Information Fusion in Logic: A Brief Overview

Laurence Cholvy[1] and Anthony Hunter[2]

(1) ONERA-CERT
2 avenue Ed Belin
31055 Toulouse, FRANCE
and
(2) Department of Computer Science
University College London
Gower Street
London WC1E 6BT, UK

Abstract. Information fusion is the process of deriving a single consistent knowledgebase from multiple knowledgebases. This process is important in many cognitive tasks such as decision-making, planning, design, and specification, that can involve collecting information from a number of potentially conflicting perspectives or sources, or participants. In this brief overview, we focus on the problem of inconsistencies arising in information fusion. In the following, we consider reasoning with inconsistencies, acting on inconsistencies, and resolving inconsistencies.

1 Introduction

Many tasks that an intelligent agent performs such as decision-making, planning, design, and specification, often involve collecting information from a number of potentially conflicting perspectives or sources, or participants with different views, and forming a single combined view or perspective — a synthesis, or consensus.

Consider requirements engineering. The development of most large and complex systems necessarily involves many people, each with their own perspectives on the system defined by their knowledge, responsibilities, and commitments. Inevitably, the different perspectives of those involved in the process intersect, giving rise to conflicts. From a logics perspective, these conflicts can be viewed as logical contradictions or inconsistencies. While classical logic is a rich and useful formalism in requirements engineering, it does not allow useful reasoning in the presence of inconsistency: the proof rules of classical logic allow any formula of the language to be inferred. Hence, classical logic does not provide a means for continued deduction in the presence of inconsistency.

Traditionally, inconsistency in logic has been viewed as a problem that requires immediate rectification. In [GH91,GH93], it was proposed that inconsistency is not necessarily a problem, as long as we know how to act in the presence of it. Indeed, inconsistencies can be viewed as being useful, since they can help to direct a cognitive activity. Certaintly, premature resolution of inconsistency can result in the loss of valuable information, and constrain an overall problem solution.

Ultimately, a single consistent view is required from a set of multiple views. Information fusion is the process of deriving this single consistent view. Whilst theoretical approaches such as belief revision [AGM85,Gar88,DP97], databases and knowledgebase updating [FKUV86,KM89,Win90,Som94], and combining knowledgebases (for example [DLP92,Mot93,BKMS92,BKMS91]) are relevant, information fusion addresses a wider range of issues raised by practical imperatives in applications such as requirements engineering.

The problem of information fusion appears in many fields, such as gathering beliefs or evidence, developing specifications, and merging regulations. The types of information to be modelled can differ, depending on the application. Frequently, they can be beliefs (describing what things are or what they are supposed to be in the real world) or they can be norms (describing how things should be in an ideal world). And of course, the type of logical formalism used to model the information depends on the type of information. Possibilities include classical logic, belief logics and deontic logics, though, potentially, any logic may be used.

But, whatever the type of the logic, the problem of information fusion raises the crucial problem of inconsistencies. Whilst the aim is to build a consistent set of information, the status of this set of information will differ, depending on the type of information. For example, in gathering beliefs, the aim of the fusion process is to build a consistent set of beliefs — a consistent representation of the real world — whereas in regulation merging, the aim is to build a consistent regulation i.e, a set of rules which consistently specifies an ideal world [Cho97].

In the following, we discuss some of the features of information fusion in logic. In particular, we focus on the problems of inconsistencies arising in information fusion, including the management of inconsistent information, reasoning with inconsistencies, and resolving inconsistencies.

2 Reasoning with inconsistencies

In practical reasoning, it is common to have "too much" information about some situation. In other words, it is common for there to be classically inconsistent information in a practical reasoning problem. The diversity of logics proposed for aspects of practical reasoning indicates the complexity of this form of reasoning. However, central to this is the need to reason with inconsistent information without the logic being trivialised.

Classical logic is trivialised because, by the definition of the logic, any inference follows from inconsistent information *(ex falso quodlibet)* as illustrated by the following example. From the set of formulae α, $\neg\alpha$, $\alpha \rightarrow \beta$, δ, reasonable inferences might include α, $\neg\alpha$, $\alpha \rightarrow \beta$, and δ by *reflexivity*; β by *modus ponens*; $\alpha \wedge \beta$ by *conjunction introduction*; $\neg\beta \rightarrow \neg\alpha$ and so on. In contrast, trivial inferences might include γ and $\gamma \wedge \neg\delta$.

