Aberystwyth University

Intelligent systems for decision support

Shen, Qiang

Published in:
Proceedings of International Joint Conference on Computational Intelligence

Publication date:
2009

Citation for published version (APA):
Shen, Q. (2009). Intelligent systems for decision support. In Proceedings of International Joint Conference on Computational Intelligence (pp. 25-36)

INTELLIGENT SYSTEMS FOR DECISION SUPPORT

Qiang Shen

Department of Computer Science, Aberystwyth University, Aberystwyth SY23 3DB, UK.
qqs@aber.ac.uk

Abstract: Decision support systems play an important role in many application domains. For instance, in the detection of serious crime, including terrorist activity, an intelligent system which is capable of automated modelling and analysis of intelligence data is of great importance. It will provide useful decision support for intelligence analysts, offering an effective means in the assessment of scenarios for possible crimes. Such systems will also facilitate rapid response in devising and deploying preventive measures. This paper describes the important challenges which face the development of intelligent decision support systems, with a focus on the problem of crime detection and prevention. It presents some recent advances in computational intelligence in general, and in fuzzy systems in particular. These advances contribute to the accomplishment of tasks essential for intelligence data monitoring (amongst other applications).

1 INTRODUCTION

Decision support systems are a specific type of computerised information system that support decision-making activities (Keen, 1978; Marakas, 1999). They are intended to assist decision makers in compiling useful information from raw data, user knowledge, and domain models in order to identify and solve problems and make decisions. As such, studies of decision support and decision support systems naturally belong to an environment with multidisciplinary foundations, including (but not exclusively) database and operations research, artificial and computational intelligence, human-computer interaction, modelling and simulation, and software engineering. In particular, the research and development of techniques that enable the construction and performance of selected cognitive decision-making functions form a key to the successful application of decision support systems.

Solving complex real-world problems often requires timely and intelligent decision-making processes, through analysis of a large volume of available information. For example, in the wake of terrorist atrocities such as September 11, 2001, and July 7,

2005, intelligence experts have commented that the failure in the detection of terrorist activity is not necessarily due to lack of data, but to difficulty in relating and interpreting the available intelligence on time. Therefore, an important and emerging area of research is the development of decision support systems that will help to establish so-called situational awareness: a deeper understanding of how the available data is related and whether or not it represents a threat.

Most criminal and terrorist organisations form flexible networks of loosely related individuals and sub-organisations. These networks are often embedded within legitimate society and remain secrete. However, organised crime and terrorist activity does leave a trail of information, such as captured communications and forensic evidence, which can be collected by police and security organisations. Whilst experienced intelligence analysts can suggest plausible scenarios, the prompt identification of potential organisations that may pose a threat, the amount of intelligence data possibly relevant may well be overwhelming for human examination. Hypothetical (re-)construction of the activities that may have generated the intelligence data obtained, therefore, presents

an important and challenging research topic for crime prevention and detection.

There have been many intelligent systems proposed in the literature which provide helpful information that may enhance efforts for crime reduction. However, their effectiveness is typically and crucially dependent upon the experience of the user/analysts since any potential threat posed by uncovered networks are identified ultimately only by the analysts. This vulnerability can be addressed by a system which automatically generates plausible scenarios when given a limited amount of real or hypothesised evidence, and which provides the user with the means to analyse such scenarios. This paper introduces a knowledge-based framework for the development of such systems, to assist (but not to replace) intelligence analysts in identifying plausible scenarios of criminal or terrorist activity, and in assessing the reliability, risk and urgency of generated hypotheses. In particular, it presents some recent work in exploiting computational intelligence techniques to build intelligent decision support systems for monitoring and interpreting intelligence data.

The rest of this paper is organised as follows. The next section presents the underlying approach adopted and describes the essential components of such a decision support system. Then, Section 3 shows particular instantiations of the techniques used to implement the key components of this framework. Essential ideas are illustrated with some simple examples. Section 4 summarises the paper and points out important further research.

2 PLAUSIBLE SCENARIO-BASED APPROACH

In order to devise a robust monitoring system that is capable of identifying many variations on a given type of terrorist activity, this work employs a model-based approach to scenario generation (Shen et al., 2006). That is, the knowledge base of such a monitoring system consists of generic and reusable component parts of plausible scenarios, called model or scenario fragments (interchangeably). Such fragments include: types of (human and material) resources required for certain classes of organised terrorist activity, ways in which such resources can be acquired and organised, and forms of evidence that may be generated (and hence acquired from intelligence databases) given certain scenarios.

Note that conventional knowledge-based systems, such as rule-based and case-based reasoners, have useful applications in the crime detection area. How-
ever, their scope is restricted to either the situations foreseen or those resulting from previously encountered cases. Yet, organised terrorist activity tends to be unique, whilst employing a relatively restricted set of methods (e.g. suicide bombing or bomb threats in public places). A model-based reasoner designed to (re-)construct likely scenarios from available evidence, as combinations of instantiated scenario fragments, seems to be ideally suited to cope with the variety of scenarios that may be encountered. Indeed, the main strength of model-based reasoning is its adaptability to scenarios that are previously unseen (Lee, 1999).

Figure 1 shows the general architecture of the approach taken in this research. Based on intelligence data gathered, the scenario generation mechanism instantiates and retrieves any relevant model fragments from the library of generic scenario fragments, and combines such fragments to form plausible explanations for the data. A description of how such intelligent decision support systems are built is given below.

