INTELLIGENTLY ADDING INTELLIGENCE

P.WILLIAMS
Atkins, UK
Paul.williams@atkinsglobal.com

ABSTRACT RÉSUMÉ

All decisions in winter service are based on input information. Decision makers trust this information whilst recognising its limitations. A weather forecast is generated using highly complex algorithms; however, human intervention is often required. Although highly reliant on model outputs, quality weather forecasts require the injection of further intelligence.

Decision making in winter service is also increasingly reliant on algorithms and the quest for higher levels of refinement comes at a cost. In the current economic climate it may not always be affordable to achieve high accuracy, therefore it is important to determine what can be done at lower cost. The fundamental questions are “How much accuracy do we need for different decision making activities?” and “Where is the intelligence best added?”.

This paper considers data use in winter service and explores how the injection of intelligence into uncertain data sets can yield useful, lower cost results. It includes an outline of an experimental, long-range forecasting decision support tool that shows promise for salt stock management, resource planning and other benefits outside of winter service.

1. INTRODUCTION

Winter service remains a high profile and politically important part of highway maintenance. This is to be expected given that transport disruption resulting from a typical winter in England was valued by the 2010 Winter Resilience Review at approximately £1bn. This is compared to the average annual spend on winter service in England of £160m. [1]

The delivery of winter service in the UK is heavily reliant upon data feeds and network based information. Detailed road weather forecasts, designated weather stations providing roadside observations and GPS tracked vehicles all contribute to a well informed operation.

Decision makers trust this information whilst recognising its limitations. A weather forecast is generated using highly complex algorithms requiring super-computing power, however, human intervention is often required. Although highly reliant on computer model outputs, high quality weather forecasts often require the addition of further human intelligence.

This paper considers data use in winter service and explores how the injection of intelligence into uncertain data sets can yield useful, low cost results. The findings to date show potential benefits to winter service managers responsible for both local and strategic road networks. In addition there is scope for benefits to be gained outside of the core winter service operation.

2. DATA PROLIFERATION

The past 10-15 years has seen an explosion in the availability and use of data in most areas of life. Highway winter service is no different with dramatic increases in the collection, storage and dissemination of data. A state of the art winter service is likely to be reliant upon highly sophisticated route based weather forecasts. There will be numerous weather stations logging and reporting actual weather conditions every few minutes, 24 hours a day. The salt spreading vehicles will be tracked and reporting back not only their position but also their actions as they proceed along their route. Decisions to treat may even be made with semi-automated decision support tools which log the what, why and when of the decision making process.

Data is becoming absolutely critical to the delivery of winter service. As more data becomes available, more sophisticated tools are created to make good use of it. This creates an ever increasing cycle driving for increased accuracy and certainty. Ultimately this is driven by the desire to deliver a safe and appropriate winter service for as little financial cost as possible. However, as sophistication increases there is a risk that we lose sight of the limitations of the data source and what we are try to achieve.

Decision support tools are developed to help to manage a defined problem. The recent proliferation of data provides the developer with more choice regarding the data set(s) they adopt and analyse.

2.1. Weather Forecasting

As a deliberately controversial statement; weather forecasts are highly educated guesses. They are a prediction of the future based upon past recorded information and some highly complex modelling. The models are so complex that they require ever increasing supercomputing power to run. However, even with all the computing power and some of the finest minds, the weather forecasters still struggle to provide accurate long range forecasts. This is to be expected as there are so many variables, dependencies and interlinks within the modelled system. The UK Met Office state in their Long Range Forecast User Guide that:

"Long-range predictions are unlike weather forecasts for the next few days. The nature of our atmosphere means it is not possible to predict the weather on a particular day months to years ahead. At this range we have to acknowledge that many outcomes remain possible, even though only one can eventually happen. Over the course of a whole season, year or decade, however, factors in the global weather system may act to make some outcomes more likely than others."[2]

This does not mean that the outputs are not useful, it is just they have in built uncertainty which needs to be understood and accommodated.

It is important to understand how the weather forecasts employed to make decisions are constructed. The following description is a generic summary of the type of process. There will be variation between different forecast organisations however it illustrates that there is uncertainty within every forecast.

The forecast model(s) rely of millions of data readings of current and past events from across the globe. These are fed into the computer and the complex algorithms run to predict the future. However there is no single 'answer'. The model may run hundreds of times with subtly different inputs resulting in a collection of 'answers'. These 'answers' can then be laid side by side and the most probable 'answer' identified. This is then output as 'the' current prediction of future weather events according to the computer. However this computer output is purely algorithm based. For short term, detailed forecasting a human weather forecaster will often use their knowledge and experience to refine the forecast. Despite the huge computing power and complexity of the algorithms there is still a need to apply some human intelligence.