For classical logic, trivialisation renders the reasoning useless, and therefore classical logic is obviously unsatisfactory for handling inconsistent information. A possible solution is to weaken classical logic by dropping some of the inferencing capability, such as for the C_ω paraconsistent logic [dC74], though this kind of weakening of the proof rules means that the connectives in the language do not behave in a classical fashion [Bes91]. For example, disjunctive syllogism does not hold, $((\alpha \vee \beta) \wedge \neg\beta) \rightarrow \alpha$, whereas modus ponens does hold. Variations on this theme include [AB75,Arr77,Bat80,PR84,PRN88,CFM91]. Alternative compromises on classical logic include three-valued logic [Lin87], four-valued logic [Bel77], quasi-classical (QC) logic [BH95], and using a form of conditioanl logic [Roo93]. Another approach is to reason with classically consistent subsets of inconsistent information. This has given rise to a number of logics (for example [MR70,Wag91,BDP93,BCD$^+$93,BDP95,EGH95]) and truth maintenance systems (for example [Doy79,Kle86,MS88]). These options from C_ω through to reasoning with maximally consistent subsets behave in different ways with data. None can be regarded as perfect for handling inconsistency in general. Rather they provide a spectrum of approaches. However, in all of them, the language is based on that of classical logic, and the aim is to preserve features of classical reasoning. For a review of these approaches see [Hun97].

Modal logics have also been developed for reasoning with inconsistent information. For instance, in [FH86], Farinas and Herzig define a modal logic for reasoning about elementary changes of beliefs which is a kind of conditional logic. The accessibility relation associated with the modality aims to capture the relation between a world and the worlds which are obtained after adding a piece of information. This logic has more recently been extended in order to take into account the dependances that may exist between the different pieces of information [FH92,FH94] and the influence these dependence links have in the updating process. Another modal conditional logic defined for belief revision is described in [Gra91]. Modal logic has also been proposed for consistent reasoning with inconsistent information by using only a consistent subset of beliefs [Lin94]. For a review of modal logics in handling inconsistent information see [MvdH97].

Problems of handling default knowledge are closely related to that of handling inconsistent information. Indeed, implicit in default, or non-monotonic, reasoning is the need to avoid trivialization due to conflicting defaults [BH97]. There are a variety of non-monotonic logics for handling default knowledge (for reviews see [Bes89,Bre91,GHR94]) with different strategies for avoiding inconsistency by selecting preferred consistent subsets of formulae. These include preference for more specific information [Poo85], ordered theory presentations [Rya92], preferred subtheories [Bre89], explicit preferences [Pra93], and prioritised syntax-based entailment [BCD$^+$93]. Non-monotonic logics therefore offer means for analysing inconsistent information and preferences over that information.

The ability to reason with inconsistencies is important, since it allows inconsistent information to be explored and analysed. However, there is also a need to act on inconsistencies, and eventually to resolve inconsistencies. We address these two topics in the next two sections.

3 Acting on inconsistencies

Immediate resolution of inconsistency by arbitrarily removing some formulae can result in the loss of valuable information. This can include loss of information that is actually correct and also loss of information that can be useful in managing conflicts. Immediate resolution can also unduely constrain cognitive activities such as problem solving and designing.

Identifying the appropriate inconsistency handling strategy depends on the kinds of inconsistency that can be detected and the degree of inconsistency tolerance that can be supported. Possible kinds of actions include:

Circumventing the inconsistent parts of the information. This can be viewed as ignoring the inconsistency, and using the rest of the information regardlessly. This may be appropriate in order to avoid inconsistent portions of the information and/or to delay resolution of the inconsistency. This includes using the logical techniques discussed in Section 2. Isolating inconsistency is a special case where the minimally inconsistent subset of the information is not used in the reasoning – it is isolated — but not deleted.

Ameliorating inconsistent situations by performing actions that "improve" these situations and increase the possibility of future resolution. This is an attractive approach in situations where complete and immediate resolution is not possible (perhaps because further information is required from elsewhere), but where some steps can be taken to "fix" part or some of the inconsistent information. This approach requires techniques for analysis and reasoning in the presence of inconsistency.

