

Abduction and Induction

APPLIED LOGIC SERIES

VOLUME 18

Managing Editor

Dov M. Gabbay, *Department of Computer Science, King's College, London, U.K.*

Co-Editor

John Barwise, *Department of Philosophy, Indiana University, Bloomington, IN, U.S.A.*

Editorial Assistant

Jane Spurr, *Department of Computer Science, King's College, London, U.K.*

SCOPE OF THE SERIES

Logic is applied in an increasingly wide variety of disciplines, from the traditional subjects of philosophy and mathematics to the more recent disciplines of cognitive science, computer science, artificial intelligence, and linguistics, leading to new vigor in this ancient subject. Kluwer, through its Applied Logic Series, seeks to provide a home for outstanding books and research monographs in applied logic, and in doing so demonstrates the underlying unity and applicability of logic.

The titles published in this series are listed at the end of this volume.

Abduction and Induction

Essays on their Relation and Integration

edited by

PETER A. FLACH

University of Bristol

and

ANTONIS C. KAKAS

University of Cyprus



SPRINGER-SCIENCE+BUSINESS MEDIA, B.V.

A C.I.P. Catalogue record for this book is available from the Library of Congress.

ISBN 978-90-481-5433-3

ISBN 978-94-017-0606-3 (eBook)

DOI 10.1007/978-94-017-0606-3

Printed on acid-free paper

All Rights Reserved

© 2000 Springer Science+Business Media Dordrecht

Originally published by Kluwer Academic Publishers in 2000

No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission from the copyright owner.

Contents

Foreword	ix
Preface	xiii
Contributing Authors	xv
1	
Abductive and inductive reasoning: background and issues	1
Peter A. Flach and Antonis C. Kakas	
1.1 Introduction	1
1.2 Abduction and induction in philosophy and logic	2
1.3 Abduction and induction in logic programming and artificial intelligence	11
1.4 Integration of abduction and induction	24
1.5 Conclusions	27
Part I The philosophy of abduction and induction	
2	
Smart inductive generalizations are abductions	31
John R. Josephson	
2.1 A distinctive pattern of inference	31
2.2 What is an explanation?	36
2.3 Smart inductive generalizations are abductions	39
2.4 Conclusion	43
3	
Abduction as epistemic change: a Peircean model in Artificial Intelligence	45
Atocha Aliseda	
3.1 Introduction	45
3.2 Abduction and induction	46
3.3 The notion of abduction	47
3.4 Epistemic change	51
3.5 Abduction as epistemic change	54
3.6 Discussion and conclusions	57
	v

4		
	Abduction: between conceptual richness and computational complexity	59
	Stathis Psillos	
4.1	Introduction	59
4.2	Ampliative reasoning	60
4.3	Explanatory reasoning: induction and hypothesis	62
4.4	Abduction	64
4.5	Abduction and computation	69
4.6	Conclusions	73
	 Part II The logic of abduction and induction	
5		
	On relationships between induction and abduction: a logical point of view	77
	Brigitte Bessant	
5.1	Introduction	77
5.2	Abduction and induction: one is an instance of the other	78
5.3	Abduction and induction: different with a common root	79
5.4	Abduction and induction: totally different	82
5.5	Abduction and induction: a logical inference	84
5.6	Conclusion	86
6		
	On the logic of hypothesis generation	89
	Peter A. Flach	
6.1	Introduction	89
6.2	Logical preliminaries	90
6.3	Explanatory reasoning	93
6.4	Confirmatory reasoning	97
6.5	Discussion	105
7		
	Abduction and induction from a non-monotonic reasoning perspective	107
	Nicolas Lachiche	
7.1	Introduction	107
7.2	Definitions	108
7.3	Abduction and explanatory induction	111
7.4	Abduction and descriptive induction	111
7.5	Discussion	114
7.6	Conclusion	115
8		
	Unified inference in extended syllogism	117
	Pei Wang	
8.1	Term logic and predicate logic	117
8.2	Extended syllogism in NARS	119
8.3	An example	123
8.4	Discussion	125

Part III The integration of abduction and induction: an Artificial Intelligence perspective

