REAL-WORLD REASONING: TOWARD SCALABLE, UNCERTAIN SPATIOTEMPORAL, CONTEXTUAL AND CAUSAL INFERENCE

BEN GOERTZEL, NIL GEISWEILLER,
LÚCIO COELHO, PREDRAG JANIČIĆ,
CASSIO PENNACHIN

ATLANTIS THINKING MACHINES
SERIES EDITOR | K.-U. KÜHNBERGER

ATLANTIS PRESS

ATLANTIS THINKING MACHINES

VOLUME 2

SERIES EDITOR: **KAI-UWE KÜHNBERGER**

Atlantis Thinking Machines

Series Editor:

Kai-Uwe Kühnberger,

Institute of Cognitive Science

University of Osnabrück, Germany

(ISSN: 1877-3273)

Aims and scope of the series

This series publishes books resulting from theoretical research on and reproductions of general Artificial Intelligence (AI).The book series focuses on the establishment of new theories and paradigms in AI. At the same time, the series aims at exploring multiple scientific angles and methodologies, including results from research in cognitive science, neuroscience, theoretical and experimental AI, biology and from innovative interdisciplinary methodologies.

For more information on this series and our other book series, please visit our website at:

www.atlantis-press.com/publications/books

ATLANTIS PRESS

AMSTERDAM – PARIS

Real-World Reasoning: Toward Scalable, Uncertain Spatiotemporal, Contextual and Causal Inference

Ben Goertzel

Novamente
LLC/Biomind LLC
Bernerd Place
MD 20851, Rockville
USA

Nil Geisweiller

yl. Kliment 4
Samokov 2000
Bulgaria

Lúcio Coelho

Rua Guajajaras
1268, apto.506
30180-101
Belo Horizonte
Brasil

Predrag Janičić

Faculty of Mathematics
Studentski trg, 16
11000 Belgrade
Serbia

Cassio Pennachin

Igenesis Av. Prof. Mário
Werneck 2170 / 501
Buritis, Belo Horizonte
30575-180 Minas Gerais
Brasil

ATLANTIS PRESS

AMSTERDAM – PARIS

Atlantis Press

8, square des Bouleaux
75019 Paris, France

For information on all Atlantis Press publications, visit our website at: *www.atlantis-press.com*

Atlantis Thinking Machines

Volume 1: Enaction, Embodiment, Evolutionary Robotics. Simulation Models for a Post-Cognitivist Science of Mind - Marieke Rohde, Ezequiel A. Di Paolo

ISBNs
Print: 978-94-91216-10-7
E-Book: 978-94-91216-11-4
ISSN: 1877-3273

Contents

Acknowledgements

This book had its origins as a technical report written in 2008 by Novamente LLC for IARPA, under a program managed by Jeffrey Morrisson with the help of Ken Kisiel. So, above all, we would like to acknowledge the vision and practical support of these two most excellent applied scientists, without whom this book would not exist. The book in is current form is not precisely identical to the report we wrote for IARPA, but there is a large amount of overlap.

We also must especially acknowledge those who played a role in the development of the PLN inference system, which arises often in these pages. PLN was created by a large number of people, but perhaps the biggest contributors (aside from Goertzel and Geisweiller who co-authored this book) were Matt Ikle', Ari Heljakka, and Izabela Freire Goertzel, the co-authors of the PLN book, and Joel Pitt who became involved later. André Senna and Welter Silva also deserve particular thanks for their leading roles in developing the Novamente Cognition Engine and OpenCog frameworks in which the current PLN implementation exists.

Thanks are also due to Jeffrey Epstein for a research grant supporting Ben Goertzel's research, during part of the time period the original IARPA report was being written, and also during part of the time period the report was being turned into this book.

And, thanks go to Zeger Karssen of Atlantis Press for clearly understanding the need for a book such as this one ... and being patient with our inordinately long delays in producing the final manuscript.

