H14-283

THE ROLE OF ATMOSPHERIC DISPERSION MODELLING IN MODERN CONSEQUENCE ASSESSMENT STUDIES

Jeffry Urban, Michael Ambroso, Keith Galvin, Steven Nunes, Dennis DeRiggi, and Nathan Platt

Institute for Defense Analyses, Alexandria, Virginia, USA

Abstract: The ability to predict the consequences of an atmospheric release of a toxic material accurately, as opposed to conservatively, is of increasing importance to decision-makers who need to plan for or respond to such incidents. A conservative estimate of casualties, for example, may involve multiplying the average population density by the area over which a conservatively-chosen threshold atmospheric concentration is exceeded on average. We have begun to investigate the requirements for producing accurate consequence estimates using modern atmospheric dispersion models, including the capabilities and limitations of different end-to-end consequence modelling approaches. Some issues of concern include the use of ensemble-average concentration predictions in conjunction with toxicity models, especially advanced models involving the toxic load; the fact that different methods exist for calculating the toxic load from time-varying exposures but none have been experimentally validated; and the potentially inapt comparison of atmospheric dispersion model predictions to the consequences of real-world historical incidents. The work presented here focuses primarily on the problems associated with the choice of toxic load model and the potential use of individual realizations of a release in place of ensemble-average concentration predictions for consequence estimation. Our study of individual releases from the FUSION Field Trial 2007 (FFT07) experiment and from VTHREAT-generated predictions of a small-scale chemical attack indicate that casualty estimates can have a substantial dependence on the choice of toxicity model. We have also probed the limitations of using ensemble-average concentration predictions for casualty estimation and the potential use of low-order concentration moments to improve those estimates.

Key words: *Hazard prediction, casualty estimation, toxic load, dose-response, probit, ensemble average plume, HPAC, VTHREAT*

INTRODUCTION

A common use of atmospheric dispersion models is to estimate the potential health hazard that would arise from an atmospheric release of a toxic material. These estimates frequently have involved calculating the number of people who may be exposed, on average, to a threshold level of the toxic material. Although these methodologies are often intended to provide conservative estimates of the number of people affected, users of dispersion modelling-based consequence assessment tools are increasingly finding themselves in need of accurate (not conservative) estimates of the range of casualties that might be expected. Accurate estimates of the range of casualties require accurate knowledge of not only the atmospheric dispersion, toxic effects, and spatial distribution of the population, but also of the statistical distribution of those quantities, such as the stochastic effects of atmospheric turbulence on dispersion or uncertainties in the distribution of the population. Furthermore, the proper end-to-end modelling of casualties requires that these calculations be combined in such a way that preserves the distributions of the random variables and correlations between them. For example, a toxicity model that is sensitive to the time dependence of the exposure may require knowledge of turbulence-induced fluctuations around the ensemble average concentration. Our recent research has indicated that certain end-to-end modelling schemes in use today actually may inadvertently underestimate casualties (Czech, C. et al., 2011).

Our current research has focused on developing a conceptual framework for an internally-consistent end-to-end calculation of casualties using the outputs of state-of-the-art atmospheric dispersion models in combination with toxicity and population models. The work presented here focuses on two areas: the differences between casualty estimates generated from predictions of chemical concentrations averaged over the ensemble of turbulent realizations of the chemical plume and those generated from predictions of individual plume realizations, and the differences between casualty estimates generated using different toxicity models that are presently in use or have been proposed to be used in hazard assessments. Other topics of investigation include the effect of uncertainties in the population distribution or in estimates of the toxicological parameters on hazard assessments and the proper way to compare the predictions of atmospheric dispersion models to the historical record of chemical releases.

