Computers and Electronics in Agriculture 88 (2012) 63–71

Contents lists available at SciVerse ScienceDirect

Computers and Electronics in Agriculture

journal homepage: www.elsevier.com/locate/compag

BREAM: A probabilistic Bystander and Resident Exposure Assessment Model of spray drift from an agricultural boom sprayer

Marc C. Kennedy [a,*], M. Clare Butler Ellis [b], Paul C.H. Miller [b]

[a] The Food and Environment Research Agency, Sand Hutton, York YO41 1LZ, UK
[b] Silsoe Spray Applications Unit, Building 42, Wrest Park, Silsoe, Bedford MK45 4HP, UK

ARTICLE INFO

Article history:
Received 25 November 2011
Received in revised form 5 July 2012
Accepted 11 July 2012

Keywords:
Uncertainty analysis
Monte-Carlo simulation
Bayesian emulator
Probabilistic sensitivity analysis
Crop spray model

ABSTRACT

Complex simulation models are available to predict the level of exposure to bystanders and residents after a crop spraying event. In this paper we consider a particle-tracking spray drift model whose input parameters define particular scenarios of interest. Model outputs based on fixed values for these inputs ignore natural random variation and therefore give no indication of realistic variation in exposures, nor do they quantify the probability of rare extreme exposures. We describe a probabilistic modelling framework that allows the effect of variability in the input parameters to be quantified. An efficient statistical method for approximating the spray drift model is used, by creating a statistical emulator. An additional statistical model is then used to link airborne spray outputs to bystander exposures based on measured data. Uncertainty and variability are quantified in this model component. Validation of our approach is considered in two stages: first the accuracy of the emulator is assessed, as a surrogate for the true spray model. Secondly, the overall probabilistic outputs are compared with corresponding field measurements. Results are presented for a selection of typical exposure risk scenarios for bystanders and residents, illustrating the potential to generate a richer source of information for decision-makers. Sensitivity analysis results suggest strategies to reduce risk, such as minimising boom height.

1. Introduction

Regulations exist to protect the health of residents and bystanders who may come into contact with agricultural pesticides during and after crop spraying. Computer simulations, provided they are sufficiently accurate, provide a cost-effective tool to supplement limited field experiments, for quantifying the risk to bystanders and residents from spray drift. They can also be used to study the influence of particular risk factors. The inputs to these models include environmental and operating conditions that define precise scenarios. In practice these are often unknown, and can also be highly variable. The results presented here relate to one-off exposures from a single spray event. Further extensions to account for possible repeated exposures, as required for a chronic risk assessment, are the subject of ongoing research.

In the UK, the existing approach for bystander exposure assessment has been empirical, based on the dataset of Lloyd and Bell (1983). This approach is of limited use where conditions are very different from those present during data collection. Modern nozzles and crop spraying practices, for example, are quite different from those of 1983. Mechanistic models provide more flexibility as they can estimate exposure under a wider range of conditions, by varying model input parameters. Examples of relevant drift models are reviewed in Butler Ellis and Miller (2010) and Teske et al. (2011). The use of conservatively parameterised drift models can be useful for initial screening within a tiered approach to risk assessment. Parameter values should correspond to a reasonable worst case risk scenario. Such models are simple to use and the conservative values allow for variability and uncertainty to be accounted for. However, when multiple parameters are given conservative values the resulting risk estimates may be unrealistic for higher tier assessments. The distribution of bystander exposures, induced by real variability in conditions, is of interest. From this the probability of exceeding any particular level of interest to risk managers can be estimated. Risk management options can then be considered for reducing this probability if necessary. The focus in this paper is the development of a probabilistic modelling framework that has 2 components to account for (i) variability in spray drift and (ii) uncertainty and variability in the relationship between spray drift and bystander exposure levels. In the first part, variability in model input parameters is propagated through a mechanistic model to produce a probability distribution of spray drift levels. It is not practical to compute the distribution directly from the spray drift model, due to the computational time required for each model run. Efficient statistical methods are therefore used to create a fast approximation to the model, allowing the

* Corresponding author. Tel.: +44 1904 462175.
E-mail address: marc.kennedy@fera.gsi.gov.uk (M.C. Kennedy).

0168-1699/$ - see front matter Crown Copyright © 2012 Published by Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.compag.2012.07.004

calculations to be performed quickly. The second model component maps airborne spray concentrations onto actual deposits on individual bystanders. A statistical model is fitted to previously measured exposure data. The model quantifies uncertainty, due to the limited data, and variability caused by natural variation in bystander behaviour and environmental conditions.

Our aim is to produce a general modelling framework for bystander and resident exposure that can be used for regulatory purposes, and so conservative assumptions are made wherever accurate information about a realistic scenario is not available. The methods presented are also designed and implemented in a way that allows alternative scenarios to be investigated quickly and easily. As far as possible, the main relevant sources of uncertainty and variability have been included, but many unquantified sources remain. To improve transparency in the modelling process, these have been listed separately in the Appendix, based on ideas developed in EFSA (2006). In the following section we introduce individual data and model components, ending in Section 2.8 with a description of the overall integrated model for exposure assessment. The integrated model is referred to using the acronym BREAM (Bystander and Resident Exposure Assessment Model). Section 3 gives a series of risk scenarios that are considered as examples and associated results. Finally, Section 4 has discussion on the model's potential for improving the management of drift and to minimise resident and bystander exposure.

2. Modelling framework

2.1. Spray drift model for airborne concentrations and ground deposits

There are three routes by which people can be exposed to pesticide spray drift: direct dermal contamination from airborne spray, inhalation of airborne spray and indirect exposure from coming into contact with contaminated surfaces. Particular subgroups of the population can experience quite different exposures due to different characteristics – adults have the greatest surface area and will therefore have higher levels of contamination than children, but they also have a greater bodyweight, and so the dose per kg bodyweight could be lower. Children are closer to the ground and therefore potentially closer to the source of exposure. Adults and children therefore need to be considered separately.