2.2. Uncertainty

All data comes with some form of uncertainty. For example, the weather conditions reported by weather stations are not 100% accurate; they are the sensor's interpretation of conditions. The accuracy varies with what is being measured and the quality of the sensor. Winter service in the UK relies upon a fleet of weather stations with good quality sensors resulting in small amounts of error. However, this may not be the case elsewhere in the world.

As has been described above, weather forecasting has uncertainty built into it. Moreover because it is a prediction of events yet to occur it is not possible to confirm if the forecast is correct at the time of production.

Decision making has to embrace the various uncertainties and define a course of action. This is something that engineers are generally well equipped to do, however, it is important to understand uncertainty when making decisions. Making a decision without this understanding is potentially high risk as the decision has to be appropriate for not just the predicted conditions but also other credible scenarios. This is why a decision to monitor through the night may be made on marginal nights. The forecast says it will not freeze however the decision maker realises that there is a chance that the conditions will vary slightly and preparations have to be made.

Dennis Lindley is a British statistician who stated in his 2006 book "Understanding Uncertainty":

"There are some things that you know to be true, and others that you know to be false; yet, despite this extensive knowledge that you have, there remain many things whose truth or falsity is not known to you. We say that you are uncertain about them. You are uncertain, to varying degrees, about everything in the future; much of the past is hidden from you; and there is a lot of the present about which you do not have full information. Uncertainty is everywhere and you cannot escape from it." [3]

On this basis it is important to understand the uncertainty and if possible make use of it.

2.3. Decision Support

Accompanied by the proliferation of data is a proliferation of decision support tools. These process the data in useful ways to provide guidance to decision makers. Some tools go as far as presenting an answer to the decision maker for acceptance. These tools are beneficial to the industry but they are not risk free.

As decision support tools become increasingly complex and widespread there is a risk that the ability to manually make the decision is lost. This would be a significant issue if the tool were to be unavailable or if the tool provided 'wrong' answers. Structural engineers are trained to carry out analysis manually so they do not have to rely on finite element analysis software and also so they can spot when the results do not look right. Winter decision making is no different.

There is also a risk that the decision support tools promote false confidence. The tools provide an answer to a problem which is helpful as long as the decision maker understands the limitations and accuracy of the information provided. Winter decision making involves analogue systems where there are ranges of possibilities. A decision support tool may only provide a yes / no answer to the most credible scenario. There will however be a range of credible scenarios which may need to be understood.

Understanding how decisions are reached, the limitations of the inputs and the ‘status’ of the outputs is critical to the safe use of decision support tools.

3. DRIVERS FOR THINKING DIFFERENTLY

In the UK, the severe winters of 2008-09, 2009-10 and 2010-11 highlighted that salt was not an endlessly available commodity. Salt is the primary de-icer in the UK and therefore any disruption to supply, during the winter season, is a significant challenge to the safe operation of the road network. There is plenty of salt available on earth however, in the UK, the issue in those winters was producing it rapidly in a form suitable for road treatment. There were numerous reviews and lesson learning exercises but one of the key themes around salt supply is being able to better predict medium / long term demand within a winter season. Strategic salt stocks work well to safeguard against future national crises however better management of operational salt stocks is an important part of the jigsaw. The answer is not simply to hold ever increasing amounts of salt in stock as this is a significant drain on available cash. Therefore a mechanism is needed to help better predict salt demand over a number of months to enable timely restocking.

Traditionally winter service decision making has focused on the ‘here and now’. Most winter service forecasts only looks five to ten days ahead. This is simply not enough time for salt suppliers to react to a sudden surge in demand.

The biggest problem faced by winter service managers was available input data. Relying on past salt use and using rolling averages has recently been proven as inadequate. Typically the calculations were done on a season by season basis so it was not refined enough to predict spikes in demand during a season. Given the variability of winter patterns in the UK using more detailed salt use figures and applying more complex statistical analyses still could not predict potential future spikes. Climate change or even the phenomenon some call “Global Weirding” make the reliance on past information an illogical solution.

The only option was to use long range weather forecasting. This could provide a salt forecast model of the coming months. However, it is well known that long range forecasting contains a high degree of uncertainty and there is no immediate prospect of significant improvements in accuracy. So the question was how can we use this ‘dirty’ data set to create a ‘clean’ answer?

3.1. Thinking in systems

The solution came as a result of a basic systems thinking approach. Systems thinking is an increasingly popular way of approaching problems. It accepts that the problem may not be a ‘component’ but is as likely to be between ‘components’. The quote below from Ballé in 1994 sums this up:

“To understand things we take them apart and study the pieces. To improve things we try to improve pieces individually. It is rather like trying to get a horse to run faster by teaching each of the legs to perform a more efficient movement ….. The systems approach focuses on the inter-relationships, how the horse’s legs relate to each other and back to the horse.” [4]

The concept that when ‘A’ happens it has an impact on ‘B’ which in turn impacts on ‘C’, ‘D’ and ‘E’ is important. This defined the approach to allow the uncertain forecast data to provide useful information. Donella Meadows provides a very relevant insight in her 2012 “Thinking in Systems: A Primer” book.