Sequencing of conflicts so that some conflicts are addressed before others. The criteria for sequencing are diverse but may include:

Granularity of inconsistency: Some conflicts are more significant than others. Furthermore, it is possible that solving a less significant conflict before a more significant conflict may unduely constrain the allowed solutions for the more significant conflict. Solving bigger or more important conflicts first means that we need to order the conflicts. For instance, when desiging a building, first solve the conflicts that are about its "function" (deciding what this building is for will determine its height, its surface...), then solve the conflicts about the "inside" of the building (this will determine the number of rooms, the exact places of the walls...), and then solve the conflicts about "decoration" (this will finally determine the style of the curtains and the colour of the walls,...). For this, the notion of topichood of information, such as in [CD89,CD92,Hun96], is potentially useful. This could allow us to say that"inconsistent set X is about topic T", and may be used in determining the significance of the inconsistency.

Temporality of inconsistency: Other temporal constraints can impose an ordering on the sequence of resolution of conflicts. For example, in a building project, delaying the resolution of conflicts about the position of a wall is temporally more sensitive than delaying the resolution of conflicts about the colour of the paint on the wall, since the construction of the wall needs to be completed before the wall is painted. Clearly the colour of the paint can be chosen before the position of the wall.

For the above examples, it can be seen that there is an overlap in granularity and temporality of inconsistencies. However, an example of equally significant inconsistencies. where the first is more temporally sensitive, is choice of paint for the ceiling of a room, and a choice of colour for the walls. And an example where some inconsistencies are far more significant than others, but where ultimately they are of equal temporal sensitivity, is in a book that is about to be published.

Resolving inconsistencies altogether by correcting any mistakes or resolving conflicts. This depends on a clear identification of the inconsistency and assumes that the actions required to fix it are known.

Circumventing, ameliorating and sequencing inconsistency all imply that the resolution of conflicts or inconsistency is delayed. In practice applications usually involve multiple conflicts of diverse kinds and significance. As a result, a combination of circumventing, ameliorating and sequencing of inconsistency is required.

4 Resolving inconsistencies

In order to resolve an inconsistency intelligently, as opposed to arbitrarily, we require appropriate information. Information that is manipulated in information fusion can be partitioned, and is often represented, in different ways. The information includes the extra information required for combination.

Object-level information: The information to be combined.

Combination information: The information used to facilitate combination. This is composed of meta-level information and domain information.

Meta-level information: Information about information. For example, information about
- The sources of the object-level information
- The reliability of sources
- Preferences about the information

Domain information: Information on the context or domain of the object-level. This is used to constrain the combination process. Examples of domain information include integrity constraints such as "everybody is either a man or a woman" and "a cube has 6 sides". Domain information can be uncertain, such as for example heuristics. Though using uncertain information significantly increases the difficulty of combining object-level information.

Combination information constitutes extra information that can be used by the combination process in order to combine the object-level information. Neither domain information nor meta-level information needs to be in the same formalism as the object-level information. The only constraint on the formalism is that it can be used by the combination process.

To illustrate these concepts, consider the following example. We have two sources of information S1 and S2 and we wish to combine their object-level information.

S1: `The colour of the object is blue.`
S2: `The colour of the object is green.`

We also have the following domain information and meta-level information.

Domain: `Green and blue are different colours.`
Meta: `The domain information is more reliable than source S1.`
Meta: `Source S1 is more reliable than source S2.`

In forming the combined information, we can accept the information from S1, because it is consistent and from the most reliable source. However, we cannot now add the information from S2 since it would cause an inconsistency.

For example, assuming an ordering over development information is reasonable in software engineering. First, different kinds of information have different likelihoods of being incorrect. For example, method rules are unlikely to be incorrect, whereas some tentative specification information is quite possibly incorrect. Second, if a specification method is used interactively, a user can be asked to order pieces of specification according to likelihood of correctness.

As a second example, assuming an order between different sources is reasonable in the case when merging beliefs or evidence provided by those different information sources [Cho94,Cho93]. In particular, it can be useful to order sources according to the topics of the information they provide — in effect adopting context-sensitive ordering over the sources [CD94,Cho95]. Indeed, assuming only one ordering over the different information sources is not very realistic, given that frequently a source can be assumed to be more reliable than a second source, on one topic, but less reliable on another topic.