The Silsoe Spray drift model (Butler Ellis and Miller, 2010) is an evolution of the earlier model of Miller and Hadfield (1989). It is a deterministic model that simulates droplet trajectories from a boom of agricultural nozzles and has been validated against experimental data for a representative flat fan nozzle operated at 3.0 bar (denoted FF/110/1.2/3.0, Doble et al., 1985). This model predicts deposits of spray onto the ground and airborne spray concentrations at specified distances downwind of a sprayed field. In order to determine potential human exposure, five specific outputs are required for each downwind distance of interest:

1. Ground deposit, ml m^{-2}, G.
2. Adult inhalation factor, ml m^{-3} s, $IF_a = \frac{q(z_a)}{Ws(z_a)}$.
3. Child inhalation factor, ml m^{-3} s, $IF_c = \frac{q(z_c)}{Ws(z_c)}$.

where $q(z)$ is the total quantity of spray liquid (ml m^{-2}) passing through unit cross sectional area at a 'breathing' height z metres above the ground and $Ws(z)$ is the wind speed (m s^{-1}) at the same height. We set $z = z_a = 1.4$ m for adults and $z = z_c = 0.7$ m for children to account for height differences. The inhalation factors are subsequently multiplied by the breathing rate, B (m^3 s^{-1}) to determine the inhaled quantity of spray, based on the following assumptions:

(i) Air is breathed from a cross-sectional area, A, m^2 (in practice this may vary depending on breathing rate and wind speed, but is assumed constant).
(ii) The total amount of air passing through the cross section A is $A.Ws$ m^3 s^{-1}.
(iii) The proportion of all droplets that could be inhaled is therefore $\frac{B}{A.Ws}$.
(iv) The spray drift model predicts total quantity of spray per unit cross sectional area at distance z above the ground, $q(z)$.
(v) Total quantity of spray passing across the breathing area, A, is $q(z).A$.
(vi) Total quantity of spray actually inhaled: $I = \frac{q(z).A.B}{A.Ws}$, i.e. $I = \frac{q(z).B}{Ws}$.

This assumes that droplets of all sizes are inhaled with equal probability – in practice only the smallest droplets will be inhaled, so this provides a conservative estimate.

4. Adult airborne spray, ml m^{-2}, AS_a the mean quantity of spray per unit area, averaged over 0–z_a m above the ground in a vertical plane representing a standing bystander, and
5. Child airborne spray ml m^{-2}, AS_c the mean quantity of spray per unit area, averaged over 0–z_c m above the ground.

The mean airborne spray is then multiplied by the height of the bystander (approximated as 2 m for adults and 1 m for children) to determine the maximum potential spray that a bystander could be exposed to. The approach taken here does not attempt to relate bystander exposure either to body surface area, as has been done previously (e.g. Vercruysse and Steurbaut, 2001), or to projected cross-sectional area, since the actual relationships are likely to be complex, and these areas are not easy to determine. For simplicity, the exposure is considered to depend only on bystander height, which is equivalent to a bystander width of 1 m, and is a conservative value. A further step (Section 2.2) is then undertaken to take account of the 'spray collection efficiency' of a person.

2.2. Bystander exposure data

Airborne spray and bystander exposure are linked via a statistical model which will be described in Section 2.7. Relevant field data were compiled from various sources and is fully described in Butler Ellis et al. (2010) and Glass et al. (2010). They include measurements of bystander contamination in the range 1–12 m from the source, with varying meteorological conditions and both children and adults (assumed to be 1 and 2 m tall respectively). As seen in Fig. 1, the measurements cover approximately 3 orders of magnitude in both airborne spray and bystander exposure values. The mean relationship appears roughly quadratic on the log–log scale, and there is variability around this mean in the bystander contamination for any given airborne spray value. This variability is due to additional bystander behaviour and other conditions as mentioned in Section 1.

2.3. Modelling of uncertainty and variability

Our analysis quantifies both uncertainty (imperfect knowledge) and variability (randomness). The main uncertainty we aim to quantify in this paper is that related to the true relationship between airborne spray and bystander contamination. The data seen in Fig. 1 shows an increasing trend in the expected contamination as a function of spray, but the data are not sufficient to precisely estimate the true parameters of a mean function fitted to these data, or the level of variability around the mean at any given spray level. The true parameters are therefore *uncertain* quantities. A Bayesian model to quantify this uncertainty is presented in Section

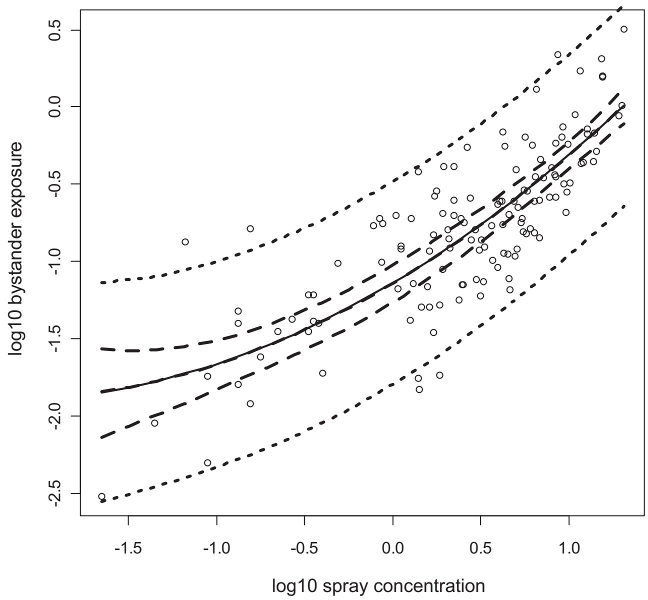

Fig. 1. Data relating airborne spray and bystander deposits, and summaries of the fitted quadratic model. The solid line is the mean relationship, the outer dotted lines correspond to the point-wise 95% probability interval of exposure for any given spray level, due to uncertainty and variability. The inner dashed lines are 95% probability interval of the fitted mean line, due to parameter uncertainty.

2.7. The Bayesian approach is a formal treatment of the process of learning from data (Gelman et al., 1995). A *prior distribution* is specified to represent an individual's subjective belief about an unknown quantity before observing some data or other information about it. After having observed the information, Bayes' Theorem is applied to update these beliefs from the prior to the *posterior distribution*. This process is applied repeatedly as more data are collected, and the weight of evidence from the combined data and subjective prior information is reflected in the final posterior distribution.

Variability is driven by real differences between spraying conditions, including climate, operator and field variables, and also between individuals' exposure to the spray. Factors influencing the latter include variations in bystander behaviour and characteristics, intervening landscape features and distance from the sprayer. We model these two types of variability sources differently as explained in Sections 2.4–2.6 and Section 2.7 respectively. Unlike with uncertainty, simply obtaining more information cannot reduce variability. An element of variability, due to short-term wind turbulence, is included within the drift model itself as a stochastic process. Our modelling approach also includes variability due to changes in the height of the boom, and variations in wind speed and direction during a spray application. Details of the probability distributions used for the inputs are given in Table 1, and the method of propagating this variability through to model outputs is explained in Sections 2.4–2.6. Other inputs that might in reality be variable were fixed to correspond to protective risk scenarios of interest (Section 3).

Separating uncertainty and variability is important as it can help to guide the risk manager in deciding whether to try and reduce the variability of effects, which would minimise the probability of extreme exposures, or to obtain more information in order to reduce uncertainty and thereby better understand the true scale of the problem before taking action. To summarise from the above discussion, fixed but uncertain parameters characterise the relationship between airborne spray and bystander deposits. Variable parameters including wind speed, wind direction and boom height. We use a 2 dimensional Monte Carlo (2DMC) implementation to separate uncertainty and variability (Cullen and Frey, 1999). The 2DMC algorithm loops through multiple realisations of the uncertain variables (the outer loop) generated from the posterior distribution of these variables. Conditional on each of these parameter values, a number of values are simulated for the variable quantities, from the assumed variability distribution (the inner loop). We used $I = 10{,}000$ outer loops and $J = 10{,}000$ inner loops. The result is a matrix of simulated exposures with I rows and J columns, covering all combinations of uncertain and variable realisations. Various summaries of this matrix are useful in highlighting the degree of uncertainty, for particular variability quantiles. Examples are presented in Section 4.