Chapter 1

Introduction

The general problem addressed in this book is how to effectively carry out reasoning, knowledge discovery and querying based on huge amounts of complex information about real-world situations. Specifically we conceive "real-world reasoning" here mainly as "massively scalable reasoning involving uncertainty, space, time, cause and context." Of course there are other important aspects to reasoning about the real world we live in, e.g. the hierarchical structure of much of the human world, and we will briefly touch on some of these here as well. But for the purposes of this book, when we mention "real-world reasoning" or RWR, we're mostly talking about uncertainty, spacetime, cause, context and scalability.

The RWR problem is critical in at least two respects: as part of the broader pursuit of artificial general intelligence (AGI) (Goertzel & Pennachin, 2006; Goertzel *et al.*, 2006a; Goertzel & Bugaj, 2008; Hart & Goertzel, 2008), and in terms of the practical information processing needs that have arisen in current society.

On the AGI side, it is obvious that every human brain ingests a huge amount of knowledge each waking hour, and somehow we manage to query and analyze our huge, dynamic internal data stores. No AGI design can possibly succeed without some way to effectively carry out intelligent judgment and discovery based on these data stores. AGI also has other aspects, e.g. procedure learning and goal refinement (to name just two), but RWR is certainly a huge part of the puzzle.

On the practical information processing side, anyone who lives in a developed country these days is aware of the tremendous amount of data continually being gathered about all manner of aspects of the human and natural worlds. Much of this data is discarded shortly after it's gathered, but much of it is retained in various repositories. However, even when the data is retained, it is rarely utilized to anywhere near the full extent possible, because our state-of-the-art technologies for storing, querying, mining and analyzing very large data stores are still very primitive and simplistic (not only compared to what is in principle possible, but compared to what we know to be possible based on contemporary mathematics and computer science).

In these pages we review a class of approaches to handling these RWR problems using uncertain, spatiotemporal, contextual and causal logic. Uncertain logic is not the only possible approach to the RWR problem, but we believe it's one very promising approach, and it's our focus here. While the first RWR-capable logic system has yet to be constructed, we make an argument, via detailed review of the literature and the state of the art and suggestion of some original ideas, that the time is ripe for their construction.

The book is intended to serve two purposes: to provide a reasonably accessible overview of the RWR problem and the available technologies and concepts for its solution; and to provide a sketch of one possible avenue toward solution.

Toward the "overview" goal, we review a number of concepts and technologies – some recently developed, some more classical – that address aspects of the RWR problem. While our treatment centers on formal logic, we also introduce material from other areas such as graph databases, probability theory, cognitive architecture and so forth as appropriate.

After reviewing a variety of other logical approaches, we present our own approach to real-world reasoning, which is based on the Probabilistic Logic Networks (PLN) framework (Goertzel *et al.*, 2008); and give some detailed suggestions regarding how one might address the scalable real-world inference problem effectively via integrating PLN with other ideas and technologies described. Our goal in this regard is not to propose a particular highly-specific technical solution, but rather to describe a class of possible solutions that might be described as "scalable spatiotemporal uncertain logic systems". In this vein, in the later chapters we give a number of detailed examples showing the kinds of results one might expect to obtain by approaching a large knowledge store containing information about everyday human activities with the Probabilistic Logic Networks inference framework that we have developed in prior publications.

1.1 The Advantages of a Logical Approach

There are many advantages to the logic-based approach relative to others, some of which will be alluded to as the text progresses, but perhaps the largest advantage is its relative representational transparency. That is, if the knowledge stored in a knowledge base, and the patterns recognized in this knowledge base, are represented in a logical format, then it is reasonably tractable for humans to inspect this knowledge and these patterns. This is a major practical advantage in terms of allowing hybridized human/artificial intelligence – and, given the comments made above about the interesting but erratic performance of AI algorithms in our domain, this seems a very important point.

Given the advantage of logic-based approaches in terms of representational transparency, the only reason to choose an opaque approach over a logic-based approach would be if the opaque approach were dramatically superior in its capabilities. However, this currently seems not to be the case: in fact the evidence so far seems to indicate that logic-based approaches are the most powerful ones in this sort of context.