2.1 Flexible Composition Modelling

As indicated above, the central idea is to establish an inference mechanism that can instantiate and then dynamically compose generic model fragments into scenario descriptions, which are plausible and may explain the available data (or evidence). A compositional modelling approach (Falkenhainer and Forbus, 1991; Keppens and Shen, 2001) will be devised for this purpose. The main potential of using this approach over conventional techniques is its ability to automatically construct many variations of a given type of scenario from a relatively small knowledge base, by combining reusable model fragments on the fly. This ensures the robustness required by the problem domain.

Essentially, a compositional modeller stores hypothetical scenarios in a hypergraph or scenario space, representing states and events within different scenarios as nodes and the (causal) relations between such nodes as directed hyperarcs. A significantly extended version of the assumption-based truth maintenance system (ATMS) (de Kleer, 1986) is developed to maintain the scenario space and to enable retrieval of partial scenarios and other useful information (e.g. potential evidence to help validate a scenario) from it efficiently.

The compositional modelling approach developed in this research differs from those in the literature in a number of ways:

1. *Scalability* to suit the requirement of generating and storing a space of plausible scenarios. This
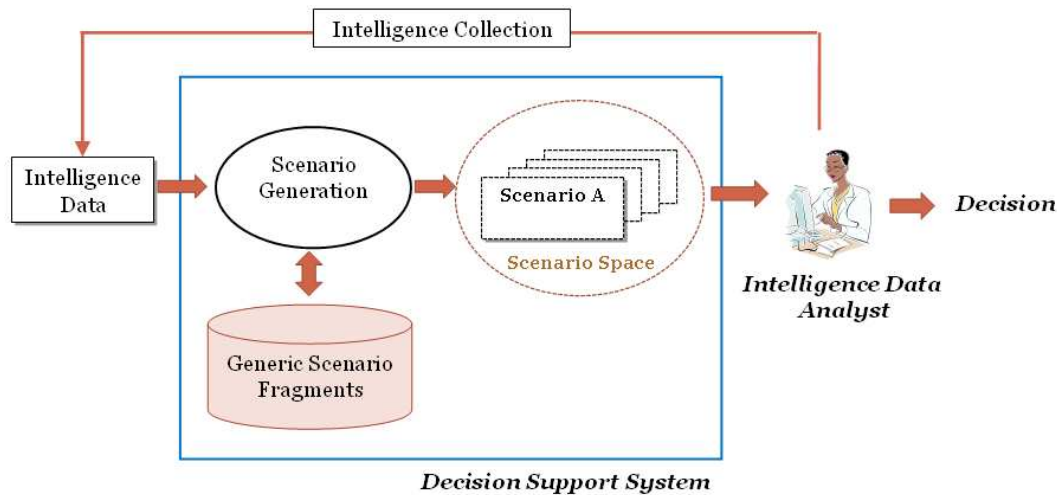
Figure 1: Architecture of Intelligent Systems for Intelligence Data Analysis

is needed because of the potential large number and variation of terrorist activity or other types of serious crime.

2. *Ability* to speculate about plausible relations between different cases. Often, intelligence information will refer to individuals and objects whose identity is only partially specified. For example, when a person is observed on a CCTV camera, some identifying information can be collected, but this may be insufficient for an exact identification. When a person with similar features has been identified elsewhere, it is important that any relation between both sightings is explored. Ideas originally developed in the area of link-based similarity analysis (Calado et al., 2006; Liben-Nowell and Kleinberg, 2007) are adapted herein for: (a) identifying similar individuals and objects in a space of plausible scenarios, and (b) supporting the generation of hypothetically combined scenarios to explore the implications of plausible matches.

3. *Coverage* to generate scenarios from a wide range of data sources, including factual data, collected intelligence, and hypothesised but unsubstantiated information. This requires matching specific data (e.g. the names of discovered chemicals) with broader (and possibly subjective) knowledge and other vague information contents. Such knowledge and information may be abstractly specified in the knowledge base, e.g. a chemical being "highly explosive". Similarly, matching attributes of partially identified objects and individuals may involve comparing vague features, such as a person's apparent height, race and age. This suggests

the use of a formal mathematical theory that is capable of capturing and representing ill-defined linguistic terms, which are common in expressing and inferring from intelligence knowledge and data. Fuzzy representation and inference methods are therefore introduced to compositional modelling, for the first time, to decide on the applicability of scenario fragments and their compositions.

2.2 Plausible Scenario-Based Intelligence Monitoring

Monitoring intelligence data for evidence of potential serious criminal activity, especially terrorist activity, is a non-trivial task. It is not known in advance what aspects of such activity will be observed, and how they will be interconnected. There are nevertheless, many different ways in which a particular type of activity may be arranged. Hence, conventional approaches to monitoring, which aim to identify pre-specified patterns of data, are difficult to adapt to this domain.

Although general and potentially suitable, the model-based approach adopted here may lead to systems that generate a large number of plausible scenarios for a given problem. It is therefore necessary for such a system to incorporate a means to sort the plausible scenarios, so that the generated information remains manageable within a certain time frame. For this purpose, generated scenarios are presented to human analysts with measurements of their reliability, risk, and urgency. The reliability of a scenario estimates the likelihood of its occurrence. The risk posed

by a scenario reflects the number of potential casualties and/or the degree of possible economic cost of failure to prevent such activity. A scenario's urgency corresponds to the time in which the suspect terrorist activity may be carried out.