2.4. Monte-Carlo uncertainty analysis

Probabilistic uncertainty analysis of a model is a process of quantifying the probability distribution of the model output that is implied by uncertainty and/or variability in model inputs (Saltelli et al., 2000).

For a given code output, the standard Monte-Carlo (MC) method proceeds as follows:

1. Simulate a large number of input sets from user-specified probability distributions for variability in individual parameters. Particular distributions we assume in the BREAM model are shown in Table 1.
2. Evaluate the output for each input set, to yield a large sample from the probability distribution of the output.
3. Use the output sample to estimate summaries of the output distribution, such as the mean or standard deviation.

MC uncertainty analysis is attractive due to its simplicity and the wide range of possible input distributions available to represent realistic variability within a current or potential future scenario. However, thousands of code runs are typically required for accurate results, so it can be extremely inefficient and not practical with complex codes. A single run of the spray drift model, for example, takes around 4 min to run so it would not practical to generate many thousands of runs as part of a regulatory tool. We next describe the creation of a fast approximation to the code, which makes the MC approach feasible.

2.5. Bayesian modelling of computer models

A cheap but accurate approximation to any code output of interest can be created, using recent developments in Bayesian analysis of computer code outputs (BACCO). BACCO methods are based around the concept of an *emulator* of the computer code. This provides an accurate statistical representation of the relationship between the inputs and outputs and is created from a series of training code runs. It provides a cheap approximate version of the code (the emulator mean function) that runs in a fraction of the time, plus a description of the uncertainty due to the approximation. When run at a training input, the emulator output equals the true code output, with zero uncertainty. At other points, the emulator interpolates the training runs, with uncertainty dependent on the distance between the emulator input and training inputs. As more points are added to the set of training runs, this uncertainty decreases and the emulator mean increasingly adapts to the true shape of the output function. Efficiency gains of up to three orders of magnitude have been reported by using an emulator compared to more conventional MC methods. This is achieved by exploiting the fact that most deterministic codes are smooth functions of their inputs. Each training run therefore provides information about the output for a region of the input space, and this is quantified in the emulator through the use of a smooth

Table 1
Inputs to the BREAM user interface and/or emulators. Note that wind angle input is not required in the user interface, but is varied in the training samples to build the emulator. The required input 4 is dependent on choices made for inputs 2 and 3.

Input	Input name	Range for emulator	Comments related to user interface/model framework
1	Nozzle type	FF 110 03 nozzle at 3 bar	Each nozzle requires a separate set of emulators
2	Bystander type		Adult or child
3	Exposure route		Inhalation or dermal
4	Breathing rate for adults (B_a) or children (B_c) or dermal absorption (Abs)		Single value input
5	Wind speed at height of 3 m (km h^{-1})	0.5–25	Normal distribution assumed, based on input mean value and standard deviation 0.185 times mean +0.0068, estimated from empirical data
6	Boom height (m)	0.1–1.5	Normal distribution assumed, based on input mean value and estimated standard deviation (0.3 times mean)
7	Number of nozzles (multiples of 6)	6–996	Single value input
8	Crop height (m)*	0.05–2.0	Single value input
9	Distance from source (m)	1–15	Single value input
10	Wind angle (°)	10–170	Normal distribution: mean fixed at 90°, standard deviation 10°, based on empirical data. The mean corresponds to wind blowing directly towards the bystander
11	Forward speed (km h)	4–25	Single value input

* Crop height is fixed at 0.1 m for the special case emulators that assume a short crop.

correlation function. Any output is therefore correlated with each of the training runs, with correlations depending on the respective distances to those points. The highest correlation of 1 occurs at training points themselves, i.e. distance of 0, whereas more distance training points are less influential. Good designs for the training runs are space-filling, such as maximin Latin hypercube designs. A more detailed introduction to emulation methods is given in O'Hagan (2006) and examples are presented in Kennedy et al. (2006).

2.6. Building and validating the emulators

Using the GEM-SA software implementation (Kennedy, 2005), we created emulators of five different outputs from the Silsoe spray drift model as listed in Section 2.1 – ground deposit (G), adult inhalation factor (IF_a), child inhalation factor (IF_c), adult airborne spray (AS_a), and child airborne spray (AS_c). Many inputs of the original drift model were assumed to be constant, so the emulators were only functions of those remaining inputs that were believed to be influential on the outputs and/or highly variable. Table 1 lists the inputs of the BREAM model. Inputs 5–11 were selected as emulator inputs. Furthermore, separate groups of emulators were created to consider both a short crop scenario (crop height fixed at 0.1 m) and a scenario in which crop height can be varied. Ground deposit output close to the sprayed area was found to be unreliable for larger crop heights, in the sense that the emulator could not reproduce the true spray drift model output accurately, so the latter scenario was not considered for ground deposits. These 5 + 4 = 9 emulators were generated for the representative FF/110/1.2/3.0 nozzle. The model was run for a range of different input configurations to provide training data to build the emulators. These training inputs were designed to give an even coverage over the range of the seven inputs of interest (inputs 5–11 in Table 1), or six inputs for the short crop scenario. Ranges were chosen to represent feasible values. A total of 198 runs were used in each training sample, to cover the 7- or 6-dimensional input space. The LP-tau design option (Sobol, 1977) was used in GEM-SA to generate these points, to give a roughly even coverage of the ranges shown in Table 1.

The combined set of emulator definition files, generated as part of this building process, was then included as data to feed into the overall integrated exposure model (Section 2.8). Each of the model outputs is positive, but has values close to zero and positively skewed distributions. The natural log of the output was instead used, in order to give output values on the real line with reduced skewness. The log-transformed outputs can be better approximated with a normal distribution. Also, by back-transforming the log-scale predictions, non-negative exposure estimates are assured.