Some theorists have argued against logic-based approaches to real-world data on the grounds that there are problems with "grounding" logical symbols in real-world data (the so-called "symbol grounding problem" (Goertzel *et al.*, 2006a)). However, these objections do not hold up to scrutiny. It is true that logic-based approaches cannot function adequately for real-world applications unless the logical symbols used are explicitly associated with observed data-patterns, but there are well-understood technologies for making such associations. Historically, many logic-based AI systems have been used in an "ungrounded" way, not containing components that directly connect the logical terms used with real-world observations – but this is a problem of poor system architecture, not a flaw of the logic-based approach in itself.

1.2 Main High-Level Conclusions

To give a small hint at what is to come, the main conclusions at the end of our investigation are that

- the logic-based approach has the in-principle power to solve the problem of querying and analyzing very large scale spatiotemporal knowledge bases, in a manner respecting the contextual and causal knowledge contained therein
- there is a significant amount of scientific and technological knowledge in the literature regarding nearly every aspect of the application of logic-based technology to this problem
- the Achilles heel of current relevant logic-based technology is scalability
- the keys to achieving scalability in this context are conceptually understood – adaptive inference control and attention allocation – but have not been explored nearly as thoroughly as they need to be

- it seems likely that special techniques may be useful for adaptively controlling real-world scalable inference as opposed to inference in other domains (e.g. mathematical theorem proving)
- one viable way to achieve scalable real-world reasoning may be to use the Probabilistic Logic Networks framework, perhaps within an integrative AGI design like OpenCog which provides flexible means for adaptive inference control

We thus suggest that a critical focus of research should be on the development of methods for *exploiting the specific statistical structure of real spatiotemporal data, to adaptively guide logical inference methods* in performing query and analytical processing.

1.3 Summary

We now briefly review the chapters to follow, summarizing the main themes and ideas to be introduced.

1.3.1 *Part I: Representations and Rules for Real-World Reasoning*

Part I of the book reviews a host of approaches described in the literature for representing and reasoning about real-world knowledge, including temporal, spatial, contextual and causal knowledge.

Chapter Two reviews many of the varieties of formal logic that have been developed during the last century, with a focus on those approaches that appear most relevant to the large-scale information-management problem. We begin with a basic review of predicate and term logic, and then move on to subtler variations such as modal logic (the logic of possibility) and deontic logic (the logic of obligation). We also discuss the methods that logic systems use to actually draw logical conclusions based on the information provided to them: forward chaining, in which information items are combined exploratorily to come to new conclusions; and backward chaining, in which a question is posed to the system and it then seeks to find the answer using multiple logical inference steps based on the information at its disposal.

Chapter Three considers various methods of handling uncertainty in formal logic, including fuzzy sets and logic, possibility theory, probability theory, and imprecise and indefinite probabilities. Uncertainty management is critical to our target application, because a great percentage of real-world data is uncertain, and most of the conclusions one can draw based on real-world data are also uncertain. So, logic systems that only deal with absolute truth or falsehood are not going to be very useful for our target application. But, the literature contains a huge number of different methods for dealing with uncertainty – and one of our conclusions is that there isn't necessarily a single best approach. Rather, a practical solution may integrate more than one approach, for instance using both fuzzy and probabilistic methods as appropriate. Figures 1.1 and 1.2 from Chapter Three illustrate several of the possible methods for representing time within logic:

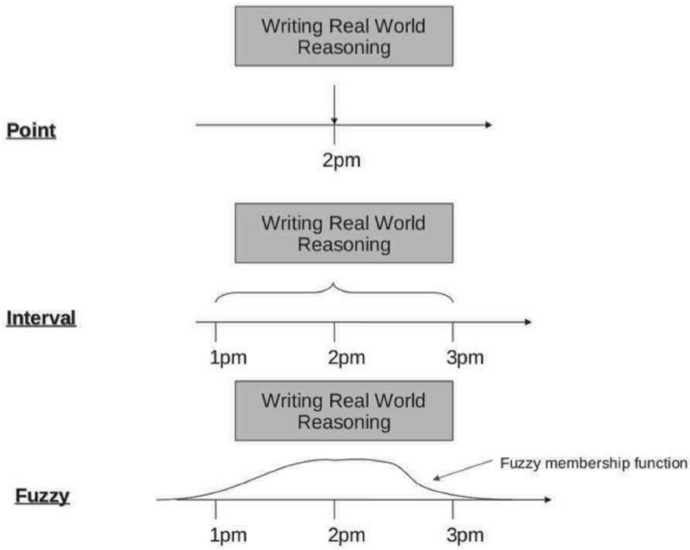

Fig. 1.1

Chapter Four grapples with the various ways logicians and computer scientists have devised to represent time within logic. This is a core issue for our current pursuit, because a large percentage of real-world knowledge involves time. The most standard method for handling time within logic is Allen's interval algebra, which treats time-intervals rather than points as the atomic temporal entities, and enumerates a set of rules for combining and reasoning about time-intervals; but it suffers the deficit of being crisp rather than explicitly handling uncertainty. So we review several methods of extending interval algebra to deal with uncertainty, including methods involving fuzziness, probability, and combinations of the two. Figure 1.3 from Chapter Four illustrates the logical relationships between time intervals specified by Allen's interval algebra:

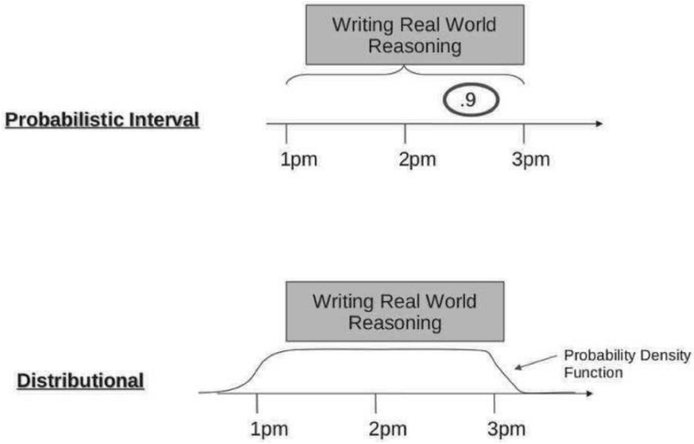

Fig. 1.2

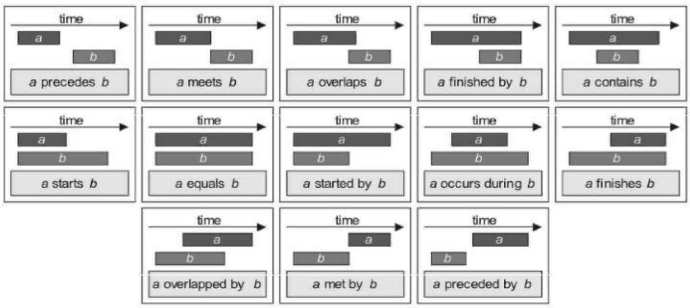

Fig. 1.3

And the Figure 1.4, also from Chapter Four, is a graphical representation of some temporal relationships between events, using a probabilistic variation of Allen's interval algebra:

Continuing the theme of its predecessor, Chapter Five deals with temporal inference, reviewing the multiple methods presented in the literature for incorporating time into logic. These include methods that simply treat time like any other logical information, and also methods that give time a special status, including reified and modal techniques. We conclude that methods giving time a special status are likely to be dramatically more efficient, and express a particular favor for reified techniques compatible with Allen's interval algebra (discussed above) and its variations. We give some concrete examples of temporal inference regarding peoples' daily activities.

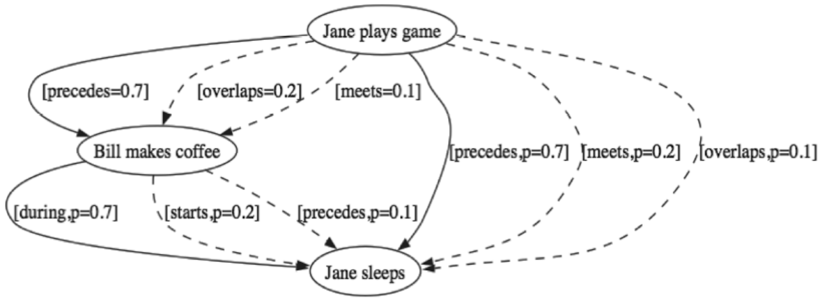

Fig. 1.4

For instance, one of the example problems we consider involves a query regarding "which people were in the same place as Jane last week," and a knowledge base with the following information:

- Susie and Jane use the same daycare center, but Jane uses it everyday, whereas Susie only uses it when she has important meetings (otherwise she works at home with her child).
- Susie sends a message stating that Tuesday she has a big meeting with a potential funder for her business.

