J1.8 REAL-TIME FUSION OF SENSOR DATA TO ACHIEVE IMPROVED SITUATIONAL AWARENESS

Ian Griffiths*, Martyn Bull and Luke Carrivick
RiskAware Ltd, Bristol, UK

1. INTRODUCTION

In hazard release incidents it is particularly important to have good situational awareness in order to make the most effective decisions and save lives. Chemical and biological (CB) and other sensors can be deployed to measure the presence of toxic airborne contaminants; however it is necessary, and a significant challenge, to make best use of the information provided to determine how a hazard will evolve in order to warn people appropriately. This has particular relevance to military and homeland security application. In the absence of any other information, the current military approach is to interpret the observation point as being the source location. In fact this is very unlikely to be the case and hence the assumption will lead to error in the hazard prediction. An additional issue is the lack of local meteorological observations or up-to-date accurate weather prediction data necessary for hazard modelling.

There has been a significant research to investigate whether computational techniques can be used to fuse sensor data and provide better hazard prediction. There are several approaches to the problem. These include fusing data to determine a source term that best matches the observations and data assimilation that uses observations to refine a modelled representation of hazard. The output of a source estimation algorithm can be used to make a hazard prediction using a standard dispersion model, whereas the data assimilation method produces a "now-cast" of the hazard.

We have developed a simple data assimilation approach that fuses real-time sensor observations to provide a CB hazard "now-cast". In the development of this Nowcast capability, we were seeking a technique that was rapid to compute even with modest computer power, while robust and providing output that would be compatible with hazard prediction models. An initial capability has been developed, which we describe in this paper. We have also carried out an evaluation of its benefits for some realistic scenarios, comparing it with the current approach used by the military.

2. TECHNICAL APPROACH

Our Nowcast approach represents the hazard cloud as a collection of Gaussian puffs, which it optimally fits to current observations producing a real-time CB hazard from the fused sensor data. It employs the expectation-maximisation (E-M) algorithm to fit the Gaussian mixture model to the observations (Dempster (1977)). In our approach the algorithm proceeds as follows:

- When sensor information is received a number of puffs are initialised with initial estimates of position and size (see figure 1a). (On the first sensor reading this will be random.)
- The sensor observations are each assigned to one of the puffs by selecting the nearest one (figure 1b).
- We then estimate values at the sensor location – the E-step – based on the Gaussian parameters.
- We update the Gaussian parameters so as to maximise the fit between the puffs and the observations – the M-step (see figure 1c).
- The E-M steps are repeated until convergence is reached (figure 1d).
- The algorithm does this for a range of different numbers of Gaussians and selects the fit that is best.

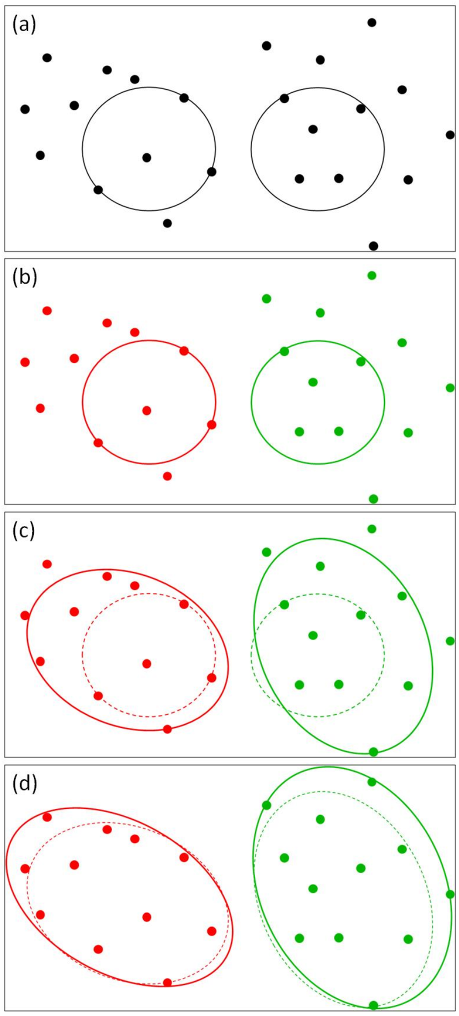

Figure 1: Fitting stages of Gaussians to sensor data using the E-M algorithm.

