate way. However, a problem with using a probabilistic approach in the winter maintenance community is the unforgiving nature of the 0°C threshold. Whilst research has shown that professionals cope admirably with threshold decisions from probabilistic forecasts [13], there is a tendency to exercise caution when it comes to winter treatment decisions. There will be a need to set a standard probability at which treatment occurs, which for winter road maintenance will *probably* be significantly below 50%. This will allow for uniform decision making across boundaries as well as providing a basis for verification. In turn, it is the verification data that will ultimately lead to the identification of a suitable probability threshold.

So far this paper has discussed the utility of probabilistic forecasting for the first 24 hours. However, with ever increasing computational power, there are significant capabilities of running models for longer time periods. Whilst there is an advantage of having such forecasts with uncertainties quantified over a 3-5 day period, could this approach be extended further? For example, a probabilistic route based forecast model could be linked in with a seasonal forecast to produce an estimate of salt usage for the forthcoming winter season. Such a model could be continually updated in real-time to ensure stocks are adequate. For longer time scales, such as those used in climate change impact assessments, the forecast model could be coupled with a weather generator. This would enable climate change scenarios to be taken into account and thus provide future guidance for medium term winter road maintenance budgets. At the moment, local authorities in the UK are under considerable pressure to demonstrate that they are achieving ‘best value’ from their winter maintenance budgets. Unfortunately, repeated budget cuts have led to complacency in winter maintenance regimes. This complacency has already resulted in significant snow related problems in winter 2009-10 and 2010[14,15]. Such an approach would articulate the long term probability and costs of severe winters to policy makers and hopefully prevent a re-run of such events in the future.

4 CONLUSIONS

The new paradigm of route based forecasting and decision support systems provide a perfect framework in which to develop probabilistic road weather forecasting. The technology to support such a move is already in place and producing a basic forecast is straightforward. However, to produce detailed ‘sharp’ forecasts is a non-trivial task and requires the integration of several emerging technologies and significant extensions to the state-of-the-art.

However, before fully embracing the potential, there needs to be clear evidence that practical benefit can be leveraged with improved probabilistic forecasts. There is evidence in the scientific literature that probabilistic information can greatly improve decision making, but this hasn’t been tested in the winter road maintenance community. There is a need for a detailed study to ensure that the added information does not reverse the advances made in reducing salt usage in recent years. If there is a 10% chance of ice on a stretch of road, does the highway engineer have a duty of care to mitigate that risk? If that is the mindset of the engineer, then an oversalted network beckons. However, detailed verification of probabilistic forecasts and thresholds will lead to the opposite and will further improve confidence in selective salting and dynamic routing practices.

Overall, probabilistic forecasting deserves to be trialed in the sector, however the inherent nervousness of practitioners is a major concern and may well render the technology obsolete even before it becomes operational. Even without probabilistic forecasts, the sector will still benefit from ever-increasing computational resource by improvements in numerical weather prediction. The development of operational mesoscale models to the sub-kilometre scale is in the not-too distant future and will allow the resolution of surface features currently only accounted for in downscaling models. Perhaps this is the true future of RWIS and where future research efforts should be focussed.

5 REFERENCES

[1] Chapman, L. & Thornes, J.E. 2011. What resolution do we need for a route-based road weather decision support system*? Theoretical & Applied Climatology* **104**:551-559

[2] Thornes, J.E. 1984. *The prediction of ice formation on motorways in Britain*. Unpublished PhD Thesis, Department of Geography, University of London, UK.

[3] Thornes, J.E. 1991. Thermal mapping and road-weather information systems for highway engineers, in Perry, A.H. & Symons, L.J. (eds) *Highway Meteorology*, E & FN Spon, London. 39-67pp

[4] Bishop, D. 2005. Nanotechnology and the end of Moore's Law? *Bell Labs Technical Journal* **10**:23-28

[5] Brown, A., Jackson, S., Murkin, P., Sheridan, P., Skea, A., Smith, S., Veal, A., Vosper, S. 2008. New techniques for route-based forecasting. *Proceedings of the 14th Standing International Road Weather Commission, Prague, Czech Republic, 14-16 May 2008 (available online)*

[6] Smeding-Zururendonk, I., Wokke, M., Wisse, J. 2010. Road surface temperature forecasting for gritting routes. *Proceedings of the 16th Standing International Road Weather Commission, Quebec, Canada, 5-7 February 2010*

[7] Chapman, L. & Thornes, J.E. 2006. A geomatics based road surface temperature prediction model. *Sci. Total Environ.* **360**:68-80

[8] Berrocal, V.J., Raftery, A.E., Gneiting, T. & Steed, R.C. 2010. Probabilistic Weather Forecasting for Winter Road Maintenance. *Journal of the American Statistical Organisation* **105**: DOI:10.1198/jasa.2009.ap07184

[9] Min, S.K., Simonis, D. & Hense, A. 2007. Probabilistic climate change predictions applying Bayesian model averaging. *Phil. Trans. R. Soc. A.* **365**:2103–2116

[10] O'Hagan, A. 2006. Bayesian analysis of computer code outputs: a tutorial. *Reliability Engineering and System Safety*, **91**: 1290–1300.

[11] Gustavsson, T. & Bogren, J. 2007. Information not data: Future development of road weather information systems. *Geogr. Ann.* **89**:263-271

[12] Joslyn, S. & Savelli, S. 2010. Communicating forecast uncertainty: Public perception of weather forecast uncertainty. Meteorological Applications 17: 180-195

[13] Joslyn, S., Karla, P., Jones, D., Pyles, J., Hunt, E. 2007. The effect of probabilistic information on threshold forecasts. *Wea. Forecasting*, **22**:804–812.

[14] Thornes, J.E., 2005. Snow and road chaos in Birmingham on 28th January, 2004. *Weather* **60**: 146-149

[15] London Assembly 2009. *Slipping up? Impact of extreme weather on London Transport.* http://www.london.gov.uk/who-runs-london/the-london-assembly/publications/transport/slipping-impact-extreme-weather-london-transport [Accessed 13/02/2012]

ACKNOWLEGEMENTS

Many thanks to Dan Cornford of Aston University for providing the initial input and ideas used in this paper.