OVERVIEW OF CASUALTY ESTIMATION

Hazard prediction models frequently are used to provide estimates of either hazardous areas or the number of casualties that may result from the release of a hazardous material. The work presented here focuses on estimating casualties resulting from toxic inhalation exposure. The traditional method of estimating such casualties uses the lognormal (probit) dose-response model, in which the probability of an individual being injured is estimated based on the inhaled dose. The inhaled dose is presumed to be proportional to the total dosage accumulated over the course of the exposure at the location of the individual. The relationship between the dosage and the probability of casualty may be stated as follows (Finney, D.J., 1947),

$$P[D(r)] = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{x_0} \exp\left[(x-\mu)^2 / (2\sigma^2) \right] dx \,, \tag{1}$$

$$x_0 = \log(D(r)) \,. \tag{2}$$

Here, $D(r)$ is the dosage at location r and x is a parameter associated with the normally-distributed random variable X, which is equated with the lognormally-distributed dosage ($X = \log(D)$). σ is the variance of the population response around the mean μ, which is usually associated with the logarithm of the median effective dose ED_{50} ($\mu = \log(ED_{50})$). When Eq. 1 is linearized through a probit transformation the toxicity relationship may be expressed as follows,

$$Y_P[D(r)] = 5 + m_P \left[\log(D(r)) - \log(ED_{50}) \right] . \tag{3}$$

Here, Y_P is the probability of casualty expressed in probit units and the probit slope is $m_P = 1/\sigma$. Casualty assessments that use the dose-response model typically assume that since it is uncertain how susceptible a given individual at location r is to

the toxic dosage $D(r)$, the probability of casualty may be equated with the fraction of individuals who would suffer casualties if exposed to a comparable dosage (as estimated, usually, from the results of animal experiments). Individual susceptibility within the population is assumed to be a lognormally-distributed function of the dosage. It may be noted that the probability of casualty is not a simple linear function of the dosage.

In the case of steady exposures, the dosage is defined as the atmospheric concentration at which the exposure occurs, $C(r)$, multiplied by the duration of the exposure, T.

$$D(r) = C(r)T \tag{4}$$

Substances whose toxic effects can be described by Eqs. 3 and 4 are said to follow Haber's Law of toxicity. Many modern atmospheric dispersion models attempt to predict the concentration history at a particular location. Real-world exposure profiles contain temporal fluctuations and even intermittency due to turbulent diffusion and plume meandering or time-varying releases. In the case of a time-varying exposure profile, the dosage is frequently generalized to the time-integrated concentration,

$$D(r) = \int_0^T c(r,t)dt \tag{5}$$

It should be noted dose-response relationships like those expressed in Eq. 3 are parameterized using data from animal experiments that involve steady exposures, so there is little empirical basis for assuming that Eq. 3 is valid when dosage is defined as in Eq. 5.

The toxic effects of certain industrial chemicals and chemical warfare agents have been found to deviate from Haber's Law. Their effects are better described by a lognormal response to a phenomenological quantity called the toxic load, which for steady exposures is defined as,

$$TL(r) = [C(r)]^n T. \tag{6}$$

The toxic load plays the role of a generalized dosage, such that $TL(r)$ may be substituted for $D(r)$ in Eqs. 1 through 3. The population response is characterized by three toxicity parameters, TL_{50}, m_P, and n. When the toxic load exponent, n, is greater than 1, the toxic effects of high-intensity concentration fluctuations are amplified relative to the effects of long exposure durations, T. There is no empirical basis for extending Eq. 6 to the case of time-varying exposure profiles and there is no accepted generalization of this equation, although modern consequence assessment models commonly – but not universally – employ a generalization suggested by ten Berge (ten Berge, W.F. and M.V. van Heemst, 1983),

$$TL(r) = \int_0^T [c(r,t)]^n dt\ . \tag{7}$$

Other generalizations of the toxic load model are discussed below. The fact that the integrand in Eq. 7 is not a linear function of the concentration when $n \neq 1$ poses some challenges for calculating toxic effects. First, the effects of biological phenomena such as respiration and tissue uptake suggest that there is some timescale τ over which the exposure profile $c(r,t)$ should be time-averaged before evaluating Eq. 7; the exponent in Eq. 7 makes the evaluation of the toxic load sensitive to this averaging time in a way that Eq. 5 is not (Sommerville, D.R. et al., 2006). It has been argued that the respiration rate places a lower bound on the value of τ, but its actual value – which may depend on the nature of the toxic substance – generally is not known. For similar reasons, Eq. 7 is particularly sensitive to the use of ensemble-averaged concentrations (a common output of modern atmospheric dispersion models) in place of individual concentration realizations; this problem is discussed below.