* *Corresponding author address*: Ian Griffiths, RiskAware Ltd, Colston Tower, Colston Street, Bristol BS1 4XE, UK; e-mail: ian.griffiths@riskaware.co.uk.

Many CB sensors provide bar reading output, which are sensor reading ranges rather than actual concentration values. These are ideally suited to the E-M algorithm. In the case that the inputs are continuous concentrations, they are first converted to bar readings before the algorithm is applied to facilitate their use by the E-M algorithm. However, if calibration data for the sensor levels is available or the inputs are continuous concentration readings this is additional information, which can be utilised. At each time step, the algorithm estimates the mass for each Gaussian using least-squares optimisation to minimise the difference between the predicted concentrations (from Nowcast) and the observed concentrations (either from the continuous sensor reading or if the calibration data is applied). This also considers the effects of sensors not alarming, which the E-M algorithm does not directly use.

As further sensor readings are received the E-M algorithm is repeated. However, the previous Gaussian fit is used as the initialisation for the algorithm, which improves the performance. Varying numbers of Gaussians are considered at each time step to ensure the best fit is obtained and also to avoid degeneracy.

The result at each time step is a set of Gaussians that yield a concentration field that provides the best fit to the observations. Because it requires no transport and dispersion modelling or other complex physics modelling, it is extremely fast providing a nowcast hazard estimate in real time. Figure 2 shows the resulting concentration fields based on the Nowcast puff fits and the input concentration data. In this example run, a large number of sensors have been simulated, although in later evaluation cases lower, more operationally realistic numbers are used.

2.1 *Hazard Prediction*

An important benefit of the E-M approach adopted is that it produces a set of Gaussians that are compatible with many hazard dispersion models, such as SCIPUFF or the Urban Dispersion Model and others. Nowcast has been interfaced to a simple rapid Gaussian puff model, which uses the output of Nowcast to automatically produce a hazard prediction.

The aim of this work is to improve situational awareness and provide improved hazard predictions. Of particular interest is investigating whether the approach being developed can support real-time hazard refinement as sensor data updates are received. This has been achieved by using the combined output of Nowcast and the linked dispersion model.

At each time step, the dosage is calculated for both the Nowcast contribution and the forward dispersion model. For Nowcast we integrate the concentration fields at each time step to produce an incremental dosage up until the current time step. We then add to this the dosage calculated using the rapid forward dispersion model. Each of these dosage contributions change each time new sensor information is received, as the Nowcast hazard fit increases the accumulated Nowcast dosage and provides a new set of inputs to the forward dispersion model.