Each of the aforementioned ranking features may be assessed by a numeric metric. However, intelligence data and hypotheses are normally too vague to produce precise estimates that are also accurate. Therefore, this work devises a novel fuzzy mechanism to provide an appropriate method of assessing and presenting the reliability, risk and urgency of generated scenarios. The framework also covers additional tools such as a facility to propose additional information sources (by exploring additional, real or hypothesised, evidence that may be generated in a given scenario).

Figure 2 shows a specification of the general framework given in Fig. 1. Technical modules include:

- Fuzzy Feature Selection carries out semantics-preserving dimensionality reduction (over nominal and real-valued data).

- Fuzzy Learning provides a knowledge modelling mechanism to generalise data with uncertain and vague information into mode fragments.

- Fuzzy Iterative Inference offers a combination of abductive and deductive inferences, capable of reasoning with uncertain assumptions.

- Flexible CSP (constraint satisfaction problem-solver) deals with uncertain and imprecise constraint satisfaction, subject to preference and priority.

- Fuzzy Interpolative Reasoning enables approximate inference over sparse knowledge base, using linear interpolation.

- Flexible ATMS is an extended truth-maintenance system that keeps track of uncertain assumption-based deduction.

- Flexible Coreference Resolution implements a link-based identity resolution approach, working with real, nominal, and order-of-magnitude valued attributes.

- Fuzzy Aggregation performs information aggregation by combining uncertain attributes as well as their values.

- Fuzzy Evidence Evaluation performs evidence assessment, including discovery of misleading information, and generates evidence-gathering proposal.

- Fuzzy Risk Assessment computes potential loss-oriented risk evaluation through fuzzy random process modelling.

This research focusses on the use of structured knowledge for monitoring intelligence data which may contain factual evidence and assumed information. Such data is assumed to be associated with standard formats available to intelligence analysts, even though they may well involve vague or even ill-defined concepts. An investigation of how model fragments may be learned from (typically high-dimensional) data is essential. Such data may be presented in a pre-specified form, e.g. in terms of properties of suspects, types of incident and classes of evidence.

Note that recent advances in semantic web research (Berners-Lee et al., 2001; Shadbolt et al., 2006) provide useful techniques for pre-processing raw data to reveal the content of source information. Use of such techniques will help to automate the pre-processing of raw information, which is currently carried out in an *ad hoc* manner, but remains outside the scope of this paper.

Systems built following the approach outlined in Fig. 2 can help to improve the likelihood of discovering any potential threat posed by criminal or terrorist organisations. In particular, the use of an automated intelligent monitoring system, whose reasoning is logical and readily interpretable by human analysts, can be very helpful in supporting human analysts when working under time constraints. For instance, this may aid in avoiding premature commitment to certain seemingly more likely but unreal scenarios, minimising the risk of producing incorrect interpretations of intelligence data. This may be of particular interest to support staff investigating cases with unfamiliar evidence. In addition, the resulting approach may be adapted to build systems that facilitate training of new intelligence analysts. This is possible because the underlying inference mechanism and the knowledge base built for intelligence data monitoring can be used to artificially synthesise various scenarios (of whatever likelihood), and to systematically examine the implications of acquiring different types of evidence.

3 ILLUSTRATIVE COMPONENT APPROACHES

As a knowledge-based approach to building decision support systems, any implementation of the framework proposed above will require a knowledge base to begin with. The first part of this section will
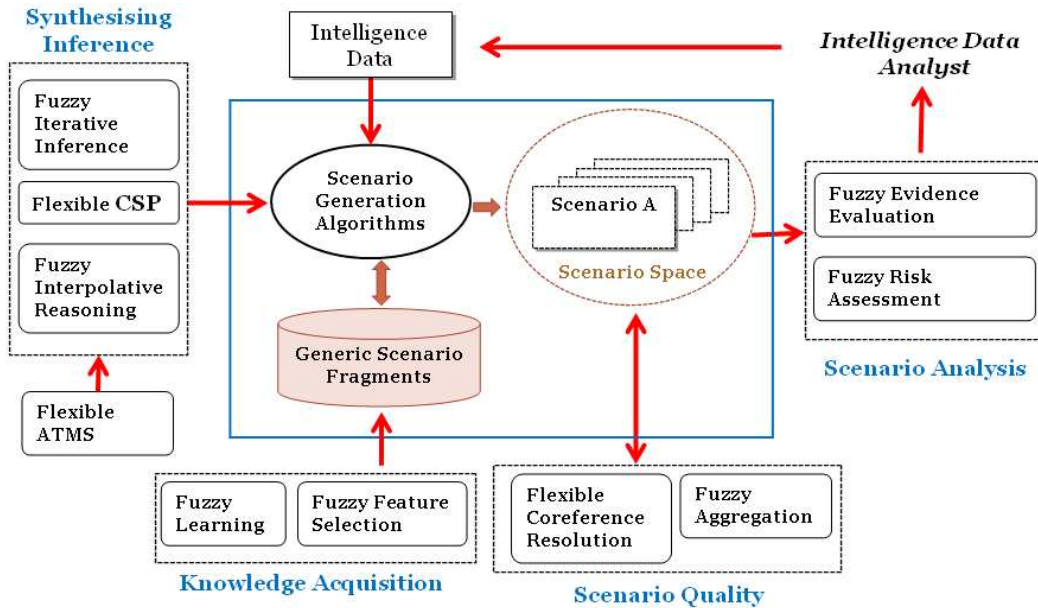
Figure 2: Instantiated Architecture

then introduce a number of recent advances in developing data-driven learning techniques that are suitable to derive such required knowledge from potentially very complex data. The second part will describe one of the key techniques that support scenario composition, especially for situations where limited domain knowledge is available. The third and final part of the section will demonstrate how risks of generated scenarios may be estimated. Figure 3 outlines a simplified version of the framework which may be implemented using the techniques described herein.