Given this information, inference is needed to figure out that on Tuesday Susie is likely to put her child in daycare, and hence (depending on the time of the meeting!) potentially to be at the same place as Jane sometime on Tuesday. To further estimate the probability of the two women being in the same place, one has to do inference based on the times Jane usually picks up and drops off her child, and the time Susie is likely to do so based on the time of her meeting. We show in detail how temporal inference methods can be used to carry out this commonsense inference, and other similar examples.

Chapter Six builds on the treatment of time and presents an analogous discussion of a more complex subject, space (critical to our core theme as a substantial percentage of real-world knowledge involves spatial as well as temporal information). We review the Region Connection Calculus, which models the logic of space in terms of a fixed set of logical relationships between logical terms that correspond to spatial regions. As this is a simple but limited technique, we then consider more complex approaches to representing space in logic, including directional calculus, and occupancy grids as utilized in robotics (which are extremely general yet also resource-intensive, and so should only be used when simpler methods fail). The following diagram, drawn from Chapter Six, depicts the relationships between various spatial regions and spatially distributed phenomena (NTPP stands for Non-Tangential Proper Part, and O stands for Overlapping; these are spatial-relationship predicates drawn from the Region Connection Calculus formalism):

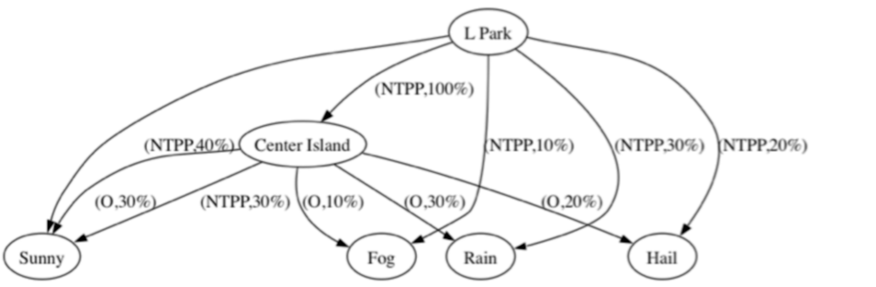

Fig. 1.5

Next, as well as time and space, another critical aspect of real-world reasoning is context. Nearly all real-world knowledge implicitly or explicitly gains its meaning from the specific context in which it is understood by human knowledge producers and consumers to exist. So if logical methods are to be applied effectively to real-world data, it is important that they explicitly represent contextuality. In Chapter Seven, we review a number of approaches to representing contextuality in logic, and give detailed examples of several. We also consider one example of context representation that is particularly acutely relevant to our application area: the use of contextual logic to handle user modeling. If different users of an information system have different biases and interests, then a logic based system can pay attention to this and give them different information via treating each user as a separate context and then doing contextually-biased reasoning.

In addition to context representation, Chapter Seven treats contextual inference, reviewing a number of techniques presented in the literature, and again finding favor in those methods that explicitly represent context as a special relationship within the base logic. We give a concrete examples of contextual inference applied to practical problems regarding people and their interrelationships. One example we consider involves the following assumptions:

- Alison is an accountant who is also a musician. Alison is emotional in the context of music, but not in the context of accounting. She frequently mentions Canadian place names in the context of music (maybe she's a Canadian music fan), but not in the context of accounting.

- Bob is in a similar situation, but he frequently mentions Canadian related stuff in both the music and accounting contexts.
- Clark is also in a similar situation, but he frequently mentions Canadian related stuff only in the accounting context, not the music context.
- People who have a lot to do with Canadian people, and a lot to do with money, have a chance of being involved in suspicious log trafficking activities.

We then show how contextual inference methods can be used to estimate the probability that Clark may be involved with log trafficking.