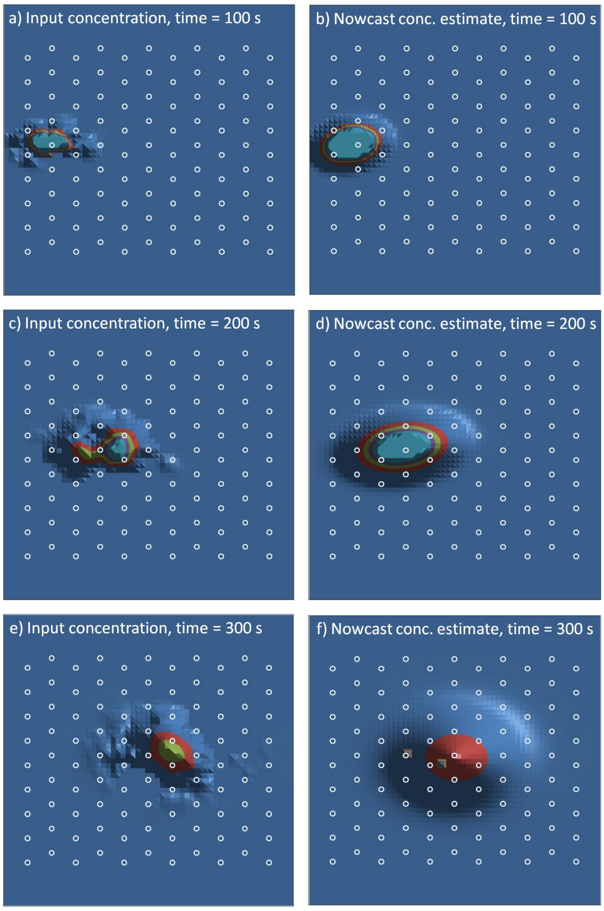

Figure 2: Comparisons of challenge concentration fields (input) shown on the left (a, c and e) and Nowcast concentration estimates (output) shown on the right (b, d and f) at times 100s (top – a and b), 200 s (middle – c and d) and 300 s (bottom – e and f) from true release time. Circles show sensor locations.

2.2 *Handling Inaccurate Meteorology*

It is notable that the Nowcast algorithm developed does not require a meteorological input in order to produce its output. A significant issue for operational hazard prediction is that the meteorology available may be from forecast data produced some time prior to the event or for observations taken a distance of several or tens of kilometres from the release. This results in meteorological inputs for the hazard prediction model that may be inaccurate by significant amounts – for example errors of 30 degrees in the wind direction are possible. Because Nowcast is providing an estimate for the evolving hazard it seems there is the potential to estimate the underlying local wind speed and direction, which can then be used by the transport and diffusion model.

There are a number of ways this could be achieved. At this time we have chosen a simple method, which we call the centre-of-mass method. At each time step, we calculate the concentration field using the Nowcast fitted Gaussian puffs. For this field we calculate a centre of mass. We then track the movement of the centre of mass over time, which provides an estimate of the wind speed and direction for the dispersing hazard cloud. This meteorological information is used by the rapid forward dispersion model.

Figure 3 shows the estimated wind speed and direction produced by Nowcast over time. The first sensor reading was at time 50 seconds. No meteorological information was provided to the algorithm, so initially there is significant instability until it has received sufficient information. From the graphs this appears at about time 150 seconds, which in this example is after about ten sets of sensors observations. The true wind speed was 5 ms^{-1} and the wind direction was from 0 degrees, so we can see that for this example the algorithm has successfully managed to estimate the local meteorological conditions within a reasonable degree of accuracy.

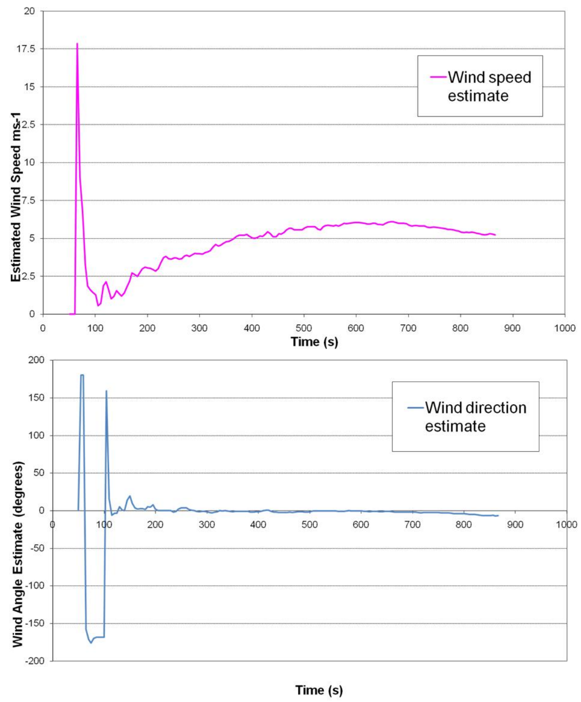

Figure 3: Inferred wind speed (top) and wind direction (bottom) estimates over time using centre-of-mass tracking.