“As our world continues to change rapidly and become more complex, systems thinking will help us manage, adapt, and see the wide range of choices we have before us. It is a way of thinking that gives us the freedom to identify root causes of problems and see new opportunities.” [5]

We understand that salt stocks diminish when road treatments occur and there is no restocking. Treatments only occur when the road surface is cold and there is moisture available to form snow, frost or ice. The road surface is cold at the same approximate times that the air is cold. Given we could obtain forecasts of likely conditions we therefore could calculate a potential salting scenario. This salting scenario will outline the likely future treatments which in turn would allow the generation of a forecast salt stock profile. This may sound trivial however it allowed a link between the available forecasts and the desired answer. It also highlighted that the tool could have potential to contribute to much wider decisions than just salt stock management.

Taking the example of the 2008-09 winter, if numerous decision makers had access to information predicting salt usage then the impending crisis could have been foreseen earlier and as orders increased the salt producers would have been able to increase production. Everyone had access to long range weather forecasts but there is little evidence to suggest that many made the link and acted upon it.

There is a direct relationship between salt usage and number of treatments. This means that the same ‘answer’ would allow forecasting of labour and plant demands allowing better management of the winter service. It also has wider potential as a forecast spike in salting activity is also likely to mean an increased number of calls to contact centres, more demand on council support services and an increased demand on the doctors due to respiratory illness. It is all part of ‘the system’.

It was apparent that there was likely to be significant benefits if the problem could be solved. The various pieces of the jigsaw existed but it now had to be put together.

4. DEVELOPING THE IDEA

Continuing the ‘systems’ based approach the first step was to identify inputs, constraints, linkages and the desired outcomes. This allowed the problem to be approached as efficiently as possible. A key point was the aim to produce a relatively low-tech output to ensure it was applicable and understandable to as wide a range of people as possible.

4.1. Forecast data input

The first activity was to examine the commercially available forecast information. This is where the recent proliferation of data helped. It is possible to purchase raw weather forecast model data over the internet for modest sums of money. There was an option to purchase higher cost, forecaster intervened data however the project was an experiment to identify the potential use of raw un-intervened model data. The option remains for the future to consider the costs and potential benefits of a move to intervened data.

Having reviewed the readily available forecast data the Global Forecasting System (GFS) model and a post processed Climate Forecast System (CFS) model were chosen. By combining these two forecast data sets we gained an accessible total forecast period of 60 days in XML format.

4.2. Constraints

The ideal data would have provided road surface temperature forecasts and an indication of the likely hazards in terms of frost, ice or snow. However this was not possible so the first constraint was the available data. A mechanism to convert the available data into more useful information was required.

The forecast data did not supply the overnight minimum temperature which is the standard decision making reference point. The forecast minimums at different times of day and night were available therefore the worst forecast condition was adopted as the minimum overnight temperature. This clearly inserts further uncertainty.

It was understood that there was uncertainty around the accuracy of the forecast data. The longer the range, the more inaccurate however there was no ready way to determine how big the error would be. It was not possible to change the accuracy levels available and therefore a way to embrace this uncertainty and drive benefit from the available data was needed.

The forecasts only contained atmospheric data and therefore could not be directly translated into winter treatments. The decision making guidance in the UK relies upon road surface temperature. It was understood that there is no direct relationship between road and air temperature therefore an assumption would be required.

The weather varies with location and at times can vary significantly within a small geographic area. Forecast information was available for numerous points around the UK however each additional forecast point added to the cost of the data and also the complexity of the tool. Ideally hundreds of forecast points would have been ingested and processed, however, this was not practical in initial development. Seven points were chosen to provide a spread of forecasting across the UK. These were in ‘average’ locations as opposed to the more unusual locations such as tops of hills or adjacent to the coast.