All of these approaches have been developed using computational intelligence techniques in general and fuzzy systems methods in particular. Introduction to these techniques will be explained at conceptual level with illustrative examples. Mathematical and computational details are omitted, but readers may find them from the relevant references.

3.1 Fuzzy Learning and Feature Selection

In general, an initial knowledge base of generic scenario fragments is built partly by generalising historical intelligence data through computer-based induction, but partly through manual analysis of past terrorist or criminal activity. This work focusses on the automated induction of model fragments. Also, in a real world setting, data may come from multiple sources and hence, may require a substantial amount of pre-

processing in order to achieve semantic interpretability. However, this consideration is beyond the scope of this paper. Any data given for learning is assumed to have been presented in a homogeneous format.

3.1.1 Fuzzy Descriptive Learning

Many real-world problems require the development and application of algorithms that automatically generate human interpretable knowledge from historical data. Such a task is clearly not just for learning model fragments.

Most of the methods for fuzzy rule induction from data have followed the so-called precise approach. Interpretability is often sacrificed, in exchange for a perceived increase in precision. In many cases, the definitions of the fuzzy sets that are intended to capture certain vague concepts are allowed to be modified such that they fit the data better. This modification comes at the cost of ruining the original meaning of the fuzzy sets and the loss of transparency of the resulting model. In other cases the algorithms themselves generate the fuzzy sets, and present them to the user. The user must then interpret these sets and the rules which employ them (e.g. a rule like: If *volume* is Tri(32.41, 38.12, 49.18), then *chance* is Tri(0.22, 0.45, 0.78), which may be learned from data presented in Fig. 4). Furthermore, in some extreme cases, each rule may have its own fuzzy set definition for every condition, thereby generating many different sets in a modest rule base. The greatest disadvan-
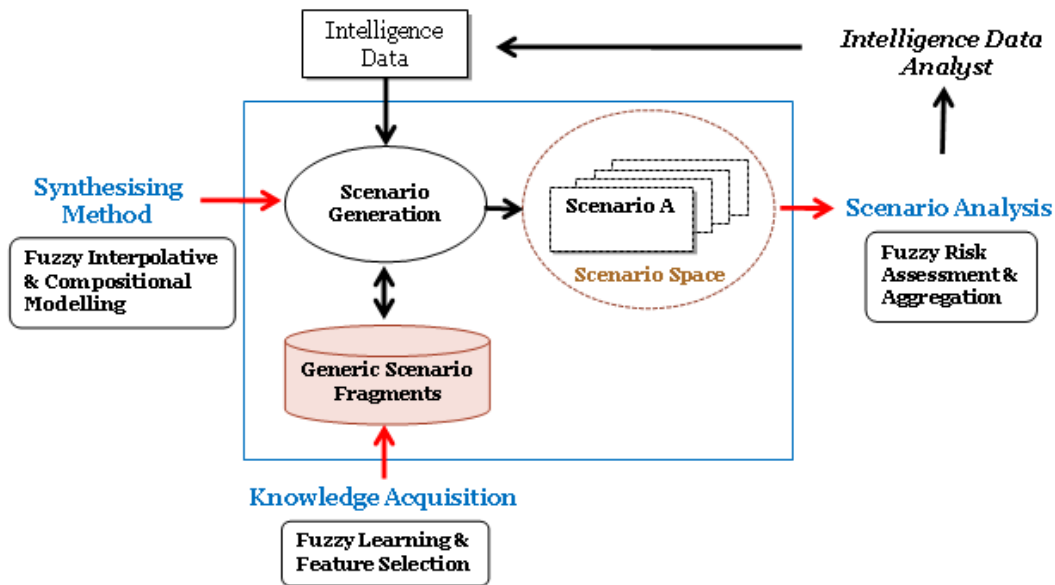
Figure 3: Focussed Illustration

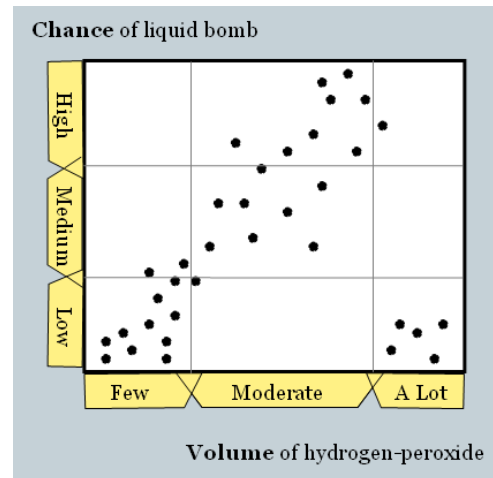Figure 4: Precise Modelling

Figure 5: Descriptive Modelling

tage of the precise approach is that the resulting sets and rules are difficult to match with human interpretation of the relevant concepts.