The accuracy associated with each emulator was checked using a cross-validation procedure described below. Results were similar across all emulated outputs, so here we report only the ground deposit output. We used the leave-one-out cross validation option built into GEM-SA, in which the emulator is used to estimate (and calculate emulator variance for) each of the training runs in turn. A slightly modified emulator is used to estimate each training point, built using all training points except the one being estimated. GEM-SA reports several diagnostic values, based on the results of the cross validation, that can be used to assess the quality of the emulator approximation and potential violations of the assumptions underlying the emulator theory. The following are computed, where n is the number of runs, y_i is the true output for the ith training run, $\hat{y}_i$, s_i are the corresponding emulator approximation and standard deviation calculated with the ith training point removed.

$$\mathrm{CVRMSE} = \sqrt{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2/n},$$
$$\mathrm{CVRMSRE} = \sqrt{\sum_{i=1}^{n}\{(y_i - \hat{y}_i)/y_i\}^2/n},$$

$$\mathrm{CVRMSSE} = \sqrt{\sum_{i=1}^{n}\{(y_i - \hat{y}_i)/s_i\}^2/n}.$$

These represent cross-validation root mean squared error, relative error, and standardised error respectively on the natural log scale for outputs. Values obtained from 198 training points were CVRMSE = 0.305, CVRMSRE = 0.839, CVRMSSE = 1.176. The standardised error is close to 1, suggesting that the emulator standard deviations are close to where they should be, given the actual emulator errors. The CVRMSE and CVRMSRE after back-transforming the training outputs and predictions to the natural exposure scale are easier to interpret. These were 0.283 and 0.395. For comparison, the raw outputs from the training runs range from 0.0107 to 9.417.

A more detailed validation study of the standardised residuals $(y_i - \hat{y}_i)/s_i$ contributing to the CVRMSSE can be carried out to identify exactly where in the input space the emulator's performance is least accurate. An example for the ground deposit output emulator is shown in Fig. 2. For a perfect emulator these would be

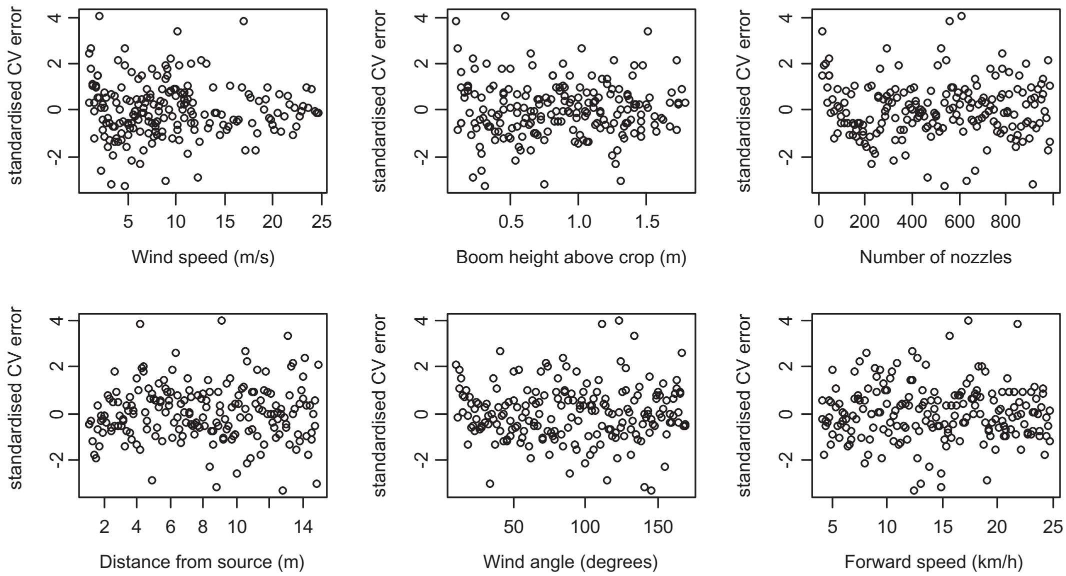

Fig. 2. Cross-validation standardised over-predictions of true training outputs on log scale, for ground deposit/FF110 03 nozzle emulator, plotted as functions of each input.

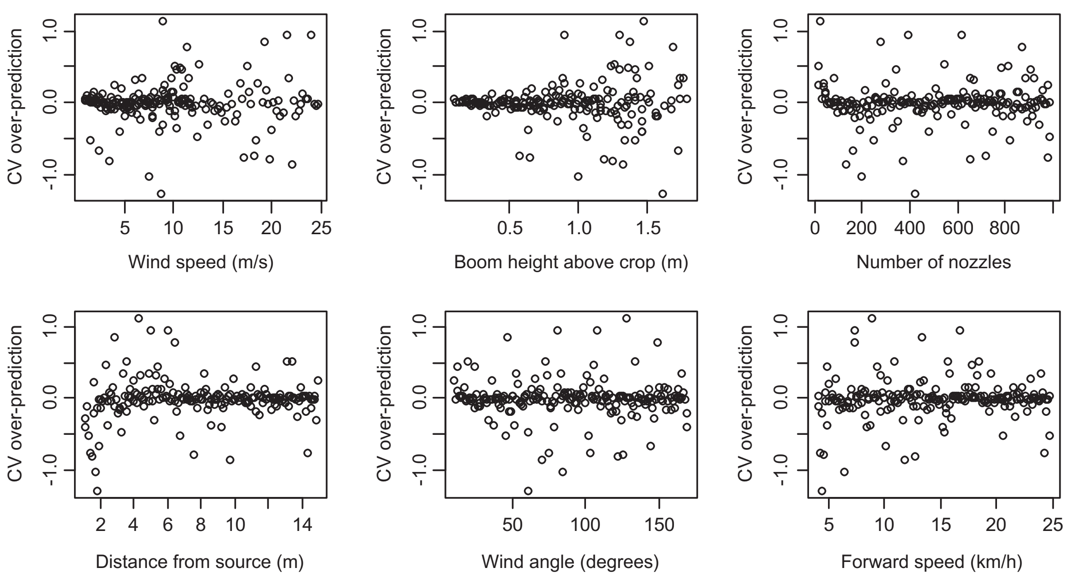

Fig. 3. Cross-validation errors on back-transformed scale, plotted against each individual input parameter.

approximately distributed as $N(0,1)$. From these plots there are no obvious regions of the input space where the emulator deviates noticeably from this distribution, although Fig. 3 suggests that the emulator is under-predicting at short distances and over-predicting when the number of nozzles is small, for example. We consider these cross-validation results are good enough to use the emulator in place of the original spray drift model, particularly if these extreme low inputs are avoided.