9		
	On the relations between abductive and inductive explanation	133
	Luca Console and Lorenza Saitta	
	9.1 Introduction	133
	9.2 Generality and informativeness	134
	9.3 A general definition of explanation	137
	9.4 Inductive and abductive explanations	140
	9.5 Analysis of inference mechanisms in the literature	145
	9.6 Related work	148
	9.7 Conclusions	151
10		
	Learning, Bayesian probability, graphical models, and abduction	153
	David Poole	
	10.1 Introduction	153
	10.2 Bayesian probability	156
	10.3 Bayesian networks	161
	10.4 Bayesian learning and logic-based abduction	162
	10.5 Combining induction and abduction	166
	10.6 Conclusion	168
11		
	On the relation between abductive and inductive hypotheses	169
	Akinori Abe	
	11.1 Introduction	169
	11.2 The relation between abduction and induction	170
	11.3 About the integration of abduction and induction	176
	11.4 Conclusion	179
12		
	Integrating abduction and induction in Machine Learning	181
	Raymond J. Mooney	
	12.1 Introduction	181
	12.2 Abduction and induction	182
	12.3 Abduction in theory refinement	183
	12.4 Induction of abductive knowledge bases	188
	12.5 Conclusions	191
	Part IV The integration of abduction and induction: a Logic Programming perspective	
13		
	Abduction and induction combined in a metalogic framework	195
	Henning Christiansen	
	13.1 Introduction	195
	13.2 A metalogic framework for models of reasoning	197
	13.3 Modelling a variety of reasoning processes	199
	13.4 Implementation of the DEMO system	204
	13.5 Summary and related work	209

14		
Learning abductive and nonmonotonic logic programs		213
Katsumi Inoue and Hiromasa Haneda		
14.1 Introduction		213
14.2 Learning nonmonotonic logic programs		216
14.3 Learning abductive logic programs		223
14.4 Related work		228
14.5 Conclusion		229
Appendix: Proof of Theorem 14.2		229
15		
Cooperation of abduction and induction in Logic Programming		233
Evelina Lamma, Paola Mello, Fabrizio Riguzzi, Floriana Esposito, Stefano Ferilli, and Giovanni Semeraro		
15.1 Introduction		233
15.2 Abductive and Inductive Logic Programming		234
15.3 An algorithm for learning abductive logic programs		239
15.4 Examples		241
15.5 Integration of abduction and induction		248
15.6 Conclusions and future work		250
Appendix: Abductive proof procedure		250
16		
Abductive generalization and specialization		253
Chiaki Sakama		
16.1 Introduction		253
16.2 Preliminaries		254
16.3 Generalizing knowledge bases through abduction		256
16.4 Specializing knowledge bases through abduction		260
16.5 Related work		264
16.6 Concluding remarks		264
17		
Using abduction for induction based on bottom generalization		267
Akihiro Yamamoto		
17.1 Introduction		267
17.2 From abduction to induction		269
17.3 SOLD-resolution		270
17.4 Finding definite clauses		273
17.5 Finding unit programs		278
17.6 Concluding remarks		279
Bibliography		281
Index		301

Foreword

Reasoning in reverse

Logic is the systematic study of cogent reasoning. The central process of reasoning studied by modern logicians is the accumulative *deduction*, usually explained semantically, as taking us from truths to further truths. But actually, this emphasis is the result of a historical contraction of the agenda for the field. Up to the 1930s, many logic textbooks still treated deduction, induction, confirmation, and various further forms of reasoning in a broader sense as part of the logical core curriculum. And moving back to the 19th century, authors like Mill or Peirce included various non-deductive modes of reasoning (induction, abduction) on a par with material that we would recognize at once as ‘modern’ concerns. Since these non-deductive styles of reasoning seemed irrelevant to foundational research in mathematics, they moved out quietly in the Golden Age of mathematical logic. But they do remain central to a logical understanding of ordinary human cognition. These days, this older broader agenda is coming back to life, mostly under the influence of Artificial Intelligence, but now pursued by more sophisticated techniques – made available, incidentally, by advances in mathematical logic...