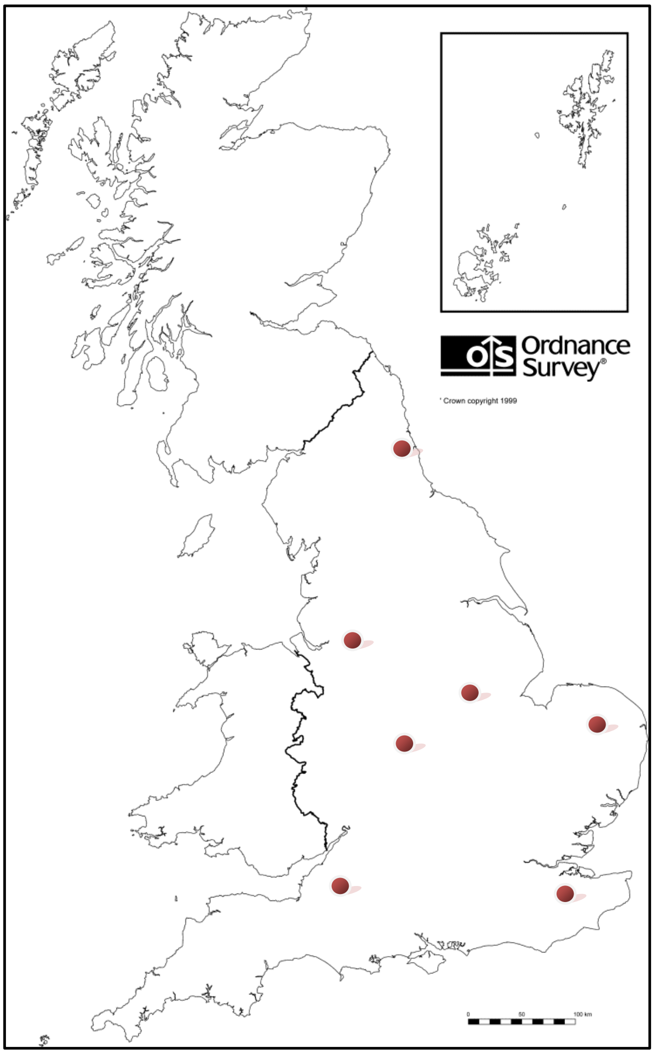

Reproduced from Ordnance Survey map data by permission of Ordnance Survey © Crown copyright 2013

Figure 1: Map of forecast points

Winter treatment decision making varies with different local authorities and agencies therefore the amount of salt used for a given set of weather conditions will vary across the UK. Whilst it may be feasible to utilise multiple different treatment guides within the decision making algorithm it would significantly increase the complexity and it is not clear if it would yield any significant benefits. A typical treatment decision guide would be needed to maintain a simple algorithm.

This pilot project was highly innovative and there was little to benchmark against. It was important to calibrate and test the output against actual data. Since the 2008-09 winter, the UK records its salt usage in more detail and therefore there was the potential of a data set to test against.

4.3. Linkages

The intention of building a tool to help predict salt usage opened up the opportunity to provide information to a wider audience than just the winter decision maker. As mentioned previously a forecast spike in salting activity is also likely to mean an increased demand on other public services.

The outputs could influence a significant range of a local authority's services. The Winter Service Manager is obviously interested in being able to manage their salt stocks, fleet and drivers. The output will assist with all of this but will impact on other areas. For example if salt suppliers see an influx of early customer orders they are likely to increase production to meet the need. Therefore the output at a Local Authority level could help the salt producers. Obviously the producers could adopt a national model to aid their planning however it does not tell them when the orders will actually be placed by customers.

The plant and labour needs of winter service will often impact on the ability of a local authority to deliver other activities. Therefore if a spike in winter treatments is expected then it is likely that other services will have to reduce or put in place plans to bring in additional resources.

Taking this one step further, there are numerous local authority services which are not directly impacted by the winter service activities, however they rely on the highway network to deliver their service. For example, carers visiting vulnerable people may struggle to get to them all if the road network is not fully available.

There are also local authority services that are impacted by the weather and can have a significant impact on the local economy. For example businesses may struggle to operate normally if their staff have childcare issues following school closures due to preventable issues such as heating failures. Forward planning may be able to prevent some of the school closures. Snow is obviously a different issue.

Finally there is the public. It would not be appropriate to disseminate the outputs of any tool based on such uncertain data to a ‘layman’. However it would allow planning of media campaigns such that as a cold spell approaches the media managers are well prepared to engage, warn and inform as is required by the UK Civil Contingencies Act: 2004.

It is therefore clear that there are numerous potential interested parties. The challenge is how to deliver usable output to meet as many needs as possible.

4.4. Outcomes

The desired outcome is simply a tool which provides guidance on the potential future winter service demands. This can then be employed in different ways to yield benefit.

It is clear from the number of linkages that there are numerous potential interested parties. The challenge is how to deliver usable output to meet as many needs as possible. Some of the intelligence from the forecasting inputs had been consciously omitted so a decision was needed to determine the best places to inject more intelligence.

The amount of uncertainty within the forecast data and the number of assumptions needed to process the data mean that a ‘single answer’ output is unlikely to be useful. The outputs will have to be potential scenarios which can then be used by a competent person to make an informed decision. In doing this some of the uncertainly is absorbed. Indeed a ‘range answer’ is only possible where there is uncertainty therefore this approach is capitalising on the uncertainty. It paves the way for techniques such as sensitivity analysis to provide confidence levels for the decision maker.

It was realised that this could only be a decision support tool and not a decision making tool.

5. PILOT IMPLEMENTATION

Having undertaken the initial research it was clear that an experimental pilot should be developed. This was against a backdrop of the second severe winter in the UK and another potential salt crisis. It was realised the tool would not be suitable for operational decision making but it may yield some useful results.

