Uncertainty in RODOS

RODOS REPORT

DECISION SUPPORT FOR NUCLEAR EMERGENCIES

RODOS(B)-RP(94)-05

Uncertainty in RODOS

RODOS(B)-RP(94)-05

G. Caminada

School of Informatics
University of Manchester
Manchester, M13 9 PL

S. French

Manchester Business School
University of Manchester
Manchester, M13 9 PL
Email: simon.french@man.ac.uk

K. Politis

School of Informatics
University of Manchester
M13 9 PL

and

J. Q. Smith

Department of Statistics
University of Warwick
Coventry CV4 7AL
Email: j.q.smith@ warwick.ac.uk

January, 2000

Management Summary

This document is a revision of RODOS(B)-RP(94)05, which was written to plan the development of data assimilation and uncertainty handling in RODOS.. It summarises the present position on these topics at the end of the Framework 5 projects and also introduces the more recent versions of the models used.

Contents

1 Introduction

The RODOS system (Real-time On-line DecisiOn Support system) has been developed to improve the emergency management of any future accidental release of radioactivity. The intention is that RODOS should be comprehensive in the analyses provided to support decision makers (DMs) and capable of finding broad application across Europe.

RODOS provides a wide range of decision support varying from the largely descriptive to providing a detailed evaluation of the benefits and disadvantages of various countermeasure strategies and ranking them according to the societal preferences as perceived by the DMs. It provides all levels of decision support (see Table 1) for all decisions ranging from consideration of the possibility of an accidental release through to long term countermeasures implemented months and years after an accident.

Although RODOS has a lot in common with traditional decision support systems (DSS), it differs in a number of distinctive ways:

- RODOS is designed to work in several different contexts, depending on where decisions need to be made. Each context will have its own particular requirements. Figure 1 outlines how RODOS will behave in an accident situation. In normal working mode RODOS will be kept up to date with geographic and demographic information as well as political constraints such as intervention levels. It will be useful as a training tool for emergency services managers and more senior DMs, and support for the development of emergency plans. If there is a threat of a release, RODOS will start working as a real-time decision aid, responding to user requests at the different stages of an accident.

Level 0: Acquisition and checking of radiological data and their presentation, directly or with minimal analysis, to DMs, along with geographic and demographic information available in a geographic information system.

Level 1: Analysis and prediction of the current time and future radiological situation (i.e. the distribution over space and time in the absence of countermeasures) based upon monitoring data, meteorological data and models.

Level 2: Simulation of potential countermeasures (e.g. sheltering, evacuation, issue of iodine tablets, food bans and relocation), in particular, determination of their feasibility and quantification of their benefits and disadvantages.

Level 3: Evaluation and ranking of alternative countermeasure strategies in the face of uncertainty by balancing their respective benefits and disadvantages (e.g. costs, averted dose, stress reduction, social and political acceptability) taking account of societal preferences as perceived by DMs.

Table 1 (after Kelly, 1994) : Levels of decision support for off-site emergency management. Decision support can be provided at various levels, here categorised into four levels. The functions provided at any level include those provided at lower levels. RODOS is unique in that it will provide support at all levels, including Level 3, for all potentially useful countermeasures at all times following an accident.

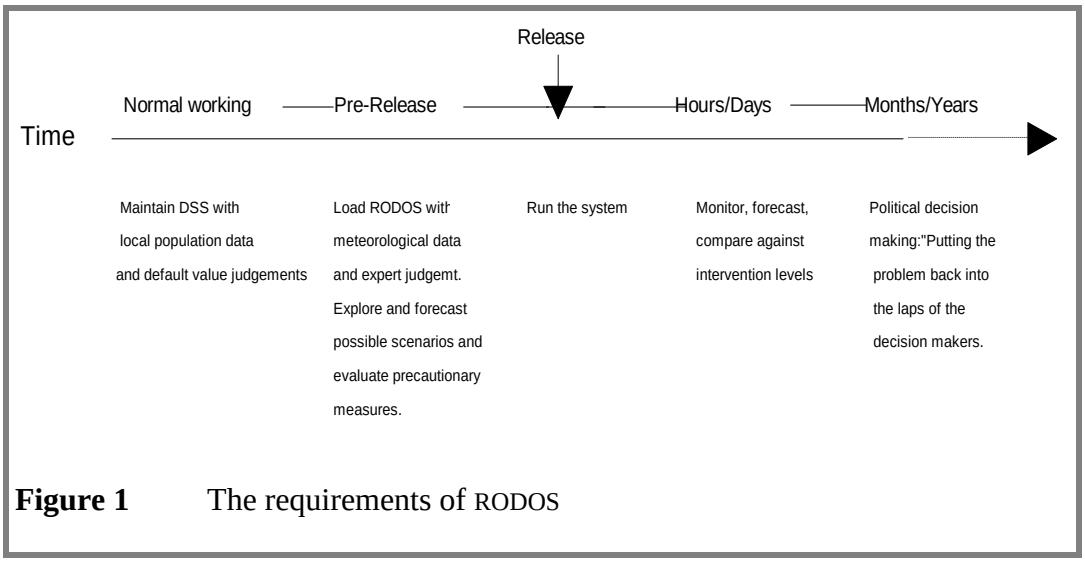

Figure 1 The requirements of RODOS

Initially, RODOS will support the decision making of the plant or site managers and local emergency services. As time progresses, decision making will become more the responsibility of regional governments and, in significant cases, national governments. RODOS will support all DMs and all decision making from initial evacuation and sheltering countermeasures, through agricultural countermeasures, to long term relocation. These support processes are described in more detail later.

- RODOS has a modular structure. The different partners within RODOS are each working on separate modules which may or may not be used depending on the circumstances surrounding a particular accident. For instance, one team may have a dispersal model which is very effective for flat, homogeneous terrain, whereas another may have a model which handles complex terrain. In any accident, the most appropriate set of modules would be used. Clearly some of these choices can be made *a priori* for each particular installation or site.

- It has a client/server architecture. It is built on a list of modules which are connected via a communication interface. A document outlining the conceptual architecture of the system is Ehrhardt (1997). In practice, each module requests a service from the system (typically this might be a request for data) and the system determines how this might be satisfied. If the data are already in the database, they can be supplied directly. Alternatively, RODOS may call another module which can calculate the required information.

- Since RODOS is to provide decision support at level 3 (see Table 1), it must explicitly address and solve the issue of uncertainty handling. The need to handle, model and represent uncertainty arises in a number of ways, as described below. It is important to incorporate all the uncertainty and to propagate it through the system in a seamless and effective manner. In French *et al* (1995b), French *et al* (1998) and French (1997a) it was argued that Bayesian statistical and decision

theory can provide such a single, consistent mechanism for handling uncertainty throughout RODOS. Since then, various attempts have been made to explore further the application of Bayesian ideas in uncertainty parameters and the communication of that uncertainty throughout the system.

- RODOS is networked firstly locally, so that advantage can be taken of the client/server architecture in supporting a variety of parallel interaction with the DMs; secondly it may be networked on much larger national and international scales to draw latest meteorological and radiological data from across Europe, and to share its prognoses with other countries.

The current document is a revised and updated version of French *et al* (1995b). Our intention is to provide a detailed exposition of how the Bayesian methodology can handle the uncertainty present at various stages of the emergency management in a unified way. We hope that our presentation demonstrates the flexibility of Bayesian approaches to cope with issues of uncertainty coherently. There is an increasing recognition of the need to ensure consistency between the management of the various phases of an accident and its aftermath, and hence a need for consistency between the modules of RODOS used at the distinct stages. The primary objectives of this report are:

- *to survey the recent progress that has been made on uncertainty handling and data assimilation within RODOS and review in some detail those developments that are directly relevant to these issues;*
- *to demonstrate that these developments have strengthened the flexibility of Bayesian statistical and decision theory to provide a single, consistent mechanism for handling uncertainty throughout RODOS; and*
- *to outline some of the major tasks and challenges in this context for the future.*

A companion report discusses the coherent support of decision making and evaluation (Papamichail and French, 1999).

To begin with, a methodology for handling uncertainty within RODOS must address the following requirements and issues:

- Data assimilation. Techniques for data assimilation must be treated as part of uncertainty handling, since the acquisition of data always changes the uncertainty in one or more quantities of interest. For instance, sudden increases in gamma dose rates at a monitoring station affect someone's uncertainty as to the spread of a plume. Also note that almost all data acquired by RODOS will be time varying.

- Source term. There may be some, probably incomplete, in-plant measurements of what is being released. As with monitoring data, these will contain measurement errors.
- Weather conditions. Wind fields used by atmospheric dispersal models are interpolated from meteorological station data or, perhaps, from an intermediate interpolation, which has been stored in a database. Thus there will be interpolation errors in addition to forecast errors. Rain conditions can also cause problems. If it is raining at one meteorological station and not at another, where does it stop raining?
- Radiation monitoring data. These can give an impression of what is happening and where radiation lies. They are typically collected non-uniformly in a both temporal and spatial sense. There can be very few data in the early phases of an accident. Monitoring data will also contain measurement errors.
- Expert judgement. At the time of a release, many parameters of the source, such as height, strength and composition, as well as of the atmospheric dispersion model, such as dispersion and deposition coefficients, are likely to be best estimates given by experts rather than 'hard' data. Such estimates are inherently uncertain and, further, subject to calibration and bias effects which give them different statistical characteristics to ordinary data.
- Model uncertainty. RODOS uses models to predict the passage and effect of the contamination. These can only ever be thought of as approximations to real world processes. It is therefore important to track, check, and make some allowance for the extent to which they diverge from reality.
- Implementing countermeasures. If countermeasures are not implemented appropriately, or people do not comply with them, this increases uncertainty in their effectiveness. Such uncertainty is called *volitional* uncertainty.
- Attitude to risk. The evaluation of equitable treatment of different population groups is a complex problem. Whereas individual attitude to risk is a well understood concept, societal risk attitude is far from clearly understood.
- Coherence between modules. The module chains and the client/server architecture used by RODOS necessitate rigid interfaces between modules. (The fact that these interfaces are realised via shared memory imposes tight size conditions as well.) More importantly though, the conceptual tools used to represent uncertainty in different modules must fit together smoothly. (Square pegs and round holes will not do!)
- Feasibility of the methods involved. When defining methods of handling uncertainty it is necessary to make sure that the proposed

solution can actually be implemented and is computationally feasible. This will almost inevitably result in the use of computational approximations. The methodology should be such that the quality of the approximations can be assessed.

- Communicating uncertainty. It is important to communicate uncertainty effectively to the DMs. After all, it would be a shame if RODOS handled all the uncertainty beautifully only to be misunderstood.

- DMs' value systems. RODOS needs to help DMs to resolve any lack of clarity in their understanding of the value systems. By their 'value systems' is meant their interpretation of the preferences and values of the society which they represent. With few exceptions such value systems are seldom fully defined. In any particular context, DMs need help to think through and refine their value systems. This need can give the DMs a feeling of uncertainty vis a vis their value systems. But note that we would argue that this is a different form of uncertainty to that related to, say, physical quantities: see French (1995). The issue of uncertainty in value judgements is addressed in Papamichail and French (1999) and not discussed further here.

One of the primary objectives of this report is to highlight the significance of a consistent machinery for handling uncertainty in the system. Although arguments favouring this view will emerge in various places in the sequel, before we discuss the capability of Bayesian methods to provide such a mechanism, some preliminary observations seem to be appropriate here, see also the recent exposition of these issues in French (1997a).

In the previous section we demonstrated the multi-dimensional character of uncertainty within RODOS at the various levels. It needs to be emphasised that this uncertainty is handled by different people with different skills and with varying backgrounds. For example, the conflicting attitudes between scientists and engineers in this context are well documented from experience, see e.g. Ranyard and Smith (1997). Scientists tend to overestimate the abilities of their models by limiting their focus on estimating the parameters of these models, while engineers on the other hand pay a lot more attention to the practical difficulties of the situation than the models themselves. Clearly a compromise has to be met, not least because eventually the main role of both scientists and engineers within RODOS is to support the DMs. No model is or ever will be perfect for dealing with uncertainty in practice, yet models (mathematical, physical or otherwise) are needed and used every day. A model which has been built to represent and/or forecast physical phenomena is not a panacea and can only be judged after experimentation. Any such judgement has to be based on the success of the model on analogous situations, or at preliminary tracer experiments. Unfortunately, due to the nature of the problem, empirical verification has proved to be a formidable task so far in RODOS -

very few complete data sets and tracer experiments exist. However this should not be interpreted as encouraging an arbitrary choice of modelling; we rather believe that it is an argument in favour of methodologies that include the possibility of modelling error in their uncertainty handling and prompts for further investigation of diagnostic tests at various levels. Further, we must not forget that the complicated structure of RODOS demands at various occasions that the output of a module has to be used as the input of another, which in turn will provide the basis for the analysis by someone else. Not only the results between different modules need to be compatible, but also their interpretation among various groups has to be consistent. After all, the success of RODOS in practice depends critically on the *actions* that will be made in a particular situation. Different data will be processed at different times possibly by different DMs. Since the aim in RODOS is to manipulate uncertainty and implement countermeasures for a lengthy period of time, the principles underpinning choice at one time should be compatible with those used at other times (French, 1997a).

Apart from these theoretical reasons that stress the need for consistency throughout the modules of RODOS, there is also the message from the already accumulated experience. Ahlbrecht *et al* (1997) have reported recently the evidence from the five elicitation workshops that have taken place so far. Among other worrying aspects, there appears to be a considerable lack of agreement for the set of criteria upon which the decision making has to be based. Further, 'the introduction of uncertainty into an exercise was very discomforting to the DMs' (Ahlbrecht et al, 1997). This appears to be the result of the somewhat vague perception of uncertainty among RODOS participants so far. In view of the above, we believe that handling and communicating uncertainty in a consistent manner seems to be one of the major challenges for future developments in RODOS.

It has been argued for some time now that the Bayesian methodology provides a unified approach to tackle uncertainty issues within RODOS. Since French and Smith (1997, chapter 1) argued for the advantages of using Bayesian analysis to encode, model and communicate uncertainty in order to aid the task of the DMs, various expositions have been made as to how a Bayesian methodology can handle all the aspects of uncertainty discussed in the previous section; see for example section 4 in French, Papamichail, Ranyard and Smith (1998). We do not wish to repeat these here. We simply remark the following key features of Bayesian ideas that assist an integrated way of handling uncertainty within RODOS:

- Bayesian analysis is developed in the context of decision making. The modelling of uncertainty and of value judgements are undertaken in compatible manners.
- Bayesian ideas recognise and accommodate two distinct cycles: (i) the continual updating of uncertainty as data arrives and (ii) the

sequential nature of decisions which fits with comprehensive support of emergency management.

- Bayesian methods combine probability as a measure of uncertainty with utility to reflect one's needs and/or preferences in a mathematically tractable and practically coherent way. Thus the demand that RODOS uses a consistent method to encode uncertainty and pass it between modules is fulfilled.

Sections 3 – 9 illustrate how these ideas can be put forward within RODOS from the early phases of an emergency until long after an accident. Notwithstanding this, it is important to be aware of the conceptual and practical difficulties in adopting a full Bayesian approach.

In principle, quantifying and communicating uncertainty within the Bayesian framework involves the subtle task of manipulating in a mathematical way subjective, and inevitably sometimes conflicting, opinions and preferences. This subjective character of uncertainty however is a mere reflection of the fact that Bayesian inference and decision theory is focused on the needs of individuals. Any decision making is inherently and unavoidably subjective, and thus it is at least implicitly based on the DM's own views about the unknown state of the world. Further, he has the flexibility to update these views constantly on the basis of new evidence, either this comes in the form of data arriving or opinions from experts (from a Bayesian perspective, there is in fact a view to treat these two pieces of information as identical, see French, 1985). Moreover, experience suggests that preferences of a single individual may not be consistent through time. It is the role of those who support the DM, hence, to ensure compatibility in his choices. The issue of handling and avoiding human preferential inconsistencies is explored in some detail in French (1986).

At an initial stage, however, before data are collected, a DM has to express his beliefs about a number of unknown parameters by assigning a prior distribution to them. The specification of a prior is an integral part of any Bayesian procedure leading to inferences and/or decision making and needs to be done with care, see for example Chapter 6 in O'Hagan (1994) for some theoretical considerations. There are a few key issues that need to be addressed here. The first is that, quite naturally, this prior assignment involves elicitation of other persons' beliefs, known as the experts. This is known as the expert problem, for a discussion of the problem and its variants see Cooke (1991), French (1985), Lindley (1985). Given the time requirements of RODOS, it is unlikely that in the emergency phase of an accident the possibility for a full consultation of the experts and elicitation of their opinions will be feasible; in such a case and in the event of absence or very vague prior information, one can code his prior beliefs by assigning a noninformative prior distribution, see Box and Tiao (1973). The role of the experts becomes a lot more significant in the later stages, when also

measurements of contamination become available. The information derived from the data needs to be combined then with the experts' beliefs, in order to form a full probability distribution for the unknown quantities. This may serve as the new prior distribution which, when more data are available, can be used for the updating. Note however that this encoding of the data and the experts' opinions in a probability distribution is a subtle task which requires careful consideration. Even with only two experts present for example, it might be not at all clear as to how one should 'average' their probabilities, see e.g. French *et al* (1991b). Moreover, it is unlikely that all the experts possess the same skills or technical knowledge, so that not all opinions need to be given equal weights. In practice, it seems advisable in many cases that a DM should himself make an assessment of the experts considering how well calibrated and informative their opinions are; see Cooke (1991) for more details. Some illustrative examples in this context, with reference to dispersion and deposition parameters of contamination after an accident, are given in Chapter 6 of Cooke *et al* (1994).

Another issue which is related to the assignment of prior distributions concerns a highly desirable, if not essential, feature of any modelling (Bayesian or not), its *robustness*. Loosely speaking, any sufficiently good model should possess the property that very small changes in its input result in minor changes in the output of the model. In a Bayesian context, this is most often interpreted as the sensitivity of the model to misspecification of the prior. Again, a number of diagnostics to examine the adequacy of the model and suggest possible modifications are available, see Chapter 7 in O'Hagan (1994).

Finally, as described above, the Bayesian methods model uncertainty and preferences for a single DM. However, they extend to support groups of DMs, although not as directly as one might think. Arrow has proved that no methodology can model ideal group decision making adequately - or rather he has shown that ideal group decision making is ill-defined; see French (1986) for details. To support groups for DMs one must rely on sensitivity analysis. Do variations in input judgements across any differences within the group lead to a change in the top ranking action? In practice, there is usually enough agreement within groups for the choice of action to be unchanged. (See the flat maximum principle in von Winterfeldt and Edwards, 1986.) Moreover, the process of sensitivity analysis helps the group towards consensus, see French (1986, 1989, 1995). We recognise that in the early stages of accident there will not be time for detailed sensitivity analysis. However, because of the urgency of decision making in the early phases, responsibility usually rests with a single emergency manager supported by a team of advisors rather than shared between a group of DMs equally. The need for true group decision support arises in later stages.

The remainder of the report is organised as follows. The need to pass measures of uncertainty seamlessly between modules together with the time varying nature of the data implies that we must consider data and control flows with the (conceptual) software architecture of RODOS. Sections 3 – 9 describe the Bayesian machinery required to propagate uncertainty through example model chains of RODOS. Appendix A contains a tutorial and survey of Bayesian methods and an understanding of these will be assumed in the main report.

2 Conceptual Structure of RODOS

Before examining in detail how uncertainty can be incorporated in each module, we look at the overall structure of RODOS (i.e. a top-down view). Our intention is to demonstrate the need for data and predictions to be stored in a (temporal) database and how assessments of uncertainty need to be passed between modules.

The conceptual RODOS architecture (not the physical software which includes, for instance, databases and graphics servers), as described in Ehrhardt *et al* (1993), has been split into three distinct subsystems:

- Analysing Subsystem (ASY) modules which process incoming data and forecast the location and quantity of contamination including temporal variation.
- Countermeasure Subsystem (CSY) modules which suggest possible countermeasures, check them for feasibility, and calculate their expected benefit in terms of a number of attributes (criteria).
- Evaluating Subsystem (ESY) modules which rank countermeasure strategies according to their potential benefit and to preference weights provided by the DMs.

Note that this conceptual structure corresponds well with the Bayesian paradigm in that, roughly speaking: ASY modules build and update appropriate probability distributions; CSY modules identify possible actions and calculate their potential consequences; and ESY modules calculate their expected utilities. However, the descriptions here of generic roles for the ASY, CSY and ESY subsystems are only indicative. There is no intention to respecify any subsystems such as LCMT and ECOAMOR.

This conceptual structure can be interpreted in too simplistic a manner. In order to cope with the diverse requirements placed on RODOS over time it needs to be placed in a temporal context (see Figure 2). As time passes, RODOS will arrive at different decision points where it must select three modules to form an ASY, CSY and ESY chain appropriate to the context of a decision. One of the basic differences between RODOS and more traditional DSS, which tend to be designed to support only one type of decision, is that RODOS provides decision support for all the different phases of a nuclear accident. During each phase, only certain types of information will be present and certain contingent decisions considered. For example, whilst the plume is still airborne above a particular area, it would not normally be sensible to consider decontaminating the roads: the critical issues would be to shelter or evacuate the population. Again actions to influence the food chain, like imposing bans on certain types of foodstuffs, would normally be taken after the plume has passed by, mainly because most measurements relevant to such bans would not have been taken yet; and the decision to ban most foodstuffs can be delayed for a short time without great risk, whereas other decisions cannot. (None the less, it should be

noted that in longer releases many of the phases of decision making may have to take place in parallel. For instance, during a long (or second) release in which the wind direction changes, evacuation decisions and food ban decisions may need to be taken simultaneously.)

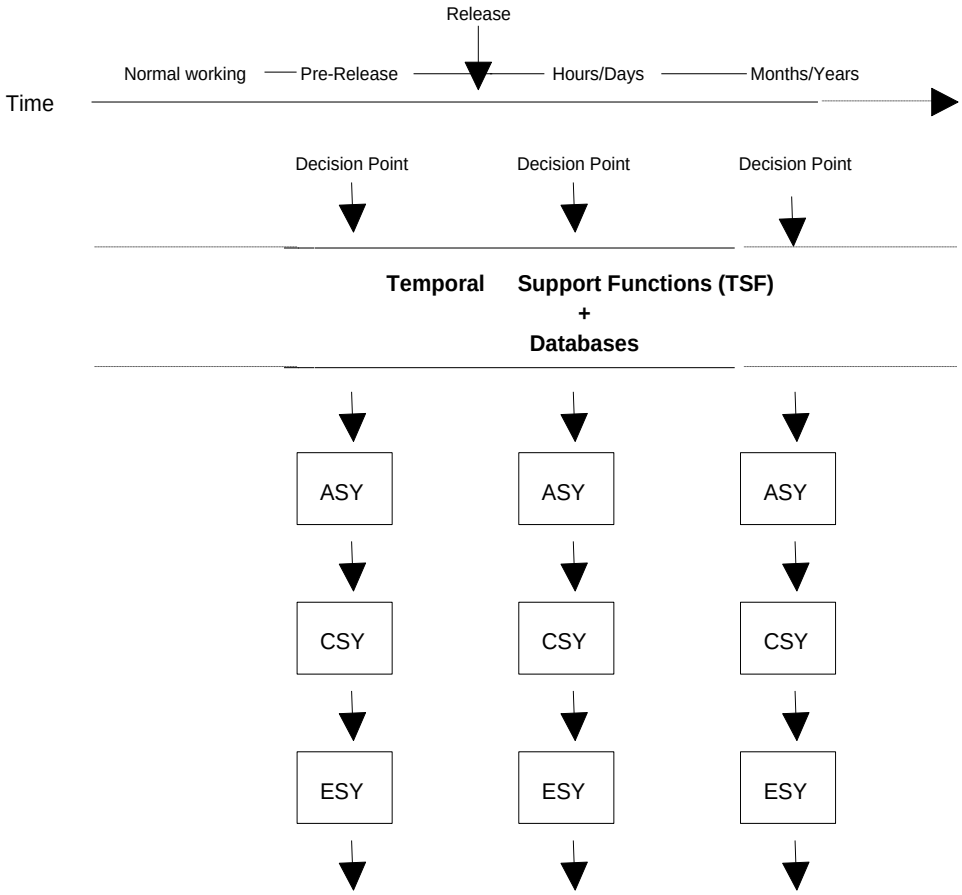

Figure 2: The temporal support functions

Thus it is helpful to recognise in the conceptualisation of RODOS the temporal control necessary to manipulate modules in response to the user's requests or some pre-programmed sequence of desired analyses. These functions are performed by the Supervising Subsystem (SSY) and the Operating Subsystem (OSY). For convenience, in the sequel we refer to these functions as the Temporal Support Functions (TSF).

The TSF would provide the flexibility within RODOS to change operating behaviour according to the current phase of the accident. Whenever RODOS receives data or input from a user (e.g. when its environment changes) it should be able to re-examine and re-evaluate whether more detailed analyses are possible or necessary. Perhaps some new piece of data allows

the use of a more appropriate model. Alternatively, maybe a user has requested some information and it must decide the best way to supply it. The TSF also need to be able to cope with data arriving out of a sequence relative to their times of observation. As an aside to our main purpose here, note that the client-server architecture of RODOS is essential in providing the mechanism for the TSF to achieve such control, initiation and scheduling of analyses.

The main tasks of the TSF include:

- Processing and storage of all incoming data in a way that is suitable, either early in the emergency or later. It is important that RODOS knows what data it has and how reliable they are. Moreover, the data have two times associated with them. The time that they were measured or observed and the time that they became available to the RODOS system. Thus the temporal database will need a fairly sophisticated functionality, see Valet and Roberts (1994). The database may also need to store intermediate stages in various analyses so that they can be rolled-back efficiently in the event of further data, which was collected earlier, becoming available.
- Managing the management of the emergency response. Throughout the incident the emergency response managers will need to gain an impression of the sort of decisions which will be needed in the near future. Thus at various times the emergency response managers will need to ask:
  - When do we next need to make a decision?
  - What class of countermeasures do we need consider at that time?
  - How detailed need analyses be to support the decision making on those alternatives?

  Initially such questions will be answered by 'quick and dirty' analyses, calculated by running a module chain without much detail. Later as further information arrives and the situation evolves, more detailed, well-calibrated forecasts can be used. By such means RODOS will help the emergency response managers identify whether a full analysis and decision of countermeasures should be made at a particular instant or whether it should or can be delayed. Some of this process management can be automatic (as in the descriptions of OSY in previous papers on the conceptual structure of RODOS, see Ehrhardt *et al* (1993); other parts will require interaction with the emergency response manager.
- Selecting which models are most suited to the particular situation and the current data available. In some cases the TSF could suggest what data would be most useful to acquire (i.e. "if we knew *this* we could explore *those* issues"). This could be used when deciding how to allocate resources for data acquisition.

- The TSF will also need to determine which previously performed analyses may be called into question by the arrival of new data. At the simplest level, the TSF could mark previous analyses as being possibly superseded by new data. Then, if the results were to be used subsequently, the analyses could be re-initiated. However, it would be better for the TSF to compare the new data with the earlier predictions. If they fitted reasonably, the TSF might re-initiate the analysis in the background in order to ensure that the predictions were based on all the available data: however, if the analyses were not to be needed in the future, this could be a waste of computing resource. The real issue would arise when new data arrived which did not fit the previous predictions. Then not only there is a case for re-initiating the analysis, but also earlier decisions may need to be revised-perhaps quickly.
- RODOS must keep a complete audit trail of the incident for later investigations. It must be possible to reconstruct any analysis using only the data available at the time it was performed. This again means that the TSF must control a very sophisticated temporal database.

In essence the TSF monitors and controls all data input to RODOS and the manipulation rules of these data. It responds to requests by the user depending on the level and type of detail required. Both the type of request and the possibility of giving an informative answer to the user will strongly depend upon the current history of the accident and the data available.

Returning to Figure 2, it is worth remarking that while logic dictates that an ASY module runs before a CSY module and that in turn before an ESY module, in practice the client/server architecture works slightly differently. According to this architecture, an ESY module would start before an ASY module. For example, through a log kept by the TSF, it might learn that the appropriate CSY had not been run and, in turn, that a preceding ASY need be run. Thus the support process would be driven by the questions which need to be answered.

The temporal functionality of the system outlined above emphasises the need for coherence between various modules in RODOS. Since the TSF together with the needs of the emergency managers may dictate that almost any ASY, CSY, ESY chain is called immediately after an analysis based on another such chain, it is essential to ensure that uncertainty is modelled consistently between different modules. Moreover, the measures of uncertainty produced by each module need be lodged in the database for subsequent recall and use by other modules.

2.1 Early Phase

When there is a warning of a potential accident, it will be necessary to consider possible scenarios and evaluate precautionary countermeasures. Figure 3 shows the model structure of RODOS in the early phases of an a accident.

At this stage and just after a release (up to 24/48 hours) the ASY would consist of an Atmospheric dispersion model such as RIMPUFF which can be used in several different ways. Firstly it can be run with very little detail (possibly deterministically) to give a quick indication of how the accident might develop and, thus, which areas are of interest. The amount of detail could be increased and, if a release actually occurs, monitoring data will be incorporated dynamically as and when a more accurate picture is required, typically when a decision is about to be made. It will also be used to forecast, based on current meteorological conditions, what the situation is likely to be say 6-12 hours later. Its output is in the form of a grid of concentrations with an associated uncertainty (probability) distribution. The Graphics Subsystem (GSY) will be able to display the concentrations via contours and, upon request, give predicted values at any given point at any given time.