As a third example, in the domain of regulation merging, assuming an ordering over regulations can be useful to consistently reason with rules [CC95]. This is related to the use of priorities in argumentation for legal reasoning, in particular [Pra93,PS95,PS96,RD96,TdT95].

There are number of ways that this approach can be developed. First, there are further intuitive ways of deriving orderings over formulae and sets of formulae. These include ordering sets of formulae according to their relative degree of contradiction [GH97]. Second, there are a number of analyses of ways of handling ordered formulae and sets of ordered formulae such as dicussed above in Section 2 on logics for inconsistent information.

5 Discussion

In general, information fusion is a difficult problem. Given the difficulty, there is the need for the following:

- Inconsistency management during information fusion, to track inconsistencies and minimize the negative ramification of inconsistency.
- A range of logics for reasoning and analysis of inconsistent information, to allow continued use of inconsistent information, and to facilitate resolution.
- Extra information, called combination information, to enable the resolution of inconsistencies during information fusion.
- Interactive support for information fusion where the support system offers suggestions but the user controls the fusion process.

Despite the difficulties, there are practical reasoning applications where information fusion is likely to be of significant import: Take for example managing inconsistencies in the development of multi-perspective software development [FGH+94,HN97], where inconsistencies can be detected using classical logic, information surrounding each inconsistency can be used to focus continued development, and actions are taken in a context-dependent way in response to inconsistency.

Acknowledgements

This work was partly funded by the CEC through the ESPRIT FUSION project.

References

[AB75] A Anderson and N Belnap. *Entailment: The Logic of Relevance and Necessity.* Princeton University Press, 1975.

[AGM85] C Alchourrón, P Gardenfors, and D Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50:513–530, 1985.

[Arr77] A Arruda. On the imaginary logic of NA Vasilev. In A Arruda, N Da Costa, and R Chuaqui, editors, *Non-classical logics, model theory and computability.* North Holland, 1977.

[Bat80] D Batens. Paraconsistent extensional propositional logics. *Logique et Analyse*, 90–91:195–234, 1980.

[BCD+93] S Benferhat, C Cayrol, D Dubois, J Lang, and H Prade. Inconsistency management and prioritized syntax-based entailment. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence.* Morgan Kaufmann, 1993.

[BDP93] S Benferhat, D Dubois, and H Prade. Argumentative inference in uncertain and inconsistent knowledge bases. In *Proceedings of Uncertainty in Artificial Intelligence*, pages 1449–1445. Morgan Kaufmann, 1993.

[BDP95] S Benferhat, D Dubois, and H Prade. A logical approach to reasoning under inconsistency in stratified knowledge bases. In *Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, volume 956 of *Lecture Notes in Computer Science*, pages 36–43. Springer, 1995.

[Bel77] N Belnap. A useful four-valued logic. In G Epstein, editor, *Modern Uses of Multiple-valued Logic*, pages 8–37. Reidel, 1977.

[Bes89] Ph Besnard. *An Introduction to Default Logic*. Springer, 1989.

[Bes91] Ph Besnard. Paraconsistent logic approach to knowledge representation. In M de Glas M and D Gabbay D, editors, *Proceedings of the First World Conference on Fundamentals of Artificial Intelligence*, pages 107–114. Angkor, 1991.

[BH95] Ph Besnard and A Hunter. Quasi-classical logic: Non-trivializable classical reasoning from inconsistent information. In C Froidevaux and J Kohlas, editors, *Symbolic and Quantitative Approaches to Uncertainty*, volume 946 of *Lecture Notes in Computer Science*, pages 44–51, 1995.

[BH97] Ph Besnard and A Hunter. Introduction to actual and potential contradictions. In *Handbook of Defeasible Reasoning and Uncertainty Management*, volume 3. Kluwer, 1997.

[BKMS91] C. Baral, S. Kraus, J. Minker, and V.S. Subrahmanian. Combining multiple knowledge bases. *IEEE Trans. on Knowledge and Data Engineering*, 3(2), 1991.

[BKMS92] C Baral, S Kraus, J Minker, and V Subrahmanian. Combining knowledge-bases of first-order theories. *Computational Intelligence*, 8:45–71, 1992.