Various sensitivity measures can be derived directly from the emulator. For example, the *main effect* of the *i*th input x_i is the expectation, with respect to the distribution of the remaining inputs, of the output as a function of x_i relative to the overall mean output level. These main effects serve as an intermediate check on the relationship between input and output variables and show where variability in individual inputs is driving variability in the output. Oakley and O'Hagan (2004) describe the use of emulators for sensitivity analysis in more detail. Fig. 4 shows the main effects associated with $\log(IF_a)$ with the overall mean level added, to show the expected output level for fixed values of individual inputs. In this case, changes to boom height have the greatest influence, although actual variation during a spray event would be small relative to the full range shown here. Emulation uncertainty shown by the interval around this main effect is small, suggesting an accurate emulator. Taller crops generally lead to lower inhalation outputs and, unsurprisingly, a wind angle of around 90° (directly towards a bystander) generates the highest inhalation factor. This sensitivity analysis is a useful check that the model behaviour is intuitively correct, and it can be applied even in the absence of validation data. In general it can identify coding errors and lead to model improvements. Note that this sensitivity analysis exercise is separate from the final use of the model. At this stage inputs are assumed to have independent uniform distributions when calculating expectations. In interpreting results, it is important to recognise that this and other underlying assumptions may be

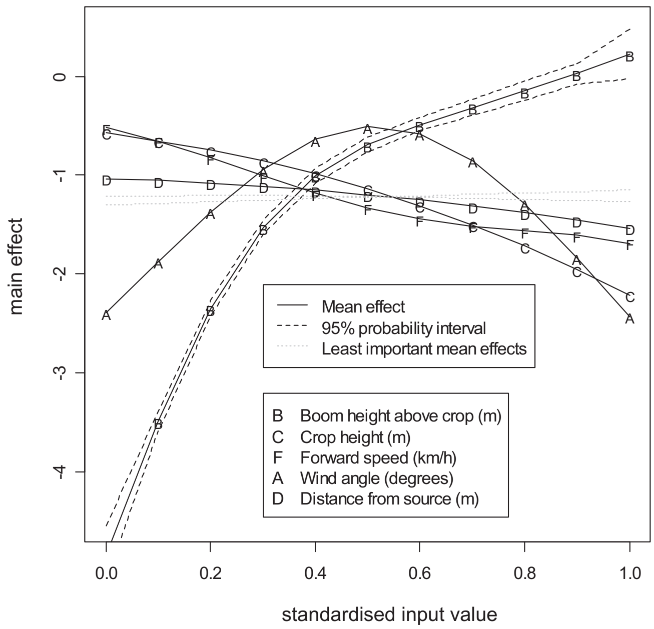

Fig. 4. Main effects for adult inhalation log(IF_a) on natural log scale (emulator assuming variable crop height). Bounds are shown for the most significant effect only, and represent uncertainty about the true code output, due to the emulator approximation. Uncertainty for the other effects is broadly similar.

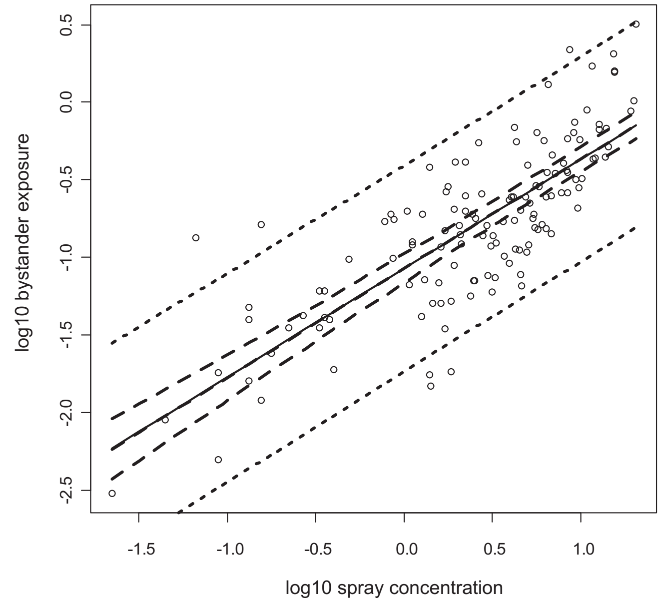

Fig. 5. Data relating airborne spray and bystander deposits, and summaries of a fitted linear model. The solid line is the mean relationship, the outer dotted lines correspond to the point-wise 95% probability interval of exposure for any given spray level, including uncertainty and variability. The inner dashed lines are 95% probability interval of the fitted mean line, due to parameter uncertainty.

unrealistic. High forward speed, for example, might be associated with higher spray rate in reality and therefore actual exposure would not necessarily be reduced when forward speed is increased.

The way in which the validated emulators are used within the overall BREAM model is explained further in Section 2.8.

2.7. Modelling bystander exposure data

Data for airborne spray versus bystander contamination was introduced in Section 2.2. For a given spray concentration AS (ml m^{-1}), we assume the quadratic regression model

$$\log_{10} BC = m(\boldsymbol{\beta}, \log_{10} AS) + \varepsilon = \beta_0 + \beta_1 \log_{10} AS + \beta_2 (\log_{10} AS)^2 + \varepsilon$$

to relate it to the corresponding bystander dose BC (ml). The ε term represents the deviation from the mean relationship on the log scale. It accounts for the combination of measurement error in the data and variability of bystander exposure for a fixed level of spray[1] Note that this variability is independent of the spray variability and will add to the overall variability in the outputs. We assume a Normal distribution $\varepsilon \sim N(0, \sigma^2)$. Thus, for a general airborne spray value AS, this model says that the expected $\log_{10}$ bystander exposure is $m(\boldsymbol{\beta}, \log_{10} AS)$ but due to variability the actual $\log_{10}$ exposure is distributed as $N(m(\boldsymbol{\beta}, \log_{10} AS), \sigma^2)$. The model parameters $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2)$ and σ^2 are unknown, and we assign the standard non-informative prior distribution expressed as $p(\boldsymbol{\beta}, \sigma^2) \propto \sigma^{-2}$. In this way, we do not introduce any subjective bias, instead allowing the data alone to derive inferences about the mean relationship and variability. Applying Bayes theorem, we obtain the posterior distribution for $\boldsymbol{\beta}, \sigma^2$ (see for example O'Hagan and Forster, 2004). A point estimate can be used to estimate the parameters, or we can generate multiple realisations from this distribution to reflect residual uncertainty, given the data. The latter is used in our 2DMC algorithm. In iteration i of the outer (uncertainty) loop of the 2DMC algorithm, a single simulated parameter set $(\beta^{(i)}, \sigma_i^2)$ is drawn from the posterior distribution. Then the corresponding inner loop randomly combines each of the spray outputs $AS_1, \ldots, AS_J$ simulated by MC (Section 2.4), with a new set of simulated bystander variability value $\varepsilon_{i1}, \ldots, \varepsilon_{iJ}$, to yield independent bystander doses $BC_{i1}, \ldots, BC_{iJ}$ using

$$\log_{10} BC_{ij} = m(\beta^{(i)}, \log_{10} AS_j) + \varepsilon_{ij}$$

where $\varepsilon_{ij} \sim N(0, \sigma_i^2)$ are simulated independently for $j = 1, \ldots, J$. Fig. 1 shows the 134 data points used to fit the model and summaries of the mean, uncertainty and variability of exposure for any given spray concentration based on the quadratic model. 1000 uncertainty simulations and 1000 variability simulations were generated to create the lines. The posterior mean for the regression parameters was estimated as (−1.1405490, 0.6796146, 0.1517974) and the mean precision was estimated as 9.435221. Attempts to fit a simpler linear model, without the quadratic term led to a smaller precision, meaning larger variability around the mean fitted regression line, as shown in Fig. 5. Notice also in Fig. 5 that at the highest observed spray concentrations the mean line is below all data points, and under-predicts bystander exposure compared to the quadratic model. Analysis of the residuals also suggested a relatively poor fit, so we decide to use the quadratic model as described.