This grid becomes one of the inputs to the CSY which at this stage would be an evacuation and sheltering simulator. In turn, the CSY calculates all possible countermeasures for each area together with an estimate of the expected dose averted if the countermeasure were implemented. It would be extremely useful to calculate an associated uncertainty with this number partially based on the uncertainty input from the ASY.

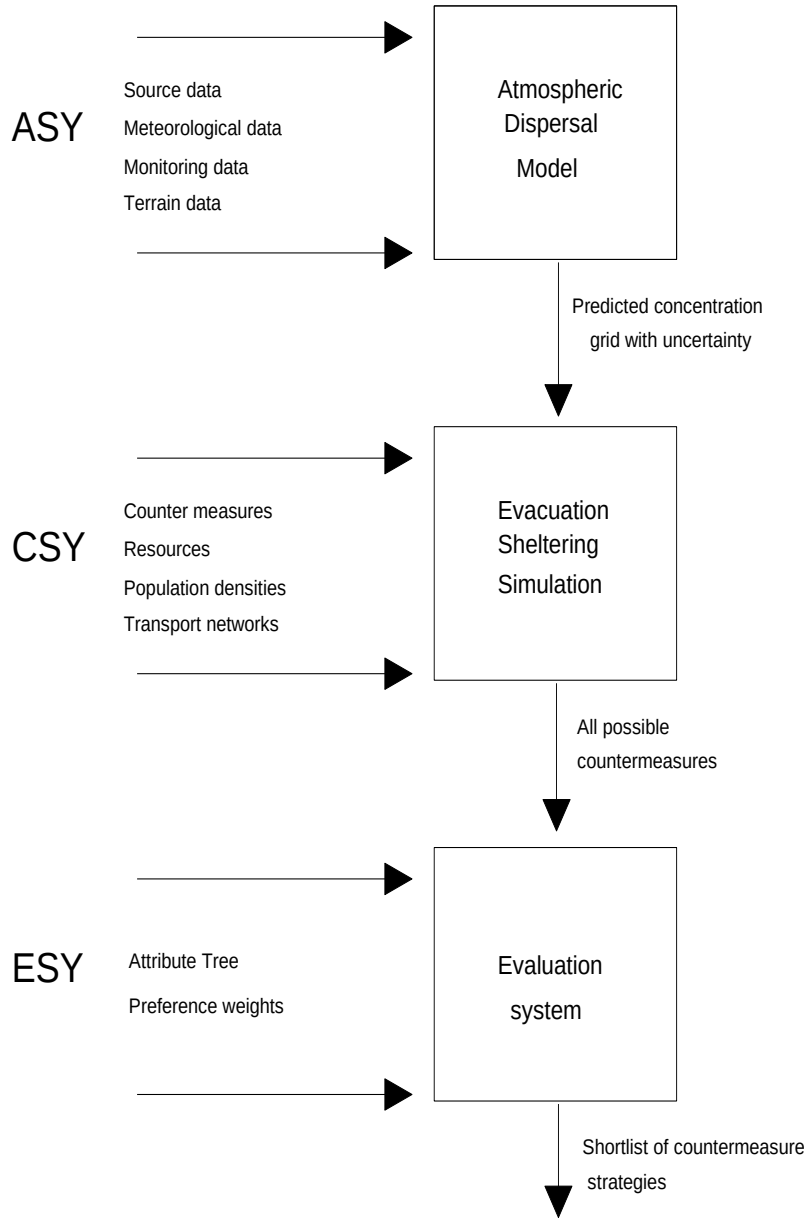

Figure 3: Early phases

The averted dose can be used as a cost function for evaluating the usefulness of a countermeasure in the ESY. The CSY then outputs a list of all possible countermeasures for input to the ESY which consists of:

- A coarse expert system filter which rejects any logically infeasible countermeasures or combinations of countermeasures (e.g. you cannot distribute iodine in an evacuated area).
- A multi-attribute value or utility (MAV/MAUT) ranking of the strategies which takes the remaining list of countermeasures and ranks them for their relative effectiveness according to previously elicited utility attributes and preference weights from the DMs. It may be necessary to revise and re-evaluate these preference weights in any given situation before a particular decision is taken.
- A fine expert system which takes the top 10-15 strategies and produces a management summary report detailing the costs and benefits of each.

The ESY is slightly different to the ASY and CSY in that, whereas for the latter pair interchangeable modules would be used depending on the time and situation of the accident, the ESY is likely to be the same software module with different attribute trees and with the preference weights changing over time.

2.2 Medium/Long Term Phases

In the medium and longer term phases (see Figure 4) there are likely to be sufficient monitoring data to stop the use of an atmospheric dispersion model. However, these data still need to be interpolated in some way. A model can be built over the data, again with an associated uncertainty, to produce a map of concentrations as output. In French *et al* (1995b) a three stage hierarchical model was considered for this implementation. Here in Section 5 we describe a similar modelling developed recently by Faria *et al* (1997). The data model will have a forecasting capability, but this time to produce information about 6-12 months time or several years into the future. The output of this model might be on a much larger geographical scale than for the early phase. However it may be more practical to consider small regions at a time and to combine these to explore the overall situation. Now the CSY would be a relocation/food bans/decontamination or hydrological pathways simulator such as LCMT or EQUAMOR. There is also a need to consider the ending of some early countermeasures: e.g. when to allow return from short term evacuation. Again a list of countermeasures is output for use in the ESY.

The ESY would work as before with the relevant attribute trees and preference weights as inputs. It may be the case that preferences will change with time. So before each major decision point the attribute trees and preference weights could be revised or re-evaluated.

So far we have considered near range countermeasures. RODOS is planned, however, to support all decision making, short and long term, near and far range. Thus RODOS will have long range dispersion models to call as ASY submodules. There is a need to consider the uncertainty in the predictions of these. Long range modelling within RODOS is being provided by an

adjoint method which is seeded with an initial run of RIMPUFF. Since adjoint methods are nearly identical to Kalman filtering used in our modifications of RIMPUFF, there is potential here to introduce automatic uncertainty handling and data assimilation.

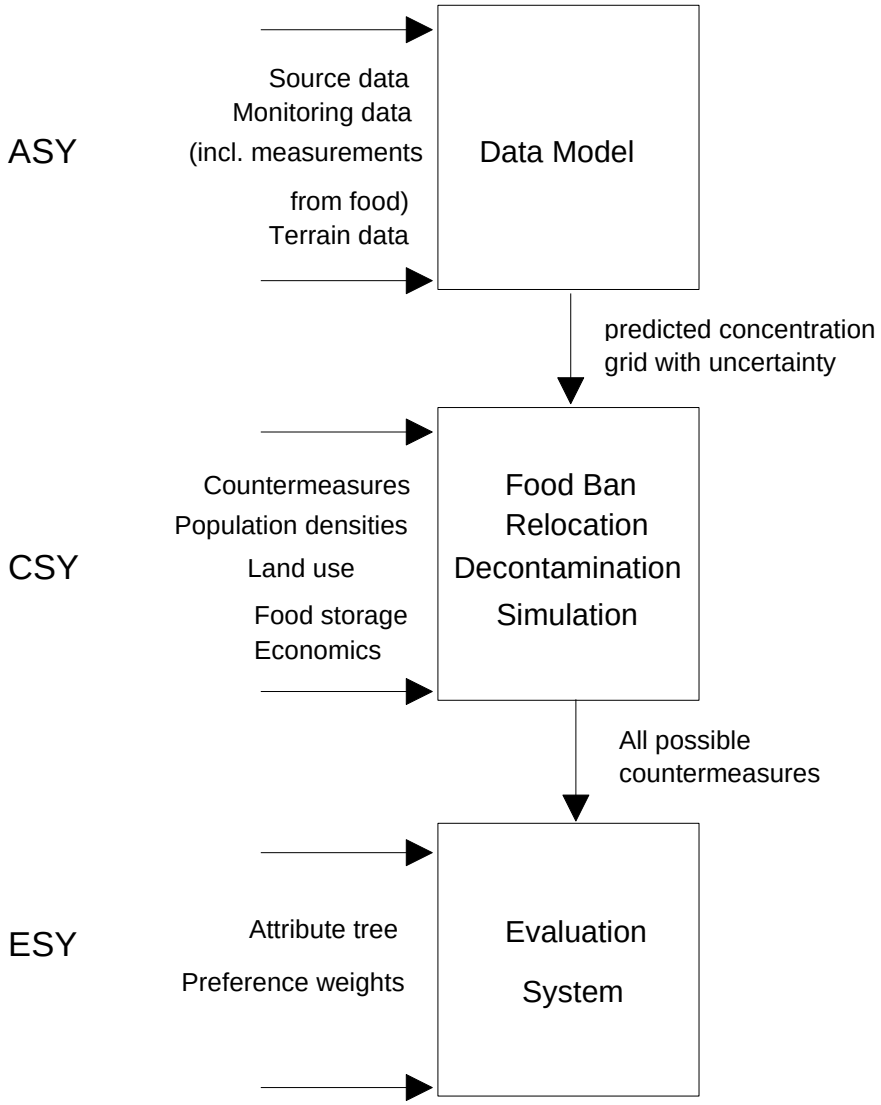

Figure 4: Medium/late phases

3 Pre-release phase calculations

In the early stages of an accident very little information will be available about the magnitude and characteristics of the release of radioactivity to the environment. This information, called the source term, will however be one of the most important inputs to a decision support system to assist in the implementation of the appropriate emergency response. The emergency response team would very much like to know an estimate of the likely release of radioactivity before it actually occurs. For this matter a Source Term Module (STM) has been developed, which can provide an indication of the likely release, based on an optimum use of data on the status of the plant, and the plant's passive and active behaviour (as a result of operator actions). The output of this module will subsequently be the input for atmospheric dispersion models, which we will examine in the next section. The STM is implemented on a PC, since that is the type of computer normally available in a control room and PCs can transmit output files to RODOS.

The requirements for the STM in RODOS can be summarised as follows:

- The module should provide an early indication of the source term based on the status of the plant.
- The module should not require expert knowledge of accident analysis or fission product behaviour. The user should be guided through the different (possible) stages.
- The software should be capable of being operated under stress and should therefore require the user to make a minimum number of decisions.
- The input data should be limited to what is easily available from plant instrumentation or observations.
- The module should not require complex computing facilities, neither require the use of complex algorithms, mainly because the response time should be rapid.
- The output of the module should be available on a time scale and level of detail required by the user. Normally the time scale will be determined by the time until release or the time until real on-site or off-site release measurements become available. It is envisaged that, following the availability of off-site measurements (and a smooth transition procedure) the source term module will not be required any more.
- The module should be interactive, so that the user can update input information as it becomes available and therefore give a better source term estimate.

- The output of the module should consist of both a range of possible source terms, with an associated probability, based on plant information and the user's response to initial key questions (including possible updates of these answers) and so-called RODOS-files carrying details of the possible STCs that we found above.

The methodology used by NNC and their RODOS partners to determine the appropriate source terms and their probabilities is based on a belief network. For those who are not acquainted with belief networks a basic introduction is given in appendix A. Belief nets (instead of fault trees) were chosen mainly for the following reasons:

- Probabilistic Safety Analyses (PSAs) have been conducted on most nuclear plants, so detailed information is available on probability ditributions.
- The structure of belief nets enables information and judgement to be elicited in a convenient form.
- Software is readily available.

The evaluation of a belief network, based on using the implicit conditional independence structure is quite straightforward (if the required conditional probabilities are given). The construction of a belief network however is a lot harder, especially in complicated cases like this. In principal we can do this in a forward direction (starting with the event node and evaluating the consequences) or in a backward direction (starting with the final result and identifying initiating events). We will not go further into this here, see for more information Smedley *et al* (1999). A RODOS-STM belief network is given in Figure 5

In this belief network there are two types of node; those representing state of a part of a plant (e.g. availability of condensers or ventilation systems) and those representing observations (data from plant instrumentation such as core exit temperature and reactor vessel pressure). The end-point of the network is the source term, which is represented here by the sum of three nodes. It should be noted however that there is normally just one dominant release route. These three source terms are:

- PC Source Term. The source term from the Reactor Building as a result of releases from the primary circuit.
- SC Source Term. The source term from the secondary circuit as a result of either steam generator tube leakage or enhanced tube leakage (SGTR).
- Auxiliary Building Source Term. The source term from the Auxiliary Building as a result of leaks from auxiliary systems.

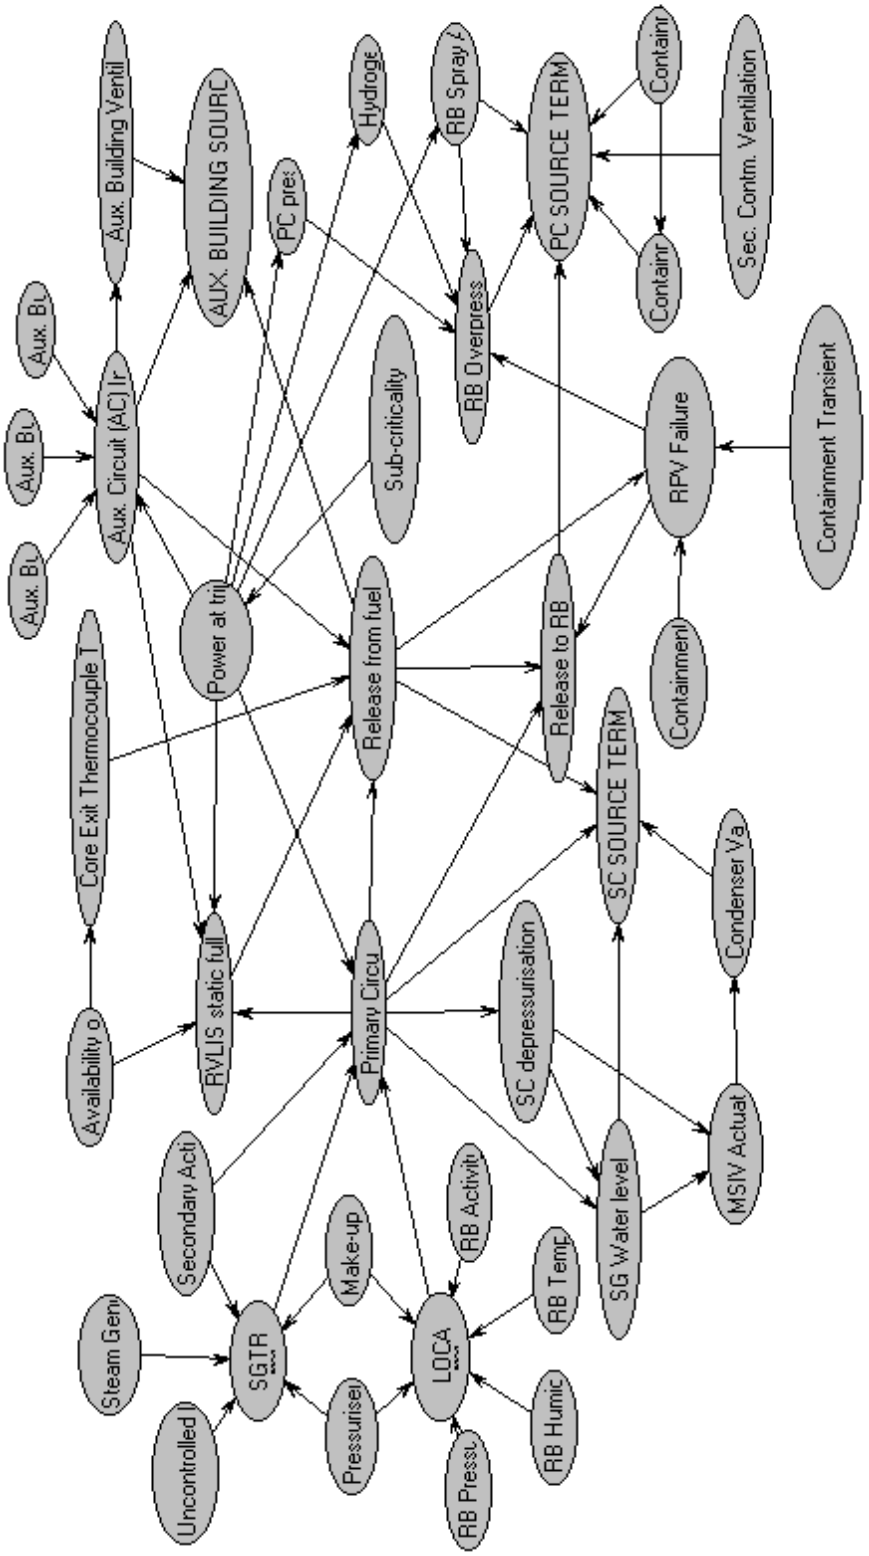

Figure 5: A belief net of the form used in RODOS_STM

By answering questions about the current state of the plant (and thus inputting evidence into the belief network!) the user is directed to the most likely fault source terms, termed Source Term Categories (STCs) and their associated probabilities. The STCs will subsequently be allocated to reduced series of Release Categories (RCs), which will give an indication of the magnitude and severity of the release. In the following section we will say a little bit more about both.

3.1 Source Term Categories and Release Categories

At a nuclear power plant there is a very large number of potential design basis and beyond design basis faults. The most straightforward way to use these data, when carrying out a safety analysis, is to categorise these faults according to certain phenomena and events that influence the magnitude of the released source term. The initiating fault events considered are:

1. Loss of Coolant Accident (LOCA) into the Reactor Building.
2. LOCA via an auxiliary circuit outside the Reactor Building.
3. Anticipated Transient Without Trip (ATWT)
4. Steam Generator Tube Rupture (SGTR)
5. Disruption of the secondary circuit causing an uncontrolled secondary circuit depressurisation.

These events will normally result in so-called ‘design basis faults’ and the primary coolant source term can be of three types:

1. the activity associated with a normal shutdown spike, denoted by ‘Spike’
2. the activity associated with the fuel/clad gap, denoted by ‘Gap’
3. the activity associated with a reactivity transient in which the fuel experiences Departure from Nucleate Boiling, denoted by ‘DNB’

If safeguard systems are unavailable however, such a design basis fault may deteriorate to a state in which fuel damage occurs; we denote this by ‘Melt’.

The release of radio activity to the environment will also depend on engineered safety features that are available during the fault sequence. Think of:

1. condition of primary containment safeguards (for releases into the Reactor Building (RB)). Can we successfully isolate the RB? Does the RB spray system work?
2. filtration units (if they are in the release path)
3. pool scrubbing (for release into water)

The combination of such an initiating event and the availability (or not) of engineered safety systems is termed a Source Term Category (STC). As

mentioned in the previous section there are basically three types of STCs (PC source term, SC source term and Auxiliary Building source term).

In order to indicate the general severity of an accident, the STCs are further grouped into Release Categories (RCs). In total there are seven of these Release Categories:

- Normal
- Design Basis
- Beyond Design Basis
- Large Release
- Extreme - Containment Intact
- Extreme - Late Failure
- Extreme - Early Failure/ Bypass

Unlike the STCs, which are defined on the basis of the plant status, RCs are defined on the basis of the maximum source terms to the environment. The series of RCs are defined to cover the radiological releases associated with a range of potential reactor faults based on the response or availability of safety systems. The RCs cover off-site consequences ranging from releases that are not expected to result in a significant number stochastic health effects in the exposed population to releases that may result in a significant number of early deaths.

These Release Categories will be of direct interest to emergency DMs, because it gives them a broad idea of the potential source term. However, they will also be interested in issues such as which places the radioactive plume will affect and how badly. In order to find that out, the release quantities of this module are used as input for atmospheric dispersion models. The next section will deal with those.

4 The Plume is spreading locally

At this point the major decisions are about evacuations, sheltering and distribution of stable iodine. For each geographic location the emergency managers need to decide which action or combination of actions is best, based on expected dose, available resources and other factors. The TSF at the same time will run a short range atmospheric dispersion deposition ASY model such as RIMPUFF. The elaboration of RIMPUFF to incorporate observational updating, started at Warwick-Leeds and now continuing by the Warwick-Manchester participants, has been one of the earliest applications of the Bayesian methodology within RODOS. Before we present the main features of this elaboration, we give some information underlying atmospheric dispersion models in this context.

A fierce requirement for any atmospheric dispersion model used in the ASY of RODOS is that the model needs to be operationally fast. Since these models are typically used for predictions within 12-24 hours after an accident, it is essential that they provide efficient predictions quickly. In the emergency stage after the release it is imperative that the model should provide an idea of the contamination profile within the next few hours, so that action can be taken. In particular, the system needs some indication of which areas are more likely to be heavily exposed to the release, so that special interest might be focused upon them. Despite their individual differences, all atmospheric models need a similar set of parameters to operate. These are:

- source term data
- dispersal patterns
- meteorological data and
- geographic information.

Thus, for example, all dispersal models assume the ability to provide certain meteorological information, such as the wind speed and direction, forecast of precipitation, information about the stability of the weather conditions and so on. More generally, for the typical assumptions underlying atmospheric dispersion models, see Ranyard and Smith (1997).

The list above ignores the fact that, in an emergency, other information, most notably monitoring data will also be available, giving still more information about the release. These data are recorded by permanent and mobile measurement stations that together provide:

- gamma dose rates;
- instantaneous air readings;
- air readings, integrated over time;
- ground readings integrated over time.

It is sensible not to rely solely on these data, just as it is sensible not to rely solely on the predictions of the model. We believe that by combining these we will get the best picture of what is actually happening and thus can best predict future contamination based on the current status.

Atmospheric dispersion models fall into two broad categories: Lagrangian and Eulerian. Lagrangian models simulate the trajectories of many emitted particles according to a stochastic differential equation, whose parameters have to be estimated. Experience shows that even simple versions of these models can not give accurately enough predictions quickly enough. Eulerian models use a diffusion/advection equation to predict dispersion. By imposing rather severe restrictions on these models (stationarity, constant windvector, homogeneous terrain, it is possible to obtain a steady state model that we can use in real time. However, this model is too simplistic to provide useful results in the early stages of an accident. A possible solution to this is to use a puff model, such as RIMPUFF.

Among the various mesoscale dispersion models that offer predictions (see, e.g. the comparative study of Päsler-Sauer, 1986), RIMPUFF is one of the very few models with the potential to provide accurate real-time predictions within the time limits of the ASY in RODOS. Models based on a a Gaussian plume (the simplistic Eulerian model) for example (Päsler-Sauer, 1986) seem to be less efficient for the short term. Further, some early trace experiments indicated that puff models, i.e. models that approximate the continuous emissions from the source by a series of small puffs (see Figure 6), are reasonably adequate in practice for quick predictions. Moreover, one of the distinguishing characteristics of RIMPUFF which fits well with the Bayesian methodology, is the process of pentification. Namely, when a puff reaches a certain diameter, it is approximated by a mixture of five smaller puffs. The diameter of each of these puffs is half that of their parent, while a crucial assumption here is that mass is preserved, More precisely, by imposing an extra condition about the second moments of the puffs on the horizontal plane, Thykier-Nielsen *et al* (1989) derived that the shape of the pentificated puff with respect to its original position and the mass allocation is as shown in Figure 7.

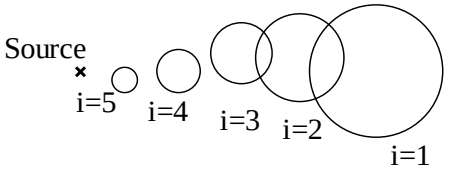

Figure 6: **A simple puff model**

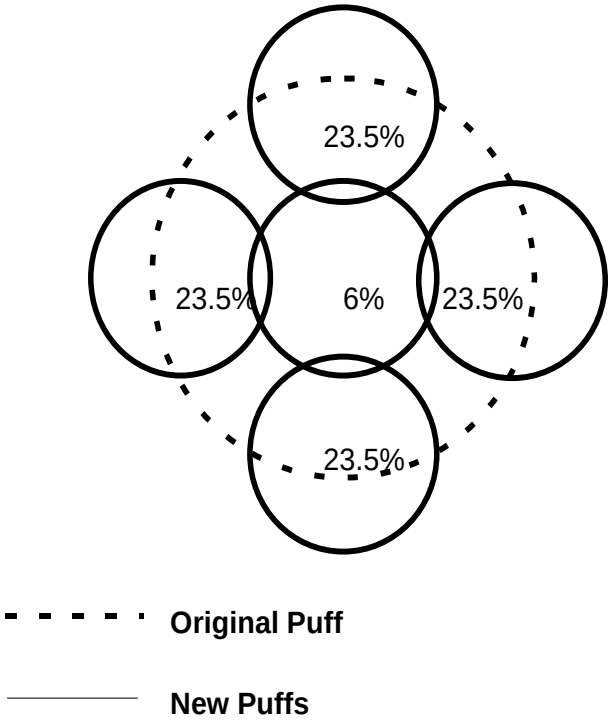

Figure 7: Puff Pentification

We do not wish to discuss further the description of the model and the accompanying code, since these can be found in many places; see, for example Mikkelsen *et al* (1984, 1987), Thykier-Nielsen *et al* (1989, 1997), and were already exhibited in detail in French *et al* (1995b).

Because of practical problems with the modification of the computer code in which RIMPUFF was written, recently a new start has been made with a different atmospheric dispersion model: ATSTEP (Pasler-Sauer, 1998). This is a segmented plume model, which makes it a little less accurate than puff models, such as RIMPUFF, but it has the advantage that it is very fast. The same updating principles that were applied to the RIMPUFF-model are currently being built into this ATSTEP-model.

Before we move on, we will say a bit more about (Lagrangian) particle models. Recently more and more voices have raised the question whether it would not be better to use these models, especially now that computational speed is rapidly increasing, which makes it possible for them to provide almost real-time predictions. We would however like to place some critical comments on the above suggestion.

First of all, when it is said that a particle model can almost run real-time, this essentially means that we can do just one run. However, this takes the possibility away to obtain answers to so-called ‘what if’-questions. With

'what if'-questions we try to find out what the results would be IF we had a stronger wind, or a different height, or merely noble gases (instead of a mixture of gases). It is not unthinkable that a decision maker would like to run the model ten times or even more! This seems to be impossible with the particle models, at least at the moment.

It seems reasonable to assume that the problem associated with multiple-run possibilities could disappear in due time, because of further increases in speed of computers. However then we're still stuck with another problem associated with the particle model which is its stochastic nature. If we run the model again with new parameter-settings, can we say then what amount of the difference between the two outcomes should be attributed to the new settings and what amount to the stochastic nature? In principle we can avoid this problem by using an enormous number of particles, but then we are confronted yet again with the problem that the model becomes too slow.

Apart from this, the stochastic nature of the model raises an issue of even more concern. In order to understand this better we need to take a look at the different time scales we are concerned with. We can distinguish:

- actual time, which is simply the time the DMs have to decide on the countermeasures that they are applying.
- data time, which is the time related to the observational data that we receive.
- model time, which is the time the model calculations have reached at the actual time.

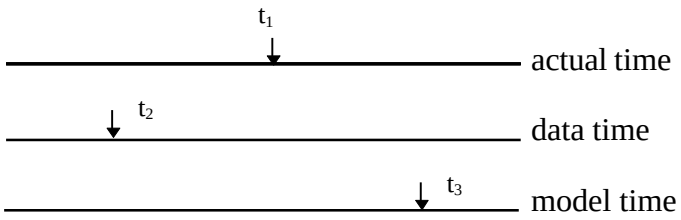

Figure 8: **Different time scales in RODOS**

Now suppose we receive data about t_2 at t_1 and that the model-calculations have progressed up till t_3 (see

Figure 8). Now we are in trouble, when we want to update if we use a particle model, since we can not 'just take back' the model from t_3 to t_2 because of its stochastic nature. Furthermore, it is not straightforward either how to make inferences about an updated distribution at t_2 to a distribution at t_3. In other words: going forward and backward in time is a problem with Lagrangian particle models that does not seem to be easy to overcome. This going forward and backward in time is, in principle at

least, possible with the puff models that we will discuss in the next sections.

4.1 Sources of uncertainty

The theme underlying the use of statistical modelling to improve the predictions of a physical dispersion model, such as RIMPUFF, is that in a data assimilation process uncertainty parameters are omnipresent. Here we indicate the statistical modelling which underlies our Bayesian version of the RIMPUFF-model. This module is designed to answer the following questions:

- What is the likely spread of contamination?
- How can this prediction be updated in the light of monitoring data?
- What are the uncertainties in the predictions?

In practice, to run an atmospheric/deposition model one needs as input:

(a) atmospheric data

(b) parameter settings.