5.1. Assumptions

The first step was to consider the assumptions needed. This was probably the single most significant injection of intelligence into the tool. Using experienced winter decision makers who had received advanced meteorology training it was possible to translate a standard treatment decision guide into a set of rules which could be used against the available forecast data.

There were actually three set of rules established to provide a high, medium and low scenario. These scenarios effectively took optimistic, pessimistic and neutral assumptions based on the weather forecast data (Assumption variance). This assumed that the forecast was accurate and the 'rule' contained all the error.

However, these scenarios could also be seen as translating the single forecast into an optimistic, pessimistic and neutral forecast (Forecast variance). This assumes the 'rules' are accurate and creates an artificial variance in the forecast.

The truth is likely to be somewhere in the middle with the forecast and assumptions both being incorrect (balanced output). This muddies the waters of uncertainty somewhat but Figure 2 shows these concepts diagrammatically. It requires the user to take a step away from pursuing the definitive answer and accepting that there is a range of answers. It does assume that the forecast and assumptions will not both be very wrong in the same direction. However, given the 'neutral forecast' and 'neutral assumptions' are the most likely results it is likely that the true output will be close to the centre of the output range.

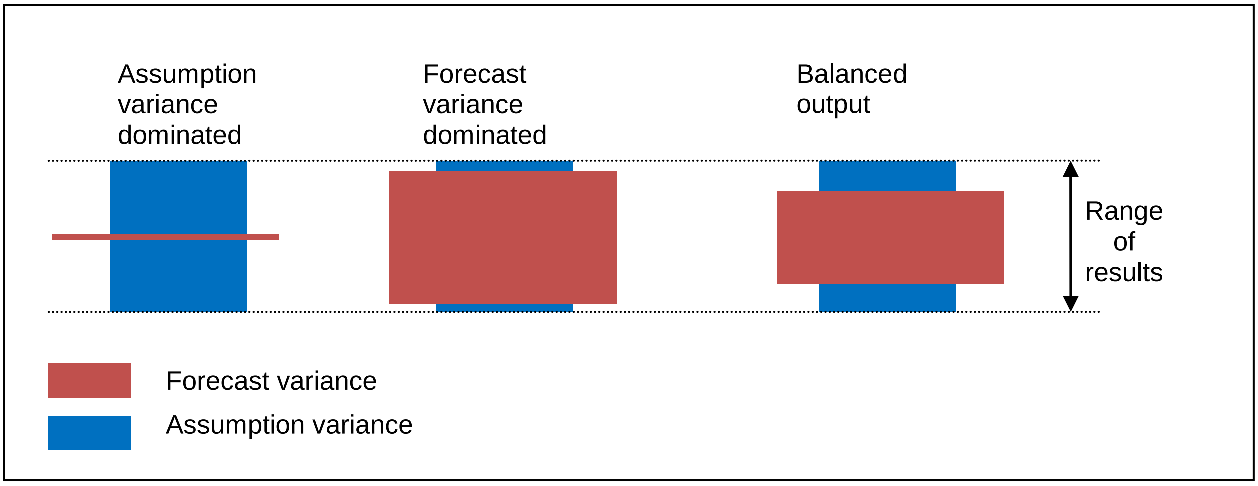

Figure 2 – Uncertainty in the model.

This approach is not overly scientific however it does only aim to give the most probable range of results. It is conceivable that the true answer will lie outside the results range however this needs to be accepted by the user. It is a salt forecast therefore may not be the perfect answer.

By varying the scenario settings it would be possible to carry out a rudimentary sensitivity analysis. This would provide an indication of the confidence the decision maker should have in the output.

The assumptions used air temperature, snow risk and precipitation amount as the main indicators to the conditions. The treatment guides were reverse engineered such that a set of conditions were matched to each treatment option. It was even possible to identify where multiple treatments could be required. This allowed, for an assumed network area, the calculation of the amount of salt required to deliver the treatments.

5.2. Data handling

Having obtained the two forecast data feeds these were initially parsed into a daily single forecast data set. This was the combined GFS and CFS model data. This data feed was processed against the rules to provide three different number of treatment and therefore salt use scenarios.

The processing in the initial pilot was carried out via a multi-sheet Microsoft Excel workbook. It used standard functions and can be performed without the need for visual basic programming. However future versions may benefit from a more sophisticated computing solution. Once the forecast data had been loaded the processing time was typically under 10 seconds.

The three scenarios are then plotted on a graph against time to provide both instantaneous and cumulative salt use. The instantaneous use graph allows ready identification of spikes of winter service activity. Given a known starting stock, the cumulative graph can be used to model the salt stock diminishing with use. This can then trigger thresholds for restocking.

5.3. Calibration

The first working edition of the tool required calibration. Assumptions had been made it was not certain if these were correct. It is akin to manufacturing a clock, setting the time for the first time and simply hoping it kept the correct time. The output required baselining against known results.