The present volume is devoted to two major varieties of non-deductive inference, namely *abduction* and *induction*, identified as logical ‘twins’ by C.S. Peirce, but discovered independently under many different names. Roughly speaking, abduction is about finding explanations for observed facts, viewed as missing premises in an argument from available background knowledge deriving those facts. Equally roughly speaking, induction is about finding general rules covering a large number of given observations. Both these phenomena have been studied by philosophers of science since the 1950s, such as Carnap (the pioneer of inductive logic) and Hempel (whose ‘logico-deductive’ model of explanation has unmistakable abductive features). Another major contribution was made by Popper. If good news travels in this forward direction, bad news travels in the opposite. Valid consequences also take us from false conclusions to falsity of at least one of the premises, allowing us to learn by revision – even though we may have some latitude in where to assign the blame. Thus, reasoning is also tied up with scientific theory change, and more generally, flux of our commonsense opinions. What the present volume shows is how these concerns are

converging with those of logicians and computer scientists, into a broader picture of what reasoning is all about.

A pervasive initial problem in this area, even if just an irritant, is *terminology*. Some people feel that ‘abduction’ and ‘induction’, once baptised, must be real phenomena. But they might be mere terms of art, still looking for a substantial denotation... Indeed, it is not easy to give a crystal-clear definition for them, either independently or in their inter-relationship. (Of course, this is not easy for ‘deduction’ either.) Fortunately, the editors do an excellent job in their introductory chapter of clearing up a number of confusions, and relating abduction and induction in a productive manner. No need to repeat that. Instead, let me highlight how the subject of this book intertwines many general features of reasoning that need to be understood – and that somehow manage to escape from the usual logical agenda. For this purpose, we must make some distinctions.

Every type of reasoning revolves around some underlying *connection*, giving us a link with a certain ‘quality’ between input (data) and output (conclusions). The classical schema for this connection is the binary format $P, Q, \dots \vdash C$. But this format leaves many ‘degrees of freedom’, that are essential to reasoning. First, the *strength* of the connection may vary. With Tarski it says “all models of the premises P, Q, \dots are also models for the conclusion C ”. But there are respectable alternatives which ask less: replacing “all” by “almost all” (as happens in some probabilistic reasoning), or by “all most preferred” (as in non-monotonic logics in AI). This variety of ‘styles of reasoning’ can be traced back to the pioneering work of Bolzano in the early 19th century, including an awareness – wide-spread now, but quite novel then – that these different styles differ, not only in the individual inferences they sanction, but also in their general structural rules, such as Monotonicity of Transitivity. Second, varieties of logical consequence multiply by the existence of very different viewpoints on the connection. All variations so far were *semantic*, in terms of models, truth, and preference. But we can also analyse cogent reasoning in a *proof-theoretic* manner (consequence as derivability), giving us options between classical, or intuitionistic, or linear logic – or a *game-theoretic* one (where valid consequence is the existence of a winning strategy for a proponent in debate), giving us yet further logical systems.

Another key dimension to reasoning is *direction*. Standard logical inference moves forward, from given premises to new conclusions. But abduction moves backwards, looking for premises that imply a given conclusion. The backwards direction is often less deterministic, as we can choose from a vast background reservoir of knowledge, prejudices, hypotheses, etc. Indeed, to put backwards reasoning in proper perspective, we need richer formats of inference. An example is Toulmin’s schema from the 1950s, where claims follow from data via a ‘warrant’, where data are backed up by evidence, and warrants by background theory. Thus, we are led naturally to a study of *theory structure*. The latter is prominent in the philosophy of science, where assertions may fall into laws, facts, and auxiliary hypotheses. This structure seems essential to logical analysis of reasoning. Thus, abduction looks for facts, while induction searches for regularities. Indeed, would the same distinctions make sense in the forward direction?

Structured theories fix different roles for assertions. A more radical view would be that such roles are not fixed globally, but just represent a focus in the current *con-*

text. What counts as a ‘relevant’ assertion, or as an ‘explanation’ for a given observation, may depend on the dynamics of some ongoing argument – and the topic of conversation. This epistemic flux reflects a general modern concern. Inferential connections and theory structures are static entities, but actual reasoning is a *dynamic process*, driven by purposes over time. This dynamics is reflected in standard accounts of abduction and induction, as being about ‘finding’ explanations or generalizations. But even standard deduction is usually dynamically goal-driven, say by a conjecture, pulling us toward intermediate results. Of course, we need to ‘compare like to like’: abduction, induction and deduction all have both static and dynamic aspects.