Although atmospheric data themselves are imperfect since they contain errors, the key idea in a statistical context is to capture the uncertainty associated with various parameters of interest, by assigning a distribution to each of these. The way that this is currently done, or can further be improved, is discussed in the next sections. However, in order to get an insight to the role of statistical models as part of the ASY and assess their specific features and attitude for prediction, one needs first to identify the sources of uncertainty in this context. Ranyard and Smith (1997) have classified the uncertainty associated both with the source term and with the dispersion of contamination into four categories:

1. First there is uncertainty associated with the actual emergency event. Thus when an accident happens, it is often clear that one of a small set of scenarios has occurred, each with an associated uncertain consequence relating to the spread of the plume. This type of uncertainty has been elicited in the previous section.
2. Once the accident has occurred, measurements begin to be taken. There are various ways in which measurement errors can happen after an accident which includes errors associated with instruments used in the right way and at the right time, but there are also various examples from previous incidents concerning misreading or misuse of the instruments.
3. A third source of uncertainty concerns the atmospheric conditions. Although dispersal models assume conditions are given at all times and positions, wind field measurements, for example, are typically interpolated from a large irregular grid and therefore only a coarse approximation of their true values can be supplied.

4. Finally, there are many different models of dispersion. The same physical experiments have been performed where the source emissions are controlled and known and meteorological conditions are known far better than they would be in a real accident. Despite this, results show that all these physical models give far from perfect results and are all different from one another.

4.2 A Bayesian Forecasting System : The implementation of Bayes_RIMPUFF.

In view of the above arguments, the use of a statistical model has been suggested as a means to model uncertainty input in a dispersal model. The statistical modelling of the input parameters provides also a solid ground for the updating of the forecasts. We have to note at this point, however, that we do not advocate assessing and inputting probability distributions on all variables which *conceptually* possess some uncertainty. Considering the potential that RIMPUFF can be used for predictions even before an accident, it is hard to envisage that any of the input variables can be assessed deterministically. This would obviously result in a statistical model with too many parameters/variables and the desired computational efficiency would be lost. Instead, we try to capture the uncertainty in the forecasts made by RIMPUFF by placing probability distributions over certain key parameters only. Essentially we model uncertainty in those variables which have a considerable effect on the predictions and take other variables as fixed. The technique of inflating certain variances to allow for model uncertainty serves also to make some allowance for the (small) effect of ignoring the uncertainty in the non-key variables.

In the light of the analysis of the model, our experience of statistical modelling and advice from atmospheric modelling experts, we have chosen as key variables:

- quantity released from the source
- wind direction at the source
- height of the source.

The major influence of these parameters in practice is also elicited in a sensitivity analysis in the Bayes_RIMPUFF validation report (Politis et al, 1999).

Note however the different qualitative character in the treatment of uncertainty among these three parameters: it is only the mass of the release in fact which is assigned a more or less full ‘continuous’ probability distribution, while the values for the other two variables are chosen from a finite number (typically three) of possible values. This set-up obviously reflects the predominant role of the released mass as a major influence for the contamination profile, but also keeps track of the sensitivity of the model’s forecasts to changes in the wind field and the height of the release. Finally the statistical model contains two types of error terms, to allow for

errors both in the model specification and for observational (measurement) errors; moreover a stochastic version of the pentification process is adopted. More details are given below.

Before we proceed to give the statistical formulation of the model as Dynamic Linear Model (DLM), we note that there are four distinct processes that need to be taken into account:

- the source emission process, representing the consecutive releases of mass from the source; this is clearly one of the most important pieces of information which will be needed to be used as a guide for very early decision making.
- during each advection step (which is mass-preserving), there will be deposition and decay which changes the masses.
- the observation process, describing the discrepancies of the measurements, taken either at the chimney stack or at various permanent and mobile detector stations at fixed time intervals, from their actual values. Note that these observations will be air and ground readings, taken either instantaneously or integrated through time, and Gamma dose rates.
- the fragmentation process which describes the pentification of the puffs after they reach a certain size, as presented above.

All these four features of the model are combined to give a mathematical formulation as a DLM in the following form.

To begin with, since we are interested in predictions in both space and time, we adopt the following notation. Let $Q(i)$ denote the mass of the ith puff emitted from the source at time $t=i$. The contribution of $Q(i)$ to the spatial concentration of the contamination at time t and location $s=(s_1,s_2,s_3)$ is given by the product $r_t(s,i)\ Q(i)$; here $r_t(s,i)$ is a function that represents how the release is distributed over time and space. It is typically a complicated function of a number of parameters, but for a given set of atmospheric input data, such as those used in the RIMPUFF code, it can be explicitly determined. Next, let $n(t)$ be the number of puffs emitted before time t. Introducing an error term to model observational errors, we can write for the observed contamination at time t and site s, $Y_t(s)$,

$$Y_t(s) = \sum_{i=1}^{n(t)} r_t(s,i)Q(i) + \varepsilon(s)\,. \qquad (6.1)$$

Note that at a given time point t, observations will arrive in one of the following two forms:

(a) a single stack observation $Y(t,0)$ of the source mass at $s=0$ which includes measurement error, *or*

(b) a vector of observations, $Y(t,\boldsymbol{s})$, where s represents the vector of the (probably, predetermined) sites where measurements are taken, again including error.

Thus writing $\boldsymbol{Y}_t=Y(t,s)$ for the observation vector at time t, and similarly $\boldsymbol{Q}_t$, and $\boldsymbol{\varepsilon}_t$ for the vectors of puff masses and observational errors respectively and R_t, = $(r_t(s,i))$ for the matrix of the dispersal parameters, we see that (6.1) can be written in matrix form as

$$\boldsymbol{Y}_t= R_t\boldsymbol{Q}_t,+ \boldsymbol{\varepsilon}_t,$$

and it is assumed that the vector $\boldsymbol{\varepsilon}_t$ has a covariance matrix V_t which is known. Thus, if one assumes in addition that the error term is Gaussian, we have that

$$\boldsymbol{Y}_t\sim N(R_t\boldsymbol{Q}_t,V_t).$$

In order to model the evolution of the spatial contamination through time as a DLM, one needs also to specify a system, or transition, equation. Since $\boldsymbol{Q}_t=(\boldsymbol{Q}_{t-1},Q(t))$, a simple such formulation can be given under the assumption that adjacent puffs are highly positively correlated, so that the emission process might be modelled as a random walk. However, expert judgement often suggests that air concentration readings increase asymptotically at an exponential rate. Therefore a modelling including a more general relationship between consecutive puffs can be adopted, see Ranyard and Smith (1997). This latter formulation though, has still some practical and technical inadequacies, see French *et al* (1995b) and Ranyard and Smith (1997).

To model the possibility that the original puff was drifting away from its predicted trajectory (e.g. due to local wind effects, plume splitting to branch around hills, or shearing), Smith and French (1993) proposed a stochastic version of the pentification process. To deal with this pentification process, we will have to introduce some new notation. Denote by $l=(t,l_1,l_2,\ldots,l_k)$ the puff fragment which is the l_k-th fragment of the l_{k-1}-th of the … of the l_1-th fragment of the puff released at time t. Further, let $Q(l)$ denote the mass of this puff and let I_t denote the set of all puff fragments appearing on or before time t. Now, since $Q(l)$ denotes the mass of puff l (either itself emitted from the source or created through pentification), the vector of masses of its children, $\boldsymbol{Q}_c(l)$, is related to $Q(l)$ by the formula

$$\boldsymbol{Q}_c(l)= \rho\, Q(l)+ w,$$

where ρ is the vector governing the mass allocation to the siblings from Mikkelsen *et al* (1984), i.e. ρ=(0.235, 0.235, 0.235, 0.235, 0.058) and w is a random vector chosen in a way that mass is conserved, i.e. $w'(1,\ldots 1)'=0$. To complete the model formulation, it remains to decide about the covariance matrix of the w's. This can be chosen so that it

represents different uncertainties in the local windfield, in particular the wind speed and direction; see Smith and French (1993) for more details.

Finally define the vector that contains both parent and children puffs as $\boldsymbol{Q}_{pc}=(Q(l),\boldsymbol{Q}_c)^T$ and revise the definition of $\boldsymbol{Q}_t$ above to be the vector of puff masses that exist on or before time t.

The last feature we have to include here is the process associated with the advection of the puffs. Although the advection process itself is mass-preserving, we also include decay and deposition in this step. And not surprisingly, there is quite some uncertainty involved in especially deposition, so we have to include that here. This is also the place where we make explicit the increasing uncertainty over time. In mathematical terms we express this in the following way:

$$\boldsymbol{Q}_{t+1} \sim N(H_t\boldsymbol{Q}_t, U_t)$$

We can also inflate the diagonal of U_t somewhat to allow for modelling error. Note that the indexed t here does not have to be the same as the indexed t earlier used for the release time in general: it is quite natural to use different time steps for release and advection.

The forecast of contamination can now take place in a straightforward manner, following the methodology given in Harrison and West (1989). Note that the computer program developed accordingly requires as a minimal input for a single run the specification of the (prior) variances of the observational and the transitional error terms and the expectation of the source release. These can be assigned after consultation from experts to reflect our prior beliefs for the parameters. Further note that we have not considered yet the uncertainty surrounding the height of the release and the windfield discussed above. As already indicated there, a natural way to incorporate this into the model is to consider a small number of possible scenarios for these quantities. Thus, instead of setting e.g. the wind direction equal to φ, we may rotate this by $\pm\theta$, and assign probabilities to each of φ, $\varphi+\theta$, $\varphi-\theta$ for the direction of the initial release. This can be thought as a three-point approximation to the probability distribution of wind direction and probably expert judgement is again needed to be implemented in practice. By doing the same for the distribution of the parameter 'height of release', we thus obtain nine combinations for the setting of these two parameters. As soon as monitoring data are collected, Bayes_RIMPUFF then finds at each time step the posterior distributions for height and wind direction given the data. Further, Bayes' theorem in a straightforward way can be employed to find the predictive distribution of contamination at a particular time and site for this mixed model, see for example Ranyard and Smith (1997). Clearly our current choice for the number of distinct prior settings for the wind direction and release height to be three is only a means to combine tractability with parsimony in the model and there is no theoretical limitation for any alternative settings. Moreover we mention that this modelling takes also account of the

sensitivity associated with the wind speed: when RIMPUFF calls LINCOM to calculate the local wind field, this produces different wind speeds at different source heights.

The formulation of the model as a discrete time series above is in accordance with RIMPUFF, which assumes that the puffs are emitted in discrete time. Although at present the model assumes that the time interval between successive emissions is constant, this restriction is easily removed if necessary, provided that the timing of the observations is known *a priori* (see Harrison and West, 1989, pp. 358-359). French and Smith (1993) have also indicated how the observations are encoded in the model if the measurements are not instantaneous, but integrated over time. Another attractive feature of the model is that at any time step, the contamination profile provides updated forecasts for the source term.

Finally it would be unfortunate, not least for its intuitive appeal, if in the absence of data the model predictions did not agree with those of the deterministic model. Mathematically this is easily taken care of: by assigning very large (theoretically infinite) variance to the observational errors, the model behaves as if no data are present. Thus the Bayesian model contains the deterministic, physical model as a special case. More generally, the initial setting of the variances in the model should be such that it reflects our degree of uncertainty about the atmospheric conditions and the spatial distribution of the plume.

4.3 Extensions of the model

This standard linear model formulation of RIMPUFF given in French and Smith (1993) and outlined in the previous section, allows for sequential predictions of contamination in the vector of sites *s* and adapts its forecasts sequentially using Bayes' theorem when more data arrive. Various practical considerations have to be taken into account however relating to the particular nature of the problem, when the model is used in conjunction with RIMPUFF in the emergency phase of an accident. Recently, a number of extensions of the model have been considered with a view to

- meet the fierce time limitations of an accident and the need for the subsequent decision making at an early stage; thus, much more efficient algorithms have to be developed
- deal with situations where the model described so far is inadequate or inappropriate. Before the use of a DLM for example, one has to be convinced that the Gaussian distribution is a reasonable description for the data observed
- explore and validate diagnostic tests for the adequacy of this, or any other subsequently developed model, for predictions.

First it was realised that in practical terms, the updating of observations using this Bayesian forecasting system had computational difficulties,

especially when the rate of pentification is fast and there are very many puffs present at the same time. In order for its computations to be feasible within RODOS, the speed of the model calculations had to be of the same order with those of the standard RIMPUFF model without updating. Technically, the computer code based on the model above was substantially slower than the original RIMPUFF code because it manipulated matrices whose number of rows and columns was several hundreds. For this reason, a much more efficient technical device was suggested by Smith *et al* (1995) ; see Appendix B in French *et al* (1995b). This is based on utilising ideas from the Bayesian belief networks, as they have been demonstrated for example in Lauritzen and Spiegelhalter (1988), and Dawid (1992). Exploiting various conditional independencies among a finite collection of random variables, it is possible to build a fast algorithm, often referred to as a probabilistic expert system, for handling the joint probability distributions of these variables. Typically, the causal probabilistic structure underlying the construction of an expert system is at first represented by a causal network, i.e. a directed graph which exhibits these causalities in a qualitative way. Manipulation of the graph involves first its triangulation by joining nodes with common descendants. Then a sequence of cliques has to be defined that satisfies the running intersection property (see Appendix A for the meaning of this). After that, connect the cliques with each other (via separators) with undirected edges. The resulting undirected graph is then known as a junction tree (see Dawid, 1992). The concept of cliques is precisely that, which allows the complex probabilistic information to be stored in an efficient manner.

These ideas have been developed and brought into a more general context in Smith *et al* (1995), with a view to be incorporated into the dynamic version of RIMPUFF. More explicitly, instead of a static relationship between the variables, it is now assumed that this relationship evolves through time. The evolution of the contamination process after an accident (source emissions/fragmentation) is now represented by a *dynamic junction tree,* whose cliques are either:

(i) the masses of two puffs emitted consecutively from the source, or

(ii) the masses of a parent puff and its associated children, provided that the parent puff has already fragmented by the time in question.

Then we define (where $\psi(t)$ is a dummy variable introduced to express prior source term information):

$C^*(t) = \{Q(t), \psi(t), Q(t+1), \psi(t+1)\}$, $1 \le t \le T\text{-}1$

and

$C(l) = \{\boldsymbol{Q}_{pc}(l)\} = \{Q(l), \boldsymbol{Q}_c(l)\}$, $l \in I_T$.

All cliques in the junction tree have one of these two forms.

Such a clique representation is shown in Figure 9 below. In this example, the cliques labelled $C^*(1)$, $C^*(2)$, $C^*(3)$, $C^*(4)$ represent cliques of variables associated with source puffs and their future neighbour. The other cliques of the form $C(l)$ represent masses of parent puffs with their children. The first puff has already pentified to give clique $C(1)$. The second, fourth and fifth fragment of this puff have in turn pentified to give rise to further cliques $C(1,2)$, $C(1,4)$ and $C(1,5)$. The fourth and fifth child in the clique $C(1,2)$ have also fragmented to give cliques $C(1,2,4)$ and $C(1,2,5)$ whilst the third child of the clique $C(1,4)$ has fragmented to give $C(1,4,3)$ and so on. The edges on this tree indicate where cliques contain a common puff or puff fragment.

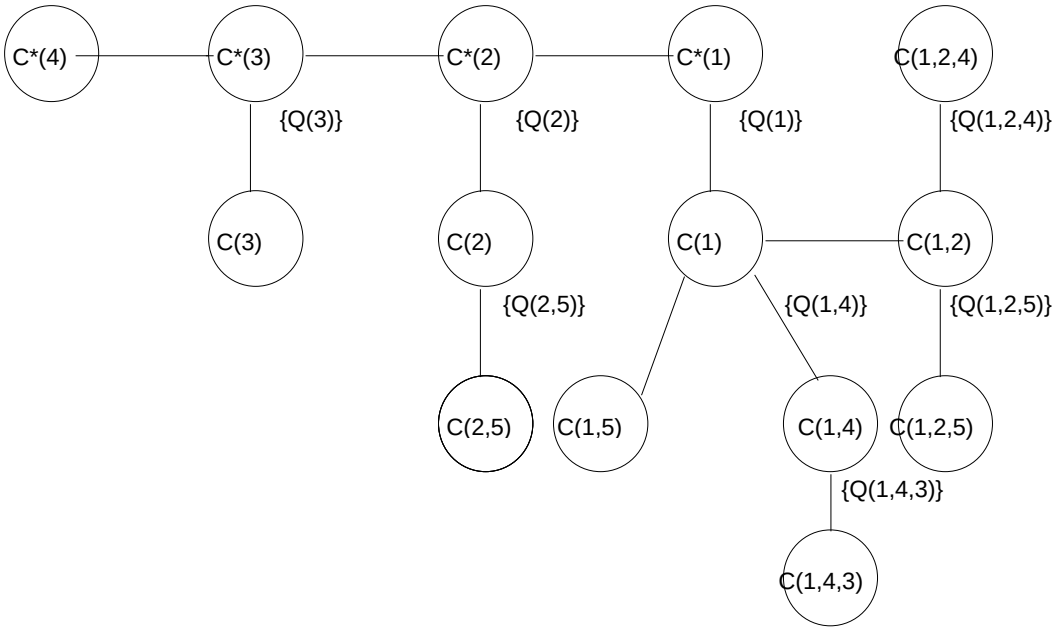

Figure 9: A Dynamic Junction Tree

In view of the systematic dependence between the masses of puff fragments and their parents, this mechanism allows for the joint density $p_T(\boldsymbol{Q}_T)$ of the vector of mass emissions and their fragments, existing on or before time T, to be expressed in terms of the marginal densities of the masses belonging to the same clique. Further, since all cliques are of small dimension, this gives an efficient propagation scheme for the updating of probabilities as new evidence arrives. Smith *et al* (1995) have demonstrated how this works by introducing a simple device of dummy variables. See that paper for details about how the system updates its forecasts when new data related to a single clique arrive. Note that we clearly keep track of the joint distribution of all puffs, even those that don't exist anymore. It is precisely this feature which allows the algorithm to

incorporate observations that were taken some time ago, but not received until now.

The above representation of the model included in the RIMPUFF code has enhanced the computational efficiency of the model significantly, so that the program can update its forecasts within the time limits of RODOS and the ASY in particular. To meet more realistic needs, however, concerning the assumptions underlying the formulation of the model, several extensions have been considered. In particular, the analysis in Smith *et al* (1995) was, apart from the Markovian property of the system, which is implicit already in most atmospheric dispersion models, such as RIMPUFF, and needed to be respected (Ranyard and Smith, 1997), also based on the following assumptions:

(i) The normality of the system, i.e. the joint distribution of observations and puff masses is Gaussian.

(ii) Any observation taken relates to the masses of puffs within a single clique.

The significance of these issues was already addressed in Smith *et al* (1995), where some directions to tackle departures from these assumptions were also given. More recently, some effort has been put towards the study and validation of models that do not depend on these two conditions.

To begin with, we note that in the emergency phase after a nuclear accident, most of the measurements collected are essentially count data; this motivates consideration of the Poisson distribution to accommodate such data. One standard way of dealing with non-normality of the contamination readings, as well as other possible discrepancies from the standard DLM framework, is to consider the more general set-up of Dynamic Generalised Linear Models (DGLM), see West and Harrison (1989). With reference to the junction tree with cliques defined above, the updating techniques for a DGLM (West and Harrison, 1989, Chapter 14) can be seen as an approximation to a full Bayesian analysis (Smith, 1992). Gargoum and Smith (1997b) have explored this by comparing the contamination predictions of a Normal model with those of one where the observations are assumed to follow the Poisson distribution using the same data set. It appears that the model using the Poisson distribution captures more quickly the shape of the contamination profile. To assess the quality of their approximations, they suggest the use of the Hellinger metric; for the use of the latter within Bayesian forecasting, see Smith (1995). An alternative criterion, based on the Kullback-Leibler separation measure, is discussed in Gargoum (1997).

A related method, based on numerical rather than algebraic computations, has been given recently in Settimi and Smith (1997). It uses the Markov Chain Monte Carlo algorithm (see, for example Gelman et al, 1995, Gilks et al, 1993, Smith and Roberts, 1993, for its use in various applications) to simulate observations from the posterior densities of the parameters.

Although the MCMC and its variants have already been established as a common practice in various contexts within the Bayesian framework, their applicability within RODOS has proved to be limited, mainly since they fail to meet the time requirements of the system. The numerical method is a lot slower compared to the algebraic approximation of Gargoum and Smith (1997b), however, Settimi and Smith (1997) advocate its use in order to assess the efficiency, in predicting future observations, of various probabilistic propagation algorithms. In the event of large Poisson counts present, the comparison of Settimi and Smith (1997) shows that predictions based on the DGLM are close to those derived from the Gibbs sampler (MCMC method). Where the efficiency of the algebraic method appears to be less satisfactory is at circumstances where small values of the Poisson distribution are observed. This is of little concern for its use when running RIMPUFF after an accident, since it is most likely that the detector points where observations are taken, will be chosen at sites with high contamination. It might be useful to note at this point that these methods extend in a theoretically straightforward manner to situations where the distribution of the observations is other than the Poisson. For instance, the approximate algorithms based on the DGLM which are discussed in Gargoum and Smith (1997b) are valid for a wide family of distributions, known in statistics as the exponential family, and in fact even the latter constraint might be unnecessary.

Another assumption which has been made so far for the clique representation of the process is that the observations at any particular site *s* taken at time *t*, *Y(t,s)*, depend on the masses of the puffs of the same clique alone. Section 4 in Smith *et al* (1995) describes how the system updates its predictions by absorbing evidence about a single clique. In cases where the wind field is turbulent, or the terrain inhomogeneous, this might not be a realistic assumption. A more general algorithm which includes the possibility that an observation which arrives depends on puffs that belong to different cliques has been developed by Smith and Papamichail (1996), see section 5 there. More explicitly, if the contamination at a particular site and time Y(t,s) depends on a number of puffs which belong to the cliques $c[1],\ldots,c[r]$, then it is possible to transform the original forest formation to a new one which has $d=\bigcup_{i=1}^{r} c[i]$ as a clique (either *d* itself is a node in the new forest or it consists of two or more nodes of finite unions of the $c[i]$) and then apply the algorithm of Lauritzen and Spiegelhalter (1988). Inevitably the propagation of probabilities using this modified algorithm is now slower than the original one, and its efficiency seems to depend critically on the normality assumption. If the system is normal, then standard matrix manipulation from multivariate normal distribution theory (see, e.g. Mardia et al, 1979) allows for efficient computations; if departure of normality is evident, then we may need an integration or summation of marginals and this can reduce the speed of the algorithm significantly. Three more algorithms for efficient propagation of probabilities are

discussed in Smith and Papamichail (1996, section 6). These again rely on transforming the original forest into a new one and they might be appropriate for updating in particular circumstances. Again the use of the Hellinger metric offers a diagnostic test for the appropriateness of these propagation methods.

4.4 Current status of the model - diagnostics

The implementation of Bayes_RIMPUFF into a fast computer algorithm, the first results of which were recorded in French *et al* (1995b), has been developed by the Warwick - Manchester group. The progress in the computer code follows, with the obvious time lag, that of the atmospheric dispersion model itself, as this is investigated at the Risø Laboratory in Denmark. At the same time, the program reflects the flexibility in the Bayesian modelling of RIMPUFF, as it has been demonstrated in the previous section, allowing for different modelling schemes as input. Currently the program operates with a single nuclide, a time-varying windfield and allows three options for the distribution of the observations: deterministic (i.e. no updating takes place), Normal and Poisson. Further it allows the possibility that the updating uses the Clique structure and the Hellinger metric as discussed in Section 4.3.

Ideally, the program should output into a temporal database positions, clique membership, means and variances of the puffs at each time step. But if the CSY and GSY subsystems are incapable of calling some RIMPUFF modules to calculate the quantities they want, we are able to produce a grid of air and ground concentrations and variances of these at each time step or at a given time step. We are also able to produce covariance matrices for any time step but this may be very time consuming. It may be better to identify default shapes for these covariance matrices *a priori* (see Smith and French, 1993) and use these together with an estimated scale factor, instead of calculating accident-specific ones.

As regards testing the adequacy of the model for particular circumstances, on a theoretical basis the use of the Hellinger metric provides approximations which work reasonably well. As indicated in Section 4.3, Bayesian computational algorithms which have been developed recently, might not be themselves at the current stage fast enough to provide future predictions, but may serve as a basis for testing the efficiency of other methods. We aim to explore and elaborate this issue further in the future.

General diagnostic tests, such as residual analysis, are available for examining the fitness of a model on a particular set of observations. Other plots which will be produced might detect noticeable departures of the observations from their predicted values or the inappropriateness of some of the model assumptions, such as normality. We note that the monitoring of diagnostic plots for the forecasting model depends largely on the abilities and the experience of the user, as it is a process that cannot be fully automated.

We further emphasise that any model output needs to be critically examined by relating predictions to observations and, maybe, by exploring sensitivity analysis on gross changes in wind direction - this would require a model rerun. Note that we can run RIMPUFF at various computational speeds and levels of approximation depending of the scale of the time step. Thus there is the facility to produce quick explanatory runs.

The main worrying aspect about the model so far, as with any other model in RODOS, is that we have not been able to test its sufficiency against *real data*. We have tested the algorithm on simulated data sets, such as Lundtofte Nord : see French and Smith (1992) and Politis *et al* (1999). The model fits the data which we have simulated from the results of these experiments reasonably well . However, while these tests are encouraging and provide us with insight and confidence in the model's behaviour, it is relatively easy to work with simulated data (i.e. data generated by the model itself) whereas a more rigorous test is to use real data from a tracer experiment. At present we have no data sets which are satisfactory for a full test of the program - nor has anyone else. Any satisfactory data set needs to record prior forecasts of wind field, source term etc. as well as the monitoring observations. It is likely that our modified RIMPUFF will be most effective when instantaneous Gamma dose rate measurements are accommodated, but this has yet to be tested and that will be a difficult practical test to arrange!

4.5 The Countermeasure Subsystem: ECOAMOR

The role of the Countermeasure subsystem at this early phase is to identify a list of possible actions to be taken, such as evacuation, sheltering, or issue of iodine tablets and quantify the consequences of each action by considering the costs and benefits of all possible countermeasures. Note that different areas will have different needs and will be exposed to varying degrees of risk, so that each of these actions (strategies) will in fact be a series of countermeasures applied to a collection of sites. In the early stages of a release, the analysis in the countermeasure modules will focus primarily on calculating the expected amount of dose averted or the collective amount of dose per individual for each strategy (i.e. a combination of countermeasures in different areas), as these are more likely to be of primary concern rather than e.g. money.

The input to the CSY modules will be predictions about the shape and spread of the plume. Other information such as forecasts for the deposition parameters and external gamma dose rates will also be available from the ASY modules.

The structure of the CSY subsystem with a description of some of the currently available modules has been given in Ehrhardt *et al* (1997), see also Päsler-Sauer and Schichtel (1997).Some of the prototype models that form part of the Early Countermeasure package, such as EVSIM or ECONOM have not as yet been integrated into the RODOS user interface and thus their output cannot be used for manipulation by the ESY modules. In principle, it is envisaged that the CSY modules will be in constant interaction with the ESY subsystem which is responsible for the assessment of the alternative strategies. Table 2 presents the data required from various CSY models to be processed by the ESY in the early phase, including the period immediately prior to an accident. Currently a number of problems remain to be resolved regarding the communication of modules in the CSY/ESY

Decision Point	Countermeasures	CSY module	Input data required
Threat phase	• Warning the public	RODOS STM	• Estimated remaining time before the release
		EVSIM / STOP	For each block: • population number and evacuation rate, or, • time needed to evacuate
Early phase	• Issue of iodine tablets • Sheltering • Evacuation	EMERSIM	• Blocks affected by the plume of radiation For each block and for each countermeasure: • average sum individual effective averted dose For each countermeasure • begin time and end time • low and upper intervention levels For each strategy: • average sum individual effective averted dose • sum collective effective averted dose • number of people involved
		HEALTH	For each strategy: • number of people with deterministic health effects • number of people with stochastic somatic health effects
		ECONOM	For each strategy: • cost (implementation + loss of productivity)

Table 2: Data required from CSY modules in early/threat phase

interface, see the discussion in French and Papamichail (1997). For example, during the analysis and evaluation of alternative strategies by the ESY using multi-attribute value/utility theory (see the next section, the ESY modules produce a list with the best 10-15 actions according to a number of specified criteria. It might then be required at this point that these strategies are fed back to the CSY modules so that a closer evaluation of their costs and benefits is available. The current functionality of the operating system of RODOS prohibits the possibility that the output of an ESY module might be used as input to a CSY module and this seems to be a desired implementation for the future.

Finally we mention that for decision making at this early stage, advice is often to use site specific intervention levels: e.g. evacuate if the measured or predicted air concentration or gamma dose rate at a certain point exceeds a nominated level. The CSY/ESY system may then simply calculate the probability of exceeding these site specific levels for each strategy. A more sophisticate analysis from the CSY modules may involve probability methods to calculate expected doses/concentrations and variances of these. Assuming that the quantities are linear functions of the mass of contamination, this involves only a straightforward modification of the current CSY modules. If linearity cannot be assumed, then approximation methods are needed ad the CSY needs to calculate in each emergency administration region/sector the reduction of the dose, costs, population size which is likely to be affected and other attributes. Currently in RODOS, the EMERSIM module identifies the areas with likely emergency actions, simulates these actions and calculates individual doses with and without countermeasures, providing also an estimate of the duration that each countermeasure needs to be in effect.