It was possible to access some Highways Agency salt usage figures which allowed calibration. The model's results were reviewed against the now known outcomes. By varying the assumptions it was possible to generate salt use forecasts that aligned to the actual results. The real test was then the tool's ability to forecast salt use into the future using the newly calibrated assumptions.

5.4. Refinement

The tool produced some interesting results over a period of time. For example during the severe weather event of December 2010 the model predicted that the UK salt use would dramatically reduce prior to salt levels falling to critical levels. The actual salt stocks did move as the tool had predicted with a sudden stop to the severe conditions that had put such strain on the winter service across the UK.

Over an extended period the settings were refined to what is thought to be a good 'all-round' set of rules for the UK as a whole. However there is further refinement to do and it is anticipated that the rules will need validating for each different network prior to any operational use.

When re-running models with slightly different rule settings it was possible to carry out a basic sensitivity analysis. Through prudent use by competent people it is possible to refine the output to provide the decision makers with a degree of confidence in the results. However, there is an element of ‘gut’ feel required. It reverts back to the concept of forecaster intervened weather forecasts. The computer models suffice the majority of the time, however, sometimes you can’t beat the human touch.

This is as far as the research has gone to date. There is a working model in existence which appears to offer promising results. Taking a national view it has produced workable results which, at times, have been shared with those responsible for supervising reserve salt stocks in the UK. The results may not have been acted directly upon but they were available if required.

6. EXAMPLE OUTPUTS

A number of output graphs have been included in the paper to provide an impression of the output that can be generated. The visualisation in the final edition is likely to change as this is currently based on a Microsoft Excel spreadsheet and charting function.

Figure 3 is a good example of the weather forecast driving likely salt use over the Christmas 2012 period. From the graph a winter service manager would be able to see that there is some activity throughout the period however the majority of the activity falls after the Christmas period. This would be an important piece of information for both salt and resource management.

Going into a little more detail the ‘neutral scenario’ line sits closer to the ‘low use scenario’ than the ‘high use scenario’. Without any further analysis this indicates that the true situation is likely to be closer to ‘low’ than ‘high’. This is very basic but it highlights the potential benefits sensitivity analysis could bring.

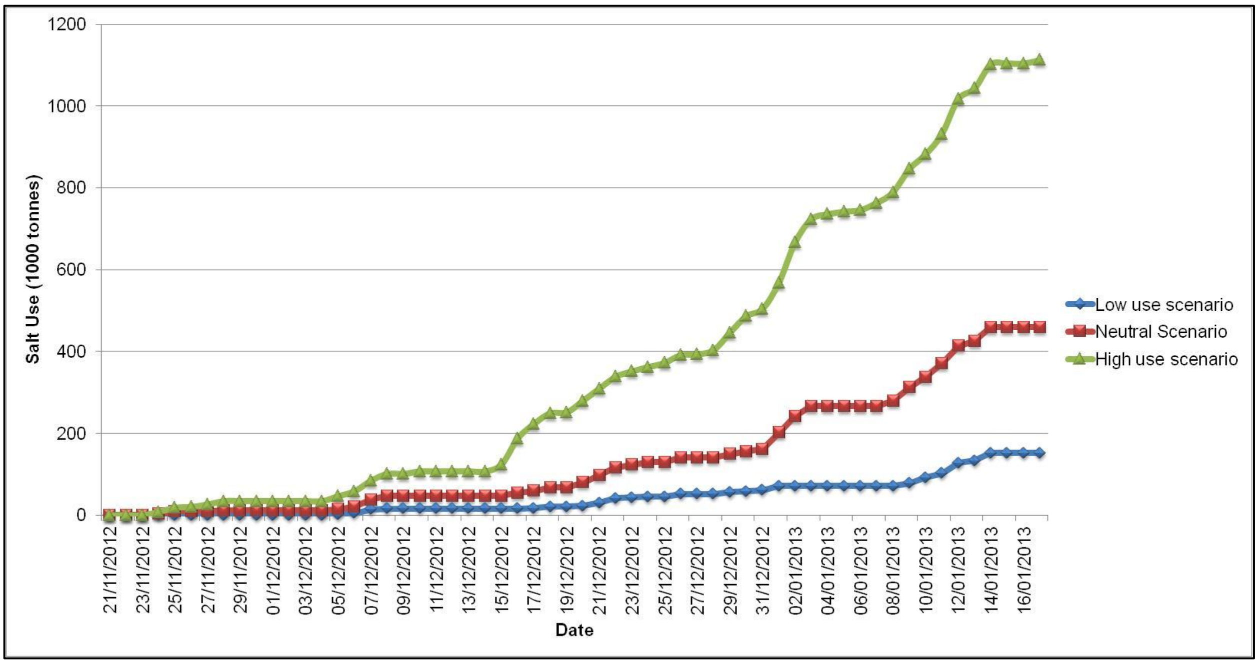

Figure 3 – Results graph from the Christmas period 2012

Long term forecasting generally gets the 'shape' of the temperature over time graph correct. The precise timing is not necessarily correct and potentially the precise temperature is also inaccurate. However, getting the shape correct is important. Figure 3 also shows that the three scenarios accentuate the same shape of curve. The 'high use scenario' amplifies the peaks more than the other lines.