It is clear from figure 3 that until the algorithm has stabilised in its meteorological output, it would not be wise to use the estimate. The algorithm currently analyses the statistics of the evolving wind speed and direction and use this to determine when the meteorological estimate appears sufficiently robust to be used. It calculates the moving average and variance of the values, and when the variance falls below a threshold, the moving average value is provided to the forward dispersion model. Prior to this point the dispersion model uses the initial meteorological input, accepting any errors in that, as it is deemed to be the best data to use until then. As would be expected, as the plume leaves the sensor array, the variance term grows again as the meteorological estimates become less stable; this is automatically detected and the dispersion model is provided with the last stable meteorological estimate. Figure 4 provides a example of the difference in the dosage calculated using a dispersion model with the original erroneous meteorology input and the Nowcast dosage which updates the meteorological estimates over time; the comparison with the true realised dosage indicates that Nowcast's dosage prediction is much closer to the true dosage and it has corrected for the input meteorology error.

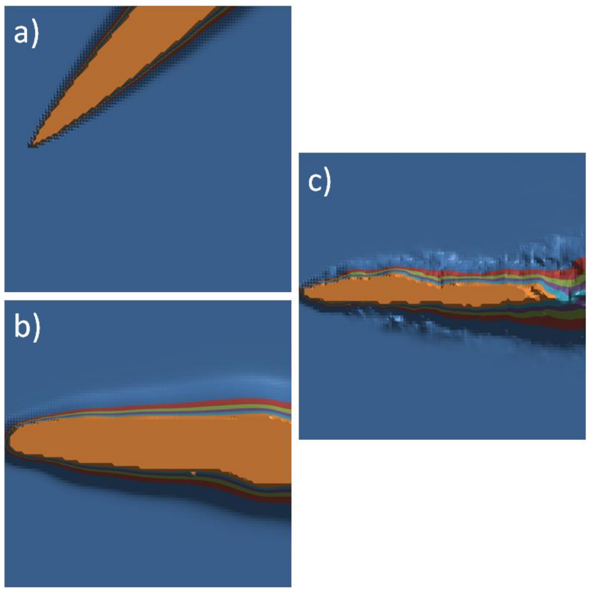

Figure 4: Dosage results when there is a error (45 degrees) in the wind direction input – a) forward prediction from first sensor reading using input wind direction; b) Nowcast estimate using inferred meteorological values; and c) the true dosage, calculated by integrating the realisation concentration fields used to generate sensor inputs.

3. STUDY DETAILS

An evaluation study has been carried out to compare the resulting hazard predictions using the current doctrinal approach of using the location of the sensor observation as the release point and the Nowcast approach. In addition we have included the hazard prediction obtained using the actual release point. These are compared to a “ground truth” hazard that is used to produce both the sensor observations and also the challenge realised dosage. This comparison has been made for cases where the meteorological conditions are known and ones where it an estimate, which includes an error, is provided to the prediction models.

3.1 *Challenge Generation*

A synthetic environment has been used to developed produce challenge data for the evaluation. As described by Bull (2009), the core component is a physics-based concentration fluctuation model that produces realistic concentration time series with the appropriate complex structure and which, importantly, is correlated in space and time. It takes as input a concentration ensemble mean, variance and length scale field (as produced by, for example, the SCIPUFF dispersion model). The synthetic environment also includes sensor models and various human effects and performance models and physiological burden models. In this way the synthetic environment can be used to create challenge data sets which can be used for test and evaluation, as in this study. Figure 5 provides an example of a simulated concentration realisation field produced by the system. By integrating these instantaneous concentration realisation fields over the simulation duration a challenge dosage can be calculated, against which modelled dosage predictions can be compared in the evaluation.

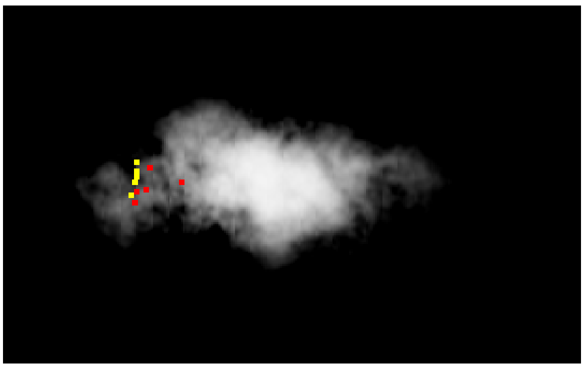

Figure 5: An example concentration realisation field with modelled sensor alarm status (yellow – non-alarmed, red – alarmed).