4.6 The Evaluation Subsystem

One of the distinctive characteristics of RODOS is its ability to offer full support at level 3 (see Table 1) to the DMs. This is implemented by the ESY modules which rank countermeasure strategies according to their potential merits and preference weights provided by the DMs. The necessity for such an implementation stems from the fact that the number of possible countermeasures which are output from the CSY can be huge. A countermeasure strategy typically applies to a region and generally specifies different protective measures for small subregions depending on the level of contamination in each of these. Thus the number of possible strategies grows combinatorially with the number of subregions and the number of possible protective actions. Further, there are often constraints on the feasibility of countermeasure strategies relating to national and international guidelines on radiation protection, practicability, and so on.

To overcome these difficulties, the use of expert system technology has been proposed; see Borzenko and French (1996), French (1996), and French and Papamichail (1997). The advantage of introducing this

machinery is that ESY modules discard unlikely strategies before a full evaluation of the remaining ones is undertaken.

The conceptual structure of the Evaluation Subsystem in RODOS has already been introduced. A more realistic picture of the process that depicts also the interactions of the subsystem with other parts of RODOS is given in Figure 10. Here we describe the functionality of each of the three component modules at the current stage and discuss their implementation.

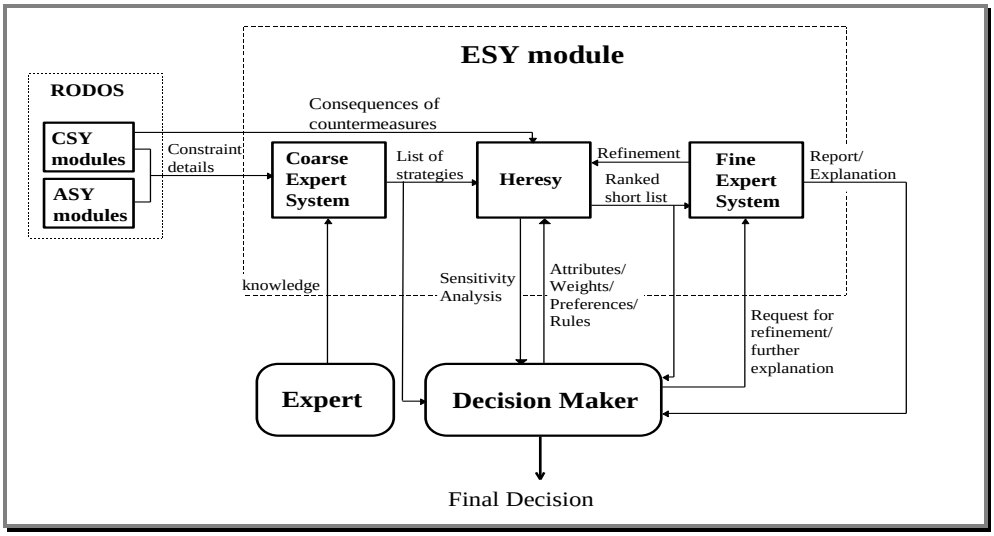

Figure 10: Interaction between DMs and ESY modules

At a first stage, strategies output from the CSY are filtered through a coarse expert system (CES). This is used to reject all those strategies that are not feasible (e.g. evacuate 100,000 people in a few minutes) or they contradict either common sense (e.g. issue iodine tablets in an empty area) or basic principles such as public acceptability, equity etc. Since the strategies remaining after the CES filtering need to be assessed as fully as possible in the subsequent stages, their number needs to be of a manageable size. In practice it has been found that although the number of possible strategies which are input in the ESY can be several billions, typically after their manipulation with the CES only a few hundred remain to be further elaborated. Moreover, the implementation of such a coarse expert system is based on constraint management satisfaction (see, e.g. Tsang, 1993) and the currently available software ensures that such a reduction can be achieved in a few seconds only. An investigation of practical aspects for this implementation was carried out by Papamichail and French (1997). In a hypothetical scenario, the area surrounding a nuclear plant was divided in 17 blocks and in any block a unified strategy has to be adopted. By considering three early phase possible countermeasures for each block, we see that there are 2^3 different strategies for each block, while the number of possible actions for the whole area rises to $2^{3.17} \cong 2.25 \times 10^{15}$. Once a number of constraints for the countermeasures is listed and input in a

constraint satisfaction software package, the program discards all the actions that fail to satisfy one or more of the desired criteria. Papamichail and French (1997) have used the ILOG Solver tool to reduce the number of possible actions with a number of different randomly generated input data. Using 1000 simulations, it was found that in almost 60% of the cases the CES produced less than 50 strategies, while 80% of the time the number of possible actions was up to 300. More details of the implementation are given in Papamichail and French (1997), where also some undesired features of its current status which need further elaboration are discussed. The two main drawbacks of the CES currently are that it does not incorporate a sensitivity analysis, i.e. if a strategy fails to satisfy a singe constraint marginally then it is rejected by the program, and perhaps more seriously there is no upper bound in the number of actions that are output by the CES. Although this number is substantially smaller than the input number of strategies, there is at present no guarantee that the system will produce a small number of actions as output in a particular release scenario. Given the time considerations in the emergency phase of an accident, it is unlikely that unlikely that several thousands of actions may be fully assessed and ranked in the following stages of the ESY in RODOS. If the system is adopted for use in RODOS, it is required to include a device that selects only (up to) a fixed number of strategies according to some generally agreed rules, e.g. the expected amount of averted dose which might be a priority at this stage.

A typical example of a CES output screen is given in Figure 11.

The second step of the ESY system consists of a ranking system which evaluates different alternatives and provides a short list of 10-15 strategies to the DMs. The first two modules that have been developed in RODOS as parts of the ESY system for this stage, HERESY and M-Crit, have been described in Borzenko and French (1996), see also French *et al* (1995b)

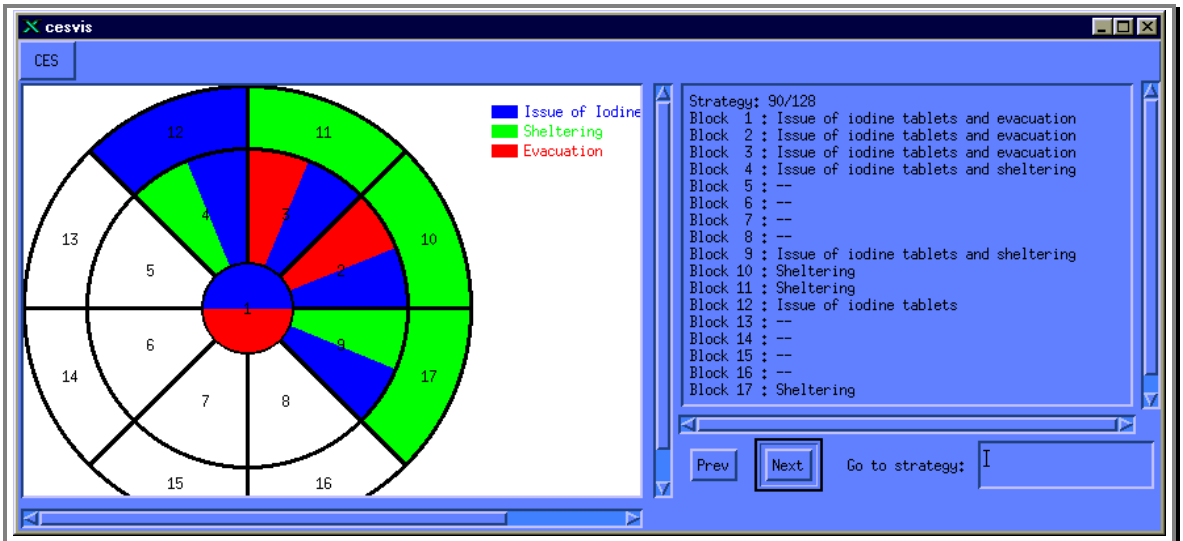

Figure 11: Example of the CES output

and Borzenko *et al* (1994). Briefly, HERESY uses MAV/MAUT theory to rank the countermeasures exploiting attribute independencies, while M-Crit exploits a numerical technique to approximate the DMs' indifferences. The current implementation of HERESY uses a linear multi-attribute value function. This requires sufficient independence among the variables (e.g. preferential independence) and any uncertainty information is ignored: value functions assume that the attribute values are known with certainty. One of the advantages of this simplification is that sensitivity analysis for the weights in the value function can be carried out in a straightforward and computationally efficient way. The main task of HERESY is thus to identify the top few, say 10, alternatives and check the sensitivity of these to the choice of the weights. On the other hand, M-Crit, which is based on much weaker assumptions, seeks to approximate more complex forms of a MAV function that allow for many dependencies in the preferences to be present. The disadvantage of this is that the method seems to be most appropriate for circumstances where the DMs (DMs) are sure of their preferences *a priori*. Further, the elicitation procedure does not include any firm consistency diagnostics, such as sensitivity analysis, to help the DMs identify the robustness of the model.

Although different in a number of characteristics, HERESY and M-Crit have a lot in common and their potential merits are such that they can be used in conjunction in RODOS. They can both be used interactively by the DMs and they are based on the same principles. For instance, they both encompass the multiplicity of criteria for the ranking of alternatives in their assessment and they model preferences as multi-attribute value functions. Further, the fact that they both use a graphical interface which facilitates communication between more than one DMs, with varying skills, background and perspective.

As soon as the countermeasure strategies have been shortlisted, the remaining 10-15 action scenarios will be passed on to a fine expert system (FES) with a sophisticated set of rules which will output to the DMs a detailed commentary on each strategy highlighting its relative strengths and weaknesses. Klein (1994) discusses a similar combination of expert system technology with MAV/MAUT ideas to provide DMs with an insight on the ranking of alternative strategies. The actual implementation of the FES within RODOS would obviously depend on the time that it is used and the particular circumstances surrounding the release. However, since this final filtering process has received little attention so far, at least the common agreed principles for its use need to be further explored and elaborated, particularly if it is to be used in the emergency phase, and this seems to deserve further consideration and discussion among RODOS participants.

More details about the use and functionality of the ESY subsystem are given in the next sections, focusing on the evaluation of strategies in the later stages of an accident. Bearing in mind the underlying temporal structure of

RODOS subsystems (see Section 2), the ESY modules will be called and function several times before, during and after an accident. The use of these modules will vary depending on the time that they are employed, since the attribute hierarchies will inevitably change through time, and taking into account the time limitations for a decision to be made. Note however that in contrast with the ASY and the CSY, the ESY may be based upon the same software module regardless of the time that it is called, with only the attribute trees, the constraints and the preference weights changing over time.

Finally, since it is envisaged that it in the event of an accident the ESY modules will be in constant interaction with the DMs, the use of the ESY modules on any particular should obviously reflect the needs of the DMs. The current evidence from the elicitation workshops that have taken place so far suggests that in the early phases of an accident, the DMs may not feel the need for support at level 3. Although this is an issue that needs to be explored further after more workshops are carried out, if it is the case that no such support is required immediately after an accident, detailed analyses about the ranking of countermeasures may be postponed until the following stages.

5 The plume is still spreading locally: food bans and other countermeasures.

5.1 ASY subsystem

While the plume is still spreading, the ASY subsystem will be RIMPUFF or some other atmospheric dispersion module. The inputs and the outputs will be as described earlier. Naturally there will be a greater number of stack and monitoring data available. As a consequence if Bayes_RIMPUFF is employed, we can get a lot more accurate forecasts than those in the emergency phase, since the parameters involved will be subject to smaller variability. Using these forecasts the ASY will output updated estimates for the areas of high contamination and the dispersion profile of the plume. Further, the greater amount of data available enables more extensive use of management by exception so that the sufficiency of the proposed modelling may be assessed more accurately. Note however that gradually the interest of the ASY modules will focus on factors such as the exposure of humans to radiation due to ingestion of contaminated foodstuff, exposure to radionuclides deposited on the ground or on skin and clothes of people, etc. Thus the output of the atmospheric dispersion model, including in particular estimates of the gamma doses, is processed by the Food and Dose Module of RODOS (FDM). The main stages of this module, as categorised in Faria *et al* (1997) are as follows:

(i) deposition calculations, where the interest is to model the capability of plants, soil and grassland in absorbing deposited radioactivity;

(ii) concentrations on plants, where from the deposited and absorbed activity diverse concentrations are calculated such as on roots, leaves and fruits of plants;

(iii) concentrations on animal products, where from plants (including grass) used as feedstuff, concentrations on animal products such as beef and milk are calculated;

(iv) concentrations on processed foodstuff; and,

(v) dose rates calculations, such as human ingested and inhaled ones.

For stage (i) above, Faria *et al* (1997) have recently considered a space-time Bayesian statistical model, the Bayes-FDM model. The model assumes that a number of fixed monitoring stations capable of monitoring gamma dose rates and aggregated ground contamination over a period of two hours are available. The success of the model seems to depend on the degree that these facilities are available; in Germany, for instance, with the densest network in Europe, there exist currently 2200 such stations, while in general other European countries have a much sparser monitoring network. Moreover the model takes account of the sensitivity of deposition parameters to the weather conditions, particularly rain intensities. The model inputs are thus not only estimates for doses/concentrations from the

atmospheric dispersion module, but also predictions about meteorological parameters and their stability. All this information is combined to form a prior distribution for the ground deposition which is in turn updated once data from the monitoring stations become available. A complicating factor in the modelling of deposition is the fact the radioactivity produced by the deposited isotope depends on the type of the radionuclide present.

At this initial investigation of Bayesian modelling for deposition, Faria *et al* (1997) have modelled ground deposition and contamination of plants. More specifically, let at given time, the total gamma dose rate at site *r* be denoted by $D_{tot}(r)$. In order to obtain an estimate for the ground deposition we decompose this into the contribution to the gamma dose by air and ground activities, $D_g(r)$ and $D_{air}(r)$ respectively. Faria *et al* (1997) quote Bleher and Jacob (1993) for the form of the ground activity which can be written as follows:

$$D_g(r) = D_{tot}(r) - D_{air}(r) = \sum_k \delta_g(k)\, A_g(k,r),$$

where the summation is taken over all nuclides *k* present, $\delta_g(k)$ is a nuclide specific factor for ground concentrations and $A_g(k,r)$ is the deposition due to ground activities only for nuclide *k.*

The calculations in the model are based on decomposing the ground deposition into its wet and dry components. Each of these can be obtained on using the gamma dose rate measurements and (measured or estimated) air concentrations for the site *r*. Further some information about the amount of rainfall is needed to obtain the wet deposition. A practical problem which arises here is that in order to derive estimates for deposition *per isotope*, we need to be able to obtain air concentrations at *r* for each nuclide. Alternatively, gamma spectrometry can be used which allows identification of the isotopes.

A convenient feature of the model for its use in conjunction with an atmospheric dispersion model, such as RIMPUFF, is that it uses the same grid of points where output predictions for the parameters are output from the dispersion model. Moreover using spatial interpolation, forecasts for ground deposition at off-grid points, such as the monitoring stations where gamma dose measurements are taken, are provided.

More explicitly, let $\underset{\sim}{s}=(s_1,s_2,\ldots s_n)$ be the spatial grid of points in the dispersion model and $\underset{\sim}{r}=(r_1,r_2,\ldots r_m)$ be the set of points for the monitoring stations. We define an $m \times n$ matrix $G(\Delta \underset{\sim}{s})$ each element g(i,j) of which represents a formula which interpolate ground depositions at a site s_i on the grid to an off-grid site r_j at a monitoring station. Then, the ground depositions $\phi(\underset{\sim}{r})$ at $\underset{\sim}{r}$ can be modelled as:

$$\phi(\underset{\sim}{r}) = G(\Delta \underset{\sim}{s})\ \theta(\underset{\sim}{s}) + \varepsilon_\varphi(\underset{\sim}{r}).$$

Here $\theta(\underset{\sim}{s})$ is a set of parameters in the model and $\varepsilon_\varphi(\underset{\sim}{r})$ is a random vector representing the interpolation errors.

Using the predictions of the atmospheric dispersion model as data and assigning a (typically, noninformative) prior distribution to the parameter vector θ, the model gives first a posterior distribution for θ given the data. This is in turn used to provide a posterior distribution for φ given the information at the monitoring stations. We mention that the variability of the updated forecasts reflects the uncertainty present in various parameters of the model such as the rainfall or the air concentrations at off-grid points

The Bayes-FDM model calculates for any point s_i in the grid, the total contamination $A(i,k,s_i)$ caused by deposition at plants of type i of isotope k. The absorbed activity $A(i,k,s_i)$ is calculated for all types of plants at each site s_i as if all types of plants were present at that site. The calculations are carried out again by decomposing the total activity into its wet and dry components.

More explicitly, if $A_{dry}(i,k,s)$ and $A_{wet}(i,k,s)$ stand for the contributions of dry and wet deposition to the total contamination at a given site s, then

$$A(i,k,s)= A_{dry}(i,k,s) + f(i,k,s)A_{wet}(i,k,s)$$

where the factor $f(i,k,s)$ is called the *interception fraction* of isotope k for plants type i and depends on the amount of rain at the site s. Further, the dry component of the contamination is given by the formula

$$A_{dry}(i,k,s)= u(i,k)\ \bar{C}_{air}(k,s).$$

Here $u(i,k)$is the deposition velocity of isotope k on plants of type i and $\bar{C}_{air}(k,s)$ denotes the time-integrated concentration of radioactivity at s.

Following the above, a suitable formulation of total contamination for isotope k on plants of type i at the whole grid of points, $A(i,k,\underset{\sim}{s})= (A(i,k,s_1),\ldots, A(i,k,s_n))$, as a linear regression model is as follows

$$A(i,k,\underset{\sim}{s}) = \mathbf{F}(i,k,\underset{\sim}{s})\ \mathbf{Z}_D(k,\underset{\sim}{s}) + \mathcal{E}_A\ (i,k,\underset{\sim}{s}). \qquad (7.1)$$

Here $\mathbf{Z}_D(k,\underset{\sim}{s})= [\bar{C}_{air}(k,\underset{\sim}{s}), A_{wet}(k,\ \underset{\sim}{s})]'$, where

$\bar{C}_{air}(k,\underset{\sim}{s})= (\bar{C}_{air}(k,s_1),\ldots,\ \bar{C}_{air}\ (k,s_n))$, $A_{wet}(k,\underset{\sim}{s})= (A_{wet}(i,k,s_1),\ldots,\ A_{wet}(i,k,s_n))$ and

$\mathbf{F}(i,k,\underset{\sim}{s})= (\underset{\sim}{u}(i,k),\ \underset{\sim}{f}\ (i,k,\underset{\sim}{s}))$

with $\underset{\sim}{u}(i,k)= (u(i,k),\ldots, u(i,k))'$, $\underset{\sim}{f}\ (i,k,\underset{\sim}{s})= f(i,k,s_1),\ldots, f(i,k,s_n))$, and

for any $i, k, \mathcal{E}_A(i,k,\underset{\sim}{s})$ is a random vector of error terms.

Note that in practice the observable quantities in obtaining the contamination $A(i,k,s)$ at a site s are the integrated air concentrations $\bar{C}_{air}(k,s_n)$ and the wet deposition $A_{wet}(k,s)$.The factors $f(i,k,s)$ and $u(i,k)$ are in fact unknown parameters which for the purposes of the modelling above are assumed to be fixed. Thus the model (7.1) is a stochastic linear approximation to the 'true' relation which has a complex non-linear form. More specifically, assuming that for a given isotope k, $\bar{C}_{air}(k,\tilde{s})$ and $A_{wet}(k,\tilde{s})$ are independent, it is straightforward to obtain this joint density given the air concentration and wet deposition factors. When gamma dose rate measurements arrive from the monitoring stations, this density is updated in a standard way using Bayes' theorem. The above regression model serves as a basis for a hierarchical model as above which is used to update the joint density of wet and dry depositions at the fixed set of sites in question. If interest lies on the posterior density of concentrations at the monitoring stations rather than the grid points, then interpolation is again needed.

Although it is conceptually possible to consider other more realistic non-linear models, the model in (7.1) simplifies the underlying mathematical machinery used for the updating and appears to be sufficient for our purposes. We stress however that the investigation of Bayesian modelling for the FDM system is still at an initial stage and we aim to explore this further in the future. For example, the assumption that air concentrations and the wet deposition factor are independent for a given set of sites may be dropped, so that the covariances of these vectors are also modelled. This should give a more realistic formulation of the problem without resulting in a too complicated mathematical model. Further, other quantities of interest for the Food and Dose Module, as outlined in the beginning of the section, such as concentrations on animal products or on processed foodstuff may be modelled in an analogous way.

Finally one of the main tasks for the future is to accommodate the above modelling in the Bayes_RIMPUFF code. This would produce a much more effective way of handling various uncertainties by a single ASY module and would facilitate the role of the CSY modules both in the threat and in the early phases of an accident.

5.2 CSY subsystem

The CSY subsystem at this stage may change to one appropriate for food, agricultural and other countermeasures, although there may still be a need for early countermeasures such as sheltering and evacuation. We will ignore the latter possibility here since it has been covered in section 4.5. We note however on passing that the incorporation of deposition modelling in the output of RIMPUFF, as indicated above, could enhance the

accuracy of simulated countermeasures, as collective doses can be assigned better estimates.

For food countermeasures, French (1996) includes an example-based analysis considering simulation of countermeasures with respect to three agricultural products: meat, milk, and vegetables.

The use of multi-attribute value theory is proposed and the relative merits of each action are determined on the basis of the attributes cost, collective averted dose and public acceptability. Using these (or any other) criteria for cost-benefit evaluation, the CSY modules should be able to produce a table of summary statistics for the cost, averted dose etc. in each area considered. From the analysis of French (1996), it emerges that often additional factors may be considered, such as the distribution of the population in each area according to age groups. This requires careful consideration for each of the attributes in question, since it may have a substantial impact on the analysis; whatever action for the milk is decided for example (e.g. decontamination, transformation to butter), it is unlikely to have the same effect on infants and adults. Thus the CSY analysis might need to accommodate the population distribution as a factor and output the results for each age group separately.

A CSY module such as LCMT can be modified to accept as input the uncertainty estimates from the ASY subsystems. The modelling in LCMT involves consideration of countermeasures such as relocation (permanent or temporary), decontamination and actions for agricultural products, e.g. food ban, processing, or changes of the crop varieties.

5.3 ESY subsystem

It has been argued in various meetings of RODOS subgroups that the evaluation of possible food and agriculture countermeasures is a non-problem for RODOS. The contention is that national and international guidance in the form of intervention levels effectively predetermines the decision. The guidance is typically given in terms of upper and lower levels such that the countermeasure is not advised if the contamination is below the lower limit and expected if it exceeds the upper level. Between the two levels the decision is left to the emergency managers. There are at least two reasons, however, why RODOS still has a problem to solve.

- Firstly, there can often e more than one countermeasure possible which can avert the dose to the population: e.g. the dose from livestock might be averted by either slaughtering the animals or by removing them from the affected region and allowing them to consume uncontaminated food for some months. There are, therefore, choices to be made between such alternative countermeasures.
- Secondly, the inevitable existence of uncertainty about key variables will cloud the issues. In most cases no forecast can state with certainty that an intervention level will be exceeded at a given region by a given

time. Normally the forecasting system will only be able to give approximate probabilities that the contamination at a certain time will be beyond/below the upper/lower level and for the event that the contamination will lie between these levels.

Consequently the emergency managers will have an explicit decision to make. This decision will be based on the information they receive when they communicate interactively with the ESY modules. It has also been discussed there that both the conceptual architecture and the software modules of the ESY subsystem will be the same whenever it is called. Thus we envisage that the output of LCMT will be screened progressively by a coarse expert system, a MAV/MAUT module and a fine expert system. In the case of food and agricultural countermeasures the constraints used to discard strategies in the CES filtering process are likely to be more complex than for the initial countermeasures for sheltering and evacuation. The development of a coarse expert system which would be designed with this more complex grid of strategies is currently under consideration and one of our main priorities for the near future. As noted in Section 4.6, modern constraint management techniques developed within the artificial intelligence and operational research communities, are well able to deal with such complexity.

HERESY or M-Crit can then be used to rank the remaining countermeasure strategies according to multi-attribute expected value theory. One of the disadvantages of this approach, however, as mentioned earlier is that it fails to capture the uncertainty about the attributes. We outline here a mechanism which would incorporate the uncertain character of attributes and appears to be better suited for practical needs. We note that the formulation of the problem given below reflects only our preliminary experimentation with this approach, as exhibited in French and Papamichail (1997),which needs to be refined further to cope with more complicated situations.

Since utility functions are used in decision making to express the DM's preferences, multi-attribute utility theory seems to be a plausible machinery to encompass the uncertainty involved with the attributes. We aim to implement a multi-attribute utility functional form into HERESY to enhance its flexibility for ranking alternatives.

More precisely, the linear multi-attribute value function which is currently used in HERESY ranks a strategy according to its 'score':

$$v\left(x_1, x_2, \ldots, x_q\right) = \sum_{i=1}^{q} w_i x_i$$

where x_i is the score of the strategy on the i^{th} attribute and w_i is the weighting factor of the i^{th} attribute. A multi-attribute *utility* function, $u(x_1, x_2, \ldots, x_q)$, would rank the strategy according to the expected value:

$$E_{X_1,X_2\ldots,X_q}\left(u\left(X_1,X_2,\ldots,X_q\right)\right)=\int u\left(x_1,x_2,\ldots,x_q\right)dx_1dx_2\ldots dx_q$$

where the expectation is taken with respect to the joint distribution of $X_1, X_2,\ldots, X_q$ as reported by the ASY and CSY modules. For illustration purposes, we may assume that this distribution is a multivariate normal distribution. This makes calculations easy and is also in accordance with other modules of RODOS, where normality is again a crucial assumption.

Moreover, for the utility function u, we choose a form which incorporates a single parameter reflecting risk aversion[1] in addition to the weights within the linear value function as already coded into HERESY. The simplest such functional form for u is the exponential transform of $v(x_1, x_2,\ldots, x_q)$ which implies constant risk aversion for the DM (Keeney and Raiffa, 1976, p167). Namely, the utility function is associated with the multi-attribute value function as follows:

$$\begin{aligned}u\left(x_1,x_2,\ldots,x_q\right)&=1-e^{-v\left(x_1,x_2,\ldots,x_q\right)/\rho}\\&=1-e^{-\sum_{i=1}^{q}w_ix_i/\rho}\end{aligned}$$

This utility function satisfies mutual utility independence between all subsets of attributes (see, for example French, 1986, Section 5.6 for various forms of 'utility independencies and their relation). It is thus equivalent to the multiplicative multi-attribute utility function studied, *inter alia*, by Keeney and Raiffa (1976, section 6.3) and used extensively in applications.

Assuming this exponential form for u, its expectation is very easy to calculate with respect to the normal distribution. Essentially one notes that the expectation has the same functional form as the normal moment generating function. If $(X_1, X_2,\ldots, X_q)^{\mathrm{T}} \sim \mathrm{N}(\boldsymbol{\mu}, \mathrm{V})$, i.e. are normal with mean vector $\boldsymbol{\mu}$ and covariance matrix V, then:

$$E_{X_1,X_2\ldots,X_q}\left(u\left(X_1,X_2,\ldots,X_q\right)\right)=1-e^{-\left(\sum w_i\mu_i/\rho-0.5\mathbf{w}^{\mathrm{T}}\mathrm{V}\mathbf{w}/\rho^2\right)}$$

where **w** is the vector formed from the weights w_i, see for example Mardia *et al* (1979).

The above (monotonic) form of the expected utility provides a straightforward criterion to rank alternatives: strategies should be ranked according to:

$$\sum_{i=1}^{q}w_i\mu_i-0.5\mathbf{w}^{\mathrm{T}}\mathrm{V}\mathbf{w}/\rho$$

In other words, strategies will be ranked by a linear function with weights w_i applied to the mean score of the strategy on the i^{th} attribute with a term

[1] For the concept of risk aversion in utility theory, see for example Chapter 4 in Keeney and Raiffa (1976).

subtracted from the total score that depends upon the covariance matrix associated with the strategy. The more uncertain one is of the scores derived from the strategy the larger the term subtracted. We also note that under the above simplifying assumptions for the form of u and the joint distribution of $(X_1, X_2, \ldots, X_q)$, there are well documented ways of assessing ρ (Keeney and Raiffa, 1976, Chapter 4).

We will develop HERESY to incorporate this multi-attribute utility ranking of alternative strategies. This will require information relevant to the covariance matrix V being passed across in the CSY/ESY interface. The means μ correspond to the data already being passed across. The risk attitude parameter *(* may be subjected to sensitivity analysis along with the w_i HERESY performs a detailed sensitivity analysis by allowing the user to vary the w_i and observe the change in ranking.

Various directions for generalising the above example might be considered. For instance, if one assumes the same exponential form for the utility function u in terms of v, the condition about the normality of the variables may be dropped; in order to calculate expected utilities, one still needs to compute the moment generating function of the corresponding multivariate distribution and this can be easily done for a rather wide family of distributions. (The resulting interpretation for ranking strategies, however, might not be as clear as above for the DMs unless the expected utility function is again monotonic). If on the other hand, one retains normality for the variables, other forms for the utility function can be considered that might in appropriate circumstances represent more accurately the needs of the DMs. The success and the simplicity of this more general formulation depends to a large extent on the amenability of the particular functional form of u for mathematical manipulation and the subsequent interpretation by the DMs.