Another dimension of fine-structure is the role of *language* in reasoning. Eventually, one cannot understand the workings of a style of inference without understanding its concomitant language design. E.g., the above distinction between individual facts and general assertions is language-dependent. (I do not know of any conclusive semantic underpinning for it.) Likewise, the choice of concrete and abstract vocabulary is essential to perspicuous theory structure, as is well-known from philosophy of science. But also, different notions of consequence suggests different vocabularies of logical operators reflecting ‘control structures’ of the process, witness classical, intuitionistic, or modal logic. Modern non-standard logics invite even more exotic new operations.

Finally, this book highlights issues of *combination*. Different forms of reasoning do not live in isolated domains, they interact. This is clear even in classical mathematics, where backwards problem Analysis lived side-by-side with forward proof Synthesis. Logic systems like semantic tableaux have this same dual character. In this book, abduction and induction occur intertwined, which raises many additional questions. This is one instance of a more general trend toward understanding the *architecture* of logical systems, and its effects on the *complexity* of their behaviour in bulk.

Reasoning is a many-dimensional process, involving a complex interplay of inferential connections, language design, changing directions, and larger-scale combinations. Naturally, the present Book does not address all this once and for all. But it does throw open windows towards understanding the true complexities of reasoning, by presenting abduction and induction intertwined as a fascinating case study for ‘real logic’.

Johan van Benthem

Preface

From the very beginning of investigation of human reasoning, philosophers had identified – along with deduction – two other forms of reasoning which we now call abduction and induction. Whereas deduction has been widely studied over the years and is now fairly well understood, these two other forms of reasoning have, until now, eluded a similar level of understanding. Their study has concentrated more on the role they play in the evolution of knowledge and the development of scientific theories.

In an attempt to increase our understanding of these two forms of non-deductive reasoning, this book presents a collection of works addressing the issues of the relation between abduction and induction, as well as their possible integration. These issues are approached sometimes from a philosophical perspective, sometimes from a (purely) logical perspective, but also from the more task-oriented perspective of Artificial Intelligence. To a certain extent, the emphasis lies with the last area of Artificial Intelligence, where abduction and induction have been more intensively studied in recent years.

This book grew out of a series of workshops on this topic. The first of these took place at the Twelfth European Conference on Artificial Intelligence (Budapest, August 1996), and concentrated on the general philosophical issues pertaining to the unification or distinction between abduction and induction. The second workshop took place at the Fifteenth International Joint Conference on Artificial Intelligence (Nagoya, August 1997), with an emphasis on the more practical issues of integration of abduction and induction. Taking place in parallel with the preparation of this book, a third workshop was held at the Thirteenth European Conference on Artificial Intelligence (Brighton, August 1998). Detailed reports on the first two workshops have been published as (Flach and Kakas, 1997a; Flach and Kakas, 1998); these reports, as well as further information about the workshops (including submitted papers), are available on-line at <http://www.cs.bris.ac.uk/~flach/abdind/>.

After the first two workshops, we invited the participants to submit a longer paper based on their workshop contribution(s), suitable for publication in an edited volume. Following a careful reviewing process, thirteen of the submitted papers were selected for publication. In addition, we invited four well-known authors to contribute a paper: John Josephson, Luca Console, Lorenza Saitta, and David Poole.

Following a general introduction into the subject, the book is structured into four main parts. The first two parts take a more theoretical perspective, while the remaining two parts address the more practical issue of integrating abduction and induction. Part 1 contains three papers addressing philosophical aspects of abduction and induction. In Part 2, four papers investigate the logical relation of the two forms of reasoning. The four papers in Part 3 deal with integration of the two forms of reasoning from the perspective of Artificial Intelligence, while the five papers that can be found in Part 4 address this problem within the more particular framework of Logic Programming. The book starts off with an introductory chapter aimed at helping the reader in two ways. It provides background material on the general subject of the book and exposes the main issues involved. At the same time it positions the other contributions in the book within the general terrain of debate.