Chapter Eight turns briefly to causal reasoning, reviewing the multiple formalisms used to represent the notion of causality, and connecting causation to probabilistic and inductive reasoning.

1.3.2 *Part II: Acquiring, Storing and Mining Logical Knowledge*

Our focus in this book is doing logical reasoning on real-world knowledge, and this is a large and critical topic – but, once one has a large store of real-world knowledge in logical format, reasoning per se is not the only thing that must be done with it. Part II, a brief interlude at the center of the book, consists of three short chapters which lightly touch three other important issues to do with large stores of logical knowledge: acquiring logical knowledge via transforming real-world data, storing and querying large volumes of logical knowledge, and mining patterns from large logical knowledge stores. Each of these topics could be a book in itself, and here we only roughly sketch the main problems involved and give some pointers into the literature.

Chapter Nine very briefly reviews existing relevant literature, discussing the use of natural language processing technology to map text and voice into sets of logical relationships; and the use of image processing and heuristic algorithms to create logical relationships out of tables, graphs and diagrams. For instance, the following diagram drawn from Chapter Six shows some logical relationships that current NLP technology can extract from the simple sentence "Gone for dinner with Bob":

Another key question that must be answered if logic-based methods are to be applied to massive real-world data stores is: how can a huge amount of logical knowledge be stored and manipulated? This is not a question about logic per se, it's a question about modern computer systems, database and database-like technologies, and so forth. In Chapter Ten,

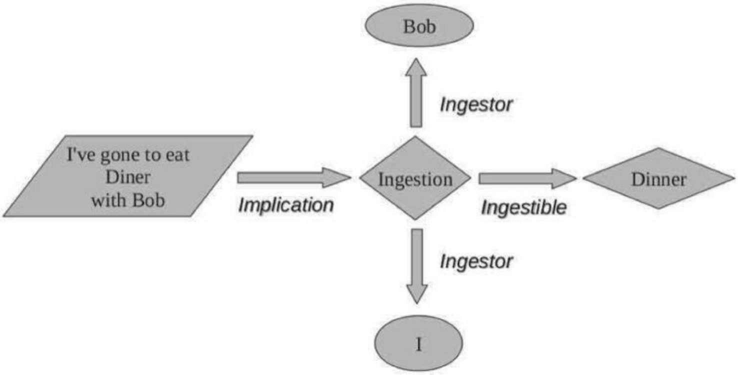

Fig. 1.6

we review a number of current technologies, including relational databases, RDF triple-stores, object databases, and hypergraph and graph databases. Our conclusion is that at present the latter form the best option, and we give some specific examples of how to translate complex logical knowledge into the specific format required for a graph database. The following table, drawn from Chapter Ten, summarizes some of our findings in more depth:

Technology	**Strengths**	**Weaknesses**
Relational DBs	• Mature, enterprise grade solutions • Ease of integration with other systems	• Poor conceptual fit for logical information storage • Inadequate model for reasoning • Complex scalability
Object-Oriented DBs	• Better conceptual fit than relational DBs (still not perfect) • Mature solutions	• Single data model • Small ecosystem • Not designed for reasoning

Continued on next page

continued from previous page

Technology	**Strengths**	**Weaknesses**
Graph DBs	• Flexible, dynamic data model • Good performance and scalability • Designed with data analysis in mind	• Less mature than competing technologies
Hypergraph DBs	• Best data model fit • Designed with reasoning and data analysis in mind	• Alpha stage technology
RDF Triplestores	• Semantic web friendly • Adequate data model for some inferences	• Less mature technology • Rigid data model
Document-oriented DBs	• Flexible data model • Performance and scalability • Rapidly maturing solutions	• Not adequate for reasoning and analysis • More work is left for application layer
Column-oriented DBs	• Very flexible, dynamic data model • Performance and scalability • Rapidly maturing solutions	• More work is left for application layer • Not designed for reasoning
Key-value DBs	• Extremely good performance and scalability • Mature and rapidly maturing solutions	• No data model, leaving most work for application layer • Not designed for reasoning

Chapter Ten turns to one of the most important applications desirable to carry out on large data stores – “data mining” (also known as “information exploitation”, “pattern discovery”, etc.). Most existing datamining techniques are either specialized for relational databases, or don't scale beyond small knowledge stores. We review here some specific datamining algorithms in depth. One conclusion drawn is that, for datamining to really be effective in this context, it will need to be hybridized with inference. Datamining technology, in itself, will always find too many potentially interesting patterns for any human user to want to explore. So logical inference technology is needed to filter the results of datamining, either via interaction with the datamining process, or via postprocessing.