6 Bayes_MATCH

In this section we will consider a long-range atmospheric dispersion model. Whereas for the short and middle range we described a (Lagrangian) puff-model, for the long range a Eulerian model seems more appropriate. One such model is MATCH (Mesoscale Transport and Chemistry Modelling System). MATCH, like many other Eulerian models, uses a transport diffusion equation to describe the contaminated plume, which is thought to be evolving through a fixed spatial grid. In contrast to what we said earlier about the complete absence of reliable data available for the validation of a short-range dispersion model, we can make a more positive statement concerning long-range model validation. Two tracer experiments on the full European scale (ETEX) have been conducted, and the results of these can be used to test the different long-range ADMs. Since MATCH has scored particularly well in these comparative studies (see, e.g. Tveten and Mikkelsen, 1995), we choose this model and construct a Bayesian version of it in order to be able to assimilate measurements in a way consistent with the rest of the system. Politis and Robertson (1999) have developed this version and this section will heavily draw upon their paper. In fact data assimilation has already been applied to MATCH before, using a variational approach. We will briefly explain this method in section 6.3.

Perhaps it is interesting first to say something about the physical model itself. The MATCH model solves the advection/diffusion equation for atmospheric tracers in a three-dimensional, Eulerian framework, which in continuous form is given by:

$$\frac{\partial \theta_i}{\partial t} = -\nabla(v\theta_i) + \nabla(K\nabla\theta_i) + q_i$$

In this equation θ_i represents the mass mixing ratio of the trace species of interest, v is the three-dimensional wind, K is the turbulent diffusion tensor and q_i represents internal sinks and sources. The approximate solution to the above equation is derived by discretisation and selection of various numerical techniques (Robertson *et al*, 1998). Atmospheric weather data are taken from an external source ('off-line'), and fed into the model at regular time intervals, usually every three to six hours. These data will have to be interpolated in time to get the hourly data we use as our timestep (normally). The model design is very flexible as far as grid-resolution is concerned. This resolution should depend on the weather conditions. The horizontal resolution is normally about 55 km and a non-uniform distribution of vertical layers has been used.

6.1 The Bayesian model

Modelling the evolution of atmospheric dispersion involves dealing with a large number of uncertainties both about the parameters that the physical

model operates on, as well as about the various external factors that influence the performance of the model. The main sources of uncertainty were already mentioned in section 4.1, but are repeated here in summarised form as: (i) the profile of the source emissions over time; (ii) the measurement errors; (iii) uncertainty surrounding meteorological forecasts; (iv) errors inherent in ADM itself. Here we would like to add two more uncertainties to these. First there are certain sub-grid scale phenomena that can not be resolved by the model. Especially in unstable meteorological conditions and when dealing with inhomogeneous topography, stochastic processes such as turbulence become a key feature. As a further uncertainty we mention interpolation errors, both in time and space. The source of time interpolation errors was already discussed above and space interpolation errors originate from the fact that MATCH operates and provides predictions for contamination on a large scale grid, whereas measurements are collected at off-grid points. Considering the three-dimensional version of the model in particular we have the additional problem of matching the contamination dispersion vertically. MATCH uses quite complicated algorithms for this, note however that, since in most cases all our measurements will be taken at (or close to the) ground level, it will be very hard to update this vertical distribution. One last feature which needs to be taken account of is the very sparse availability of data after a nuclear accident. Moreover, not only will there be very few data available in the first few hours of the accident, but these data may also arrive a lot later than when they were observed.

We have chosen the DLM approach here (just as we did in section 4) to describe the temporal evolution of the spatial field of contamination. A particular feature that is automatically captured within such a model is that it is equally easy to 'backcast' as it is to forecast. This fulfils the desire within the RODOS project not only to forecast the contamination profile forwards in time, but also to reconstruct the (time-varying) source emissions over time. In view of this, we include a discretised version of the emission profile as one of the state parameters in the DLM. We note however that, although we will show below that our model can indeed give an estimate of the source term, it is questionable whether this should really be the task of a long-range dispersion model, especially since there are already quite a few modules in RODOS dedicated to this task. Nevertheless, the formulas below show that it is at least *possible* to do this.

To write down the equations of the DLM, we first need to introduce some notation. Let $\mathbf{s} = (s_1, s_2, ..., s_n)$ be the vector of grid points and $\mathbf{r} = (r_1, r_2, ..., r_m)$ the vector of detector points. We assume that at time $t = 0$ the (true) contamination at the grid cells is described by an initial state $\theta_0(\mathbf{s})$, while $\theta_t(\mathbf{s})$ is the contamination profile at time $t = 1,2,.....$ We can now express our DLM equations as:

$$\theta_t(\mathbf{s}) = G_t \theta_{t-1}(\mathbf{s}) + q_t(\mathbf{s}) + w_t(\mathbf{s}) \qquad (1)$$

$$Y_t(\mathbf{r}) = F\theta_t(\mathbf{s}) + v_t(\mathbf{r}) \quad (2)$$

The meaning of all the other variables: (i) $Y_t(\mathbf{r})$ is the observed contamination at the detector points; (ii) G_t is the system matrix representing advection and diffusion processes (often time-varying); (iii) F is the interpolation matrix, which is assumed constant here, although under very unstable meteorological conditions this might be not a very realistic assumption; (iv) $q_t(\mathbf{s})$ is the source emission at time t; (v) $w_t(\mathbf{s})$ is the error vector that accounts for errors in the physical model itself and errors in the emission profile (and, in principle also time interpolation errors); (vi) $v_t(\mathbf{r})$ accounts for spatial interpolation errors and measurement errors. If there is not a direct correspondence between the time resolution of the model and the measurements we have to use an average θ_t^* instead of θ_t in (2). We will assume throughout that both errors are constant and normally distributed. These assumptions are not critical though and the model can be extended to include more general conditions. The errors can thus be written as $w_t(\mathbf{s}) \sim N(0,W)$ and $v_t(\mathbf{r}) \sim N(0,V)$ and we assume both V and W known. Keeping in mind the size of the grid is the model covers, it is unlikely that two measurement stations will be (too) close. The matrix V is therefore assumed diagonal; we do however assume correlation between the model error terms.

Politis and Robertson (1999) argue that expert judgement suggests that a temporal resolution of one hour can adequately model the continuous emissions from the source and we have used this as the time step in the above DLM. The resolution of the emissions (and other processes) as treated by MATCH has a much smaller temporal resolution (typically 5 minutes), so that the parameters q_t in the model are essentially hourly averages of the parameters in this more refined discretisation scheme of the physical model. Therefore we can take these emissions that are not very close in time to be independent of each other, which simplifies the formulas below to a great extent. Further, we also assume that the contamination at the grid at time t, $\theta_t(\mathbf{s})$, is independent of future emissions $q_\tau(\mathbf{s})$, for $\tau > t$.

Our prior knowledge of the emission profile and the initial state of the system at time $t = 0$ are expressed by

$$q_t(\mathbf{s}) \sim N(m_t, C_t)$$

$$\theta_0(\mathbf{s}) \sim N(p_0, C_0)$$

By iterating (1) and taking the expectation and variance of (2) respectively we obtain the following expressions for the moments of the vector $Y_t(\mathbf{r})$:

$$E(Y_t(\mathbf{r})) = FG_tG_{t-1}\ldots G_1p_0 + FG_tG_{t-1}\ldots G_2m_1 + \ldots\ldots + FG_tm_{t-1} + Fm_t$$

$$Var(Y_t(\mathbf{r})) = (FG_tG_{t-1}...G_1)C_0(FG_tG_{t-1}...G_1)^T + \\ +(FG_tG_{t-1}...G_2)C_1(FG_tG_{t-1}...G_2)^T + \\ +(FG_tG_{t-1}...G_2)W(FG_tG_{t-1}...G_2)^T + \\ +.....+FC_tF^T + FWF^T + V$$

Denote the vector and matrix on the righthandside of these two expressions respectively as M_t and $\sum_t$. Finally, for the covariance matrix between $Y_t(\mathbf{r})$ and the emissions we find (because of the before mentioned independency assumptions) that for any $\tau = 1,2,....,t$

$$Cov(q_\tau(\mathbf{s}),Y_t(\mathbf{r})) = (FG_tG_{t-1}.....G_{\tau+1}C)^T = A_\tau$$

while for all $\tau > t$, $Cov(q_\tau(\mathbf{s}),Y_t(\mathbf{r})) = 0$, since measurements at time t do not give us any information about subsequent emissions.

Using standard normal theory (see e.g. the appendix of West and Harrison (1989) for an abstract of the general formulas, and Mardia et al (1979) for the proof of all the theorems), we can combine the above results to obtain a multi-normal distribution. We will not write down this transitional step, but instead (again using standard normal theory) directly give the posterior mean and (co)variance of the emissions at time τ for $\tau = 1,2,.....,t$. Before we do this however, we want to stress once more that if the distributions are not normal, the derivation of these formulas is a lot more awkward, but the principles remain valid. The posterior moments of q_τ are:

$$E(q_\tau(\mathbf{s})|Y_t(r_k) = y_t(r_k)) = m_\tau + A_\tau^k\sigma_k^{-1}(y_t(r_k) - M_{t,k}) := m_\tau'$$

$$Var(q_\tau(\mathbf{s})|Y_t(r_k) = y_t(r_k)) = C_\tau - A_\tau^k\sigma_k^{-1}(A_\tau^k)^T := C_\tau'$$

where σ_k is the k^{th} diagonal element of $\sum_t$ and A_τ^k is the k^{th} column of the matrix A_τ. We have constructed an iterative procedure now, since we can use the same algorithm again if we replace m_τ and C_τ by m_τ' and C_τ' respectively when new measurements arrive. If we would like to obtain the posterior distribution of θ_t (instead of focusing on the source term, which has been done above), then we obtain formulas completely analogous to the above: we just have to replace m_τ, C_τ and A_τ^k by p_0, C_0 and A_0^k respectively to get the posterior moments for θ_0 . From there it is straightforward to obtain forecasts for some *t_ahead* for both θ_{t_ahead} and Y_{t_ahead} (see e.g. West and Harrison (1989) for the appropriate formulae).

Politis (1999) argues that for an ADM always the main emphasis should be on accurate and quick updating of contamination forecasts and for this matter it seems logical that a long-range ADM such as MATCH takes its inputs from a short-range model such as Bayes_RIMPUFF. This however introduces its own class of problems. The paper by Brandt et al (1996) discusses the advantage of combining a short- and long-range ADM and also the problems associated with this combining procedure. Brandt argues that an Eulerian model can not provide good predictions for the short

range. We note that MATCH can in principle bypass this problem itself by means of its own particle initialisation procedure, but despite the problems associated with it, it would seem better to use the outcomes of a separate short-range dispersion model (which has already updated predictions) to form the initial state of a long-range ADM. For the modelling in this case we refer to Politis (1998).

6.2 Computational aspects

Even with all the simplifying assumptions made for our DLM, the model will always need a long computer calculation time, simply because there are so many variables involved in the problem. To get an idea of this problem we note that the number of grid points in MATCH is typically of order 10^6, which raises the serious concern for the computer time needed at each time step involving the covariance matrices of the model. Fortunately enough, in many (simple) cases all the matrices involved in the model are sparse, which makes considerable savings possible.

In most cases we can assume that the initial emission occurs from a single grid point in the horizontal plane. Thus only a few elements of m_t are non-zero normally. Moreover, if there is a substantial degree of certainty about the location of emissions, we only have to update grid cells with non-zero prior mean and variance. A problem that remains though, is to describe the vertical distribution of the emissions in an appropriate way. We already said that MATCH deals with this problem by applying complicated algorithms and we also said that it will be very hard to update this distribution. We would like to add here though that we can put all expert knowledge we have in the prior of the distribution. Moreover, we inflate the variance of the grid points relative to their height level. But, as long as we do not take measurements at different heights this difficulty can not be solved easily. However, as a contrast to this we note that, if there is uncertainty about the emission profile in the *horizontal* plane, then our model addresses this uncertainty directly and can in fact detect prior misspecifications.

Another problem which we have to address seems less obvious, but is in fact even more problematic than the previous one. When we apply the updating procedures for our DLM, the matrices F, G_t, V and W are all assumed known at every point. And in principle it should be possible to read the matrix G_t somewhere from the MATCH model code, given the contamination at two successive time points. However, although most physical processes are linear, MATCH does not itself calculate the weights $g_{ij,t}$. Building in such a routine in the program would not only require a total restructuring of the program, but would also slow it down dramatically. Therefore we have to find another solution to this problem. Politis and Robertson (1999) suggest an approximation by means of the adjoint technique. By applying this technique, we can update directly the vectors we are interested in (rows of $FG_t....G_i$ for all i), which typically

contain many zero entries, instead of calculating all the elements of the matrices G_t explicitly. Because we use the variational method here and because it offers an alternative as a completely different method of updating, we will go a bit deeper into the principles of this method.

6.3 The variational method and some possible extensions of the model

The variational technique for data assimilation is quite popular among physicists. We will try to explain here why this is the case, and why these arguments do not apply to the purposes we want to use our ADM for. Talagrand and Courtier (1987) summarise the history of the development of the method and also give a very general description of the method itself as it could be applied to a meteorological model. They describe the method as an improvement over the Kalman filter method, attacking as the main flaw of this method the demand that the evolution of the forecast error has to be known beforehand. The assignment of these values may indeed not be easy, but is in fact a natural part of the way the uncertainty of the model itself is dealt with. The variational method actually excludes modelling error. Whereas this may not be a problem when we are just interested in best estimates (and not the uncertainty of these estimates), in our application the propagation of uncertainty is absolutely essential! However, there is one feature of the variational method that is very interesting and that is its general applicability. Whereas the Kalman filter only works for a *linear* model, this is not a necessary requirement for the variational method to work. We do note though, that the linearity of our model seems quite justifiable (we only assume that the contamination at a point at a certain timestep is a linear combination of the contamination at all the grid points at the previous time step).

The variational problem can be stated in the following terms: find the initial state such that the corresponding model solution minimises the scalar function measuring the distance to the observations. As such a scalar function (usually called penalty function) we choose here (observation at time $t = t'$):

$$J = (F\theta_{t'}(\mathbf{s}) - Y_{t'}(\mathbf{r}))^T V^{-1} (F\theta_{t'}(\mathbf{s}) - Y_{t'}(\mathbf{r}))$$

After making a first guess for the initial state we are interested in (e.g. the source term), we can use the gradient of the above penalty function with respect to the initial state to update this guess iteratively in the right direction (normally by using a quasi-Newton iteration method). In this case we obtain:

$$q_t^{n+1}(\mathbf{s}) = q_t^n(\mathbf{s}) - \rho \nabla_{q_t^n} J$$

where n is the iteration number and ρ is the approximation of the Hessian matrix (see Gilbert and Maréchal, 1989). It can now be shown that by using adjoint equations (see Talagrand and Courtier(1987) and Robertson

and Persson(1993) for details (the latter handles the specific application to MATCH)) we can obtain the following expression for the gradient of *J*:

$$\nabla_{q_t^n} J = G_t^T G_{t+1}^T .... G_{t'}^T F^T (F\theta_{t'}^n(\mathbf{s}) - Y_{t'}(\mathbf{r}))$$

As we compare the two equations above to the ones earlier found for the Kalman filter we do in fact see quite some similarity, the difference between the methods just being the exclusion of the modelling error in the case of the variational method.

Finally we would like to suggest a few extensions that might be made to the model we described here. First of all, just as we did for the Bayes_RIMPUFF model we can let the observations follow a Poisson distribution (instead of a normal distribution). Just as we did there we could use dynamic belief networks and maybe Monte Carlo Markov Chain (MCMC) methods to update the distributions in this case. As a second possible improvement we mention covariance learning procedures (see e.g. Le and Zidek (1992, 1994)) that could be incorporated in the model. Probably this will slow down the calculations too much, but this will have to be investigated. Another logical extension seems to be the assignment of a prior distribution to important external parameters in the model such as the wind direction, which would be able to detect misspecifications. We successfully applied a 3-way distribution to the height and wind direction in the Bayes_RIMPUFF model and there is no reason to assume why it would not work here. It at least seems to deserve further investigation. Finally, as a last possible extension we mention the split-up of measurement and interpolation errors. To do this we just need to introduce an extra parameter (and an extra equation). We actually made this extension for the modelling of the food chain contamination, which we will discuss in the next section. The arguments applied there could easily be transferred to this section.

7 The Bayes_DEMM

In this section we will describe the Bayes_DEMM (DEposition Monitoring Module), a space-time Bayesian method proposed for the statistical modelling of radioactivity deposition in the food chain contamination of the RODOS system. It is supposed to integrate the outputs of an ADM with inputs of the deposition module of the food chain model based on ECOSYS-87, also incorporating assimilation of different kinds of monitoring data. It handles uncertainties associated with the physical model in a natural and structured way, consistent with the Bayesian methodology throughout the project. A complicating factor in the modelling of deposition is the fact that the radioactivity produced depends very much on the deposited *type* of radio-nuclide. The Bayes_DEMM should therefore be able to handle multiple nuclides and also distinguish between wet and dry deposition because of the enormous influence of rain on the deposition process. It is mainly because of these features that some new ideas will be introduced in this section. We will follow largely the set-up of the paper by Faria and Smith (1999), apart from a few slight emendations (renaming the variables amongst other things).

7.1 The deposition module of ECOSYS

We will first give some characteristics of the physical model we will have to deal with: the deposition module of ECOSYS, which is used as the initiating module of the food chain contamination. In its deposition module, the ECOSYS calculates for each location l of interest the total radioactivity $A(i,j,l)$ produced by the deposited isotope j which is intercepted by the plant type i. The calculations are performed by decomposing $A(i,j,l)$ into wet and dry components.

$$A(i,j,l) = A_{dry}(i,j,l) + A_{wet}(i,j,l)$$

where

$$A_{wet}(i,j,l) = f_w(i,j,l)A_{wet}(j,l)$$

In this last equation $f_w(i,j,l)$ stands for the interception fraction of isotope j for plants type i. The dry component depends on the air concentration (C_{air}), whereas the wet component depends on both air concentration and the amount of rainfall ($R(l)$). The dependence of these wet and dry components on the air concentration and rainfall is shown in Figure 12 and can be represented by physical laws (and approximations) which are not of much interest to us here, but can be found in Faria and Smith (1999). Using these formulas it is possible to write $A(i,j,l)$ as the following function depending (non-linearly) on C_{air} and $R(l)$:

$$A(i,j,l) = (v_g(i,j) + b(i,j)R^*(i,j,l))C_{air}(j,l)$$

In this equation $v_g(i,j)$ is the so-called deposition velocity, $b(i,j)$ is a complicated (but known) function and R^* is a function, depending

(amongst other things) on the amount of rainfall $R(l)$. If we assume the in principle unknown parameters $v_g(i,j)$ and $f_w(i,j,l)$ to be constant and fixed in this physical model, we can achieve a reasonable linear approximation of the above equation, which is suitable for the decision support purposes of our system.

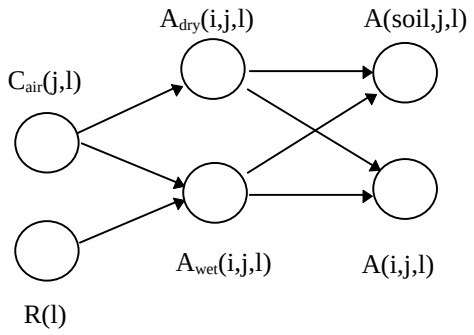

Figure 12: the deposition module of ECOSYS

7.2 Definitions and transformations

Next we introduce some notation of variables that we are going to use here. We will discretize the geographical region by laying over it a regular grid of size n_1 whose points $s_i = (x_i,y_i)$ represent the sites for which the ADM will generate predictions. Off-grid points $r_j = (x_j,y_j)$ (size n_2) will represent the monitoring station network (size m). Unmonitored sites of interest, like e.g. cities, should also be incorporated in this latter vector. We will use $l_a = (x,y)$ as a general representation for a location.

The model as we have constructed it has three main components:

1. total ground activity depositions
2. radio-nuclide specific proportions of the total ground depositions
3. wet deposited radioactivity fractions of the total ground depositions

Associated with these components are the data coming from measurement sites:

- instantaneous air concentrations (c_{air}) and gamma-dose rates (GDR) which are informative about the time (and nuclide) aggregated total depositions
- gamma-spectrometry (γ-Spec), informative about the proportions of different radio-nuclides of the total depositions
- rainfall intensity r, giving information about the proportion of wet deposition relative to the total.

Finally, before we introduce the variables, we note that we have used the subscript P when the variables are associated with the ADM predictions.

These are basically Pseudo-observations (which is what the letter *P* stands for), since we essentially treat them as observations in our model, whereas *M* is the subscript (and stands) for the Measuring stations. The subscript *e* is used to denote either of them. We will not give the definition of all the variables explicitly here, since they are all defined in similar ways and on the basis of the above it should hopefully be clear what they stand for. For the full definitions of the variables used we refer again to Faria and Smith (1999).

$\theta^*_{e,\ tot\_m} = (\theta^*_{e,\ tot\_m}(l_1),.....,\theta^*_{e,\ tot\_m}(l_v))^{'}$, where $\theta^*_{e,\ tot\_m}(l_a)$ are either the ADM's pseudo-observations ($e = P$, $l = s$, $v = n_1$ and $\dim(\theta^*_{P,\ tot\_m}) = n_1$ x 1 or the measured ($e = M$, $l = r$, $v = n_2$ and $\dim(\theta^*_{M,\ tot\_m}) = n_2$ x 1 total radioactivity deposited at l_a (a = 1,....v). The small *m* stand here for (pseudo-)Measured.

Further $\theta^*_{e,\ tot\_r}$ stands for the Real total deposition to which the (pseudo)-observations $\theta^*_{e,\ tot\_m}$ correspond. In similar ways we can now define the variables $\theta^*_{e,\ nuc\_m}$ ((pseudo-)Measured nuclide specific proportions), $\theta^*_{e,\ nuc\_r}$ (Real nuclide specific proportions), $\theta^*_{e,\ wet\_m}$ ((pseudo-)Measured wet deposition fractions) and $\theta^*_{e,\ wet\_r}$ (Real wet deposition fraction).

If we want to apply simple linear algorithms, which we do, then it would be quite convenient if the variables we defined are (approximately) normally distributed. However, it is quite obvious that this is not the case. The contamination values will not be strictly Gaussian, since we can not have negative values of contamination, and the proportion values which all range between 0 and 1 are even more obviously non-Gaussian. To force the variables of our model to have an approximately normal distribution, we have to transform the above defined variables. We note beforehand that we will only transform non-zero variables, leaving the zeros as they are.

For the total deposition variables $\theta^*_{e,\ tot\_m}$ and $\theta^*_{e,\ tot\_r}$ the transformation is quite simple: to take the natural logarithm of these variable at each location seems appropriate. We define thus:

$$\theta_{e,tot\_m} = \ln(\theta^*_{e,tot\_m})$$

$$\theta_{e,tot\_r} = \ln(\theta^*_{e,tot\_r})$$

For the wet deposition ratios, which ranged originally all between 0 and 1, the following transformation is used at each location:

$$\theta_{e,wet\_m} = \ln(\frac{\theta^*_{e,wet\_m}}{1-\theta^*_{e,wet\_m}})$$

$$\theta_{e,wet\_r} = \ln(\frac{\theta^*_{e,wet\_r}}{1-\theta^*_{e,wet\_r}})$$

An extra complication arises with the transformation of the nuclide-specific proportions. The problem is that all these proportions must add to

unity at each location. To solve this problem a so-called *neutral to the right* (NTTR) process (Walker and Muliere, 1998) is used. We first note that the dimensions for $\theta^*_{e,\ nuc\_m}$ and $\theta^*_{e,\ nuc\_r}$ defined above are n_1k x 1 or n_2k x 1, where k is the total number of isotopes we distinguish in our model, so if we still keep the location vector intact, we can split these variables up in ($\theta^*_{e,\ nuc\_m}(1)$, $\theta^*_{e,\ nuc\_m}(2)$,....., $\theta^*_{e,\ nuc\_m}(k)$) and ($\theta^*_{e,\ nuc\_r}(1)$, $\theta^*_{e,\ nuc\_r}(2)$,....., $\theta^*_{e,\ nuc\_r}(k)$) respectively. We now define

$$\theta^{**}_{e,nuc\_m}(1) = \theta^{*}_{e,nuc\_m}(1)$$

$$\theta^{**}_{e,nuc\_m}(j) = \frac{\theta^{*}_{e,nuc\_m}(i)}{1-\sum_{i=1}^{j-1}\theta^{*}_{e,nuc\_m}(i)} \qquad (j = 2,.....,k\text{-}1)$$

This transformation guarantees the independence of the $\theta^{**}_{e,\ nuc\_m}(j)$ for all j. The variables $\theta^{**}_{e,\ nuc\_m}(j)$ can now be transformed in the same manner as we transformed the wet deposition variables. We thus define:

$$\theta_{e,nuc\_m}(j) = \ln\left(\frac{\theta^{**}_{e,nuc\_m}(j)}{1-\theta^{**}_{e,nuc\_m}(j)}\right)$$

In a similar way we perform a double transformation to find $\theta_{e,\ nuc\_r}(j)$. Note that the vectors $\theta_{e,\ nuc\_m}(j)$ and $\theta_{e,\ nuc\_r}(j)$ do not have the same distribution as their untransformed equivalents: they have dimension $n_1(k\text{-}1)$ x 1 or $n_2(k\text{-}1)$ x 1.

7.3 The basic state space equations and some parameter settings

With the information we acquired so far we can draw the influence diagram as depicted in Figure 13. Note that we omitted some of the above defined variables. This was done for reasons of clarity.

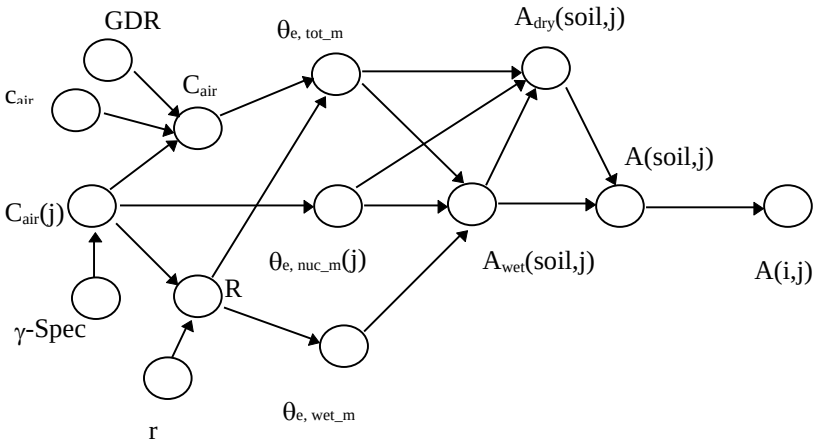

Figure 13: the Bayes-DEMM influence diagram

Two important conditional independencies can be stated from this diagram:

- $\theta_{e,\,tot\_m}$ is independent of $\theta_{e,\,nuc\_m}(j)$ and $\theta_{e,\,wet\_m}$ given C_{air} and R;
- $A(soil,j)$ (and $A(i,j)$) is (are) independent of $\theta_{e,\,tot\_m}$, $\theta_{e,\,nuc\_m}(j)$ and $\theta_{e,\,wet\_m}$ given $A_{wet}(soil,j)$ and $A_{dry}(soil,j)$.

We will now give the state-space equations for our preliminary model:

(pseudo-)Observation-equations:

$$\begin{pmatrix} \theta_{P,tot\_m} \\ \theta_{P,nuc\_m} \\ \theta_{P,wet\_m} \\ \theta_{M,tot\_m} \\ \theta_{M,nuc\_m} \\ \theta_{M,wet\_m} \end{pmatrix} = \begin{pmatrix} \theta_{P,tot\_r} \\ \theta_{P,nuc\_r} \\ \theta_{P,wet\_r} \\ \theta_{M,tot\_r} \\ \theta_{M,nuc\_r} r \\ \theta_{M,wet\_r} r \end{pmatrix} + \begin{pmatrix} \varepsilon_{P,O\_tot} \\ \varepsilon_{P,O\_nuc} \\ \varepsilon_{P,O\_wet} \\ \varepsilon_{M,O\_tot} \\ \varepsilon_{M,O\_nuc} \\ \varepsilon_{M,O\_wet} \end{pmatrix}$$

Interpolation equations:

$$\begin{pmatrix} \theta_{M,tot\_r} \\ \theta_{M,nuc\_r} \\ \theta_{M,wet\_r} \end{pmatrix} = G \begin{pmatrix} \theta_{P,tot\_r} \\ \theta_{P,nuc\_r} \\ \theta_{P,wet\_r} \end{pmatrix} + \begin{pmatrix} \varepsilon_{I\_tot} \\ \varepsilon_{I\_nuc} \\ \varepsilon_{I\_wet} \end{pmatrix}$$

In these equations $G = \mathrm{diag}(G_{tot}, G_{nuc}, G_{wet})$, where these sub-matrices are the interpolation matrices for the respective components. The ε_I-variables stand of course for the interpolation errors. All errors are considered mutually independent. Furthermore, the three components (total deposition and nuclide specific and wet fractions) are assumed mutually independent, given the total air concentration and the rainfall intensity (see also the conditional independencies stated from the above influence diagram).