THE CHOICE OF TOXIC LOAD MODEL

The ten Berge generalization of the toxic load equation (Eq. 7) is not the only method that has been proposed for calculating toxic loads from time-varying exposure profiles. Other methods that have been proposed (Czech, C. et al., 2011 and Sommerville, D.R., 2006) include equating the steady concentration in Eq. 4 with the time average of the concentration over the exposure interval ("average concentration method", Eq. 8) (Hilderman, T.L. et al., 1999); equating the steady concentration with the peak concentration over the exposure interval ("peak concentration method", Eq. 9) (Stage, S.A., 2003); using a definition of the toxic load that additionally uses information about the concentration variance (we refer to one such method that has been proposed for use with the HPAC modelling system developed by the US Defense Threat Reduction Agency (DTRA) as the "concentration intensity method" (Sykes, R.I. et al., 2007), which we have modified to apply to individual plume realizations (Eqs. 10-11); and using a definition of the toxic load that accounts for concentration intermittency (we consider one such method called the Griffiths and Megson model, Eq. 12 (Griffiths, R.F. and L.C. Megson, 1984, Griffiths, R.F. and A.S Harper, 1985)).

$$TL_{AverConc}(r) = \left(\frac{1}{T}\int_0^T c(r,t)dt\right)^n T \tag{8}$$

$$TL_{PeakConc}(r) = \frac{\int_0^T c(r,t)dt}{c_{Peak}^{1-n}(r)} \tag{9}$$

$$TL_{ConcIntens}(r) = \left(\frac{1}{T_{ConcIntens}(r)}\int_0^T c(r,t)dt\right)^n T_{ConcIntens}(r) \tag{10}$$

$$T_{ConcIntens}(r) = \frac{\left(\int_0^T c(r,t)dt\right)^2}{\int_0^T c^2(r,t)dt} \tag{11}$$

$$TL_{GriffithsMegson}(r) = \left(\frac{1}{1-(T-T_+)/T} \right)^{n-1} \left(\frac{1}{T}\int_0^T c(r,t)dt \right)^n T \qquad (12)$$

In Eq. 12, T_+ is the time during which $c(r,t) > 0$ over the exposure interval T. Since no method of calculating the toxic load has been validated for use with time-varying exposure profiles, it is important to be able to understand the effect of the method of calculating the toxic load on casualty estimates. Our research has been able to place some theoretical bounds on the problem (Czech, C. et al., 2011), but the results of applying these different toxic load models to real-world atmospheric dispersion data is also illuminating. Figure 1 shows the variability in the toxic load calculated concentration measurements made during the FUSION Field Trial 2007 (FFT07) short-range atmospheric dispersion experiment (Storwold, D.P., 2007). This experiment involved a number of different types of releases (continuous vs. puff, single vs. multiple, day vs. night, etc.). The differences between toxic load models can be substantial; for example, the difference between toxic loads calculated using the peak concentration model and the average concentration model can differ by several orders of magnitude. For the toxic load exponent chosen in this comparison, $n = 2.1$, the concentration intensity and ten Berge models nearly agree; this is an artefact of the choice of toxic load exponent, as we have demonstrated that these two models should agree exactly when $n = 2$. These results indicate that care should be taken when considering toxic load-based casualty estimates since none of these models have been validated using data from toxicological experiments. The exact impact on casualty estimates depends, of course, on whether the toxic load values fall near the TL_{50} value or far from it, where "near" is defined by the probit slope.