2.8. Overall modelling framework

Based on the model components described above, two types of results were generated. The first quantified the effect of variability in boom height, wind speed and wind angle using Monte Carlo simulation. We call this the V-model (variability). Note that the distributions assumed (Table 1) are distinct from the earlier sensitivity analysis, which was carried out as a separate exercise. Uncertainty in the parameters for the relationship between airborne spray and bystander contamination was not considered in the V-model, as we fixed the parameters at their best-fit values and assumed them to be known. A user interface has been created to allow selection of the output of interest, and display the appropriate

[1] If we had replicate bystander measurements we could in principle separately estimate and subtract the measurement error, as when we predict we are only interested in the variability. However, by assuming it is all variability we will present more conservative estimates of bystander exposure at the prediction stage.

results. In this case the V-model is most suitable for computational speed. Secondly, we used the 2DMC algorithm to consider both uncertainty and variability, for the case of adult inhalation. The computational burden is much higher for this 2DMC approach, so is not practical for routine use and for multiple 'what-if' scenarios. We refer to this as the UV-model (uncertainty and variability).

The overall framework for the V-model is summarised in the following five steps:

- *Step 1*: (User input). The user selects appropriate values for inputs (1–9 and 11 in Table 1) for the required exposure assessment. For inputs 5 and 6 (wind speed and boom height), the selected value is used as an expected value. The other inputs are constant and can be set to define or compare alternative scenarios. A user may not know a precise value but can easily check the impact of alternative assumptions by running multiple calculations with best guess and worst-case conditions, for example. Wind angle is assumed to be distributed around 90°, as a conservative choice. This is because we do not expect more detailed information on wind angle to be available, and are considering the worst case where the bystander is directly downwind of the sprayer.
- *Step 2*: A Monte Carlo algorithm is used, first by simulating a large sample of values for inputs 5, 6, and 10 from their probability distributions (Section 2.3), then computing the corresponding set of emulator outputs (with other inputs fixed at the values chosen in Step 1). Each emulator output is computed using the appropriate emulator definition files created by GEM-SA (Section 2.6). For each of the 5 outputs (G, IF_a, IF_c, AS_a, AS_c), the sample of emulator runs is a Monte Carlo sample from the probability distribution of that output (Section 2.4). Summaries of these can be presented separately in Step 5, or used in subsequent probabilistic modelling in steps 3–4.
- *Step 3*: The statistical model described in Section 2.7 is used to relate measurements of airborne spray and bystander contamination. The model is fit using the data described in Section 2.2, and includes a measure of variability in contamination for any given spray level (i.e. the spread about the mean line fitted to the data). For the V-model, the posterior mean values of $(\boldsymbol{\beta}, \sigma^2)$ are used, whereas the UV-model repeats the process using multiple realisations generated from the posterior distribution. As we have a large sample of AS_a or AS_c points from Step 2, the result is also a sample from the distribution of the appropriate output BC_a for adult or BC_c for child.
- *Step 4*: Calculate: Inhalation for adult and child ($I_a = IF_a \times B_a$, $I_c = IF_c \times B_c$) where B_a, B_c are breathing rates of adults and children respectively; Dermal exposure D_a for adult and D_c for child ($D_a = BC_a \times Abs$, $D_c = BC_c \times Abs$) where Abs is a user specified clothing absorption factor; Total exposure for adult and child ($TE_a = I_a + D_a$, $TE_c = I_c + D_c$). Again, a sample of points is produced for each calculation, using simulated points from earlier steps.
- *Step 5*: From the generated samples, the mean and 95th percentiles of distributions for G, I_a, I_c, D_a, D_c, TE_a, TE_c are calculated and reported, as point estimates and high quantiles resulting from variability.

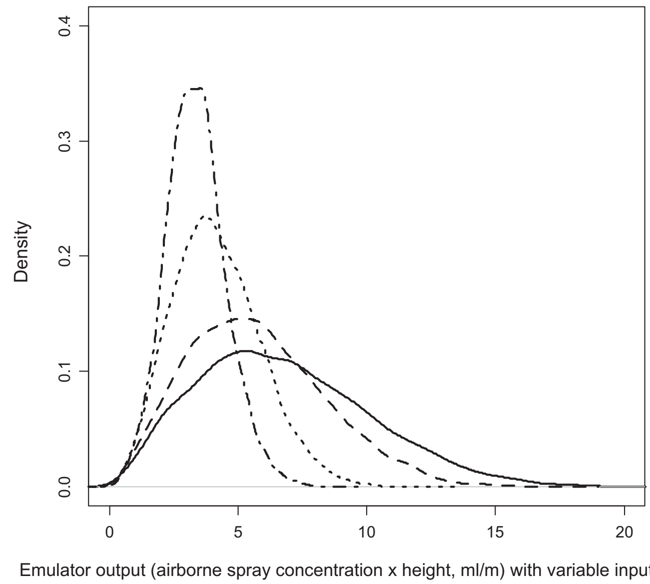

Fig. 7. Monte Carlo uncertainty analysis to propagate input variability through spray model. The 4 lines represent probability distributions of concentration at 1 m (solid), 2 m (dashed), 4 m (dotted) and 10 m (dash-dot) downwind from the sprayer.

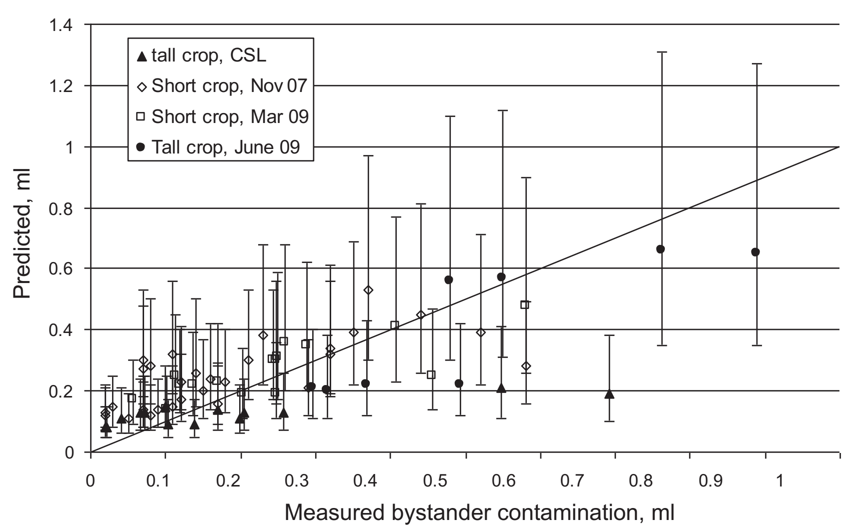

Fig. 6. Comparison between predicted and measured bystander contamination for four experimental data sets. Bars relate to the predicted 25th and 75th percentiles; data points relate to the predicted mean.