We will now suggest some plausible settings or at least usual forms for the fixed matrices in the model. First of all we take a look at the matrix G. As mentioned earlier, this matrix G can be written as $G = \mathrm{diag}(G_{tot}, G_{nuc}, G_{wet})$ and we will try to give every component an appropriate interpolation sub-matrix. An obvious choice of interpolation-type *seems* to be linear interpolation. However, this appears to result in rather too smooth predictions for this application. A solution to this problem is to use exponential interpolation functions, which are very flexible in allowing the choice of strength of spatial influence. This choice corresponds to the sub-matrices of G having elements of the form

$$g(i,j) = \begin{cases} \alpha_{ij} \exp(-\beta_{ij} d_{ij}) & \text{if } i \in Ne(j) \\ 0 & \text{otherwise} \end{cases}$$

where α_{ij} is a proportionality constant to ensure that the sum of all elements of a row of G add to one, d_{ij} is the distance from (off-grid point) i to (grid point) j and $\beta_{ij} > 0$ is the rate of exponential decrease. Further, Ne(j) stands for the set of nearest adjacent neighbouring grid points.

There are good scientific reasons to assume that the nuclide specific proportions are not expected to vary as much from one site to another as the total contamination. This should be reflected in our choice of β_{ij} for the sub-matrix G_{nuc} as compared to our β_{ij}-choice for the sub-matrix G_{tot}. This parameter will be assigned a much smaller value (returning almost to linear interpolation) in the former case.

We will now investigate the error-terms in our model. They all are assumed to be normally distributed with expectation 0 and thus we can write e.g.

$$\varepsilon_{P,O\_tot} \sim (0, \Sigma_{\varepsilon_{P,O\_tot}})$$

The other errors have similar distributions. It is now useful to split these covariance matrices up in respectively:

$$\Sigma_P = diag(\Sigma_{\varepsilon_{P,O\_tot}}, \Sigma_{\varepsilon_{P,O\_nuc}}, \Sigma_{\varepsilon_{P,O\_wet}})$$
$$\Sigma_M = diag(\Sigma_{\varepsilon_{M,O\_tot}}, \Sigma_{\varepsilon_{M,O\_nuc}}, \Sigma_{\varepsilon_{M,O\_wet}})$$
$$\Sigma_I = diag(\Sigma_{\varepsilon_{I\_tot}}, \Sigma_{\varepsilon_{I\_nuc}}, \Sigma_{\varepsilon_{I\_wet}})$$

We will now discuss possible settings for these covariance matrices. Before we do this though, we want to make sure that our settings result in *valid* covariance matrices. One possible form which guarantees this is the exponential form (Cressie, 1993) with the diagonal elements containing the variances σ_i^2 and the off-diagonal elements containing the covariances cov_{ij}, such that

$$\mathrm{cov}_{ij} = \begin{cases} \sigma_i^2 \exp(-\gamma_{ij} d_{ij}) & \text{for } j \in Ne(i) \\ 0 & \text{otherwise} \end{cases}$$

where d_{ij} is the distance again and $\gamma_{ij} > 0$ a scalar representing the strength of spatial decay in correlation. Note that the meaning of Ne(i) is slightly different from the one given above for interpolation: here we should interpret it as adjacent gridpoints of a certain gridpoint i. Moreover, within this definition of Ne(i) it is free to give your own interpretation of adjacent points: you can choose whether opposite diagonal points fall within this range or not.

For the setting of Σ_M the above is not relevant however, since we assume that the measurement errors in one site are independent of measurement errors at another site. Further, we will assume independence of interpolation errors, resulting also in diagonal sub-matrices of Σ_I. We do assume correlation between pseudo-observational errors though (leading to

a non-diagonal Σ_P). It seems unreasonable to say that the ADM's prediction errors will be site-independent. We will thus assume spatial covariance of the form discussed above, choosing appropriate values σ_i^2 and γ_{ij} for the respective total, nuclide-specific and wet components.

There is one last setting that we still have to make and that is the setting of the prior means and covariances of θ_{tot_r}, θ_{nuc_r} and θ_{wet_r}. It seems best to set relatively uninformative values for these, where the prior covariance matrix should have the form of the corresponding pseudo-observational variance.

Now that we have fixed all the necessary matrices, we can perform data assimilation in the way we have done throughout this entire document, that is by using the Kalman filter. In this way we can obtain formulas for the posterior means and variances assimilating only ADM's predictions, only measurements or both. We will not give the results of these (somewhat extensive) formulas, but the interested reader can find them in Faria and Smith (1999).

7.4 Incorporating the time factor in the model

Although the Bayes_DEMM assumes that deposition occurs only during the release phase (which simplifies things quite nicely for us), this does not imply that our deposition components do not have a certain behaviour in time. Although the wet and nuclide-specific proportions will not change much after the release phase, the total contamination will decrease, due to decay. For that reason we include an extra equation in our model, describing this behaviour with an evolution matrix H_t. Further all our previously introduced variables remain the same, the only change being that a time index is added to all of them.

Our extended model can now be written as (note that the three components have been merged into one vector):

$$\begin{pmatrix} \theta^t{}_{P,m} \\ \theta^t{}_{M,m} \end{pmatrix} = \begin{pmatrix} \theta^t{}_{P,r} \\ \theta^t{}_{M,r} \end{pmatrix} + \begin{pmatrix} \varepsilon^t{}_{P,O} \\ \varepsilon^t{}_{M,O} \end{pmatrix}$$

$$\theta^t{}_{M,r} = G\theta^t{}_{P,r} + \varepsilon^t{}_I$$

$$\theta^t{}_{P,r} = H^t\theta^{t-1}{}_{P,r} + \varepsilon^t{}_{P,D}$$

On the basis of these we can now use the temporal Kalman filter, which we already used so many times in previous sections. After having updated all the information up to *t*-1, we can obtain a posterior distribution for that time; we can then easily find a prior distribution for time *t*. We further write down an expression for the one-step ahead forecast. Combining these two (in the usual way) gives the posterior distribution for time *t*.

7.5 Translating results back to physical reality

When we started to build the model we found out that it was necessary to transform the variables to obtain approximately normally distributed variables. This is the moment to transform these variables back to give them some physical meaning again. For the total depositions this is straightforward: an inverse transformation shows directly that the original variables are log-normally distributed, a distribution with well-known attributes. For the proportion variables this inverse transformation is not so easy however. We can only make approximations of the asymptotic distributions. For expressions of these approximations see (again) Faria and Smith (1999).

In the end, what we are really interested in, is the contamination at each location. To calculate this value we decomposed the total deposition into wet and dry components. With the results of our model we can now directly see the results for these components:

$$\theta_{wet}(j,l) = (\theta^*_{e,tot\_r})(\theta^*_{e,nuc\_r})(\theta^*_{e,wet\_r})$$
$$\theta_{dry}(j,l) = (\theta^*_{e,tot\_r})(\theta^*_{e,nuc\_r})(1-(\theta^*_{e,wet\_r}))$$

The mean of these variables follows now in a straightforward manner and for the variance of these variables we use a standard formula for the variance of a product of variables which can be found in (for instance) Mardia et al (1979).

8 The plume has stopped spreading: the contamination is on the ground with little suspension

From this point on the ASY cannot use RIMPUFF or any other ECOAMORS. Intuitively we want the system to move smoothly from the predictions of the atmospheric dispersion model to averages of ground contamination measurements as more and more of these are taken over the region.

In French *et al* (1995b) we demonstrated that this can be achieved by using a three-stage Bayesian hierarchical model (see Section 8 and Appendix A there). Faria *et al* (1997), in their modelling of ground contamination, adopt a similar approach to the one which we outlined in Section 5.1. This model can also be used to provide forecasts for the deposition after the passage of the plume and in fact it can be adapted rather more easily for this case. For example, when the plume has stopped spreading, the contribution of air activities to the total gamma dose at any site is zero: with the notation of Section 5.1, $D_{tot}(r)=0$. As the number of ground contamination measurements increase with time, the weight placed upon them with the Bayesian model increases so that with sufficient data collected the weight placed upon the atmospheric dispersion model's predictions will diminish to (effectively) zero.

Note that the Bayesian hierarchical modelling of Faria *et al* (1997) is currently developing within the Food and Dose Module of RODOS as its endpoint will form the input for considering food bans and agricultural countermeasures. For the food chain modelling, within the system they are currently considered 22 feedstuffs (17 based on plants, 4 based on animal products and feeding water) and 35 foodstuffs: 17 plant products, 17 animal products and drinking water (Ehrhardt *et al* , 1997). Note that this corresponds to a huge number of countermeasures being processed by the CSY modules. Although at this stage the time pressures will have been considerably relaxed so that more detailed analyses are anticipated by a CSY module like LCMT, it is difficult to envisage that hundreds of billion of strategies can be analysed simultaneously and sufficient output such as summary statistics etc. will be produce for each of them. This would require a huge database which apart from being awkward to manipulate, would include a lot of information which is useless. In order to maintain a manageable size of data for the countermeasure strategies, it seems necessary to build an CSY/ESY interface which would function along the principles that were discussed in earlier sections. Thus a coarse expert system, which reduces the number of possible implementations to a small fraction of feasible ones, may provide a very satisfactory solution to the problem. LCMT then may only consider these strategies and provide detailed information about the cost-benefit forecasts of each strategy to the ESY. Moreover there might be sufficient time available to produce an elaborate sensitivity analysis both in the coarse expert system (an additional feature of the CES which has to be implemented) and the MAUT

module, HERESY. There is also sufficient time for the DMs to discuss with a variety of experts and seek advice on the feasibility and the potential consequences of particular strategies.

9 Decision making for the longer term

Decision making in the longer term, to include such issues as whether permanent relocation is desirable or necessary, will begin several days after the accident and involve local and probably national government. The format of such decision making is likely to be like a decision conference. French *et al* (1992) and French (1996) discuss the purposes and the general principles of such a process and reflect upon the experience from the conferences held as part of the International Chernobyl Project.

One of the primary tasks of the participants in a decision conference would be to agree on a common hierarchy of attributes. This may not only serve as a guideline for the DMs, but also it can assist to a substantial degree the analysis and the evaluation of countermeasures. In particular if multi-attribute value analysis is employed, the attribute hierarchy is essential to form the basis for the determination of the weights used in the value functions. The use of utility theory in decision conferences seems to be unnecessary and simple additive value decision models are usually sufficient (French, 1996). The current experience suggests that the use of software, including graphical displays and sensitivity analyses is useful for the participants. Thus a decision conference in RODOS might be supported by HERESY or some other standard decision analytic software such as VISA or DPL. A detailed discussion about the proposed format of a decision conference commenting upon the Chernobyl and the BER studies is given in French (1996).

10 References

M. Ahlbrecht, R. Borodin, V. Borzenko, J. Ehrhardt, S. French, S. Shershakov, A. Sohier, E. Trakhtengerts, and E. Verbruggen (1997) . Decision Support Issues in RODOS: the needs of DMs. In *Proceedings of Fourth Intern. Workshop on Dec. Supp. Systems for off-site Emergency Management*, Aronborg 6-11 October. To appear in Radiation Protection Dosimetry.

M. Aoki (1966) *Optimisaton of Stochastic Systems*. Academic Press, New York.

V. Barnett (1973) *Comparative Statistical Inference.* John Wiley and Sons, Chichester.

T. Bedford, R.M. Cooke (1997) *Mathematical tools for Probabilistic Risk Analysis.* Lecture notes, TU Delft

J.O. Berger (1985) *Statistical Decision Theory and Bayesian Analysis.* Second Edition, Springer-Verlag, New York.

J.M. Bernardo, and A. F. M. Smith (1994) *Bayesian Theory*. Wiley, Chichester.

D.A. Berry (1996) *Statistics: a Bayesian Perspective.* Duxbury Press, Belmont.

V. Borzenko, S. French (1996). Decision analytic methods in RODOS. Proceedings of Intern. *Conference on the Consequences of the Chernobyl Accident*, Minsk, Belarus 18-22 March 1996, Report EUR 16544 (EC, Office of Publications, Luxembourg), pp.1149-1158.

V. Borzenko, E. Trakhtengerts, R. Borodin, and V. M. Shershakov (1994) m-crit: A Multi-Criteria Decision Support Subsystem for rodos Prototype 2. Scientific Production Association "Typhoon", Emergency Centre, Obninsk, Russia.

G. E. P. Box, and G. C. Tiao (1973) *Bayesian Inference in Statistical Analysis*. Addison Wesley, Massachusetts, London.

J. Brandt, T. Mikkelsen, S. Thykier-Nielsen, Z. Zlatev (1996) *Using a Combination of Two Models in Tracer Simulations*. Mathl. Comput. Modelling Vol. 23, No. 10, pp 99-115.

J. Brown, K. P. Smith, P. Mansfield, and J. Smith (1996) Models for decontamination, relocation, and agricultural countermeasures. *Proceedings Fourth International Workshop on Real-Time Computing of the Environmental Consequences of an Accidental Release from a Nuclear Installation*, Aronsborg 7-11 October 1996.

J. Brown, L. H. J. Goosens, B. Kraan, R. M. Cooke, J. A. Jones, F. T. Harper, F. E. Haskin, M. L. Abott, M. L. Young, S. C. Hora, and A. Rood (1997) *Probabilistic Accident Consequence Uncertainty Analysis: Food Chain Uncertainty Assesment*. NUREG/CR-6523, Publication no. EUR 16771 EN of the Comision of the European Communities, Brussels, Luxemburg.

R. M. Cooke (1991) *Experts in Uncertainty*: *Opinion and Subjective Probability in Science*. Oxford University Press.

N. A. C. Cressie (1993) *Statistics for Spatial Data* (*revised edition*). John Wiley and Sons, New York.

A. P. Dawid (1992) *Applications of a general propagation algorithm for probabilistic expert systems*. Statistics and Computing, **2**, 25-36.

M. H. De Groot (1970) *Optimal Statistical Decisions*. Mc-Graw Hill, New York.

J. Ehrhardt, J. O. Schule, G. Benz, M. Rafat, and J. Richter (1993) *Development of RODOS, a comprehensive decision support system for nuclear emergencies in Europe - an overview*. Radiation Prot. Dosim.

J. Ehrhardt, J. Brown, S. French, G. N. Kelly, T. Mikkelsen, and H. Müller (1997) *RODOS: Decision-making support for off-site emergency management after nuclear accidents*. Kerntechnik, **62**, 122-128.

A. E. Faria Jr., J.Q. Smith (1999) Bayes-DEMM: *A Bayesian Model for Multiple Nuclides and Wet Deposition in the Food Chain Contamination*. Dep. of Stat., University of Warwick.

A.E. Faria Jr., J. Q. Smith, F. Gering, and S. Hering (1997) *Bayes-FDM model for contamination of plants*. RODOS document (in preparation).

T.L. Fine (1973) *Theories of Probability.* Academic Press, New York.

P. C. Fishburn (1970) *Utility Theory for Decision Making*. Wiley, New York.

P.C. Fishburn (1982*) The Foundations of Expected Utility*. D. Reidel Publ. Co., Dordrecht, London.

S. French (1985) *Group consensus probability distributions: a critical survey*. In, J.M. Bernardo, M.H. DeGroot, D.V. Lindley and A.F. M. Smith (eds) Bayesian Statistics II, North Holland Publishing Company, Dordrecht, pp. 183-202.

S. French (1986) *Decision Theory: an introduction to the mathematics of rationality*. Ellis Horwood, Chichester.

S. French (ed.) (1989) *Readings in Decision Analysis.* Chapman and Hall, London.

S. French (1995) *Uncertainty and Imprecision: Modelling and Analysis*. J. Oper. Res. Soc., 46, 70-79.

S. French (1996*) Multi-Attribute Decision Support in the Event of a Nuclear Accident*. J. of Multi-Criteria Decision Analysis, **5**, 39-57.

S. French (1997a) *Uncertainty Modelling, Data Assimilation and Decision Support for Management of Off-site Nuclear Emergencies* Radiation Protection Dosimetry **73**, 11-15.

S. French (1997b) *Statistical decision Theory*. Kendall's Advanced Theory of Statistics Library. Edward Arnold, London (to appear).

S. French, R. M. Cooke, and F. Vogt (1991a) *The use of expert judgement in the context of a postulated mathematical model*. In, J.M. Bernardo, J. M. Berger, A. P. Dawid, and A.F. M. Smith (eds) Bayesian Statistics IV, Oxford University Press, 617-624.

S. French, R. M. Cooke, and M. Wiper (1991b) *The Use of Expert Judgement in Risk Assessment*. Bull. IMA, **27**, 36-40.

S. French, and J.Q. Smith (1992) *The use of Model Predictions and Monitoring Data in Decision making about off-site Emergency Actions in event of an Accident*, Report for CEC. Grant no B170075GB.

S. French, N. Kelly, and M. Morey (1992) *Decision Conferencing and the International Chernobyl Project*. J. Radiol Prot., 12, 17-28.

S. French, and J.Q. Smith (1993) *Using Monitoring Data to Update Atmospheric Dispersion Models with an Application to the RIMPUFF Model*. Radiation Prot. Dosimetry, **50**, 317-320.

S. French, D. Ranyard, and J. Q. Smith (1993). *Models, data and expert judgement in decision support*. International Conference on Modelling and Simulation for the Nuclear Industry, Glasgow, September 1993.

S. French, E. Halls, and D. Ranyard (1995a) *Equity and MCDA in the event of a nuclear accident*. School of Computer Studies, University of Leeds Research Report. RODOS(B)-RP(95)03.

S. French, D. Ranyard, and J. Q. Smith (1995b) *Uncertainty in RODOS*. School of Computer Studies, University of Leeds Research Report. RODOS(B)-RP(94)05.

S. French, D. C. Ranyard, and M. T. Harrison (1997a). Event conditional attribute modelling in decision making when there is a threat of a nuclear accident rodos(b) - rp(95)04. School of Informatics, University of Manchester. In S. French and J. Q. Smith (1997) (eds) Bayesian Analysis in Practice. Edward Arnold, London, pp. 131-149.

S. French, and K. N. Papamichail (1997b) *Further Specification of RODOS ESY*. RODOS(WG5)-TN(97)-06.

S. French and J.Q. Smith (1997c) The Practice of Bayesian Analysis. Edward Arnold, London.

S. French, K.N. Papamichail, D.C. Ranyard and J.Q. Smith (1998) Design of a decision support system for use in the event of a radiation accident. F. Javier Girón and M. L. Martínez (eds) *Decision Analysis Applications* Kluwer Academic Publishers, Dordrecht. 3-18.

A. Gargoum (1997) Issues in Bayesian *Forecasting of Dispersal after a Nuclear Accident*. PhD Thesis, University of Warwick.

A. Gargoum and J.Q. Smith (1993) "*Approximation schemes for efficient probability propagation in evolving high dimensional Gaussian processes*" Dept. of Statistics, University of Warwick.

A. Gargoum, and J. Q. Smith (1994) Approximating Dynamic Gaussian junction trees. University of Warwick Technical Report.

A. Gargoum, and J.Q. Smith (1997a) *Dynamic stochastic approximations of distributions in dispersal models based on a modified Kullback-Leibler divergence measure*. University of Warwick Technical Report.

A. Gargoum, and J.Q. Smith (1997b) Dynamic Generalised Linear Probability Propagation. University of Warwick, Technical Report no 315.

A. Gelman, J. B. Carlin, H. S. Stern, and R. D. Rubin (1995) *Bayesian Data Analysis*. Chapman and Hall, London.

J. Ch. Gilbert, C. Marechal (1989) *Some numerical experiments with a variable storage quasi-Newton algorithm*. Mathematical programming, **45**, pp. 407-435.

W. R. Gilks, D. G. Clayton, D. J. Spiegelhalter, N. G. Best, A. J. McNeil, L. D. Sharples, and A. J. Kirby (1993) *Modelling complexity: applications of Gibbs sampling in medicine*. J. R. Stat. Soc. B, **55**, 39-52.

F. T. Harper, S. C. Hora, M. L. Young, L. A. Miller, C. H. Lui, J. C. Helton, L. H. J. Goosens, R. M. Cooke, J Päsler-Sauer, B. Kraan, and J. A. Jones (1994) Probabilistic Accident Consequence Uncertainty Analysis: Dispersion and Deposition Uncertainty Assesment. nureg/cr-6244, Publication no. eur 15855 en of the Comision of the European Communities, Luxemburg.

F.V. Jensen (1996) *An introduction to Bayesian Networks.* UCL press, London.

G. N. Kelly (1994) *RODOS - A Decision Support System for off-site emergency management background and future perspectives*. Technical Report.

R. L. Keeney (1992) *Value-Focused Thinking: A Path to Creative Decision Making.* Harvard University Press, Cambridge, MA.

R. L. Keeney, and H. Raiffa (1976) *Decisions with Multiple Objectives: Preferences and Value Trade-offs*. Wiley, New York.

R. L. Keeney, and R.L. Winkler (1985*) Evaluating decision strategies for equity of public risks*. Operations research, **33**, 955-970.

D. A. Klein, (1994) *Decision Analytic Intelligent Systems - Automated Explanation and Knowledge Acquisition.* Lawrence Erlbaum Associates, New Jersey.

S.L. Lauritzen (1996) *Graphical Models*. Oxford University Press.

S. L. Lauritzen, and D. J. Spiegelhalter (1988) *Local computations with probabilities on graphical structures and their application to expert systems*. J. R. Stat. Soc. B, **50**, 157-224.

P. Lee (1989) *Bayesian Statistics: an Introduction.* Oxford University Press.

K. V. Mardia, J. T. Kent, and J. M. Bibby (1979) *Multivariate Analysis*. Academic Press.

T. Mikkelsen, S.E. Larsen, and S. Thykier Nielsen (1984) *Description of the Risø Puff Dispersion Model*. Nuclear Safety, **67**, 56-65.

T. Mikkelsen, S. E. Larsen, and H. L. Pecsli (1987) Diffusion of Gaussian puffs. Q. Journal R. Meteorol. Soc., 113, pp.81-105.

A. O'Hagan (1988) *Probability: Methods and Measurement*

A. O'Hagan (1994) *Bayesian Inference*. Kendall's Advanced Theory of Statistics Volume IIB. Edward Arnold, London.

A. O'Hagan (1996) First Bayes. Software for teaching Bayesian Inference. Available from http:\\www.maths.nott.ac.uk\personal\aoh\.

N. D. Le, J. Zidek (1992) Interpolation with Uncertain Spatial Covariances: a Bayesian Alternative to Krigging. J. Multivariate Analysis, **43**, 351-374

N. D. Le, J. Zidek (1994) *Network Designs for Monitoring Multivariate Random Spatial Fields*. Recent adv. in Stat. and Prob., J. P. Vilaplana and M. L. Puri (eds), pp. 191-206.

K. N. Papamichail, and S. French (1997) *Screening Strategies in Nuclear Emergencies*. RODOS(WG5)-TN(97)-01.

K. N. Papamichail, and S. French (1999) "Generating Feasible Strategies in Nuclear Emergencies - A Constraint Satisfaction Problem", *Journal of the Operational Research Society* **50**, 617-626.

J. Päsler-Sauer (1986), *Comparative Calculations and Validation Studies with Atmospheric Dispersion Models*. Kernforchungszentrum, Karlsuhe.

J. Päsler-Sauer, and T. Schichtel (1997). *The simulation of early emergency actions in RODOS. Proceedings Fourth International Workshop on Real-Time Computing of the Environmental Consequences of an Accidental Release from a Nuclear Installation*, Aronsborg 7-11 October 1996. To appear in Radiation Protection Dosimetry.

K. Politis (1998) *Some thoughts on applying data assimilation to a long range atmospheric dispersion model*. Internal note, School of Informatics, University of Manchester.

K. Politis and L. Robertson (1999) *Bayesian Updating of Atmospheric Dispersion After a Nuclear Accident*. In preparation

H. Raiffa (1968) *Decision Analysis*. Addison Wesley, Reading, Massachusetts.

D.C. Ranyard (1996) *Decision Support for Nuclear Emergency Response.* PhD Thesis, university of Leeds.

D. Ranyard, and J. Q. Smith (1997) Building a Bayesian model in a scientific environment: managing uncertainty after an accident. In S. French and J. Q. Smith (1997) (eds) Bayesian Analysis in Practice. Edward Arnold, London, pp. 245-258.

L. Robertson, C. Persson (1993) *Attempts to apply four dimensional data-assimilation of radiological data using the adjoint technique*. Radiation Protection Dosimetry, Vol. 50, Nos 2-4, pp. 333-337, Nuclear Technology Publishing.

L. Robertson, J. Langner, M. Engart (1998) An Eulerian limited area atmospheric transport model. To appear in J. Appl. Meteor.

L. J. Savage (1954) *Foundations of Statistics*. First Edition, Wiley, New York. Second Edition (1972), Dover, New York.

R. Settimi, and J. Q. Smith (1997) A comparison of approximate Bayesian forecasting methods for non Gaussian time series. Preprint.

C. Smedley, L.M.C. Dutton, D. Clarke, D.A. Millington, E. Grindon (1999) *RODOS Source Term Module (*stm*)* RODOS(WG10)-TN(98)-08

A. F. M. Smith, and G. O. Roberts (1993) *Bayesian computation via the Gibbs sampler and related Markov chain Monte Carlo methods (with discussion).* J. R. Stat. Soc. B, **55**, 3-23.

J. Q. Smith (1992) *A comparison of the characteristics of some Bayesian forecasting models.* Intern. Stat. Review, **1**, 75-87.

J.Q. Smith (1988) *Decision Analysis: a Bayesian approach*. Chapman and Hall, London.

J.Q. Smith (1989) "*Influence diagrams for Bayesian decision analysis*", Eur.J.Oper.Res., **40**, 363-376, 1989.

J. Q. Smith (1995) *Bayesian approximations and the Hellinger metric*. University of Warwick, Technical Report no 279.

J. Q. Smith, S. French, and D. C. Ranyard (1995) An efficient graphical algorithm for updating estimates of the dispersal of gaseous waste after an accidental release. In A. Gammerman (ed) Probabilistic reasoning and Bayesian Belief Networks, Alfred Waller, Henley-on-Thames, pp. 125-144.

O. Talagrand, P. Courtier (1987) Variational Assimilation of Meteorological Observations with the Adjoint Vorticity Equation, I: theory. Q. J. Meteorol. Soc., 113, 1311-1328.

S. Thykier-Nielsen, and T. Mikkelsen (1991). *Rimpuff Users Guide:Version 30*. Risø National Laboratory, Denmark.

S. Thykier-Nielsen, T. Mikkelsen, S. E. Larsen, J. Troen, A. F. de Baos, R. Kamada, C. Skupniewicz, and G. Schacher (1989*) A model for Accidental Releases in Complex Terrain.* Air Pollution and its Applications VII, pp. 65-76 (Plenum).

S. Thykier-Nielsen, S. Deme, and Lang (1993) *Calculation Methods for gamma-Dose rates for Spherical Puffs*. Risø National Laboratory, Risø-R-692(EN), Denmark.

S. Thykier-Nielsen, S. Deme, and Lang (1995) *Calculation Methods for gamma-Dose rates for Gaussian Puffs*. Risø National Laboratory, Risø-R-775(EN), Denmark.

S. Thykier-Nielsen, S. Deme, and T. Mikkelsen (1997) Rimpuff: Version RIMDOS 7 Users Guide. RODOS(A)-TN(95)3.

E. Tsang (1993) *Foundations of Constraint Satisfaction*. Academic Press.

U. Tveten, T. Mikkelsen (1995) (eds) *Report on the Nordic Dispersion - Trajectory Model Comparison with the ETEX-I Fullscale Experiment*. Riso National Laboratory, Roskilde, Denmark, Report no. R-847(EN).

T. Valet, and S. A. Roberts (1994) Identifying the requirement for history of time-varying objects during an object oriented analysis. School of Computer Studies, University of Leeds Technical Report.

D.B. Vleeshhouwer (1996) An application of Belief Networks for Source Term Estimation. Dissertation, TU Delft

D. von Winterfeldt, and W. Edwards (1986) Decision Analysis and behavioural Research. Cambridge University Press.

S. Walker, P. Muliere (1998) *Beta-Stacy Processes and a Generalisation of the Polya-urn scheme*. The annals of statistics, vol. 25, 1762-1780.

M. West, and J. Harrison (1989) *Bayesian Forecasting and Dynamic Models*. Springer-Verlag,. London, New York.

Appendix A: A Tutorial and Survey of Bayesian Methods

Tea, Haydn and a right Tosser

Let's begin with an example. It is based on a set of remarks of L.J. Savage[2], one of the modern founders of Bayesian statistics.

Consider three similar scenarios:

1. A friend of yours who is a connoisseur of tea claims that he can tell whether the milk or tea was put into the cup first just from one sip of tea. You set up a randomised trial. 10 cups of tea will be made behind a screen. The choice of whether milk or tea is put in the cup first is made at random for each cup by using a random number generator on a calculator using (50:50) odds. Your friend is allowed one sip and has to say which was put in the cup first: tea or milk. He is right 10 times out of 10. What is the probability that when he is offered an eleventh cup of tea, he will be right?