As an alternative to the precise approach, there exist proposals that follow the descriptive (or linguistic) approach. In such work no changes are made to human defined fuzzy sets. The rules must use the (fuzzy) words provided by the user without modifying them in any way. One of the main difficulties with this type of approach is that the possible rules available are predetermined, equivalently speaking. This is be-

cause the fuzzy sets can not be modified, and only a small number of them are typically available. Although there can be many of these rules they are not very flexible and in many cases they may not necessarily fit the data well (e.g. a rule like: If *volume* is Moderate, then *chance* is High, which may be learned from data and predefined fuzzy sets given in Fig. 5). In order to address this problem, or at least partially, linguistic hedges (aka. fuzzy quantifiers) can be employed.

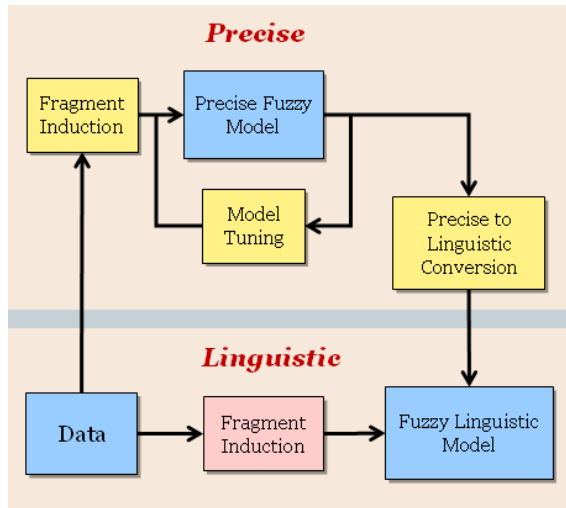The concept of hedges has been proposed quite early on in fuzzy systems research (Zadeh, 1975). A

Figure 6: Two-Step Learning of Descriptive Models

Figure 7: Feature Selection Process

linguistic hedge produces a new fuzzy set by changing the original fuzzy set, in a fixed and interpretable manner. The interpretation of the resultant set emanates from the original fuzzy set and a specific transformation that the hedge implies. In so doing, the original fuzzy sets are not changed, but the hedged fuzzy sets provide modifiable means of modelling a given problem and therefore, more freedom in representing knowledge in the domain.

This research adopts the seminal work of (Marín-Blázquez and Shen, 2002) which champions this approach. As shown in Fig. 6, this technique produces descriptive fuzzy system models with a two-step mechanism. The first is to use a precise method to create accurate rules and the second to convert the resulting precise rules to descriptive ones. The conversion is, in general, one-to-many. It is implemented by using a heuristic algorithm that derives potentially useful translations and then, by employing evolutionary computation to perform a fine tuning of these translations. Both steps are computationally efficient. The resultant descriptive model is ready to be directly applied for inference; no precise rules are needed in runtime.

Note that Fig. 6 shows the learning of a "model" in a general sense. Such a model may be a set of conventional production fuzzy if-then rules, or one or more generic model fragments which involve not only standard conditions but also assumptions or hypotheses that must be made in order to draw conclusions.
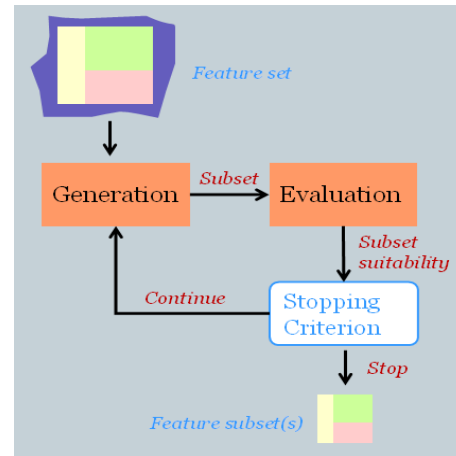
3.1.2   Fuzzy-Rough Feature Selection

Feature selection (Liu and Motoda, 1998; Jensen and Shen, 2008) addresses the problem of selecting those characteristic descriptors of a domain that are most informative. Figure 7 shows the basic procedures involved in a feature selection process. It is a problem encountered in many areas of computational intelligence. Unlike other dimensionality-reduction methods, feature selectors preserve the original meaning of the features after reduction. This has been applied to perform tasks that involve datasets containing huge numbers of features (in the order of tens of thousands) which, for some learning algorithms, may be otherwise impossible to process further.

There are often many features involved in intelligence data, and combinatorially large numbers of feature combinations, to select from. It might be expected that the inclusion of an increasing number of features would increase the likelihood of including enough information to distinguish between classes. Unfortunately, this is not necessarily true if the size of the training dataset does not also increase rapidly with each additional feature included. A high-dimensional dataset increases the chances that a learning algorithm will find spurious patterns that are not valid in general. More features may introduce more measurement noise and, hence, reduce model accuracy (Jensen and Shen, 2009a).

Recently, there have been significant advances in developing methodologies that are capable of minimising feature subsets in an imprecise and uncertain environment. In particular, a resounding amount of research currently being done utilises fuzzy and rough sets (Jensen and Shen, 2004; MacParthalain and Shen, 2009; MacParthalain et al., 2009; Shen

Figure 8: Fuzzy-Rough Feature Selection

Figure 9: Spare Knowledge Base

and Jensen, 2004; Tsang et al., 2008). The success of rough set theory is due in part to the following two aspects: (a) only the facts hidden in data are analyzed, and (b) no additional information about the data is required, such as thresholds or expert knowledge on a particular domain. However, it handles only one type of imperfection found in data, it is complementary to other concepts for this purpose, e.g. fuzzy set theory. The two fields may be considered analogous in the sense that both can tolerate inconsistency and uncertainty. The difference rests in the type of uncertainty and their approach to it; fuzzy sets are concerned with vagueness, and rough sets are concerned with indiscernibility. Therefore, it is desirable to extend and hybridise the underlying concepts to deal with additional aspects of data imperfection.