3.2 *Comparison Technique*

For this evaluation we have selected the Measure of Effectiveness (MOE) approach, described by Warner (2004). This is well suited to comparing contour areas and determining the degree of under- and over-prediction. The MOE is calculated as a coordinate using the expression:

$$MOE = \left(\frac{A_{OV}}{A_{OB}}, \frac{A_{OV}}{A_{PR}}\right) \qquad (1)$$

where A_{OB} is the area of the challenge dosage (i.e. observation) above a threshold value, A_{PR} is the area of the predicted dosage above the threshold and A_{OV} is the area in which both dosages are above the threshold. Figure 6 shows a schematic for challenge dosage plume (the ground truth or observation) and the predicted dosage plume, and Figure 7 provides an interpretation of the MOE coordinate. In the evaluation we also consider the distance of MOE coordinate to the coordinate (1, 1), which represents the perfect match, which we have called the MOE distance to facilitate comparisons between the approaches.

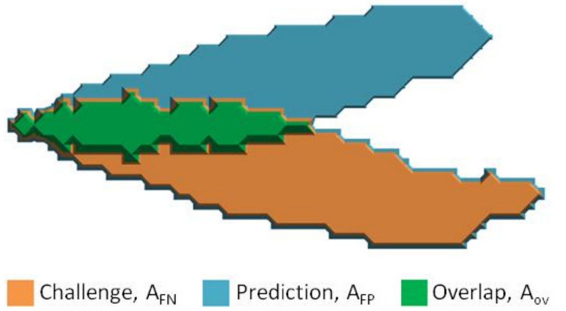

Figure 6: Schematic of plumes used in MOE calculation – orange represents area of challenge dosage (observations) only i.e. false negatives, A_{FN}; blue the predicted (modelled) dosage only, i.e. false positive, A_{FP}; and green is the overlap area of the two, A_{OV}.

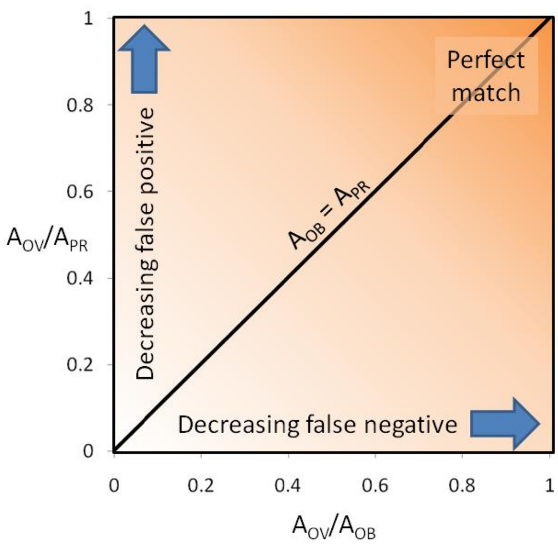

Figure 7: Graph depicting interpretation of MOE values.

3.3 *Description of Evaluation Scenarios*

The scenarios used for evaluation are based in Bristol, UK. They are based on a 5 km square domain. The Evaluation System was used to generate concentration challenge realisations that were used to provide the sensor input and the challenge dosage. The releases were 5 kg of a chemical agent with an initial source size of 1 m (all three Gaussian puff spread parameters were set to this). The challenges were Monte Carlo sampled from:

- Four threat location distributions – one focused on a open space in front of the local government headquarters (A), another in a square in a commercial area (B), a third widely dispersed but centred in the city centre (C) and the fourth along the river route into the city (D).
- Five different meteorological conditions based on annual meteorology for Bristol, two climatology data sets for two separate months, a three day weather forecast and a uniform meteorological distribution.

For each of these 20 combinations, ten separate samples were taken, providing a total of 200 different challenges. Figure 8 shows the release locations for all 200 releases.

The sensors used in the study represented typical chemical sensors and they were placed using the SPARTA automated sensor placement tool, described by Griffiths (2010). Placements were made for each of 20 threat-meteorological combinations. A total of 60 modelled sensors were placed, although on average less than ten of the instruments measured a detectable amount of agent in each simulation.