The present book is one of the first books to address explicitly the problem of understanding the relation and interaction between abduction and induction in the various fields of study where these two forms of reasoning appear. As such, it should be relevant to a variety of students and researchers from these different areas of study, such as philosophers, logicians, and people working in Artificial Intelligence and Computer Science more generally.

Acknowledgments

We would like to thank all persons who helped in one way or another with the preparation of this book.

These are all those involved in the organisation of the three workshops on the subject at ECAI'96, IJCAI'97 and ECAI'98, where much of the groundwork for this book has been done. In particular, we would like to thank the other members of the organising committees of these workshops: Henning Christiansen, Luca Console, Marc De-necker, Luc De Raedt, Randy Goebel, Katsumi Inoue, John Josephson, Ray Mooney and Chiaki Sakama. A special thanks goes to the three invited speakers at these workshops, John Josephson, David Poole and Murray Shanahan. And of course we are grateful to all the participants, for the pleasant atmosphere and lively discussions during the workshops. Finally, we would like to thank the two European networks of excellence, Compulog-Net and ML-Net, for their financial support in organising these workshops.

We thank everybody who submitted a paper to this book; those who helped reviewing the submissions; the invited authors for their marvellous contributions; and Johan van Benthem for his beautiful and thought-provoking foreword.

Part of this work falls under the workplan of the ESPRIT project ILP2: Inductive Logic Programming. We wish to thank the other partners of the project for their help and valuable discussions on the subject of the book. We also thank the Universities of Cyprus, Tilburg and Bristol for providing the opportunities to prepare this book.

Special thanks go to Kim and Nada for their patient understanding and support with all the rest of life's necessities, thus allowing us the selfish pleasure of concentrating on research and other academic matters such as putting this book together.

Peter Flach and Antonis Kakas

Contributing Authors

Akinori Abe (ave@cslab.kecl.ntt.co.jp) is a Senior Research Scientist at NTT Communication Science laboratories. He obtained his Doctor of Engineering degree from University of Tokyo in 1991, with a thesis entitled *A Fast Hypothetical Reasoning System using Analogical Case*. His main research interests are abduction (hypothetical reasoning), analogical reasoning and language sense processing. He is a member of the Planning Committee of *New Generation Computing*.

Atocha Aliseda (atocha@filosoficas.unam.mx) is an Associate Professor at the Institute for Philosophical Research of the National Autonomous University of Mexico. She obtained her PhD from Stanford University in 1997, with a thesis entitled *Seeking Explanations: Abduction in Logic, Philosophy of Science and Artificial Intelligence*, which was also published by the Institute for Logic, Language and Computation (ILLC) of the University of Amsterdam, 1997. Her main research interests are abductive logic, heuristic reasoning and the connection between philosophy of science and artificial intelligence. Her homepage is at <http://www.filosoficas.unam.mx/~atocha/home.html>.

Brigitte Bessant (bessant@cmi.univ-mrs.fr) is a Lecturer in computer science at University of Artois in France. She obtained her PhD in 1999, with a thesis entitled *Contributions to techniques of belief revision in artificial intelligence: semantic and practical aspects*. Her main research interests are nonmonotonic logics, commonsense reasoning, belief revision, update and machine learning.

Henning Christiansen (henning@ruc.dk) is an Associate Professor at the Computer Science Department of Roskilde University, Denmark. He obtained his PhD from Roskilde University in 1988, with a thesis entitled *Programming as language development*. His main research interests are logic programming with emphasis on metaprogramming, constraints, and abduction, and query-answering systems. He is co-editor of books on *Flexible Query-Answering Systems* (Kluwer, 1997; Springer, 1998). His homepage is at <http://www.dat.ruc.dk/~henning/>.

Luca Console (Luca.Console@di.unito.it) is an Associate Professor of Computer Science at the Dipartimento di Informatica of the Università di Torino. His main research interests regard reasoning mechanism, with specific attention to model-based reasoning and diagnosis, temporal reasoning, abductive reasoning, adaptive systems. He is author of several papers, of introductory books and editor of collections on model-based diagnosis.