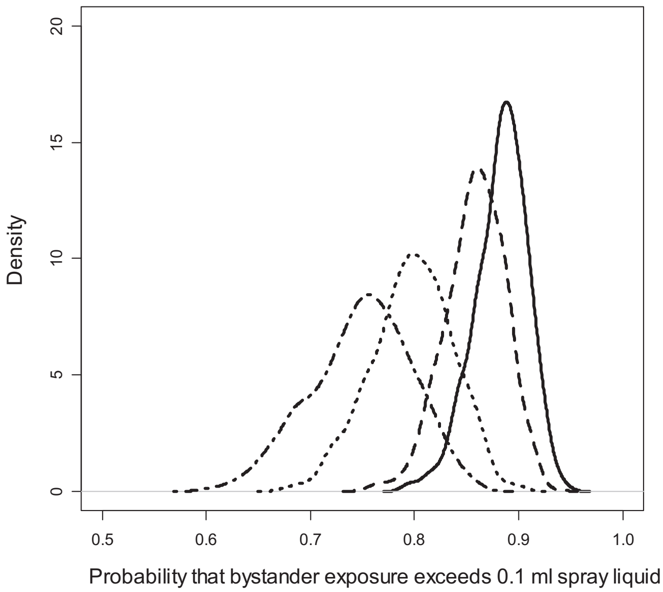

Fig. 8. Uncertainty distributions for the probability of dermal exposure exceeding a threshold of 0.1 ml spray liquid, with an adult bystander standing at 1 m (solid line), 2 m (dashed), 4 m (dotted) and 10 m (dot-dash).

The UV-model produces an $I \times J$ matrix of exposure values $\log_{10}BC_{ij}$ values as described in Section 2.7. More flexible inferences are possible using these outputs, beyond those listed in steps 1–5 above. As an example of a quantity that a risk manager may be interested in, for each $i = 1, \ldots, I$ an estimate $\hat{p}_i$ of the probability that bystander contamination exceeds a particular limit (e.g. 0.1 ml spray liquid) was calculated as the proportion of simulated BC_{ij} values exceeding this limit. The sample $\hat{p}_1, \ldots \hat{p}_I$, was then used to estimate the distribution (uncertainty) of the true probability of exceedances. To summarise, we estimate a true probability that is due to *variability*, but this probability is estimated from limited data and therefore has a probability distribution due to *uncertainty*. The latter should be reduced if we were to obtain more bystander data. The V-model can be used to generate a single $\hat{p}$ estimate without uncertainty.

3. Example BREAM results using test scenarios

The BREAM model was tested in a range of realistic but relatively high risk scenarios. First, the V-model was tested and prediction results were compared against field data reported by Glass et al. (2002, 2010) and Butler Ellis et al. (2010). These data were generated using a range of wind speeds, boom heights and forward speeds, and relate to spraying a single pass with a boom of width 20 or 24 m. Fig. 6 shows the results of the comparison, using the V-model and inputs matching the experimental conditions. For each data point we set the fixed BREAM inputs to match the corresponding field data conditions, and then applied the method of Section 2.8 to investigate the effect of variation in inputs and map the resulting distribution of spray onto a distribution of bystander contamination (Section 2.8). Variability distributions assumed for wind speed, boom height and wind angle around the stated values are listed in Table 1. A user-friendly interface has been created to generate outputs from the V-model, with user-defined inputs defining specific scenarios. The limits of possible inputs are defined by the ranges used to build the emulators (Table 1).

Figs. 7 and 8 illustrate some probabilistic outputs available from the V-model and UV-model respectively. Shown here are distributions for airborne spray for the intermediate spray model output AS_a (Fig. 7) and the associated probability that bystander exposure exceeds 0.1 ml spray liquid (Fig. 8) for downwind distances of 1 m, 2 m, 4 m and 10 m. This case is for an adult bystander and assumes a short crop, mean wind speed of 3 m s^{-1}, mean boom height 0.7 m, 48 nozzles (i.e. 24 m boom), and forward speed 12 km h^{-1}. These variables are realistic typical values for a spray application. The distributions clearly show that the probability of exceeding the exposure threshold reduces as the bystander moves further from the sprayer, but the distributions overlap due to uncertainty and at 10 m from the source this probability might be between 0.6 and 0.9. Similar plots could be produced to show the impact of varying forward and wind speed or the difference between adult and child exposures. As an illustration of the efficiency gains possible using the emulator, each of these distributions required 10,000 emulator runs and completed in 18 s. It was noted earlier that a single run of the original code takes around 4 min.

No comparisons were possible for inhaled spray because, as found from earlier work, measurements were generally below

Table 2
Uncertainty/variability table for all identified sources of uncertainty that were not included in the model. The ± symbols indicate whether the feature has the potential to make our exposure estimates higher (+) or lower (−) than the true exposures.

Source of uncertainty/variability	Direction and magnitude: over (+) or under (−) prediction
Bystander movement is not accounted for. Duration of exposure is probably worst case, inhalation includes all droplet sizes	±
The spray drift model is an approximation to reality. Model inadequacy is not accounted for but comparison of the model with real code outputs for specific inputs gives an indication of where the model under- or over-predicts.	±
The real variability distributions for wind speed, angle and boom height are unknown, and will also vary between spray events. Only relatively short term turbulence variability included. We also assume entire boom oscillates, rather than pivoting.	±
Measurement uncertainty in the bystander data is not addressed	±
The relationship between spray and bystander contamination may not be linear or quadratic	±
The different data sources are modelled as if they are from a common underlying 'population'. Differences in conditions and measurement methods will in reality lead to systematic differences.	±
Emulation uncertainty is not quantified in the final risk estimates. The cross-validation suggests that this is small, and errors tend to cancel out in the outputs generated by MC simulation.	±
Downwind structures not accounted for. CFD work showed highest factor of 3 increase, but other structures could reduce exposure	±
Vegetation effects not included. Limited evidence suggests that vegetation filters more than model predicts.	+
Overall conclusion: By leaving these aspects out of the analysis we expect our overall exposure estimates to be larger than the true values.	+

the level of detection (Lloyd and Bell, 1983). This was supported by model predictions which also were very low.

4. Conclusions and discussion

We have presented a general framework for estimating bystander and residue exposure resulting from spray drift. Particular attention has been given to the effects of variability due to boom height and wind turbulence, which is an improvement on the conventional use of deterministic point estimates. Fig. 6 shows that for available validation data, the resulting uncertainty intervals generally cover the true measured bystander contamination values, with a few exceptions. Most of these exceptions relate to measurements made when spraying over a tall (around 0.6 m) crop, where the version of the spray drift model used to develop the emulator was known to be less reliable. It should be noted that because the quantiles are 25th and 75th percentiles, we would expect 25% of actual residues to be below the lower quantile and 25% to be above the higher quantile. We have identified some additional uncertainties and variabilities that have not been quantified in the model, and these are listed in Table 2. Conservative assumptions have been built into the model so that taking a high percentile of exposure predictions should in most cases result in a (conservative) overestimate of actual bystander and resident exposure. The kind of outputs illustrated in Fig. 7 should be useful for risk managers to investigate potential risks and some factors that might be targeted to minimise this risk, such as distance between spraying areas and public spaces/residential areas. The sensitivity analysis results seen in Fig. 4 suggest that reducing boom height, and variations in boom height, also has potential to reduce exposure.