Using some Highways Agency actual treatment information it was possible to draw an approximate comparison to the forecast model.

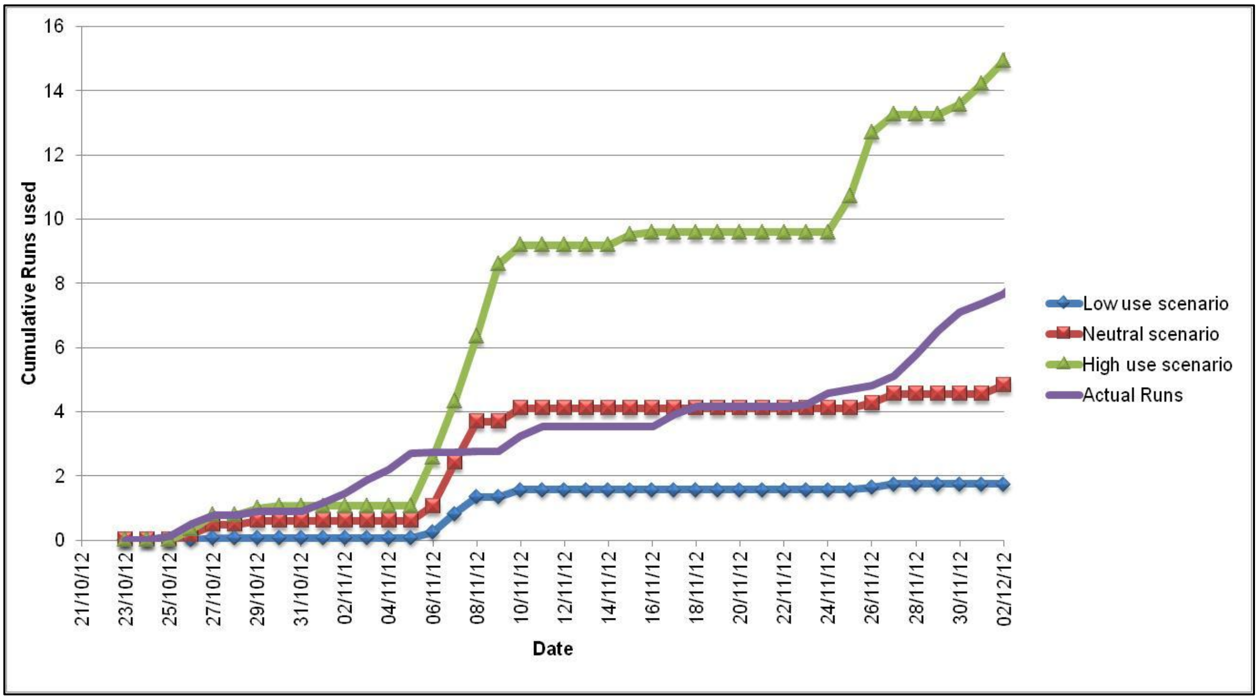

Figure 4 – Results Comparison

Figure 4 shows the average number of treatments delivered and it is clear that the model is capable of following the overall 'shape' of the service. At times the timing is slightly incorrect, however, considering the level of data being input and the range of assumptions being made the results are encouraging.

A visualisation was developed to help Winter Service Managers to forecast at county level. Figure 5 is an example output showing how salt stocks may vary over the coming period.