Floriana Esposito (esposito@di.uniba.it) is a Professor of Computer Science at DIB, University of Bari. She graduated in Electronic Physics at University of Bari in 1970 and since 1974 was Assistant Professor of Computer Science. Her research interests are artificial intelligence, machine learning, programming languages and symbolic computation.

Stefano Ferilli (ferilli@di.uniba.it) is currently a PhD student in Computer Science at DIB, University of Bari. He graduated in Computer Science at University of Bari in 1996. His research interests include logic programming, machine learning and theory revision.

Peter A. Flach (Peter.Flach@bristol.ac.uk) is a Lecturer at the Computer Science Department of the University of Bristol. He obtained his PhD from Tilburg University in 1995, with a thesis entitled *Conjectures: an inquiry concerning the logic of induction*. His main research interests are inductive logic programming, intelligent reasoning, and philosophy of artificial intelligence. He is author of the textbook *Simply Logical – intelligent reasoning by example* (John Wiley, 1994). He is academic coordinator of ILPnet2: the European Network of Excellence in Inductive Logic Programming. His homepage is at <http://www.cs.bris.ac.uk/~flach/>.

Hiromasa Haneda (haneda@kobe-u.ac.jp) is a Professor at the Department of Electrical and Electronics Engineering at Kobe University, Japan. He obtained his PhD from the University of California, Berkeley in 1972 in the area of computer-aided analysis of electronic circuits and systems. His main research interests are machine learning as applied to computer-aided design and analysis of industrial systems.

Katsumi Inoue (inoue@eedept.kobe-u.ac.jp) is an Associate Professor at the Department of Electrical and Electronics Engineering at Kobe University, Japan. He obtained a Doctor of Engineering from Kyoto University in 1993 with a thesis entitled *Studies on abductive and nonmonotonic reasoning*. His main research interests are automated reasoning, knowledge representation, machine learning, and logic programming. His homepage is at <http://cslab.eedept.kobe-u.ac.jp/~inoue/>.

John R. Josephson (jj@cis.ohio-state.edu) is a Research Scientist and the Associate Director of the Laboratory for AI Research (LAIR) in the Department of Computer and Information Science at the Ohio State University. He received his Ph.D. in Philosophy (of science) from Ohio State in 1982; he also holds B.S. and M.S. degrees in Mathematics from Ohio State. His primary research interests are artificial intelligence, knowledge-based systems, abductive inference, causal reasoning, theory formation, perception, diagnosis, the logic of investigation, and the foundations of science. He has worked in several application domains including: medical diagnosis, diagnosis of engineered systems, logistics planning, speech recognition, genetics, molecular biology, and design of electro-mechanical systems. He is the co-editor, with Susan G. Josephson, of *Abductive Inference* (Cambridge University Press, 1994, 1996). His homepage is at <http://www.cis.ohio-state.edu/~jj/>.

Antonis C. Kakas (antonis@ucy.ac.cy) is an Associate Professor at the Computer Science Department of the University of Cyprus. He obtained his PhD in Theoretical Physics from Imperial College, London, in 1984. In 1989 he started working in Computational Logic and Artificial Intelligence. His main research interests are abduction, with specific interest in the integration of abductive, inductive and constraint logic programming and its applications in the areas of planning and information integration, argumentation and the theory of actions and change. He is the editor in chief of the magazine *Computational Logic* published by *Compulog-Net: the European Network of Excellence in Computational logic*.

Nicolas Lachiche (lachiche@iutsud.u-strasbg.fr) is a Lecturer in Computer Science at the University Robert Schuman (Strasbourg, France). He was previously a Research Associate at the Computer Science Department of the University of Bristol. He obtained his PhD from the University of Nancy in 1997, with a thesis focusing on classification and descriptive induction, and their relations. His research mainly concerns data mining and machine learning, from both supervised and unsupervised perspectives, in either an attribute-value representation or a first-order logic language.

Evelina Lamma (elamma@deis.unibo.it) is an Associate Professor at DEIS, University of Bologna. She graduated in Electrical Engineering in 1985 at the University of Bologna and obtained her PhD in Computer Science in 1990. Her main research interests are artificial intelligence, and extensions of logic programming in particular.