2. A music expert claims that by looking at a page from a score he can determine whether the music was written by Haydn or a contemporary composer. You assemble a set of scores from Haydn and his contemporaries and choosing at random (50:50) select pages from the Haydn pile and the contemporary pile. The expert is right 10 times out of 10. What is the probability he is right on an eleventh?

3. You are in a pub idly tossing a coin. Heads and tails are coming up seemingly randomly, and idly over 100 tosses, you have noted that there were 53 heads and 47 tails. You believe it to be a fair coin. A drunk sits next to you and claims he can predict the next 10 tosses accurately. He does! What is the probability that his prediction is right next time on an eleventh toss?

Suppose in each scenario we set up two hypothesis – for such is the way of classical statistics as taught in most universities:

H_0 - the expert is right no more than would be expected by chance.

H_1 - the expert performs better than chance.

[Strictly, there are more hypotheses that we might make, but these will be sufficient for the points we wish to make here.]

In each case, the evidence under H_0 has probability $(1/2)^{10}$. By any standards in classical statistics this is significant evidence against H_0. Thus in all cases we should believe the 'expert' is making valuable predictions. But a question or, to be precise, three questions stand unanswered: what is the probability that the expert is right in eleventh test? There are several ways that classical statistics might approach these questions, but the point to note here that in each case *the approach would be the same*. The objective evidence in each case is the same giving a likelihood ratio of 2^{-10}. Yet, as Savage originally noted, any normal person – and, indeed, any *scientist* – would react differently in each case. Some people have remarkable palates: it may be that connoisseur tea drinkers can tell whether milk was added first even if others cannot. In this case the evidence is persuasive. Music experts should be able to recognise a Haydn score. In this case the evidence confirms one's prior expectations. One would have been surprised, to some extent, if he hadn't assigned the scores correctly. In the third case, would one really be impressed by a drunk's run of luck? Thus one's instinct would be to predict that the tea drinker would be correct on the

[2] See Barnett (1973, page 12) for a history of the example

eleventh test with quite a high probability; that the music expert would be right with a very high probability; and that the drunk would do no better than a chance guess.

This example is important in two respects for our discussion. Firstly, it emphasises that one cannot learn form data independently of one's prior expectations. To claim that the 'objective' evidence is identical in each case and, therefore, one should reach identical conclusions on the validity of the tea connoisseur's, music expert's and drunk's claims simply does not accord with the way that our instincts – our sensible instincts – tell us.

Secondly, it begins to point to some difficulty with prediction. If one accepts H_0, then the probability of the expert being right on the eleventh trial would be 0.5; and if one accepts H_1 the probability is ill defined. Of course, one might choose another hypothesis as H_1: maybe the expert is right 95% of the time. If one accepted H_1 (which is not quite the same as rejecting H_0, we admit), then the probability would be 0.95. But why choose 95%. A more general H_1, such as the expert is right greater than 50% of the time does not lead to a straightforward predictive probability.

Later in the tutorial, we see how a Bayesian analysis would proceed in each of the three cases.

Probability and Bayes' Theorem: some basics

Subjective versus objective interpretations of probability

The key to the Bayesian approach is that we use probability to represent uncertainty[3]. The probability of a future event is the (observer's judgement of the) likelihood of that event occurring. There is great debate in literature extending from statistics through economics and management to philosophy as to whether probability should have a subjective or objective interpretation: see e.g. Barnett (1982), Cooke (1991), Fine (1973), French (1986), O'Hagan (1988), and Wright and Ayton (1994). Essentially, this comes down to whether we should model beliefs or frequencies. The probability that a parameter takes a particular value is an interesting philosophical beast. It is difficult to interpret this as an objective probability because parameters are subjective modelling constructs. Thus while for the present we shall take a pragmatic view of the debates between objective and subjective interpretations, we shall lean pretty far towards the subjective.

Laws of Probability

Mathematically, the laws of probability are simple and well known:

- Probabilities are non-negative.

[3] Not strictly true: sometimes a decision maker uses the expression "I am uncertain …" to mean "I haven't thought very deeply about and clarified my thinking …". Bayesians do not attempt to model such uncertainty. Rather they use various forms of discussion to encourage the decision maker(s) to explore the issues more thoroughly. (French, 1995)

- The probability of an impossible event is zero. (Use with care in modelling! Once something is declared impossible and modelled with a probability of 0, the multiplicative properties of zero imply that the assertion of impossibility cannot be withdrawn. This advice has been dubbed *Cromwell's Law* by Dennis Lindley.)
- Probabilities of disjoint events add up: *viz.*

  $P(A \cup B) = P(A) + P(B) \text{ if } A \cap B = \varnothing$

Usually this works over a countably infinite[4] set of events and probability is said to be countably additive. The total probability over all possibilities (probability of certain event) is 1.

$$\sum_{\text{All possible events}} P(\text{event}) = 1$$

Conditional probabilities and independence

The *conditional probability* of event A given that event B is known (or assumed) to have happened is:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

i.e. the ratio of the probabilities of *A* and *B* happening together to that of *B* happening with or without *A*. Note $P(B|B) = 1$, i.e. a certainty always has the probability 1.

A and *B* are *independent* if learning *B* tells you nothing about *A*: *viz.*

$$P(A|B) = P(A)$$

An equivalent definition says that *A* and *B* are independent if

$$P(A \cap B) = P(A) \times P(B)$$

Independence is important for building models that are computationally tractable as we shall see over the tutorial, but dependence is equally important. Without it we cannot learn. Thus there is a balance to be drawn which we shall return to at several points over the course.

Bayes' Theorem

We remarked above that *dependence* (i.e. the absence of independence) gives the opportunity of learning. Let's consider how this might happen through a simple example.

> By consulting your seaweed, which is damp but not sopping wet and looking at the sky, you come to the conclusion that the chances of rain tomorrow are about 60:40. You then ask a friendly meteorologist (and the RODOS project has plenty of those!) and he says that it is going to rain tomorrow. Call him 'M'. Now you have asked M

[4] If words such as "countably infinite" and some of the later mathematics mean little or nothing to you, pass them by. The aim of this tutorial is to get the basic ideas over.

for his forecasts on many occasions in the past and have noticed that he either says "it is going to rain" or "it is not going to rain" (meteorologists clearly have a limited vocabulary). Moreover, you have noticed that M is not always right (See: this is a realistic example!). On the occasions that it has rained, he has predicted rain about 90% of the times, but 10% of the times he hasn't. And on those occasions that it has not rained, he has predicted no rain 80% of the times, but rain 20% of the times.

We can illustrate these numbers on a diagram: see Figure 14. We have drawn a unit square, which we shall use as a Venn diagram, and remembering that probabilities sum to 1, we can make areas in this square correspond to probabilities. Your probability of 0.6 of rain is represented by dividing the square vertically in the ratio 0.6:0.4. Next we divide each vertical column according to the proportion of times that M says it will rain and it will not. The column which corresponds to your likelihood of rain will, therefore, be divided so that 90% of it corresponds to the times M says it will rain and the remaining 10% to the times that he says it will not. The column corresponding to your likelihood of no rain is divided similarly, but in the ratio 0.2:0.8. Now he has said it will rain. So your attention is confined to the shaded area, the two regions corresponding to M's statement "it is going to rain". The larger area to the left corresponds to your likelihood of rain; that to the right to your likelihood of no rain. Thus your probability of rain *after* hearing the M's view is:

$$P(\text{rain}|\text{M says "rain"}) = 0.9\times0.6/(0.9\times0.6 + 0.2\times0.4)$$
$$= 54/62 = 0.87$$

Your probability of rain has increased from 60% to 87% on hearing and assimilating M's statement. Translating the numbers above into probabilities,

$$P(\text{rain}|\text{M says "rain"}) = P(\text{M says "rain"}|\text{rain}) \times P(\text{rain})/P(\text{M says "rain"})$$

This is the simplest example of the application of Bayes' theorem:

Bayes' Theorem[5]

$$P(B|A) = \frac{P(A \cap B)}{P(A)} = \frac{P(A|B)P(B)}{P(A)} \propto P(A|B)P(B)$$

The probability *a posteriori* of *B* after you have learnt *A* is the probability of *B a priori* multiplied by the conditional probability of *A* if you assume *B*, known as the *likelihood*[6].

posterior probability $\propto$ likelihood $\times$ prior probability

[5] In the example event *B* corresponds to 'rain' and event *A* to M's statement "It is going to rain".

[6] In statistics, the likelihood function is the probability of the data actually observed given the unknown parameters *considered as a function of those parameters*.

Bayes' theorem underpins Bayesian approaches to statistics and decision making in the light of evidence. Note that in the example there is dependence: the probability that M says "It is going to rain" varies with the occurrence of rain. It is this dependence that allows you to update your belief in rain.

Discrete and Continuous Distributions

All the above were stated in terms of discrete events, but most of the theory of probability applies with continuous variables or quantities. Probabilities are given by integrals of density functions; and generally in the laws of probability integrals replace sums. Ideas of *conditionality* carry through with suitable care about dividing by zero. Independence means that the joint distribution is the product of the *marginals*. To be more precise, for continuous variables *X* and *Y*:

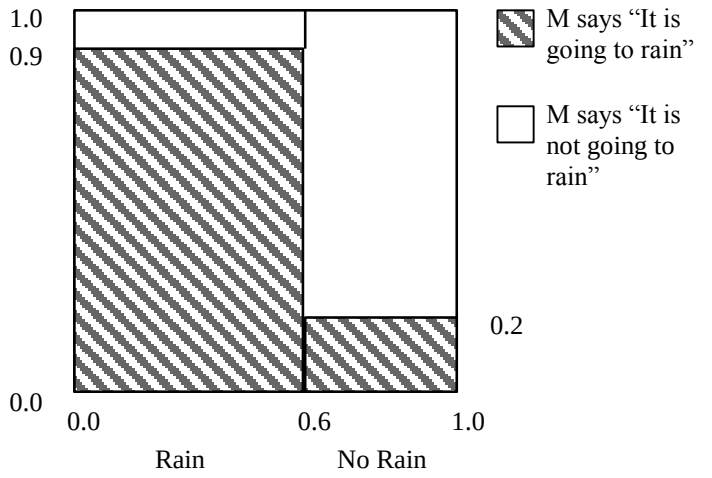

Figure 14: The Rain example

$$P_{X,Y}\left(X \le x, Y \le y\right) = \int_{-\infty}^{x}\int_{-\infty}^{y} p_{X,Y}\left(x', y'\right)dy'dx' \quad (1)$$

$$P_X\left(X \le x\right) = \int_{-\infty}^{x} p_X\left(x'\right)dx' \quad (2)$$

$$p_X\left(x\right) = \int_{-\infty}^{\infty} p_{X,Y}\left(x, y'\right)dy' \quad (3)$$

$$p_X\left(x|Y = y\right) = p_{X,Y}\left(x, y\right) \Big/ \int_{-\infty}^{\infty} p_{X,Y}\left(x', y\right)dx' \quad (4)$$

where:

1. $P(.)$ represents a *distribution function*, that is the probability that uncertain quantities are less than or equal to given values and $p(.)$ represents the density function.
2. We have used subscripts to indicate the uncertain quantities concerned.
3. Equation (1) relates the joint distribution function of X, Y to the joint density.
4. Equation (2) relates the marginal distribution function of X to its marginal density function.
5. Equation (3) shows that the marginal density of X is found by integrating out Y from the joint density.
6. Equation (4) shows that the conditional density of X given that $Y = y$ is given by dividing the joint density of X, Y at (x, y) by the marginal density of Y at $Y = y$, assuming that this does not involve a division by zero.

Two quantities are probabilistically independent when their joint density is the product of their marginals:

$$p_{X,Y}\left(x, y\right) = p_X\left(x\right) \times p_Y\left(y\right)$$

For density functions of continuous quantities, Bayes' Theorem becomes:

$$p_X(x|\ Y = y) \propto_x p_Y(y|x) \times p_X(x) \quad (5)$$

where $\propto_x$ means is proportional to as a function of x. As before we may write:

posterior density $\propto$ likelihood $\times$ prior density

Note that the use of the word 'likelihood' here corresponds precisely to the likelihood function used in non-Bayesian statistics. In fact, Bayesian analysis necessarily obeys the likelihood principle: namely, Bayesian inferences depend on the data only through the likelihood function.

The use of proportionality here is an important Bayesian technique because it allows one to forget many difficult calculations during an analysis and then mop things up at the end by recognising that probability densities must integrate to one. Let's look at a simple application of Bayes' Theorem to normal distributions.

Normal priors and likelihoods

A common assumption in statistics is that the distributions involved are normal (Gaussian). This assumption is made for a number of reasons. The first is important but unconvincing, so the fact that other more convincing reasons exist is comforting!

1. Tractability: the calculations are feasible for normal distributions but not for non-normal ones. This was particularly true in the forties and fifties before computers were readily available to statisticians. Normal distributions have the sample mean and variance as sufficient statistics which ensures very straightforward calculations in both Bayesian and non-Bayesian analysis alike.
2. The Central Limit Theorem shows that if the uncertainty encoded by a probability distribution arises from a number of independent sources which each obey certain reasonable regularity assumptions, then normality is a good approximation.
3. The use of Taylor series approximations is familiar to many users of mathematics. We know how to use this to give first order, second order, third order, etc. approximations to functions. There is a similar concept of an Edgeworth expansion of a probability distribution and – lo and behold – the normal distribution with suitable mean and variance gives the 'first order' approximation[7].

Suppose that we are interested in some physical quantity θ. Our prior knowledge of this suggests that we expect this to be, say, 10.5 ± 2.0, where we interpret '$\pm$ 2.0' to mean that we are about 95% sure that the 'true' value of θ is within the range [8.5, 12.5]. We may model this by a normal distribution with standard deviation 1.0:

$$X|\theta \sim \mathrm{N}(10.5, 1.0)$$

where we have used the notation $\mathrm{N}(\mu, \sigma^2)$ for the normal distribution with mean μ and variance σ^2, and '~' to mean 'is distributed as'.

Next assume that we may make an observation X on θ such that:

$$X \sim \mathrm{N}(\theta, (0.1)^2).$$

Note that this notation states that the variance of the observation process is $(0.1)^2$. Suppose that we observe $X = 10.23$. The distribution $P_\theta(\theta \mid X = 10.23)$ will represent the resulting uncertainty on θ. To apply Bayes' Theorem with a normal prior and normal likelihood we need:

Lemma 1

Suppose that:

[7] Note that some of the work done at the University of Warwick on the Bayes_RIMPUFF model and algorithm has investigated the sensitivity to the assumption of normality in this context and the effect of using non-normal distributions such as the Poisson (Settimi and Smith, 1998).

$$\theta \sim \mathrm{N}\left(\mu, \sigma^2\right)$$
$$X|\theta \sim \mathrm{N}\left(\theta, \tau^2\right)$$

Then:

$$\theta|x \sim \mathrm{N}\left(\mu', \sigma'^2\right)$$

where

$$\mu' = \left(w_0\mu + w_1 x\right)\Big/\left(w_0 + w_1\right)$$
$$\sigma'^2 = 1\Big/\left(w_0 + w_1\right)$$

and $w_0 = 1/\sigma^2$ and $w_1 = 1/\tau^2$.

Proof

[This is offered *only* for those who are *really* interested! It does demonstrate the power of working with proportionality: the algebra, although not fun, is much, much simpler than if we kept the constants of proportionality throughout.]

Applying Bayes' Theorem:

$$p_\theta(\theta|x) \propto_\theta p_X(x|\theta) \times p_\theta(\theta)$$

Inserting the densities for the normal distributions:

$$p_\theta\left(\theta|x\right) \propto_\theta \frac{1}{\sqrt{2\pi\tau^2}}\exp\left(-\frac{\left(x-\theta\right)^2}{2\tau^2}\right) \times \frac{1}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{\left(\theta-\mu\right)^2}{2\sigma^2}\right)$$
$$\propto_\theta \exp\left(-\frac{\sigma^2\left(x^2 - 2x\theta + \theta^2\right) + \tau^2\left(\theta^2 - 2\mu\theta + \mu^2\right)}{2\tau^2\sigma^2}\right)$$
$$\propto_\theta \exp\left(-\frac{\left(\tau^2 + \sigma^2\right)\theta^2 - 2\left(\tau^2\mu + \sigma^2 x\right)\theta}{2\tau^2\sigma^2}\right)$$

$$\propto_\theta \exp\left(-\frac{\theta^2 - 2\left(w_0\mu + w_1 x\right)\theta\Big/\left(w_0 + w_1\right)}{2\,\tau^2\sigma^2\Big/\left(\tau^2 + \sigma^2\right)}\right)$$
$$\propto_\theta \exp\left(-\frac{\left(\theta-\mu'\right)^2}{2\sigma'^2}\right)$$

The above derivation follows through on remembering throughout to cancel all multiplicative quantities that are independent of θ and on 'completing the square' in the exponential. *QED.*

Note that μ' is a weighted average of the prior mean μ and the observation x, which seems sensible. The weights are related to the ratio of the prior

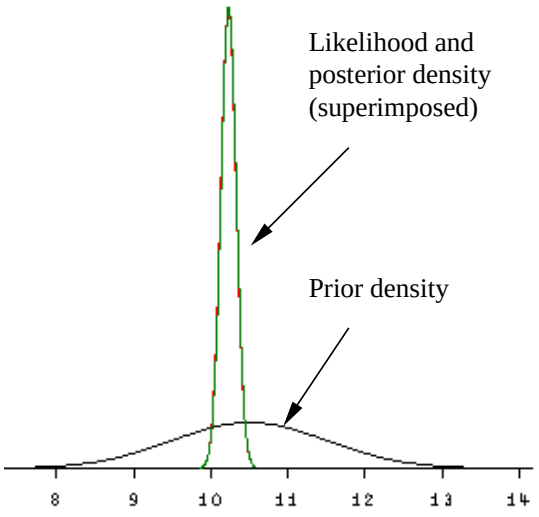

Figure 15: Prior and posterior densities for the normal example

variance to the observation variance, which again seems sensible. This is a very common property of Bayesian inference: the posterior mean is a sensible interpolate between the prior mean and the observation.

In the numerical example, we have:

$$\mu' = (1\times 10.5 + 100\times 10.23)/101 \approx 10.23$$
$$\sigma^2 = 1/(1/1 + 1/(0.1)^2) = 1/101 \approx (0.1)^2$$

In other words, since the observation was so informative, the posterior has moved to be effectively the likelihood. This is precisely what would happen in classical statistical estimation of θ.

Figure 15 illustrates this application of Bayes' Theorem. The likelihood function and the posterior are almost identical and so are superimposed on the diagram.

Returning to Lemma 1, it is interesting to rewrite the formula for μ':

$$\mu' = \left(w_0\mu + w_1 x\right)\Big/\left(w_0 + w_1\right)$$
$$= \mu + w_1\left(x - \mu\right)\Big/\left(w_0 + w_1\right)$$

This says that the posterior mean is obtained by adjusting the prior mean by a linear multiple of the 'error' (x - μ). This is a very simplified form of a *Kalman filter*, The multiplier $w_1/(w_0 + w_1)$ is known as the Kalman filter because it 'filters out' the noise in the 'error signal' (x - μ) to adjust the prior mean. This result due to R. Kalman is heavily used in control theory (Aoki, 1966). This is developed in appendix B.

Tea, Haydn and the drunk revisited

The Beta-Binomial model

Let us set the probability that the expert (be he a tea connoisseur, a musicologist or a drunk) is right in a single test to be π. Then in n trials the probability of x correct and (n - x) incorrect predictions is

$$p_X\left(x|\pi\right) = \binom{n}{x}\pi^x(1-\pi)^{n-x}, \qquad (6)$$

i.e. a binomial probability. This gives the likelihood function. For our data π^{10}; but more generally (6) considered as a function of π for fixed n and x. What should we use as a prior? We consider a Beta distribution:

$$p_\Pi(\pi) = \frac{\Gamma(\alpha)\Gamma\left(\beta\right)}{\Gamma\left(\alpha+\beta\right)}\pi^{\alpha-1}(1-\pi)^{\beta-1} \qquad (7)$$

for $0 \leq \pi \leq 1$. At first this looks a completely arbitrary choice, but remember Bayes' theorem is applied using proportionality. Multiplying (6) by (7) and remembering that we are interested in the result as a function of π:

$$p_\Pi(\pi|x) \propto_\pi \pi^{x+\alpha-1}(1-\pi)^{n-x+\beta-1}. \qquad (8)$$

This has the same functional form as (7), *viz.* a beta distribution with parameters (x+α-1) and (n-x+β-1). In other words, the data updates the parameters of the distribution for π, but the form of the distribution does not change. The prior and posterior have the same form, as will the next posterior if further data become available. Because of this we say that the beta and the binomial form a conjugate pair. Looking back at Lemma 1, note that in our consideration of a normal model above, the same property occurred: if the prior is normal with known variance and the likelihood is normal with known variance, then the posterior is normal with variance that is known (i.e. can be constructed from the variances of the prior and likelihood).

It would be nice if all distributions could be put into conjugate pairs, but few can. Thus, if the prior and the likelihood are constructed to model the DM's beliefs, it may well be that they are not conjugate distributions and hence the application of Bayes' Theorem is non-trivial. Usually it requires multi-dimensional numerical integration. For this reason, conjugacy is important; but less so than it was. Until the advent of really high powered computing, conjugacy gave the only straightforward way of conducting Bayesian analyses. It made calculations possible. Now there are very powerful Monte Carlo Markov Chain (MCMC) methods available (see, e.g., Gelman *et al*, 1995). This does not mean that conjugacy can be forgotten: it provides a very good method of making approximate calculations and, pedagogically speaking, it provides lecturers with nice examples! In fact, in the RODOS project we are using conjugate prior-posterior pairs in order to achieve very fast computation, faster than possible even with MCMC. However, we are also using MCMC methods to validate the conjugate methods (see Settimi and Smith, 1997).

But back to the example: the mean of the beta distribution (7) is $\alpha/(\alpha+\beta)$ and the variance is $\alpha\beta/[(\alpha+\beta)^2(\alpha+\beta+1)]$. Thus the mean of the posterior (8) is $(\alpha+x)/(\alpha+\beta+n)$.

The tea connoisseur

Let us suppose that you have an open mind about the ability of a tea connoisseur to tell whether milk is put into the cup before or after the tea. If we take $\alpha = 1$ and $\beta = 1$ in (7), the resulting prior is the uniform distribution on the unit interval, $0 \leq \pi \leq 1$; i.e. it does not favour any value of π over any other. This seems to express your open mind on his ability. Now observing 10 out of 10 successes will change your posterior mean for π to $(1+10)/(2+10) = 11/12$ which seems a reasonably high probability given the data. The variance is correspondingly small to reflect your growing certainty that he can tell.

Figure 16 shows the prior, likelihood and the posterior for this case. It shows how the evidence is winning over your open mind and reflecting the data which suggests that the connoisseur can tell whether milk was added first.

The music expert

Now let us consider the musicologist. We would be surprised if he could *not* tell Haydn from other composers. So our prior for π would give most of its weight to the region near 1. For instance, if we choose α=10, β = 1, our prior mean is 10/11 and it has variance 0.006 $\Rightarrow$ standard deviation of 0.08. Observing 10 out of 10 successes changes our mean for π to 20/21 with standard deviation 0.04. We are becoming much more convinced that π is close to 1.

Figure 17 shows the prior, likelihood and posterior for this example. Note that the prior and likelihood are pretty much the same. In the case of the tea expert, the data greatly outweighed our prior expectations; we had an open mind. Here we were pretty sure that the musicologist would identify the scores correctly and we were confirmed in our beliefs.

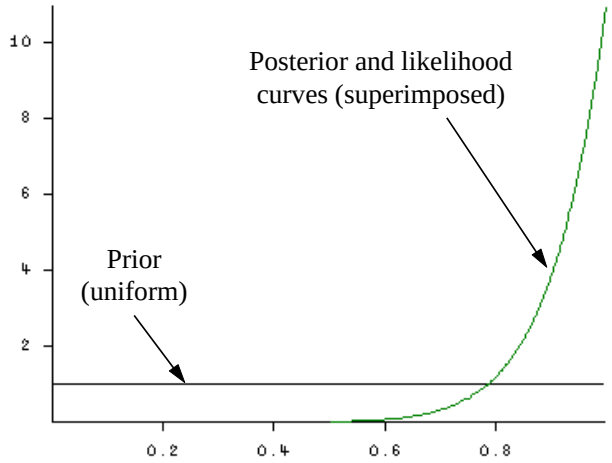

Figure 16: Prior, likelihood and posterior for the tea connoisseur

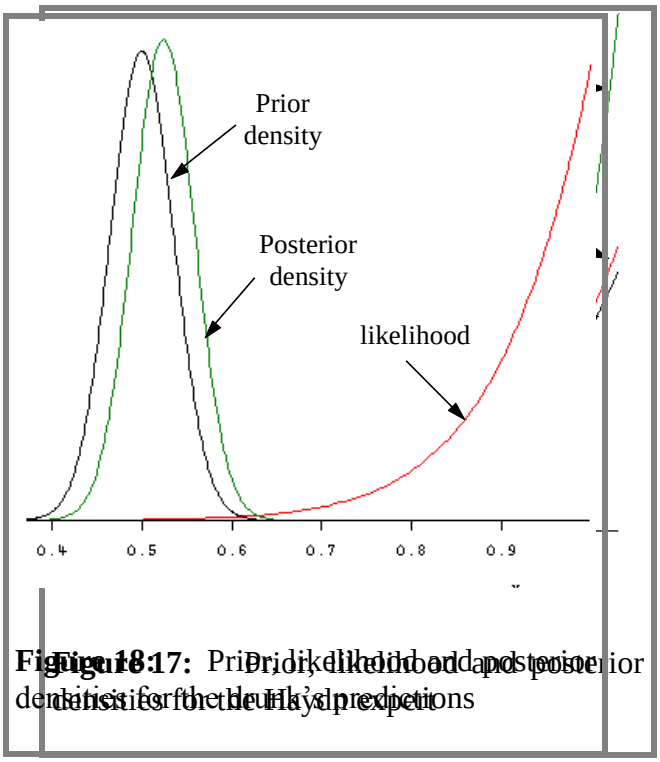

Figure 18.17: Prior, likelihood and posterior densities for the drunk's predictions

The drunk's run of luck ….

Now let us turn to the drunk predicting our coin. What should we set as our prior? Well we have seen 53 heads and 47 tails in idly tossing the coin. We could presume that we started out with little view on whether the coin was fair ($\alpha = 1$, $\beta = 1$). This would give a Beta posterior with ($\alpha = 54$, $\beta = 48$). But although this is a sensible way of constructing a prior on the fairness of the coin, it does not really represent our disbelief in the drunk's ability to predict a patently random mechanism with roughly equal chances of heads or tails. Thus we choose to inflate these values of α and β to 100 each. This gives a very firm prediction that he is equally likely to be right as to be wrong. Observing ten successes out of ten gives a posterior with $\alpha = 110$ and $\beta = 100$. This hardly changes our views: see Figure 18.

Predictive probabilities

The actual problem was to predict the outcome of the eleventh trial in each case. In this example, the posterior mean is the value we need. Respectively:

Tea Connoisseur	0.917
Musicologist	0.952
Drunk	0.523

These values seem to accord well with common sense.

Belief nets

The main part of this section has been drawn from Bedford and Cooke, 1997. There it is assumed that all random variables are discrete. However, we say here without proof that it is possible to extend the definition of a belief net to use continuous variables as in the case of RIMPUFF. Further we'll assume throughout that the numbering of the nodes is consistent with the hierarchy of the network (thus if X_i and X_j, with $i < j$, denote nodes in the network, then X_i is higher in the hierarchy than X_j).

Bayesian belief networks are graphical tools used to represent a high-dimensional probability distribution. They are convenient tools for making Bayesian inferences about uncertain states which is especially helpful when limited information is available. Belief nets have more and more applications these days as the computational demands of Bayesian inference are not as major a problem as they used to be. Belief nets are already very heavily used in many application areas and as a specific example we mention the use of a belief net within the help function of the Windows printer system.

A belief net is a directed acyclic graph (DAG) whose nodes, called chance nodes represent random variables (or groups of random variables). An arrow from one node to another represents probabilistic influence. Figure 19 shows an extremely simple belief net, which represents that one variable, *X*, influences another, *Y*. The probabilistic specifications in this belief net are represented by the marginal distribution of *X*, and the conditional distribution of *Y* given *X*, for every value of *X*. This specifies the full joint distribution as

$$p(x,y) = p(x)\,p(y|x)\,.$$

More generally, for *n* variables $X_1,\ldots,X_n$ one can always decompose the joint distribution by

$$p(x_1,\ldots,x_n) = p(x_1)\,p(x_2|x_1)\,p(x_3|x_1,x_2)\ldots p(x_n|x_1,\ldots,x_{n-1}).$$

The righthandside of this expression is not usually a compact way to write down a joint distribution. Under certain assumptions however, the expression can be simplified. If the X_i form a Markov chain for example, then for each *i*, X_i is independent of $X_1,\ldots,X_{i-2}$ given X_{i-1}, so that we can write

$$p(x_i|x_1,\ldots,x_{i-1}) = p(x_i|x_{i-1}) \qquad (9)$$

and thus:

$$p(x_1,\ldots,x_n) = p(x_1)\,p(x_2|x_1)\,p(x_3|x_2)\ldots p(x_n|x_{n-1}).$$

So a product of an *n*-dimensional function with an (*n*-1)-dimensional function with an (*n*-2)-dimensional function with an …….with a 1-dimensional function has simplified in the product of a 1-dimensional function and an (*n*-1)-fold product of 2-dimensional functions.