Figure 8: Bristol showing release locations (in red) and threat areas: A – College Green, B – Queen's Square, C – dispersed threat area, and D – River Avon.

4. RESULTS

For each of the 200 different challenge scenarios, the following runs were made for the comparison:

1. A release from the true source location and time – the True Release approach.
2. A release from the location and time of the first sensor to raise its alarm status – this is the Doctrine approach.
3. The Nowcast assimilation method. As this provides updated estimates over time, a hazard dosage prediction was selected for a time between the time of the first and last sensor observation.

In each case the models were provided with the actual meteorological conditions used in the challenge simulation.

The MOE results for each of these runs are shown in Figure 9. This shows that all three approaches are clustered near to the (1, 1) perfect match. The following observations can be made:

- The True Release dosages are not a perfect match to the Challenge dosages. This is because the Challenge is a single realisation whereas the true release prediction represents an ensemble average, and so the dosages would be expected to be different. The graph shows that the True Release cases have a tendency to provide predictions with false positives, i.e. over predict the hazard area, which is what would be expected.

- The Doctrine approach (assuming a release point at the first sensor alarm) does reasonably well. In these scenarios, this is not altogether surprising as the Sparta sensor placement capability has optimised the sensor locations and minimised the distance between the releases and the upwind sensors. However, it is discernable that it performs the worst of three methods.

- The Nowcast performs well. If one considers the average distance from the ideal (1, 1) coordinate, we find that the distance is 0.196 for the True Release approach, 0.216 for the Doctrinal approach and 0.166 for the Nowcast method. In fact in 144 of the 200 cases, Nowcast provides the best estimate

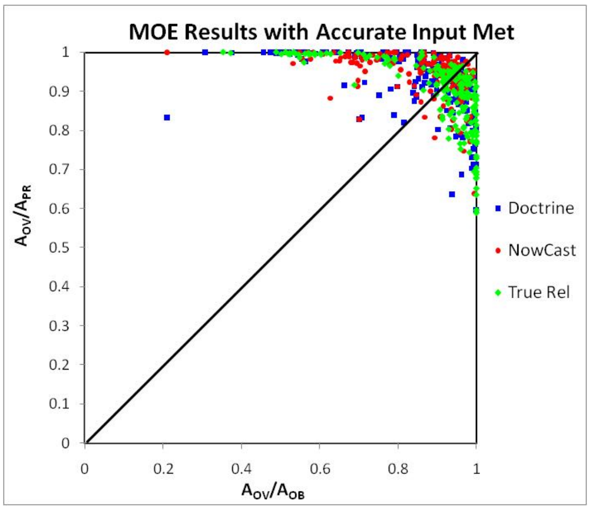

Figure 9: MOE results for Doctrine, Nowcast and True Release dosage predictions compared to Challenge Dosage, and where accurate input meteorology is provided to the predictive models.

in terms of closeness of match to the challenge. This suggests that Nowcast is performing better than the Doctrinal approach and better even than the True Release approach, which is somewhat surprising as the True Release results were originally the target for the Nowcast performance. It is likely that the Nowcast algorithm is refining the prediction based on the data from the challenge realisation to an extent that it out performs the True Release modelling.

Of particular interest in this work was whether the Nowcast data assimilation approach could compensate for errors in the meteorological input, which is a significant operational issue. The scenarios were rerun but this time with a random error between 10 and 30 degrees introduced into the meteorological input provided to the modelling. The results are shown in Figure 10. As would be expected the overall results when we provide erroneous meteorological input deteriorate substantially, however it is clear that the Nowcast results are significantly better. In particular:

- The True Release and Doctrine results are poor and reflect the error in the input wind direction. The average distance to the (1, 1) ideal coordinate is 0.831 for the True Release approach and 0.771 for the Doctrine approach. This suggests that the Doctrine based approach is slightly better, which is not unexpected as its use of the first sensor alarm location guarantees some degree of overlap with the Challenge plume.