Paola Mello (pmello@deis.unibo.it) is a Professor at DEIS, University of Bologna. She graduated in Electrical Engineering in 1982 at the University of Bologna and obtained her PhD in Computer Science in 1989. Her main research interests are artificial intelligence, and extensions of logic programming in particular.

Raymond J. Mooney (mooney@cs.utexas.edu) is an Associate Professor in the Department of Computer Sciences at the University of Texas at Austin. He received his Ph.D. in 1988 from the University of Illinois at Urbana/Champaign with a thesis on explanation-based learning. He is an editor for the journal *Machine Learning* where he recently co-edited a special issue on natural language learning. His current research interests include natural language learning, knowledge-base refinement, inductive logic programming, and learning for text categorization and recommender systems. His homepage is at <http://www.cs.utexas.edu/users/mooney/>.

David Poole (poole@cs.ubc.ca) is a Professor of Computer Science at the University of British Columbia. He received his Ph.D. from the Australian National University in 1984. He is known for his work on knowledge representation, default reasoning, assumption-based reasoning, diagnosis, reasoning under uncertainty, and automated decision making. He is a co-author of a recent AI textbook, *Computational Intelligence: A Logical Perspective* (Oxford University Press, 1998), co-editor of the *Proceedings of the Tenth Conference in Uncertainty in Artificial Intelligence* (Morgan Kaufmann, 1994), serves on the editorial board of the Journal of AI research, and is a principal investigator in the Institute for Robotics and Intelligent Systems. His homepage is at <http://www.cs.ubc.ca/spider/poole/>.

Stathis Psillos (psillos@netplan.gr) is a Lecturer at the Department of Philosophy and History of Science, University of Athens. Between 1995-1998 he was a British Academy Postdoctoral Fellow at the London School of Economics. He completed his Ph.D in 1994, at King's College London. His book *Scientific Realism: How Science Tracks Truth* is due to appear in November 1999 by Routledge.

Fabrizio Riguzzi (friguzzi@deis.unibo.it) is currently affiliated to DEIS, University of Bologna. He obtained his PhD from the University of Bologna in 1999, with a thesis entitled *Extensions of Logic Programming as Representation Languages for Machine Learning*. His main research interests are logic programming, machine learning and inductive logic programming in particular.

Chiaki Sakama (sakama@sys.wakayama-u.ac.jp) is an Associate Professor at the Department of Computer and Communication Sciences of Wakayama University. He obtained his Doctor of Engineering degree from Kyoto University in 1995, with a thesis entitled *Studies on Disjunctive Logic Programming*. His research interests include abductive/inductive logic programming, nonmonotonic reasoning, and belief revision. His homepage is at <http://www.sys.wakayama-u.ac.jp/~sakama/>.

Lorenza Saitta (saitta@di.unito.it) is a Professor of Computer Science at the Universita' del Piemonte Orientale, Alessandria, Italy. Her main research interests are in Machine Learning, specifically learning relations, multistrategy learning, and

complexity issues. Recently, she became also interested in Genetic Algorithms and Cognitive Sciences. She has been the Chairperson of the *International Conference on Machine Learning* in 1996.

Giovanni Semeraro (semeraro@di.uniba.it) is an Associate Professor at DIB, University of Bari. He graduated in Computer Science in 1988. He joined the University of Bari in 1991. His research interests are centered on the logical and algebraic foundations of inductive inference, document classification and understanding, multi-strategy learning, theory revision and intelligent digital libraries.

Pei Wang (pwang@cogsci.indiana.edu) is the Director of Artificial Intelligence at IntelliGenesis Corporation, and an Adjunct Researcher at the Center for Research on Concepts and Cognition, Indiana University. He obtained his BS and MS in Computer Science from Peking University, and his PhD in Computer Science and Cognitive Science from Indiana University. His main research interests are the foundation of intelligence, reasoning with uncertainty, learning and adaptation, and decision making under time pressure. His publications are available at <http://www.cogsci.indiana.edu/farg/peiwang/papers.html>.