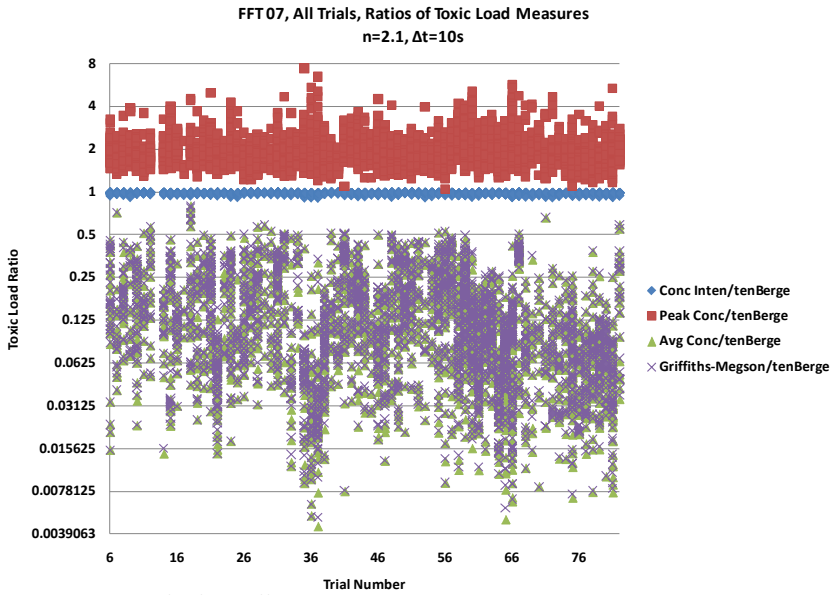

Figure 19. Ratios of toxic load values, for four different toxic load models compared to the ten Berge toxic load model, calculated from FFT07 field experiment data. The ratios are ordered by release (trial) number and each point represents data taken from an individual concentration sampling instrument. The toxic load exponent was taken to be $n = 2.1$ and a 10-second time-averaging was applied to each concentration time series prior to calculating the toxic load.

THE USE OF ENSEMBLE-AVERAGED PLUMES IN PLACE OF INDIVIDUAL PLUME REALIZATIONS

The hazardous plume resulting from the release of a toxic material will develop in an unpredictable manner due to random atmospheric turbulence. If the atmospheric turbulence is characterized statistically, however, it is possible to predict the development of the plume in an average sense or to predict "typical" (not "actual") realizations of the plume. Most atmospheric dispersion models that are presently used for consequence assessment modelling predict only the low-order statistical moments of the distribution of the spatiotemporal concentration field $c(r,t)$ (here referred to as the "plume") over the ensemble of possible turbulent realizations of the plume. Most models predict only the lowest-order concentration moment, $\overline{c(r,t)}$, or the lowest-order dosage moment, $\overline{D(r)}$, where the overbar indicates an average over all members of the ensemble. Others, such as the SPICPUFF model that drives the HPAC modelling system developed by the US Defense Threat Reduction Agency (DTRA) (Sykes, R.I. et al., 2007), additionally predict the second-order moments $\overline{[c(r,t)]^2}$ and $\overline{[D(r)]^2}$. It should be noted that the shape of the distribution of concentrations at a particular location may, in general, differ from the shape of the distribution of dosages.

Since the output of a typical atmospheric dispersion model is an ensemble-averaged concentration field, casualties are usually estimated using Eqs. 3 and 5 (or the corresponding equations for the toxic load) with $\overline{c(r,t)}$ in place of $c(r,t)$. Therefore casualties are estimated from the ensemble-averaged plume, which is an entity that is never observed in nature. What is actually of interest in casualty estimation is the statistical distribution of casualties over the ensemble of turbulent plume realizations – perhaps described by low-order moments such as the mean number of casualties and their variance – rather than the number of casualties resulting from the lowest-order moment of the concentration (the ensemble-average plume). The typical method of calculating casualties using atmospheric dispersion models is summarized on the left of Eqs. 8 and 9, whereas the actual quantity of interest is shown on the right,