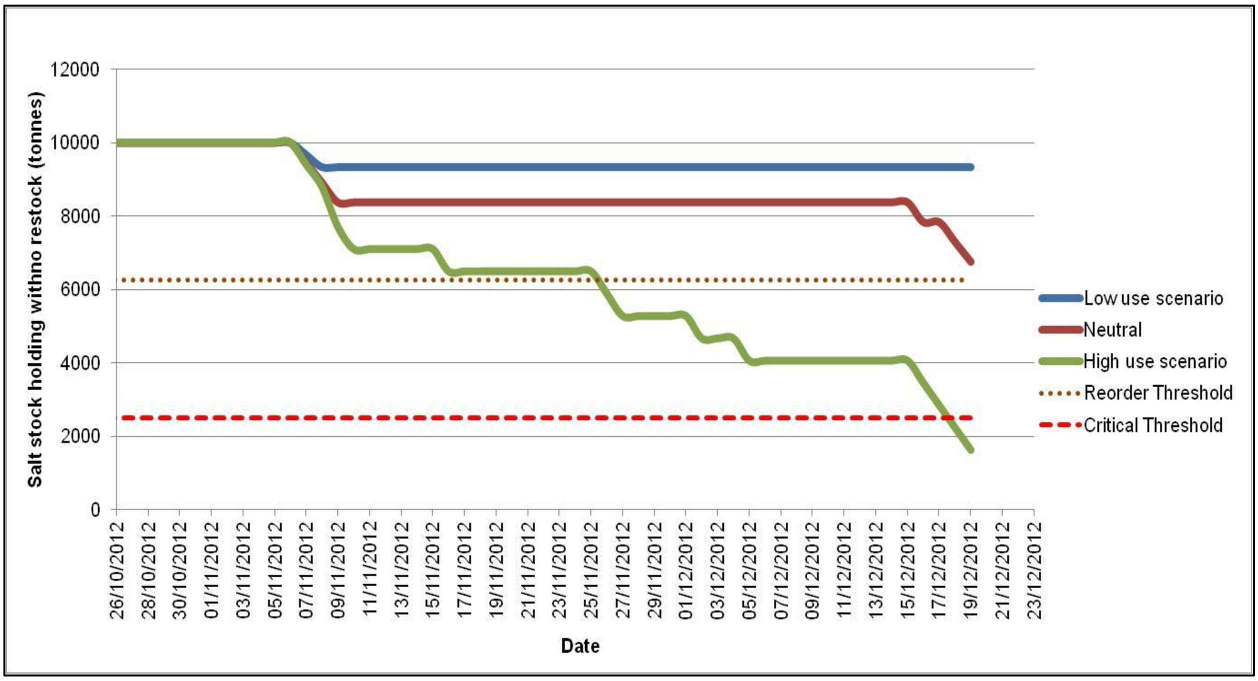

Figure 5 – Salt stock model

Figure 5 shows how the previous two outputs could be translated into a forecast stock holding. The re-order threshold is purely as an example which shows that the point where re-ordering should occur could be as early as mid November but is more likely to be mid December. This should aid decision making in terms of timing of salt ordering. This example also shows that there is a possibility that the service could become compromised as early as mid December if restocking does not occur.

7. LESSONS AND THE FUTURE

The pilot project was created out of a need for longer term decision guidance and a curiosity to see what can be achieved at low cost. Along the route a great deal has been learnt not only about weather but also about how data can be used. It has enabled the identification of ways to take the pilot forward.

7.1. Lessons learned

The concept of uncertainty is an important one in long term planning. It is simply not possible to create certainty over an extended period, especially if natural systems are implicitly linked to the activity. Rather than expend significant efforts failing to achieve the impossible the pilot embraced the uncertainty.

The decision not to use forecaster intervened weather forecasts did increase uncertainty however the amount of benefit did not outweigh the costs of the more detailed forecasts. The key to the decision tool was translating weather data into treatment information and ultimately salt use forecasts. By injecting the intelligence into the right part of the system it has shown that usable results can be achieved even using low-cost ‘dirty’ data. However, currently it does require an intelligent user to ensure the results are appropriate.

Decision support tools are a useful and inevitable part of winter service. However, it is important to question the inbuilt uncertainty around the inputs and the algorithms. If basing decisions wholly on the outputs of tools it is important that the output is truly understood.

When considering decisions and decision support it is important to consider not only the activity that the tool directly supports but the subsequent activities it influences. The systems thinking approach helps identify linkages and drive benefit from the tools you have available.

7.2. Planning for the future.

The pilot project has used coarse data inputs and tested against national figures. It appears possible to forecast salt use down to a regional level, however, with an increased data set it is thought that this could be further refined to provide more detailed salt forecasting. This would then benefit individual local authorities in their salt stock management.

The forecast data we ingested works well however this is low cost data and with the advances over the recent years in long term forecasting a new and improved data set may now be available.

By reviewing historic long term weather forecasts against an authority’s salt use of the same period, the model could be calibrated and validated to a specific network. It is hoped that operational trials with highway authorities can be established to test the usefulness of the tool.

Further work is also needed to identify scope to integrate the tool, or similar tools, into the wider operational planning of local authority services. The tool currently serves tactical level managers in the highway department but it has potential to provide guidance at a more strategic level to other departments. Emergency planning, social services, education and health are all areas where long term identification of severe weather is important. They are also areas of budgetary constraint.

It is important in times of economic restraint to maximise the benefits from any investment you make. A business case is needed to identify the core and peripheral benefits, attempt to monetise these and then offset against the cost of implementing the tool. It is hoped that this business case can be developed with ‘live’ operational partners in the near future.

REFERENCES RÉFÉRENCES

1. Quarmby, D (2010) The Resilience of England’s Transport Systems in Winter, Para 109
2. http://www.metoffice.gov.uk/research/climate/seasonal-to-decadal/long-range/user-guide
3. Lindley, D. (2006). Understanding Uncertainty, Preface pp xi
4. Ballé, M (1994) Managing with Systems Thinking. pp35
5. Meadows, D (2009). Thinking in Systems: A Primer. pp17