Fuzzy-rough feature selection (Jensen and Shen, 2007; Jensen and Shen, 2009b) provides a means by which discrete or real-valued noisy data (or a mixture of both) can be effectively reduced without the need for user-supplied information. Additionally, this technique can be applied to data with continuous or nominal decision attributes, and as such is suitable for the nature of intelligence data. A particular implementation is done via hill-climbing search, as shown in Fig. 8. It employs the fuzzy-rough dependency function, which is derived from the notion of fuzzy lower approximation, to choose those attributes that add to the current candidate feature subset in a best-first fashion. The algorithm terminates when the addition of any remaining attribute does not result in an increase in the dependency.

Note that as the measure of fuzzy-rough dependency degree is nonmonotonic, it is possible that the hill-climbing search terminates having reached only a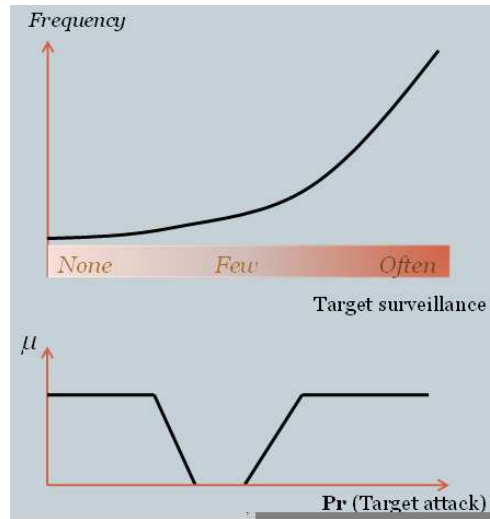 local optimum. The global optimum may lie else-where in the search space. Nevertheless, it is one of the most developed methods for approximate feature selection and hence is adopted in this research for semantics-preserving dimensionality reduction.

3.2 Fuzzy Interpolative Reasoning

In conventional approaches to compositional modelling, the completeness of a scenario space depends upon two factors: (a) the knowledge base must cover all essential scenario fragments relevant to the data, and (b) the inference mechanism must be able to synthesise and store all combinations of instances of such fragments that constitute a consistent scenario. However, in the real-world, especially for the problem domain concerned here, it is difficult, if not impossible, to obtain a complete library of model fragments. Figure 9 shows an example, where the following two simplified model fragments (i.e. two simple if-then rules in this case) are given:

$Rule_i$: If *frequency* is None then *attack* is No
$Rule_j$: If *frequency* is Often then *attack* is Yes

Then, with an observation that states "*frequency* is Few", no answer can be found to the question of "Will there be an attack"? A popular tool to deal with this type of problem is fuzzy interpretative reasoning (Baranyi et al., 2004; Tikk and Baranyi, 2000). In this work, the transformation-based approach as proposed in (Huang and Shen, 2006; Huang and Shen, 2008) is employed to support model composition, when given an initial sparse knowledge base.

The need for a fuzzy approach to interpolation is due to the fact that the precision degree of the available intelligence data can be variable. Finding a

match between the given data and the (already sparse) knowledge base cannot in general be achieved precisely. The potential sources of variability in precision include vaguely defined concepts (e.g. materials that constitute a "high explosive", certain organisations that are deemed "extremist"), quantities (e.g. a "substantial" amount of explosives, "many" people) and specifications of importance and certainty (e.g. in order to deploy a radiological dispersal device, the perpetrator "must" have access to "radioactive material and "should" have an ideological or financial incentive). Therefore, the approach adopted must be able to describe and reason with knowledge and data at varying degrees of precision. Fuzzy interpolation works well in this regard.

Figure 10 illustrates the basic ideas of fuzzy interpolative reasoning. It works through a two-step process: (a) computationally constructing a new inference rule (or model fragment in the present context) via manipulating two given adjacent rules (or related fragments), and (b) using scale and move transformations to convert the intermediate inference results into the final derived conclusions.

3.3 Fuzzy Risk Assessment

In developing intelligent decision support systems for intelligence data monitoring, there is often a trade-off that must be considered. That is, between the completeness of the scenario space generated and the potential efficiency in subsequent examination of the resultant space. On the one hand, it is important not to miss out any potentially significant scenarios that may explain the observed evidence. On the other hand, too many unsorted and especially, spurious scenarios produced may confuse human analysts. Thus, for intelligence data modelling and analysis, it is desirable to be able to filter the created scenario space with respect to certain objective measures of the quality of the generated scenario descriptions. Fortunately, as indicated previously, preferences over different hypothetical scenarios can be determined on the basis of reliability, risk and urgency of each scenario.