A belief net is a way to graphically represent a certain type of Markov property enabling a similar simplification to that given in equation 9.

It is important to distinguish belief nets from influence diagrams. The main emphasis in an influence diagram is on determining the optimal decision, whereas the emphasis in belief nets is on Bayesian inference. An influence diagram contains a decision node as well as chance nodes. The decision node represents the various decision alternatives available to the DM. Also there is very often a value node which represents the value or utility given the values of the other (random and non-random) variables in the problem.

The Markov property used in both belief nets and influence diagrams is expressed using the concept of conditional independence:

Suppose that *X*, *Y*, *Z* are vectors of random variables. We say that *X* and *Y* are *conditionally independent* given *Z*, if, after learning *Z*, no information given about *Y* would change our uncertainty about *X*. Or equivalently: if given any value z_0 of *Z* the joint probability density function of (*X*, *Y*, *Z*) factorizes as:

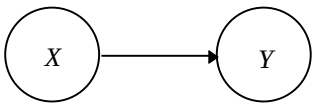

Figure 19: An extremely simple belief net

$$f_{(X,Y,Z)}(x,y,z_0) = g(x)h(y)\,,$$

where the functions *g* and *h* may depend on z_0.

Imagine for example the following events:

- The presence of a (big, hairy, dirty) rat
- We see the rat walking in the room
- A Greek[8] girl is not present.

[8] The only reason we mention Greek here, is that research has proved that Greeks hate rats significantly more than other Europeans. Unfortunately we can't go further into the underlying Freudian principles here.

These events can all be in the states true or false. Given that there is a rat present, the events that the Greek girl is not present and that we see the rat walking in the room might reasonably be considered to be independent of one another. However, they are not unconditionally independent, since if we see a rat walking in the room, the chance of a rat being present indeed increases (a lot) and as a consequence, the chance that the Greek girl has felt his presence too (and is therefore not around!) has considerably increased.

The arrows drawn between nodes represent qualitative influences which must be quantified by the model builder. This basically means assigning a probability distribution for the value a node can take, conditional on the value-distribution of the parent(s) node(s). It can be shown that these specifications uniquely determine a probability distribution *P* over the nodes if we assume the conditional independence Markov-property.

Belief nets are extremely useful for constructing a model. It does this in an intuitive and very straight-forward manner. When you have built your graphical model for example, it is actually possible to find the conditional independencies that are implied in the graph structure (by a process called *d-separation*) and by making sure if these are valid or not you are already doing some sort of validation.

When you have finished building your model it is quite easy to calculate the conditional probabilities that you are possibly looking for. Some standard graph operations like ‘arrow reversal’ and ‘node removal’ are provided for this matter. More information about this can be found in Smith (1985).

Another way of evaluating belief nets uses the method of junction trees. Since that is what is central to some applications in this report, we will focus a little bit more on that. The advantage of this way of representing the belief net is that we can propagate evidence through the network in a very efficient way. The first thing that should be mentioned is that there are (at least) two possible ways of so-called *clique-forming* (defined below) underlying this method. The methods of on the one hand Smith and on the other hand Lauritzen and Spiegelhalter differ slightly. Basically, Lauritzen and Spiegelhalter always use an intermediate stage to get their cliques, called a *decomposable graph*, whereas Smith tries to form cliques directly on basis of parent-children relations to see if this gives an appropriate clique-formation. However, if the belief net has a very simple structure in which children have only one parent, then both methods result in the same cliques and the transformation to the junction tree is then very straightforward indeed. Our belief net application in Bayes_RIMPUFF has this simple structure! However, we will show how to do the transformation for a slightly more complicated belief net; to get the idea over better. For the ideas behind transformation techniques for complicated belief nets we refer to e.g. Jensen (1996) and Smith (1995). From here on we will use the method of Smith.

A clique C_k is (according to Smith) a set consisting of a node x_k and its parents, fulfilling the condition that it is not a complete subset of another clique.

Now take a look at Figure 20:

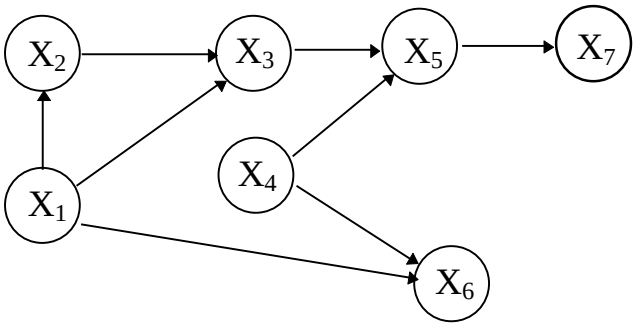

Figure 20: The original belief net

The cliques that can be found from this structure are:

$C_1 = \{x_1, x_2, x_3\}$

$C_2 = \{x_3, x_4, x_5\}$

$C_3 = \{x_1, x_4, x_6\}$

$C_4 = \{x_5, x_7\}$

The cliques $\{x_1\}$, $\{x_1, x_2\}$ and $\{x_4\}$ are omitted because they are subsets of other cliques.

The cliques constructed above now must be placed in a *junction graph*, but how do we connect the cliques? For that matter we define *separators* S_k as the set of nodes contained in $C_k \cap (C_1 \cup \ldots \cup C_{k-1})$, fulfilling the condition that they are not empty and not a duplicate set of another separator.

According to this definition the separators of Figure 20 are:

$S_1 = \{x_3\}$

$S_2 = \{x_1, x_4\}$

$S_3 = \{x_5\}$

This results in the following junction graph:

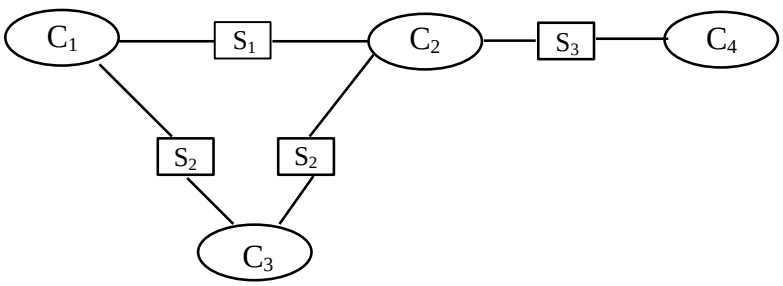

Figure 21: The junction graph

The updating algorithm that we used for Bayes_RIMPUFF is based on a junction *tree* (a graph without cycles) that fulfills the *running intersection property*, which essentially states that you can define some ordering of the cliques in which the nodes of a clique C_k that are also contained in previous cliques $\{C_1, C_2,\ldots,C_{k-1}\}$ are all members of *one* previous clique (the parent clique). The graph of our example does *not* fulfill this property (and it is not a tree either). That the running intersection property is not fulfilled can easily be verified by noticing that the separator S_2 is only part of the *union* of C_1 and C_2, and is not a subset of any of these cliques alone. We could have achieved this by some technical operations not further explained here. What is most important though is that the junction tree in Bayes_RIMPUFF *does* (automatically) fulfill this property (and is a tree). The running intersection property (combined with the tree property) ensures that when evidence arrives at a node in the graph it is only necessary to update probabilities in one clique which that node belongs to. These updates are propagated around the clique graph via the shared nodes, using a simple adjustment algorithm (see Smith, 1995).

In the case of Bayes_RIMPUFF the implication of all this is that we can partition the puffs into cliques such that:

1. observations update the estimates of mass of contaminant of puffs within a clique in a straightforward manner.
2. simple algorithms exist for propagating the information from the observations first into neighbouring cliques, then into more distant ones.

Bayesian Decision Analysis

Subjective probability and utility

So far we have been talking about Bayesian statistics. We need to note that Bayesian analysis extends in a straightforward manner to decision analysis; and, indeed, that is why the methodology has been adopted within

RODOS. Bayesian methods have developed over the last 40 years or so from the conceptual framework laid down by Savage (1954, 1972) into a subtle methodology designed to help DMs explore issues, communicate their concerns and uncertainties, investigate the balance of their conflicting for uncertain consequences and come to a decision on a course of action in the light of a greater and shared understanding of the problem before them (French, 1989). Briefly, Savage (1954) built a normative model of decision making based upon several axioms of rationality, i.e. principles of good decision making: see also French (1986; Smith, 1988). The model assumes that a DM faces a choice between strategies. The consequences of any chosen strategy will be determined by its interaction with an unknown *state* of the world: i.e.

$$\text{strategy} \oplus \text{state} \rightarrow \text{consequence}.$$

where the '$\oplus$' indicates the interaction (which is almost certainly non-additive!)

The problem facing a DM is that he wishes to construct a ranking of the actions which reflects his preferences between the consequences taking into account his beliefs about the unknown state. However, not only does he not know how to rank the possible actions, but also he may not have thought through his preferences and beliefs clearly. Decision analysis seeks to help a DM think through, explore and evolve his beliefs and preferences systematically. Bayesian decision analysis suggests that a DM's (final) ranking of the actions should be related to his beliefs and preferences through two functions.

- A *subjective probability distribution*, $P(.)$, which represents his beliefs about the unknown state of the world. The earlier sections of this tutorial have presented how this may be formed and updated in the light of data.
- A *utility function*, $u(.)$, which represents his preferences.

These obey the following three properties.

1. The subjective probability distribution represents his beliefs in the sense that:

$$P(s) > P(s')$$

if and only if, after due reflection he believes state s to be more likely to occur than state s'.

2. The utility function represents his preferences in the sense that:

$$u(x) > u(x')$$

if and only if, after due reflection he strictly prefers consequence x to consequence x'.

3. Bayesian decision analysis asserts that to combine his beliefs and preferences coherently in order to rank the actions the DM should form expected utilities:

$$Eu[a] = \int_S u(x(a,s)).dP(s)$$

where

S - is the set of possible states of the world;
$x(a,s)$- is the consequence which arises from taking action a when the state of the world is $s \in S$;
$dP(s)$- signifies that the expectation is taken with respect to the subjective probability distribution $P(.)$.

Then he should rank a_i above a_k if and only if:

$$Eu[a_i] > Eu[a_k].$$

Detailed expositions of the Bayesian paradigm for decision making can be found in French (1986) and Smith (1988): see also French and Smith (1997). It should be emphasised that $P(.)$ and $u(.)$ are not just any probability distribution nor any function which models a preference ranking. They need to be assessed in such a way as to ensure that they encode their DMs' judgements and attitudes to risk.

Note that the Bayesian paradigm separates two aspects of decision making (see Figure 22): $P(.)$ encodes knowledge and beliefs about what is happening or may happen; $u(.)$ encodes value judgements about the relative worth of different outcomes. This separation is particularly useful in building decision support systems: within RODOS much of the development of $P(.)$ is accomplished by the ASY modules, while the ESY modules focus on the development and use of $u(.)$.

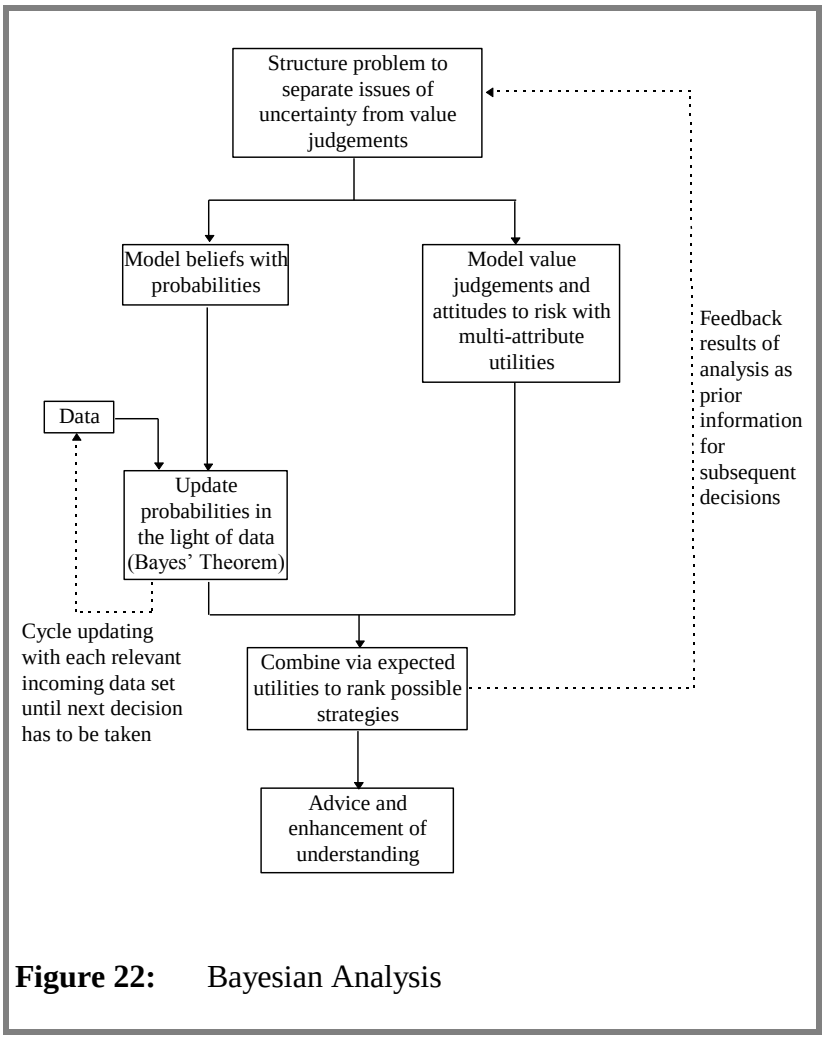

Figure 22: Bayesian Analysis

Prediction versus explanation

As scientists we have been educated to explain the past, to estimate parameters, and to develop and test models. Science progresses by gathering data in controlled conditions, exploring that data to discover any patterns and then using those patterns – or, as we call them, *models* – to summarise the data and increase our understanding of the world. That potted description of the scientific process, of course, begs a multitude of questions, but it should be sufficient to evoke in your mind the context that is needed for this discussion. Scientific modelling seeks to fit models to previously collected data sets and when it succeeds in fitting the trends, averages and (some of) the detailed wrinkles in those data in a satisfying way.

Decision support requires that we deploy these scientific models to predict the future and that requires a different emphasis and viewpoint. In simple terms, the future is seldom a perfect replication of the past and so a model that has been fitted to a previous data set is biased to specifics within those data. Hence it does not give as good a prediction of the future as one might hope. Obviously if the data set is very large and the model

extremely well tested, then the bias is likely to be negligibly small. In such cases the future will in all likelihood replicate aspects present in previously observed data. Thus there is little danger in using Newton's Laws of Motion to predict the dynamics of small objects moving a low speeds. The Laws have been very well verified against enormous numbers of data. But in most cases the data sets collected in the past may be of a similar size to the sets of events and quantities to be predicted and observed in the future. When this is the case, there is much evidence in the statistical literature of the dangers of over-fitting the past to predict the future. The advice is to use simple models in forecasting. Simple models are far less likely to be biased by a particular wrinkle in past data than complex models which can adjust through the choice of parameters to fit every nook and cranny. All this is not new advice, of course: Ockham[9] advised on the benefits of simplicity in the 14th century.

There is another reason to seek simplicity: modelling for prediction is quite simply a different task to modelling for explanation. It may require a quite different perspective. For instance, in nuclear emergency management, many talk of back-fitting or reconstructing the source term. This reflects a scientific interest to understand the scale of an accident and in a sense this is perfectly reasonable: to predict the spread of the contamination we need to have some knowledge of the source term. But the 'best' estimate of the source term may not give the 'best' prediction of the spread of health effects arising from the plume. It is not what goes up that matters so much as what comes down! Consider, for example, an atmospheric dispersion model which is good for identifying the spread of the plume and the timescales over which it spreads, but which tends to underestimate the air concentrations and gamma dose rates at ground level. If this model is used in emergency management, one would need over estimates of the source term to give realistic estimates of the doses downwind. Accurate estimation of the source term would not serve the emergency managers well *in this respect at least*.

Guide to the literature

Introductions to Bayesian statistics may be found in Berry (1996), Lee (1989), and Smith (1988). More advanced presentations and current surveys of the literature are available in Berger (1985), Bernardo and Smith (1994), Berry *et al* (1996), Gelman *et al* (1995), and O'Hagan (1994). There are many other texts: two classic texts are Box and Taio (1973) and DeGroot (1970). French and Smith (1997) provide a recent collection of readings on the practice of Bayesian analysis. Downloading

[9] "Entities must not be multiplied unnecessarily" William of Ockham (c. 1290-1349). He is quoted with this remark, but none of his writings contain this precise formulation. None the less, he did undoubtedly favour the use of as simple a model as possible: "Nature abhors the pomp of superfluous causes" quoted in A. Flew (1971) *An Introduction to Western Philosophy*. Thames and Hudson, Lond

the *First Bayes* software package (O'Hagan, 1996) and palying with that may also provide insight.

Acknowledgements

Some of the calculations and plots were performed using First Bayes (O'Hagan, 1996).

APPENDIX B : Kalman filtering equations

Introduction: The basic two stage normal model

All the following is based upon the *Bayesian two stage normal model:*

Stage 1 (Prior) $\boldsymbol{\theta} \sim \mathbf{N}(A\boldsymbol{\mu}, W)$ (1)

Stage 2 (Observation) $\mathbf{y}|\boldsymbol{\theta} \sim \mathbf{N}(B\boldsymbol{\theta}, V)$ (2)

Interpret this as: $\boldsymbol{\theta}$ is an unknown parameter vector[10], our prior knowledge of which is represented by a normal distribution with mean $A\boldsymbol{\mu}$ and covariance matrix W. We may observe $\mathbf{y}$, the distribution of which depends on $\boldsymbol{\theta}$. Specifically $\mathbf{y}$ is distributed normally with mean $B\boldsymbol{\theta}$ and covariance matrix V. Suppose now that we observe $\mathbf{y}$. The posterior distribution of $\boldsymbol{\theta}$ is:

$$\boldsymbol{\theta}|\mathbf{y} \sim \mathrm{N}(\boldsymbol{\mu}^*, W^*) \quad (3)$$

where

$$\boldsymbol{\mu}^* = (W^{-1}+B^{\mathrm{T}}V^{-1}B)^{-1}(W^{-1}A\boldsymbol{\mu}+B^{\mathrm{T}}V^{-1}\mathbf{y}) \quad (4)$$

$$W^* = (W^{-1}+B^{T}V^{1}B)^{-1} \quad (5)$$

Proof

Throughout the following remember that to form $p_{\boldsymbol{\theta}}(\boldsymbol{\theta}|\mathbf{y})$, all multiplicative factors which do not depend on $\boldsymbol{\theta}$ may be dropped. At the end of the manipulation the posterior may be formed by normalising the distribution to integrate to unity.

$$p_{\mathbf{y}}(\mathbf{y}|\boldsymbol{\theta}) = k_1 \, exp(-\tfrac{1}{2}(\mathbf{y}-B\boldsymbol{\theta})^{\mathrm{T}}V^{-1}(\mathbf{y}-B\boldsymbol{\theta}))$$

$$p_{\boldsymbol{\theta}}(\boldsymbol{\theta}) = k_2 \, exp\,(-\tfrac{1}{2}(\boldsymbol{\theta}-A\boldsymbol{\mu})^{\mathrm{T}}\mathrm{W}^{-1}(\boldsymbol{\theta}-A\boldsymbol{\mu}))$$

$$\Rightarrow p_{\boldsymbol{\theta}}(\boldsymbol{\theta}|\mathbf{y}) \underset{\boldsymbol{\theta}}{\propto} p_{\mathbf{y}}(\mathbf{y}|\boldsymbol{\theta}) \times p_{\boldsymbol{\theta}}(\boldsymbol{\theta})$$

$$\underset{\boldsymbol{\theta}}{\propto} \exp\,(-\tfrac{1}{2}(\boldsymbol{\theta}^{\mathrm{T}}(B^{\mathrm{T}}\mathrm{V}^{-1}B+W^{-1})\boldsymbol{\theta}-\mathbf{y}^{\mathrm{T}}V^{-1}B\boldsymbol{\theta} -$$

$$\boldsymbol{\mu}^{\mathrm{T}}A^{\mathrm{T}}W^{-1}\boldsymbol{\theta} - \boldsymbol{\theta}^{\mathrm{T}}B^{\mathrm{T}}V^{1}\mathbf{y}-\boldsymbol{\theta}^{\mathrm{T}}W^{1}A\boldsymbol{\mu}))$$

$$\underset{\boldsymbol{\theta}}{\propto} \exp\,(-\tfrac{1}{2}(\boldsymbol{\theta}-\boldsymbol{\mu}^*)^{\mathrm{T}}W^{*-1}(\boldsymbol{\theta}-\boldsymbol{\mu}^*))$$

[10] Note That we have moved from the single variable case of lemma 1 in Appendix A to a vector case.

To show the final step, introduce the identity matrix $I = W^*(W^*)^{-1}$ into, e.g., $\boldsymbol{\theta}^T(B^TV^{-1}\mathbf{y} + W^{-1}A\boldsymbol{\mu}) = \boldsymbol{\theta}^T I(B^TV^{-1}\mathbf{y} + W^{-1}A\boldsymbol{\mu})$**.** Thus the posterior is normal with mean $\boldsymbol{\mu}^*$ and covariance matrix W^* as claimed.

Next we note some alternative expressions for $\boldsymbol{\mu}^*$ and W^*.

$$\boldsymbol{\mu}^* = A\boldsymbol{\mu} + K(\mathbf{y} - BA\boldsymbol{\mu}) \qquad (6)$$

$$K = W^*B^TV^{-1} \qquad (7)$$

Standard Normal theory says that combining (1) and (2) gives

$$\left.\begin{array}{l}\boldsymbol{\theta} \sim N(A\boldsymbol{\mu}, W) \\ \mathbf{y}|\boldsymbol{\theta} \sim N(B\boldsymbol{\theta}, V)\end{array}\right\} \mathbf{y} \sim N(BA\boldsymbol{\mu}, V + BWB^T) \qquad (8)$$

So *a priori* our expectation for $\mathbf{y}$ is $BA\boldsymbol{\mu}$. Thus (6) corrects the prior expectation for $\boldsymbol{\theta}$**,** *viz.* $A\boldsymbol{\mu}$, with a factor that depends on the difference between the actual observation and its prior expectation, *viz.* $BA\boldsymbol{\mu}$.

K is known as the *Kalman filter*

Proof of (6)

substituting for *K* in (6)

$$\begin{aligned}\boldsymbol{\mu}^* &= A\boldsymbol{\mu} + W^*B^TV^{-1}(\mathbf{y} - BA\boldsymbol{\mu}) \\ &= (\mathrm{I} - (W^{-1}+B^TV^{-1}B)^{-1}B^TV^{-1}\mathrm{B})A\boldsymbol{\mu} + (W^{-1}+B^TV^{-1}B)^{-1}B^TV^{-1}\mathbf{y} \\ &= (W^{-1}+B^TV^{-1}B)^{-1}((W^{-1}+B^TV^{-1}B-\mathrm{B}^TV^{-1}B)A\boldsymbol{\mu} + B^TV^{-1}\mathbf{y}) \\ &= (W^{-1}+B^TV^{-1}B)^{-1}(\mathrm{W}^{-1}A\boldsymbol{\mu}+B^TV^{-1}\mathbf{y}) \qquad \text{as in (4)}\end{aligned}$$

The dimensions of the matrices

Note that the dimension of *V* is the dimension of $\mathbf{y}$, the observation. Whereas the dimension of *W* is the dimension of $\boldsymbol{\theta}$, the parameter. Expression (5) implies that to form W^* and hence *K* we have to invert a matrix of the dimension of the number of parameters. However there is an alternative expression for W^*:

$$W^* = (W^{-1}+B^TV^{-1}B)^{-1} = W - WB^T(V+BWB^T)^{-1}BW \qquad (9)$$

Proof

$$\begin{aligned}&(W^{-1} + B^TV^{-1}B)(W - WB^T(V+BWB^T)^{-1}BW) \\ =\ & I - B^T(V+BWB^T)^{-1}BW - B^TV^{-1}BWB^T(V+BWB^T)^{-1}BW + B^TV^{-1}BW \\ =\ & I - B^TV^{-1}(V+BWB^T)(V+BWB^T)^{-1}BW + B^TV^{-1}BW \\ =\ & I\end{aligned}$$

So expression (9) may replace expression (5). Note that the dimension of the matrix being inverted in the right hand side of (9) is the dimension of the observation **y**.

Thus, if we form $\boldsymbol{\mu}^*$, W^* using (4) and (5), *three* matrix inversions are required: namely W^{-1},V^{-1} and $(W^{-1}+B^TV^{-1}B)^{-1}$. Two of these inversions involve the dimension of the parameter, $\boldsymbol{\theta}$.

If we form $\boldsymbol{\mu}^*$, W^* using (6), (7) and (9) *two* matrix inversions are needed, both of dimension of the observation **y**.

There is yet a further equivalent recursion which involves a single matrix inversion of dimension of the observation **y**.

From (7) and (9)

$$K = W^*B^TV^{-1}$$
$$= (W - WB^T(V+BWB^T)^{-1}BW)B^TV^{-1}$$
$$= WB^TV^{-1} - WB^T(V + BWB^T)^{-1}(V + BWB^T)V^{-1} + WB^T(V + BWB^T)^{-1}$$

Thus

$$K = WB^T(V + BWB^T)^{-1} \quad (10)$$

There is now only one matrix inversion, *viz*. $(V + BWB^T)^{-1}$.

Finally note that from (9)

$$W^* = W - WB^T(V + BWB^T)^{-1}BW$$

So substituting (10)

$$W^* = (I - KB)\, W \quad (11)$$

In summary to form the posterior of a two-stage normal model:

i) calculate $(V+BWB^T)^{-1}$; (12a)

ii) form K using expression (10); (12b)

iii) form W^* using expression (11) (12c)

iv) form $\boldsymbol{\mu}^*$ using expression (6) (12d)

This assumes that the dimension of the observations **y** is (substantially) less than that of the parameters $\boldsymbol{\theta}$. Given that for independent observation errors, it is possible to partition the observations into small sets and take the sets in sequentially, this can usually be made the case. If the dimension of the parameter space, however, is small then *any* of the equivalent sets of recursions may be used.

Two building block modules.

To implement the algorithms below two generic modules are required, at least in the prototyping phase. For the final implementation there is likely

to be considerable advantage in writing specific versions of the modules to take advantage of sparseness patterns in specific matrices.

The first module simply convolutes two normal distributions:

$\boldsymbol{\eta} \sim \mathrm{N}(\mathbf{m}, R)$

$\mathbf{x}|\boldsymbol{\eta} \sim \mathrm{N}(E\boldsymbol{\eta}, S)$

$\Rightarrow \mathbf{x} \sim \mathrm{N}(E\mathbf{m}, ERE^{\mathrm{T}} + S)$

Thus we need

`normal_convolute`

`inputs:` $\mathbf{m}, R, E, S$

`outputs:` $\boldsymbol{\gamma}, T$

`defined by:` $\boldsymbol{\gamma} = E\,\mathbf{m}$

$T = ERE^{\mathrm{T}} + S$

Sometimes we will need a variant on this module

Suppose $\boldsymbol{\eta} \sim \mathrm{N}(\mathbf{m}, R)$

$\mathbf{x}|\boldsymbol{\eta} \sim \mathrm{N}(E\boldsymbol{\eta}+\boldsymbol{\upsilon}, S)$

$\Rightarrow \mathbf{x} \sim \mathrm{N}(E\mathbf{m}+\boldsymbol{\upsilon}, ERE^{\mathrm{T}} + S)$

Thus we need a modified

`normal_convolute_shift`

`inputs:` $\mathbf{m}, \boldsymbol{\upsilon}, R, E, S$

`outputs:` $\boldsymbol{\gamma}, T$

`defined by:` $\boldsymbol{\gamma} = E\,\mathbf{m} + \boldsymbol{\upsilon}$

$T = ERE^{\mathrm{T}} + S$

Secondly, we need a module to update a normal two stage model (1) and (2) by a Kalman filter.

`Kalman_update`

`inputs:` $A, \boldsymbol{\mu}, W, B, V$

`outputs:` $\boldsymbol{\mu}^*$, W*

`defined by: expressions (12a-d) or some equivalent equations.`

Document History

Document Title:	Uncertainty in RODOS
RODOS number:	RODOS(B)-RP(94)-05
Version and status:	Version 2.1 (draft)
Authors/Editors:	G. Caminada, S. French, K. Politis and J.Q. Smith
Address:	School of Informatics, Coupland 1, University of Manchester, Manchester M13 9PL
Issued by:	Simon French
History:	Version 1.0 Final: Issued: January 1995 Version 2.0 Draft: Issued: October 1997 Version 2.1 Draft: Issued: January 2000
Date of Issue:	January, 2000
Circulation:	Authors and others in RODOS project for comment.
File Name:	Uncertainty_in_RODOS
Date of print:	December 1, 2011