Akihiro Yamamoto (yamamoto@meme.hokudai.ac.jp) is an Associate Professor of the Division of Electronics and Information Engineering at Hokkaido University. He is also a researcher of Precursory Research for Embryonic Science and Technology (PRESTO) at Japan Science and Technology Corporation (JST). He obtained the Dr. Sci. degree from Kyushu University in 1990, with a thesis entitled *Studies on Unification in Logic Programming*. His main research interests are logic programming and its applications to abduction, inductive inference and computer networks.

1 ABDUCTIVE AND INDUCTIVE REASONING: BACKGROUND AND ISSUES

Peter A. Flach and Antonis C. Kakas

1.1 INTRODUCTION

This collection is devoted to the analysis and application of abductive and inductive reasoning in a common context, studying their relation and possible ways for integration. There are several reasons for doing so. One reason is practical, and based on the expectation that abduction and induction are sufficiently similar to allow for a tight integration in practical systems, yet sufficiently complementary for this integration to be useful and productive.

Our interest in combining abduction and induction is not purely practical, however. Conceptually, the relation between abduction and induction is not well understood. More precisely, there are several, mutually incompatible ways to perceive this relation. For instance, Josephson writes that ‘it is possible to treat every good (...) inductive generalisation as an instance of abduction’ (Josephson, 1994, p.19), while Michalski has it that ‘inductive inference was defined as a process of generating descriptions that imply original facts in the context of background knowledge. Such a general definition includes inductive generalisation and abduction as special cases’ (Michalski, 1987, p.188).

One can argue that such incompatible viewpoints indicate that abduction and induction themselves are not well-defined. Once their definitions have been fixed, studying their relation becomes a technical rather than a conceptual matter. However, it is not self-evident why there should exist absolute, Platonic ideals of abduction and induction, waiting to be discovered and captured once and for all by an appropriate defini-

tion. As with most theoretical notions, it is more a matter of pragmatics, of how useful a particular definition is going to be in a particular context.

A more relativistic viewpoint is often more productive in these matters, looking at situations where it might be more appropriate to distinguish between abduction and induction, and also at cases where it seems more useful to unify them. Sometimes we want to stress that abduction and induction spring from a common root (say hypothetical or non-deductive reasoning), and sometimes we want to take a finer grained perspective by looking at what distinguishes them (e.g. the way in which the hypothesis extends our knowledge). The following questions will therefore be our guidelines:

- When and how will it be useful to unify, or distinguish, abduction and induction?
- How can abduction and induction be usefully integrated?

Here and elsewhere, by *unification* we mean considering them as part of a common framework, while by *integration* we mean employing them together, in some mutually enhancing way, for a practical purpose.

The current state of affairs with regard to these issues is perhaps most adequately described as an ongoing debate, and the reader should look upon the following chapters as representing a range of possible positions in this debate. One of our aims in this introductory chapter is to chart the terrain where the debate is taking place, and to position the contributions to this volume within the terrain. We will retrace some of the main issues in this debate to their historical background. We will also attempt a synthesis of some of these issues, primarily motivated by work in artificial intelligence, sometimes taking positions that may not be shared by every author in this volume.

The outline of this chapter is as follows. In Section 1.2 we discuss the philosophical and logical origins of abduction and induction. In Section 1.3 we analyse previous work on abduction and induction in the context of logic programming and artificial intelligence, and attempt a (partial) synthesis of this work. Section 1.4 considers the integration of abduction and induction in artificial intelligence, and Section 1.5 concludes.

Before we embark on this, let us express our sincere thanks for all authors contributing to this volume, without whom we couldn't have written this introduction – indeed, some of the viewpoints we're advocating have been strongly influenced by the other contributions. Wherever possible we have tried to indicate the original source of a viewpoint we discuss, but we apologise in advance for any omissions in this respect.

1.2 ABDUCTION AND INDUCTION IN PHILOSOPHY AND LOGIC

In this section we discuss various possible viewpoints on abduction and induction that can be found in the philosophical and logical literature. The philosophical issue is mainly one of categorisation (which forms of reasoning exist?), while the logical issue is one of formalisation.

As far as categorisation is concerned, it seems uncontroversial that deduction should be singled out as a separate reasoning form which is fundamentally different from any other form of reasoning by virtue of its truth-preserving nature. The question, then, is how non-deductive reasoning should be mapped out. One school of thought holds