The *reliability* of a generated scenario may be affected by several distinct factors: the given intelligence data (e.g. the reliability of an informant), the inferences made to abduce plausible scenarios (e.g. the probability that a given money transfer is part of an illegitimate transaction), and the default assumptions adopted (e.g. the likelihood that a person seen on CCTV footage is identified positively). The *urgency* of a scenario is inversely proportional to the expected time to completion of a particular terrorist/criminal activity. Therefore, an assessment of ur-

gency requires a (partial) scenario to be described using the scenario's possible consequences and information on additional actions required to achieve completion. The *risk* posed by a particular scenario is determined by its potential consequences (e.g. damage to people and property). Whilst these are very different aspects that may be used to differentiate and prioritise scenarios composed by the compositional modeller, the underlying approaches to assess them are very similar. Thus, in this paper, only the scenario risk aspect is discussed.

Risk assessment helps to efficiently devise and deploy counter measures, including further evidence gathering of any threat posed by the scenario concerned. However, estimating the risk of a plausible event requires consideration variables exhibiting both randomness and fuzziness, due to the inherent nature of intelligence data (and knowledge also). Having identified this, in the present work, risk is estimated as the mean chance of a fuzzy random event (Halliwell and Shen, 2009; Shen et al., 2008) over a pre-defined confidence level, for an individual type of loss. In particular, plausible occurrence of an event is considered random, while the potential loss due to such an event is expressed as values of a fuzzy random variable (as it is typically judged linguistically). In implementation, loss caused by an event is modelled by a function mapping from a boolean sample space of {Success, Failure} onto a set of nonnegative fuzzy variables. Here, success or failure is judged from the criminal's viewpoint, of course, in terms of whether they have carried out a certain activity or not.

Risks estimated over different types of loss (e.g. range of geometric destruction and number of casualties) can be aggregated. Also, assessments obtained using different criteria (e.g. resource and urgency) may be integrated to form an overall situation risk. To generalise this approach further, order-of-magnitude representation (Parsons, 2003; Raiman, 1991) may be introduced to describe various cost estimations. Figure 11 shows such an example.

Incidentally, expert intelligence analysts have commented that the estimation in this figure matches the real dataset used in the corresponding set of terrorist attacks. However, they also commented that in general, it is not necessarily always the case that failure in carrying out a certain terrorist attack (or the successful prevention of plausible terrorist attack from the counter-terrorism perspective) would incur lower cost to the public. The ability of a decision support system to correctly capture such cases clearly deserves more through investigation in future. A possible approach to addressing this issue is to utilise the measures of risk, urgency and reliability as flexible
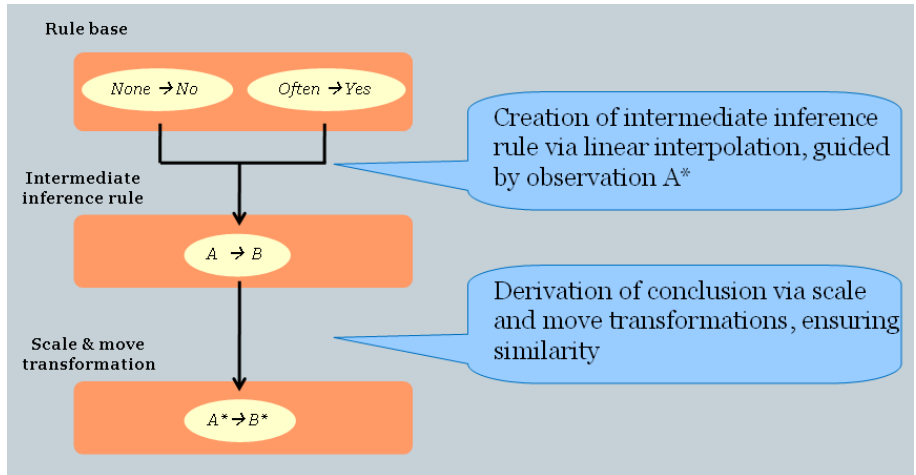
Figure 10: Transformation-Based Fuzzy Interpolation

Figure 11: Risk Assessment

constraints imposed over the planning process of police resource deployment. This will help to minimise the cost of successful surveillance, for example. Techniques reported in (Miguel and Shen, 2003) may be used to automate such planning.

4 CONCLUSION

This paper has presented a novel framework upon which to develop intelligent decision support systems, with a focussed application to intelligence data modelling and monitoring. It has proposed methods which can aid intelligence analysts in considering as widely as possible the range of emerging scenarios that may reflect organised criminal activities. The resulting approach has the ability to link seemingly distinct and unrelated intelligence data, associating and prioritising such data with logically inferred and justified possible scenarios.

In short, this work has demonstrated that computational intelligence in general, and fuzzy systems in particular can provide useful means to capture, learn and reason from (intelligence data under) uncertainty. It has also shown that evidence-driven plausible scenario synthesis is helpful for (intelligence monitoring) decision support. Furthermore, this research has outlined, though very briefly, some of the aspects that fuzzy techniques can be very successful for:

- Fragment induction
- Feature selection
- Interpolative reasoning
- Model composition
- Constraint satisfaction
- Truth maintenance
- Co-reference resolution
- Information aggregation
- Evidence evaluation
- Risk assessment

However, important research remains. In addition to what has been mentioned in the paper, the following lists a number of further issues (amongst possibly many others) that are worthy of investigation and/or development in order to reinforce the potential of this work:

- Learning hierarchical model fragments
- Hierarchical and ensemble feature selection
- Unification of scenario generation algorithms
- Dynamic coreference resolution and information fusion

- Evidence-driven risk-guided scenario generation
- Reconstruction of reasoning process
- Discovery of rare cases
- Meta-feature learning and selection for scenario synthesis

Such further studies may also bring up fresh challenges to computational intelligence research and hence new technologies for building intelligent decision support systems. It will help in consolidating and broadening the scope of its applications. In particular, the proposed framework itself may be adapted to suit tasks such as: investigator training, policy formulation, multi-modal profiling, and to address problems in domains such as academic performance evaluation and financial situation forecasting.