1.3.3 *Part III: Real World Reasoning Using Probabilistic Logic Networks*

The second major of the book provides a detailed exploration of the applicability of one particular logical framework, Probabilistic Logic Networks, to real-world reasoning problems. This part is different from the previous ones, in that it comprises primarily original work, rather than literature survey and summary.

Chapters Twelve and Thirteen summarize Probabilistic Logic Networks (PLN), the particular uncertain logic system called which several of the authors (Goertzel and Pennnachin and Geisweiller) and their colleagues have developed over the last years (and published extensively on elsewhere). We outline the basic mechanisms via which PLN deals with a variety of aspects of inference, including term and predicate logic, extensional and intensional inference, and contextual, causal, spatial and temporal inference.

Chapter Fourteen turns to the specific problem of inference about changes in large knowledge bases. We consider several concrete examples including the following causal inference scenario:

- Before March 2007, Bob never had any Canadian friends except those who were also friends of his wife.
- After March 2007, Bob started acquiring Canadian friends who were not friends of his wife.
- In late 2006, Bob started collecting Pokemon cards. Most of the new Canadian friends Bob made between March 2007 and Late 2007 are associated with Pokemon cards
- In late 2006, Bob started learning French. Most of the new Canadian friends Bob made between March 2007 and Late 2007 are Quebecois.

We show in detail how a PLN inference engine, combining temporal inference with causal inference and numerous other aspects, can attempt to answer the question: What is the probable cause of Bob acquiring new Canadian friends who are not also friends of his wife?

Chapter Fourteen also considers spatial inference in the context of change analysis, giving particular attention to the incorporation of the Region Connection Calculus (RCC) into PLN. It is shown how a fuzzy/probabilistic version of RCC may be used together with a fuzzy/probabilistic version of Allen's interval algebra to carry out commonsense inferences about the causes of peoples' activities and relationships, based on knowledge involving time and space. To exemplify the practical use of these ideas, we extend the example of Bob and his Pokemon cards, from the previous chapter, to include the case where some of Bob's friends live near Canada but not actually in Canada, and the inference system has to deal with the notion of "fuzzy Canadian-ness" as related to spatial geometry. The following figure illustrates the fuzzy spatial membership function corresponding to Canada, used in the example inference:

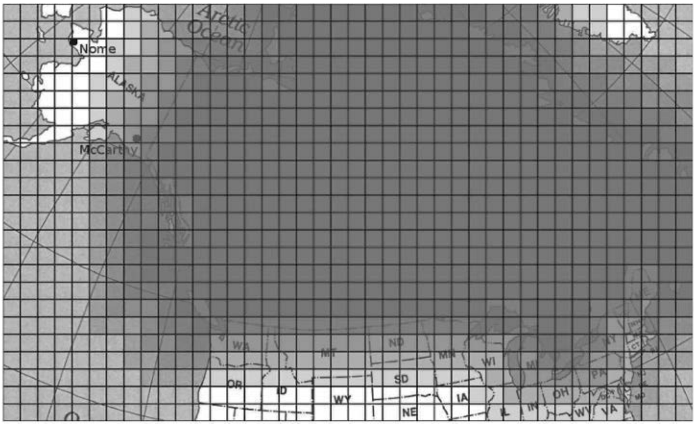

Fig. 1.7

Finally (before Chapter Sixteen which is a brief conclusion), Chapter Fifteen confronts the thorny conceptual and algorithmic issue of inference control: determining which inference steps to take, in which order, in order to answer a question, filter a datamining results list, or carry out an analysis. Far from being "merely an efficiency issue," inference control actually hits many of the deepest issues of AI, including the "frame problem" (briefly stated, that AI systems tend to lack tacit background knowledge about what questions not to bother asking because their answers are supposed to be obvious, or are irrelevant). We discuss a number of specific techniques that may be able to achieve effective inference control in the context of inference on large stores of spatiotemporal logical knowledge, including techniques that hybridize logic with other AI methods such as activation spreading. Here the discussion broadens from logic per se to the topic of “cognitive architectures” and general AI systems, the point being made that the integrative architectures underlying many such systems exist largely in order to provide effective, scalable inference control. As an example, the OpenCog cognitive architecture in which the PLN inference system is embedded is briefly considered.