- The Nowcast results are much closer to the perfect match (1, 1) coordinate. The average distance of the Nowcast MOE points to (1, 1) is 0.278, which is much lower than both the True Release and Doctrine approaches, and Nowcast provides the best estimate in 193 of the 200 cases in terms of closeness to exactly matching the Challenge dosage. This demonstrates that Nowcast is assimilating the sensor data and providing an improved meteorological estimate as this is used by the coupled dispersion model.

- In some cases Nowcast is not performing as well. Although further analysis is required it is likely that this is due to a small number of sensors producing a concentration reading; in almost a third of the scenarios, five or less sensors measure a non-zero concentration. This provides little information for the Nowcast algorithm, although as can be seen in the previous point, Nowcast still out performs the other approaches in over 95% of the scenarios.

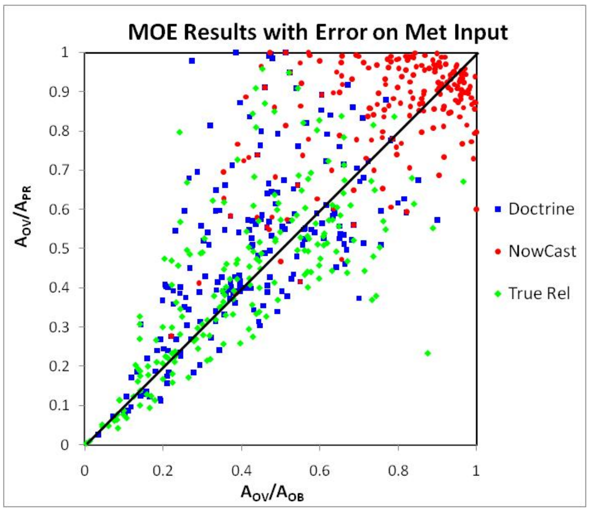

Figure 10: MOE results for Doctrine, Nowcast and True Release dosage predictions compared to Challenge Dosage, and where there is a random error of between 10^o and 30^o in the meteorological input provided to the predictive models.

5. CONCLUSIONS

A sensor data assimilation capability has been developed that uses the E-M algorithm. It fuses CB sensor observations in real time to produce immediate now-cast estimates of the current hazard. This has been linked to a rapid dispersion model to provide hazard predictions that update over time as additional sensor information is received. The algorithm also estimates the current local meteorology and this can be supplied to the dispersion model.

Comparisons with two other methods – the current standard doctrinal approach used by the military and using the true release location – have been made using a synthetic challenge generator. This has demonstrated that the Nowcast approach provides improved hazard predictions compared to the other approaches, even when these are provided with information that would not generally be available: the exact release time and location (for the true release approach), the mass of release and also accurate local meteorology. Despite this, Nowcast outperforms these approaches, even with relatively modest numbers of sensors (typically less than ten).

In the situation where the local meteorology is not known exactly, Nowcast is seen to provide significant benefits over the other two approaches. It is able to estimate the local meteorological conditions and enable the dispersion model to use this to provide improved hazard prediction. In these more operationally realistic circumstances, Nowcast provides hazard estimates that are much closer to the true hazard. This indicates that the Nowcast approach is resulting in improved situational awareness.

6. REFERENCES

Bull, M.D., and Griffiths, I.H., 2009: A Flexible CB Evaluation System. Presented at *13th Annual George*

Mason University Conference on Atmospheric Transport and Dispersion Modelling, 14-16 July 2009.

Dempster, A.P., Laird, N.M., and Rubin, D.B., 1977: Maximum Likelihood from Incomplete Data via the EM Algorithm. *J. of Royal Statistical Society*, Series B (Methodological) 39 (1), 1–38.

Griffiths, I.H., Bull, M.D., and Bush, I, 2010: Evaluation of Sensor Placement Techniques. Presented at *90th AMS Annual Conference*, Atlanta, Georgia, 17-21 Jan 2010.

Warner, S., Platt, N., and Heagy, J.F., 2004: User-Oriented Two-Dimensional Measure of Effectiveness for the Evaluation of Transport and Dispersion Models. *J. Appl. Meteor.*, 43, 58-73.