ACKNOWLEDGEMENTS

REFERENCES

Baranyi, P., Koczy, L., and Gedeon, T. (2004). A generalized concept for in fuzzy rule interpolation. *IEEE Transactions on Fuzzy Systems*, 12(6):820–837.

Berners-Lee, T., Hendler, J., and Lassila, O. (2001). The semantic web. *Scientific American*, pages 34–43.

Calado, P., Cristo, M., Goncalves, M., de Moura, E., Ribeiro-Neto, E., and Ziviani, N. (2006). Link based similarity measures for the classification of web documents. *Journal of American Society for Information Science and Technology*, 57(2):208–221.

de Kleer, J. (1986). An assumption-based TMS. 28:127–162.

Falkenhainer, B. and Forbus, K. (1991). Compositional modeling: finding the right model for the job. 51:95–143.

Halliwell, J. and Shen, Q. (2009). Linguistic probabilities: theory and application. *Soft Computing*, 13(2):169–183.

Huang, Z. and Shen, Q. (2006). Fuzzy interpolative reasoning via scale and move transformation. *IEEE Transactions on Fuzzy Systems*, 14(2):340–359.

Huang, Z. and Shen, Q. (2008). Fuzzy interpolative and extrapolative reasoning: a practical approach. *IEEE Transactions on Fuzzy Systems*, 16(1):13–28.

Jensen, R. and Shen, Q. (2004). Semantics-preserving dimensionality reduction: Rough and fuzzy-rough approaches. *EEE Transactions on Knowledge and Data Engineering*, 16(12):1457–1471.

Jensen, R. and Shen, Q. (2007). Fuzzy-rough sets assisted attribute selection. *IEEE Transactions on Fuzzy Systems*, 15(1):73–89.

Jensen, R. and Shen, Q. (2008). *Computational Intelligence and Feature Selection: Rough and Fuzzy Approaches*. IEEE and Wiley.

Jensen, R. and Shen, Q. (2009a). Are more features better? *IEEE Transactions on Fuzzy Systems*.

Jensen, R. and Shen, Q. (2009b). New approaches to fuzzy-rough feature selection. *IEEE Transactions on Fuzzy Systems*, 17(4):824–838.

Keen, P. (1978). *Decision support systems: an organizational perspective*. Addison-Wesley.

Keppens, J. and Shen, Q. (2001). On compositional modelling. *Knowledge Engineering Review*, 16(2):157–200.

Lee, M. (1999). On models, modelling and the distinctive nature of model-based reasoning. *AI Communications*, 12:127–137.

Liben-Nowell, D. and Kleinberg, J. (2007). The link-prediction problem for social networks. *Journal of American Society for Information Science and Technology*, 58(7):1019–1031.

Liu, H. and Motoda, H. (1998). *Feature Selection for Knowledge Discovery and Data Mining*. Springer.

MacParthalain, N. and Shen, Q. (2009). Exploring the boundary region of tolerance rough sets for feature selection. *Pattern Recognition*, 42(5):655–667.

MacParthalain, N., Shen, Q., and Jensen, R. (2009). A distance measure approach to exploring the rough set boundary region for attribute reduction. *IEEE Transactions on Knowledge and Data Engineering*.

Marakas, G. (1999). *Decision support systems in the twenty-first century*. Prentice Hall.

Marín-Blázquez, J. G. and Shen, Q. (2002). From approximative to descriptive fuzzy classifiers. *IEEE Trans. on Fuzzy Systems*, 10:484–497.

Miguel, I. and Shen, Q. (2003). Fuzzy rrdfcsp and planning. *Artificial Intelligence*, 148:11–52.

Parsons, S. (2003). Qualitative probability and order of magnitude reasoning. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 11(3):373–390.

Raiman, O. (1991). Order-of-magnitude reasoning. *Artificial Intelligence*, 51:11–38.

Shadbolt, N., Hall, W., and Berners-Lee, T. (2006). The semantic web revisited. *IEEE Intelligent Systems*, pages 96–101.

Shen, Q. and Jensen, R. (2004). Selecting informative features with fuzzy-rough sets and its application for complex systems monitoring. *Pattern Recognition*, 37(7):1351–1363.

Shen, Q., Keppens, J., Aitken, C., Schafer, B., and Lee, M. (2006). A scenario driven decision support system for serious crime investigation. *Law, Probability and Risk*, 2(5):87–117.

Shen, Q., Zhao, R., and Tang, W. (2008). Modelling random fuzzy renewal reward processes. *IEEE Transactions on Fuzzy Systems*, 16(5):1379–1385.

Tikk, D. and Baranyi, P. (2000). Comprehensive analysis of a new fuzzy rule interpolation method. *IEEE Transactions on Fuzzy Systems*, 8(3):281–296.

Tsang, E., Chen, D., Yeung, D., Wang, X., and Lee, J. (2008). Attributes reduction using fuzzy rough sets. *IEEE Transactions on Fuzzy Systems*, 16(5):1130–1141.

Zadeh, L. A. (1975). The concept of a linguistic variable and its application to approximate reasoning I. *Information Sciences*, 8:199–249.