$$\overline{c(r,t)} \rightarrow D[\overline{c(r,t)}] \rightarrow P\{D[\overline{c(r,t)}]\} \rightarrow Cas[P\{D[\overline{c(r,t)}]\}] \neq \overline{Cas[P\{D[c(r,t)]\}]}\,, \quad (8)$$

$$\overline{c(r,t)} \rightarrow TL[\overline{c(r,t)}] \rightarrow P\{TL[\overline{c(r,t)}]\} \rightarrow Cas[P\{TL[\overline{c(r,t)}]\}] \neq \overline{Cas[P\{TL[c(r,t)]\}]}\,. \quad (9)$$

Here, *Cas*[...] represents the integral of the probability of casualty at location *r*, *P*(*r*), over the spatial population distribution to estimate the total number of casualties. It should be noted that if there is population movement during the exposure, the spatial population distribution needs to be accounted for in an earlier stage of the calculation. In general, the full statistical distribution of casualties cannot be calculated unless some assumptions are made about the statistical distribution of the concentration, dosage, or toxic load. Our research has indicated that if the statistical distribution of concentrations is well-described by the two lowest-order concentration moments along with a simple assumption about temporal correlations, then the statistical distribution of casualties may be calculated in a straightforward manner when Haber's Law holds (Ambroso, M. et al., 2011). The calculation is not so straightforward when toxic load modelling is employed, particularly in the case of some toxic load models that are more complicated than the ten Berge model (Eq. 7); we have yet to explore the full implications of toxic load-based casualty estimation using statistical distributions of the concentration field.

Our previous research has indicated that using the ensemble average plume (described by $\overline{c(r,t)}$) to calculate the toxic load may significantly underestimate the spatial extent of a hazard compared to using the ensemble-averaged toxic load calculated from the toxic loads for individual plume realizations (each described by a unique $c(r,t)$) to estimate the hazard area in the case where $n > 1$ (Czech, C. et al., 2011). This is because the ensemble averaging smoothes out the plume in time and space, reducing the high-intensity concentration fluctuations that contribute significantly to the toxic load in toxicity models such as the ten Berge model (Eq. 7). This is a case in which using a more sophisticated toxicity model (the toxic load in place of the dosage) actually may be capable of producing worse casualty estimates if used improperly (i.e., when applied to the ensemble-average concentration field predicted by many modern atmospheric dispersion models).

Figure 2 illustrates the variability in hazard area estimation that can arise when applying toxic load models to individual concentration realizations generated using the US National Center for Atmospheric Research (NCAR) VTHREAT computational fluid dynamics (CFD)-based simulation environment (Bieberbach, G. et al., 2010). VTHREAT was used to simulate the simultaneous release of 18 instantaneous puffs of 1.6 kilograms of a neutrally-buoyant gas under convective atmospheric conditions. The 18 puff sources were spread over an area of approximately 100 meters by 200 meters. Each point in Figure 2 represents, for one of the 18 sources, the ratio of the area exceeding a particular toxic load value calculated using one of three toxic load models to the area exceeding the threshold calculated using the ten Berge model. The colour of each point indicates which method of calculating the toxic load is being compared to the ten Berge calculation. The area exceeding the threshold is meant to represent the area over which an inhalation hazard might arise from a particular source. The horizontal axis of the figure corresponds to different choices of the threshold toxic load, with lower toxic load thresholds (larger absolute hazard areas) on the left and higher toxic load thresholds (smaller absolute hazard areas) on the right. The coloured vertical lines indicate the range of toxic loads over which the hazard is expected to be lethal (TL_1 to TL_{99}) if the source mass is scaled to the values indicated (300 lbs., 1 ton, and 10 tons) and the toxicity parameters (TL_{50}, m_P, and n) are taken to be those of chlorine.