4.1. Further possible refinements

We have found that the initial emulators described above do not accurately reproduce the true outputs of the original spray model for all input values. This is not surprising, as a dense coverage of the 6- or 7-dimensional input space is difficult to achieve with a few hundred training code runs. The cross-validation diagnostics we have demonstrated give information about where additional runs should be focused in order to improve the emulator. Regions of the input space which are more important for the risk assessment will also be given priority in the next iteration of this process. A further possible direction for improving the emulation process is to combine the multiple emulators (e.g. for multiple crop heights and adult, child categories) within a treed Gaussian Process framework, as described in Gramacy and Lee (2008) and the extensions to categorical variables implemented in the tgp software of Gramacy and Taddy (2010). This should allow for sharing of information between the different sets of training runs, but is beyond the scope of the current study.

As the spray model is developed further, for example in the recently commenced BROWSE project (www.browseproject.eu), new emulators will be created as surrogates for revised versions. Examples of important updates could include longer term exposure estimation that takes account of multiple spray applications over time. Similarly, any new bystander data will be added and the Bayesian model of Section 2.7 updated accordingly. This should reduce uncertainty about the true probability of exposure exceeding a threshold dose. Further modelling might also be considered to associate ground deposits with bystander and resident exposure.

It was assumed that the fitted variability distributions for the model inputs are known precisely. In a more realistic analysis, we could acknowledge that these are also uncertain, and apply a Bayesian analysis to include uncertainty in the parameters of the variability distributions in a similar way to the uncertainty analysis described in Section 2.7. This process could combine limited data on the inputs with subjective expert opinions.

Acknowledgements

This work was funded by the UK Health and Safety Executive's Chemicals Regulation Directorate (Defra project No. PS2005). We are grateful to Paul Hamey for his advice and support. Helen Owen of the Food and Environment Research Agency coded the user interface, which was used to generate Fig. 6.

Appendix A. Qualitative assessment of the impact of features not included in the model

The inherently variable processes involved in spray drift, caused by random fluctuations in the emission of droplets and by natural air turbulence, mean that a model output can never fully capture all variability. Many uncertainties also remain, and while the uncertainties could be reduced by further work, this natural variability may still dominate.

References

Butler Ellis, M.C., Lane, A.G., O'Sullivan, C.M., Miller, P.C.H., Glass, C.R., 2010. Bystander exposure to pesticide spray drift: new data for model development and validation. Biosystems Engineering 107, 162–168.
Butler Ellis, M.C., Miller, P.C.H., 2010. The Silsoe Spray Drift Model: a model of spray drift for the assessment of non-target exposures to pesticides. Biosystems Engineering 107, 169–177.
Cullen, C.A., Frey, H.C., 1999. Probabilistic Techniques in Exposure Assessment: A Handbook for Dealing with Variability and Uncertainty in Models and Inputs. Plenum press, New York.
Doble, S.J., Matthews, G.A., Rutherford, I., Southcombe, E.S.E., 1985. A system for classifying hydraulic nozzles and other atomisers into categories of spray quality. Proceedings, Brighton Crop Protection Conference – Weeds, pp. 1125–1153.
EFSA, 2006. Guidance of the scientific committee on a request from EFSA related to uncertainties in dietary exposure assessment. The EFSA Journal 438, 1–54, http://www.efsa.europa.eu/en/science/sc_commitee/sc_opinions/uncertainty_exp.html.
Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 1995. Bayesian Data Analysis. Chapman and Hall, London.
Glass, C.R., Mathers, J.J., Harrington, P., Miller, P.C.H., Butler Ellis, C., Lane, A., O'Sullivan, C., Ferreira, M.C., 2010. Generation of field data for bystander exposure and spray drift with arable sprayers. Aspects of Applied Biology 99, International Advances in Pesticide Application, pp. 271–276.
Glass, C.R., Mathers, J.J., Harrington, P., Gilbert, A.J., Smith, S., 2002. Field Validation of LERAP and Assessment of Bystander Contamination. Central Science Laboratory Report Number 1728/02/02.
Gramacy, R.B., Lee, H.K.H., 2008. Bayesian treed Gaussian process models with an application to computer modeling. Journal of the American Statistical Association 103 (483), 1119–1130.
Gramacy, R.B., Taddy, M.A., 2010. Categorical inputs, sensitivity analysis, optimization and importance tempering with TGP Version 2, an R Package for Treed Gaussian Process Models. Journal of Statistical Software 33 (6).
Kennedy, M.C., 2005. Gaussian Emulation Machine for Sensitivity Analysis (GEM-SA). <http://www.tonyohagan.co.uk/academic/GEM/index.html>.
Kennedy, M.C., Anderson, C.W., Conti, S., O'Hagan, 2006. Case studies in Gaussian process modelling of computer codes. Reliability Engineering and System Safety 91, 1301–1309.
Lloyd, G.A., Bell, G.J., 1983. Hydraulic Nozzles: A Comparative Drift Study. MAFF report SC7004.
Miller, P.C.H., Hadfield, D.J., 1989. A simulation model of the spray drift from hydraulic nozzles. Journal of Agricultural Engineering Research 42, 161–170.
Oakley, J.E., O'Hagan, A., 2004. Probabilistic sensitivity analysis of complex models: a Bayesian approach. Journal of the Royal Statistical Society Series B 66, 751–769.
O'Hagan, A., 2006. Bayesian analysis of computer code outputs: a tutorial. Reliability Engineering and System Safety 91, 1290–1300.
O.Hagan, A., Forster, J.J., 2004. Kendall's Advanced Theory of Statistics, vol. 2B,, .. Bayesian Inference, second ed. Edward Arnold, London.
Saltelli, A., Chan, K., Scott, E.M., 2000. Sensitivity Analysis. Wiley, New York.
Sobol, I.M., 1977. Uniformly distributed sequences with an additional uniform property. USSR Computational Mathematics and Mathematical Physics 16 (1977), 236–242.
Teske, M.E., Thistle, H.W., Schou, W.C., Miller, P.C.H., Strager, J.M., Richardson, B., Butler Ellis, M.C., Barry, J.W., Twardus, D.B., Thompson, D.G., 2011. A review of computer models for pesticide deposition prediction. Transactions of the ASABE 54 (3), 789–801.
Vercruysse, F., Steurbaut, W., 2001. On-farm exposure to pesticides. Parasitica 57, 39–50.