[Bre89] G Brewka. Preferred subtheories: An extended logical framework for default reasoning. In *Proceedings of the Eleventh International Conference on Artificial Intelligence*, pages 1043–1048, 1989.

[Bre91] G Brewka. *Common-sense Reasoning*. Cambridge University Press, 1991.

[CC95] L. Cholvy and F. Cuppens. Solving normative conflicts by merging roles. In *Proceedings of the fifth International Conference on Artificial Intelligence and Law*, Washington, May 1995.

[CD89] F. Cuppens and R. Demolombe. How to recognize interesting topics to provide cooperative answering. *Information Systems*, 14(2), 1989.

[CD92] S. Cazalens and R. Demolombe. Intelligent access to data and knowledge bases via users' topics of interest. In *Proceedings of IFIP Conference*, pages 245–251, 1992.

[CD94] L. Cholvy and R. Demolombe. Reasoning with information sources ordered by topics. In *Proceedings of Artificial Intelligence : Methods, Systems and Applications (AIMSA)*. World Scientific, Sofia, september 1994.

[CFM91] W Carnielli, L Farinas, and M Marques. Contextual negations and reasoning with contradictions. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'91)*, 1991.

[Cho93] L. Cholvy. Proving theorems in a multi-sources environment. In *Proceedings of IJCAI*, pages 66–71, 1993.

[Cho94] L. Cholvy. A logical approach to multi-sources reasoning. In *Proceedings of the Applied Logic Conference*, number 808 in Lecture notes in Artificial Intelligence. Springer-Verlag, 1994.

[Cho95] L. Cholvy. Automated reasoning with merged contradictory information whose reliability depends on topics. In *Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU)*, Fribourg, July 1995.

[Cho97] L. Cholvy. Reasoning about merged information. In *Handbook of Defeasible Reasoning and Uncertainty Management*, volume 1. Kluwer, 1997.

[dC74] N C da Costa. On the theory of inconsistent formal systems. *Notre Dame Journal of Formal Logic*, 15:497–510, 1974.

[DLP92] D. Dubois, J. Lang, and H. Prade. Dealing with multi-source information in possibilistic logic. In *Proceedings of ECAI*, pages 38–42, 1992.

[Doy79] J Doyle. A truth maintenance system. *Artificial Intelligence*, 12:231–272, 1979.

[DP97] D Dubois and H Prade. *Handbook of Defeasible Reasoning and Uncertainty Management*, volume 1. Kluwer, 1997.

[EGH95] M Elvang-Goransson and A Hunter. Argumentative logics: Reasoning from classically inconsistent information. *Data and Knowledge Engineering Journal*, 16:125–145, 1995.

[FGH+94] A Finkelstein, D Gabbay, A Hunter, J Kramer, and B Nuseibeh. Inconsistency handling in multi-perspective specifications. *IEEE Transactions on Software Engineering*, 20(8):569–578, 1994.

[FH86] L. Farinas and A. Herzig. Reasoning about database updates. In Jack Minker, editor, *Workshop of Foundations of deductive databases and logic programming*, 1986.

[FH92] L. Farinas and A. Herzig. Revisions, updates and interference. In A. Fuhrmann and Rott H, editors, *Proceedings of the Konstanz colloquium in logic and information (LogIn-92.* DeGruyter Publishers, 1992.

[FH94] L. Farinas and A. Herzig. Interference logic = conditional logic + frame axiom. *International JOurnal of Intelligent Systems*, 9(1):119–130, 1994.

[FKUV86] R Fagin, G Kuper, J Ullman, and M Vardi. Updating logical databases. *Advances in Computing Research*, 3:1–18, 1986.

[Gar88] P Gardenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States.* MIT Press, 1988.

[GH91] D Gabbay and A Hunter. Making inconsistency respectable 1: A logical framework for inconsistency in reasoning. In *Fundamentals of Artificial Intelligence*, volume 535 of *Lecture Notes in Computer Science*, pages 19–32. Springer, 1991.

[GH93] D Gabbay and A Hunter. Making inconsistency respectable 2: Meta-level handling of inconsistent data. In *Symbolic and Qualitative Approaches to Reasoning and Uncertainty (ECSQARU'93)*, volume 747 of *Lecture Notes in Computer Science*, pages 129–136. Springer, 1993.