Figure 2 indicates that the size of the lethal hazard area can sometimes vary significantly when comparing closely-spaced identical puff sources. A model that predicts the ensemble-average concentration would not capture this variation; each source would yield identical average concentration predictions if the wind field were uniform. Also, one can again see that toxic load-based hazard predictions can vary significantly among the different methods of calculating the toxic load, with the peak concentration method tending to yield the largest toxic load hazard areas and the average concentration method tending to yield the smallest hazard areas. The variation in hazard area size among toxic load models also has a dependence on the release mass.

CONCLUSIONS AND FUTURE WORK

Many consequence assessment studies result either in a single point estimate of casualties or a range of estimates based on different release or meteorological conditions. We have suggested, based on theoretical considerations and CFD simulations of individual realizations of chemical releases, that turbulence-induced variability in atmospheric concentrations may be an important source of variability in casualty estimates. Particularly problematic is the use of atmospheric dispersion models that predict only ensemble average concentrations in conjunction with toxicity models. This methodology is incorrect even for dosage-based casualty estimates and the problem is amplified in the case of toxic load-based casualty estimates; counter-intuitively, the "better" toxicity model (toxic load) has the potential to produce worse casualty estimates for highly fluctuating or intermittent exposures. We are presently investigating the degree to which predictions of the statistical distribution of casualties can be improved by including predictions of higher-order concentration moments or information about spatiotemporal correlations. We also have investigated another issue of concern with toxic load modelling, which is that no validated and accepted toxic load model exists for the case of time-varying exposures. The possibly-chemical-dependent concentration averaging time is not known experimentally and different proposed methods of calculating the toxic load can lead to considerable differences in casualty estimates.

It may be useful to consider other sources of uncertainty or variability in casualty estimates. For example, the toxicity parameters used in casualty estimation (ED_{50}, TL_{50}, m_P, and n) are often associated with considerable experimental uncertainty that could propagate into significant uncertainty in casualty estimates. Another source of uncertainty is the population distribution; in particular, high spatial resolution atmospheric dispersion modelling (such as in urban areas) could provide misleading casualty estimates if they were applied to the use of coarse-grained (low spatial resolution) population distributions. One other important area of research is the comparison of atmospheric dispersion model predictions to the

historical record of toxic chemical releases (deliberate or accidental). Atmospheric concentrations typically are not measured during real-world incidents, so such comparisons usually rely on comparing model estimates of the lethal area to observations of where people actually died. Such comparisons are fraught with difficulties, not least among them being the potential need to model the effects of high spatial resolution terrain elevation, land cover, and population distributions to capture adequately the effects of gravity-driven dense gas dispersion, vegetative filtration (possibly coupled with surface reactivity), or heterogeneous population distributions (possibly including the effects of self- or directed-evacuation and sheltering or other protective measures). Future studies of the sensitivity of consequence assessment modelling to these modelling assumptions and model inputs may be able to provide bounds on the confidence with which casualty estimates can be stated.