PART I

Representations and Rules for Real-World Reasoning

Chapter 2

Knowledge Representation Using Formal Logic

Now we begin to dig into the nitty-gritty of our subject matter. Before discussing querying and analysis of complex, heterogeneous spatiotemporal and contextual knowledge, we must discuss representation of temporal knowledge (as well as, to a certain extent, spatial knowledge) ... and before that, we must address knowledge representation in general.

In the course of our investigation we must work through a number of difficult questions regarding knowledge representation, including:

- Which of the many species of uncertain logic to use as the basis for our knowledge representation
- How specifically to represent temporal knowledge?
- How specifically to represent spatial knowledge?
- What is the best low-level (e.g. graph) representation of logical knowledge for efficient storage and processing?

"Logic" itself is not a monolithic entity; it comes in many different flavors. At the highest level, there is the dichotomy between predicate logic and term logic (and there are also systems that hybridize the two, e.g. (Goertzel *et al.*, 2008; Wang, 2006a)). There are also many types of logical system within each of these broad categories, some of which will be reviewed later on.

The material in this chapter becomes somewhat formal and technical, for which we apologize to the reader who lacks the relevant taste or experience; but which unfortunately seems unavoidable if we are to give a serious treatment of our topic. The reader lacking appropriate expertise may either consult relevant background material (Copi & Cohen, 1998), or less ideally, skim this material and proceed to the later chapters, some of which will be quite clearly comprehensible without grasp of these preliminaries, some less so.

2.1 Basic Concepts of Term and Predicate Logic

Term logic, or *traditional logic*, was founded by Aristotle and was the dominating logical framework until the late nineteen century. Term logic uses subject-predicate statements of the form "S is P" (for instance, "Socrates is a man"). There are singular and universal terms (the former correspond to unique subjects). There are just four forms of propositions in term logic:

- Universal and affirmative (e.g. "All men are mortal")
- Particular and affirmative (e.g. "Some men are philosophers")
- Universal and negative (e.g. "No philosophers are rich")
- Particular and negative (e.g. "Some men are not philosophers").

New conclusions are derived from premises by *syllogisms*. Aristotle introduced fourteen syllogisms, of which we will give just two here for illustrative purposes:

- (*Barbara*) If every M is L, and if every S is M, then every S is L. (for instance, "if every man is mortal, and if every philosopher is a man, then every philosopher is mortal")
- *(Celarent)* If no M is L, and if every S is M, then no S is L. (for instance, "if no philosopher is rich and if every poet is a philosopher, then no poet is rich").

Syllogisms provide a method for deduction – deriving new facts from already proved facts.

In addition there are rules for induction and abduction:

- *(Induction)* If every M is L, and if every M is S, then every S is L. (for instance, "if every poet is mortal, and if every poet is a philosopher, then every philosopher is mortal")
- *(Abduction)* If every L is M, and if every S is M, then every S is L. (for instance, "if every poet is mortal, and if every philosopher is mortal, then every philosopher is poet")

Notice that the induction and abduction rules do not neccesarily derive true statements. Nevertheless these are important forms of inference in the face of insufficient evidence, in modern AI reasoning systems as well as in classical Aristotelian term logic (Dimopoulos & Kakas, 1996). Induction and abduction are omnipresent in human commonsense inference.

Put simply, induction aims at *generalization*. In the above example ("if every poet is mortal, and if every poet is a philosopher, then every philosopher is mortal"), the first premise yields that all philosophers that are also poets are mortal, but then it is generalized to conclude that all philosophers are mortal. Yet, it is possible that there are some philoso-