[GH97] D Gabbay and A Hunter. Negation and contradiction. In *What is negation?* Kluwer, 1997.

[GHR94] D Gabbay, C Hogger, and J Robinson. *Handbook of Artificial Intelligence and Logic Programming*, volume 3. Oxford University Press, 1994.

[Gra91] G. Grahne. A modal analysis of subjonctive queries. In R. demolombe, L. farinas, and T. Imielinski, editors, *Workshop on nonstandard queries and answers*, Toulouse, 1991.

[HN97] A Hunter and B Nuseibeh. Analysing inconsistent specifications. In *Proceedings of 3rd International Symposium on Requirements Engineering*, pages 78–86. IEEE Computer Society Press, 1997.

[Hun96] A Hunter. Intelligent text handling using default logic. In *Proceedings of the Eighth IEEE International Conference on Tools with Artificial Intelligence (TAI'96)*, pages 34–40. IEEE Computer Society Press, 1996.

[Hun97] A Hunter. Paraconsistent logics. In *Handbook of Defeasible Reasoning and Uncertainty Management*. Kluwer, 1997.

[Kle86] J De Kleer. An assumption-based TMS. *Artificial Intelligence*, 28:127–162, 1986.

[KM89] H Katsuno and A Medelzon. A unified view of propositional knowledgebase updates. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 1989.

[Lin87] F Lin. Reasoning in the presence of inconsistency. In *Proceedings of the National Conference on Artificial Intelligence (AAAI'87)*, 1987.

[Lin94] J Lin. A logic for reasoning consistently in the presence of inconsistency. In *Proceedings of the Fifth Conference on Theoretical Aspects of Reasoning about Knowledge*. Morgan Kaufmann, 1994.

[Mot93] A. Motro. A formal framework for integrating inconsistent answers from multiple information sources. Technical Report ISSE-TR-93-106, George Mason University, 1993.

[MR70] R Manor and N Rescher. On inferences from inconsistent information. *Theory and Decision*, 1:179–219, 1970.

[MS88] J Martins and S Shapiro. A model of belief revision. *Artificial Intelligence*, 35:25–79, 1988.

[MvdH97] J Meyer and W van der Hoek. Modal logics for representing incoherent knowledge. In *Handbook of Defeasible Reasoning and Uncertainty Management, Volume 3*. Kluwer, 1997.

[Poo85] D Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36:27–47, 1985.

[PR84] G Priest and R Routley. Introduction: Paraconsistent logics. *Studia Logica*, 43:3–16, 1984.

[Pra93] H Prakken. An argument framework for default reasoning. In *Annals of mathematics and artificial intelligence*, volume 9, 1993.

[PRN88] G Priest, R Routley, and J Norman. *Paraconsistent logic*. Philosophia, 1988.

[PS95] H. Prakken and G. Sartor. On the relation between legal language and legal argument : assumptions, applicability and dynamic priorities. In *Proc. Fifth Conference on Artificial Intelligence and Law*, University of Maryland, May, 1995.

[PS96] H. Prakken and G. Sartor. A system for defeasible argumentation with defeasible prorities. In *Proc. of FAPR'96*, May, 1996.

[RD96] L. Royakkers and F. Dignum. Defeasible reasoning with legal rules. In *Proc of DEON'96, Sesimbra*. Springer, 1996.

[Roo93] N Roos. A logic for reasoning with inconsistent knowledge. *Artificial Intelligence*, 57(1):69–104, 1993.

[Rya92] M Ryan. Representing defaults as sentences with reduced priority. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference*. Morgan Kaufmann, 1992.

[Som94] Leá Sombé. *Revision and updating in knowledgebases*. Wiley, 1994.

[TdT95] Y Tan and L Van der Torre. Why defeasible deontic logic needs a multi preference semantics. In Ch. froidevaux and J. Kohlas, editors, *Quantitative and Qualitative Approches to Reasoning and Uncertainty*, number 946 in Lectures notes in Artificial Intelligence. Springer, 1995.

[Wag91] G Wagner. Ex contradictione nihil sequitur. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'91)*, 1991.

[Win90] M Winslett. *Updating logical databases*. Cambridge University Press, 1990.