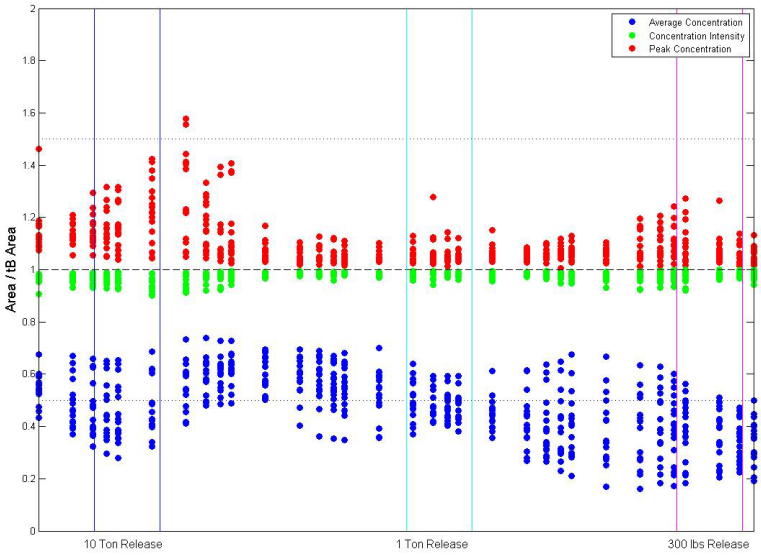

Figure 2. Ratios of the areas exceeding a given toxic load threshold (toxic load hazard areas), for three different toxic load models compared to the ten Berge toxic load model, calculated using VTHREAT simulation data of an 18 instantaneous puff source chemical attack under convective atmospheric conditions. Each point represents a ratio of toxic load hazard areas for one of the 18 sources. The toxic load threshold is varied from low to high across the horizontal axis and vertical lines indicate the range of lethal toxic loads for scaled-up releases of a neutrally-buoyant chemical with the same toxicity as chlorine.

REFERENCES

Ambroso, M., D. DeRiggi, N. Platt and J. Urban, 2011: Consistent Consequence Assessment in HPAC, presentation at the 15th Annual George Mason University Conference on Atmospheric Transport and Dispersion Modeling, Fairfax, Virginia, USA.

Bieberbach, G., Bieringer, P.E., Wyszogrodzki, A., Weil, J., Cabell, R., Hurst, J., and J. Hannan, 2010: Virtual chemical and biological (CB) agent data set generation to support the evaluation of CB contamination avoidance systems, The Fifth Symposium on Computational Wind Engineering (CWE 2010), Chapel Hill, North Carolina, USA.

Czech, C., Nathan Platt, Jeffry Urban, Paul Bieringer, George Bieberbach, Andrzej Wyszogrodzki and Jeffrey Weil, 2011: A comparison of hazard area predictions based on the ensemble-mean plume versus individual plume realizations using different toxic load models, Paper 2.5 at the Special Symposium on Applications of Air Pollution Meteorology, 91st Annual American Meteorological Society Meeting, Seattle, Washington, USA, January 2011.

Finney, D.J., 1947: Probit Analysis: A Statistical Treatment of the Sigmoid Response Curve, Cambridge University Press, London, United Kingdom.

Griffiths, R.F. and A.S Harper, 1985: A speculation on the importance of concentration fluctuations in the estimation of toxic response to irritant gases, J. Hazardous Mater., 11, 369-372.

Griffiths, R.F. and L.C. Megson, 1984: The effects of uncertainties in human toxic response on hazard range estimation for ammonia and chlorine, Atmos. Environ., 18, 1195-1206.

Hilderman, T.L., S.E. Hrudey and D.J. Wilson, 1999: A model for effective toxic load from fluctuating gas concentrations, J. Hazardous Mater., 64, 115-134.

Sommerville, D.R., Park, K.Y., Kierzewski, M.O., Dunkel, M.D., Hutton, M.I. and N.A. Pinto, 2006: Toxic Load Modeling, in Inhalation Toxicology, Second Edition, edited by H. Salem and S.A. Katz, Boca Raton, FL, USA, CRC Press, pp. 137-158.

Stage, S.A., 2004: Determination of Acute Exposure Guideline Levels in a Dispersion Model, Journal of the Air & Waste Management Association, 54, pp. 49–59.

Storwold, D.P., 2007: Detailed test plan for the Fusing Sensor Information from Observing Networks (FUSION) field trial 2007 (FFT-07), West Desert Test Center report, US Army Dugway Proving Ground, WDTC-TP-07-078.

Sykes, R.I., Parker, S.F., Henn, D.S., and B. Chowdhury, 2007. SCIPUFF Version 2.3 Technical Documentation, L-3 Titan Technical Report.

ten Berge, W.F., and M.V. van Heemst, 1983: Validity and accuracy of a commonly used toxicity-assessment model in risk analysis, Fourth International Symposium on Loss Prevention and Safety Promotion in Process Industries, Vol. 1, Institute of Chemical Engineers, Rugby, United Kingdom, pp. 11-112.