Abduction and Induction in AI: Report of the IJCAI'97 Workshop

Peter A. Flach
Dept. of Computer Science, University of Bristol
Merchant Venturers Building, Woodland Road, Bristol BS8 1UB, United Kingdom
Peter.Flach@cs.bris.ac.uk

Antonis Kakas
Dept. of Computer Science, University of Cyprus
POBox 537, CY-1678 Nicosia, Cyprus
antonis@turing.cs.ucy.ac.cy

January 6, 1998

1 Introduction

Over twenty researchers participated in this IJCAI'97 workshop to dicuss the general relationship of abduction and induction. The main aim of the workshop was set out to address the issue of the relation and integration of Abdcution and Induction in the context of (practical) AI problems. This focus was summarized by the following questions:

- What (if anything) distinguishes/characterizes each form of reasoning and their corresponding computational models?
- How can we characterize different prototypical AI tasks for which it is appropriate to use one of these two forms of reasoning?
- How can Abduction and Induction be integrated in the context of Artificial Intelligence problems? For example,
  - How do we learn abductive theories?
  - How do we use abduction in machine learning problems?

The workshop's organising committee consisted of Peter Flach (then at Tilburg University, Netherlands), Antonis Kakas (University of Cyprus), Raymond Mooney (University of Texas at Austin, USA) and Chiaki Sakama (Wakayama University, Japan). The submitted papers were reviewed and selected by a program committee additionally including Randy Goebel (University of Alberta, Canada), Katsumi Inoue (Kobe University, Japan) and John Josephson (Ohio State University, USA). The workshop assumed the same format as the the preceding ECAI'96 workshop on Abductive and Inductive Reasoning [13], devoting ample time to plenary discussions.

However, this workshop had a different emphasis from its predecessor, which indicates that a certain amount of progress has been made. Whereas the discusson at the ECAI'96 workshop revolved almost exclusively around the question whether and how abduction and induction are different forms of reasoning, the discussion at the IJCAI'97 workshop took, for most of its part, a more practical stance, concentrating on how to integrate them within an Artificial Intelligence context. Among the 12 accepted papers the program committee selected 9 for presentation at the workshop. The presentations were evenly divided over 3 sessions, addressing the issues pertaining to the integration of abduction and induction in a bottom-up fashion starting from specific practical problem of their integration and ending up with theoretical issues. These sessions are reviewed in more detail below.

2 Invited talk: David Poole

David Poole (Universiy of British Columbia, Canada) gave an invited talk entitled *Learning, Bayesian Probability, Graphical Models, and Abduction*.[1] In his talk he tried to tie together logic and probabilistic approaches to induction in terms of belief networks and probabilistic Horn abduction. Belief (Bayesian) networks are a graphical representation of independence and provide a way to structure knowledge and to exploit the structure for computational gain. Poole pointed out the relationship between belief networks and logic-based (abductive) representations for evidential reasoning. If we want to do evidential reasoning (from effects to causes) without knowing the underlying process, there is a choice between causal modelling (learning cause→effect rules) and evidential modelling (learning effect→cause rules); we can model causally and use abduction for evidential reasoning (as do abductive diagnosis and belief networks) or model evidentially and use deduction for evidential reasoning (as in neural networks and consistency-based diagnosis). Poole overviewed the tradeoffs in this choice.

One of the most interesting aspects of Poole's talk was that it approached the issue of abduction vs. induction from both a probabilistic and a logical perspective. For instance, he argued that Bayesian conditioning (the kind of evidential reasoning done in Bayesian networks to find causes for observed effects) can be seen as abduction in probabilistic logic programs, where we try to explain the evidence conditioning on all the knowledge obtained since the knowledge base was built. He also provided a view of Bayesian learning as essentially abductive, where the abductive step is finding the right parameter values. From this viewpoint, then, the conclusion that there is no essential difference between abduction and induction seems inescapable.

In his conclusion, Poole argued that the logical and the probabilistic approaches can learn from each other. Bayesians have good methods to handle noise and avoid overfitting, a universal method for finding explanations (conditioning), and good algorithms for exploiting sparse structures. They lack however the rich representations used by the logicians.

3 Discussion panel 1: Abduction and Inductive Logic Programming

Chiaki Sakama (Wakayama University, Japan) discussed ways to use abduction in the process of induction [10]. If a clause can only entail an example by adopting an abductive explanation, a generalisation is obtained by dropping the abducible from the clause. Conversely, a knowledge base that entails a negative example can be specialised by weakening the causes for the entailment by disjoining them with new abnormality abducibles.

Takashi Kanai and Susumu Kunifuji (Japan Advanced Institute of Science and Technology) proposed a new integrated method of inductive generalisation and abductive reasoning [6]. Abduction is used to supplement incomplete background knowledge. The approach is a variant of FOIL [14], including an improved information gain heuristic to deal with the cost of abductive explanations.

Akihiro Yamamoto (visiting Technische Hochschule Darmstadt, Germany; now at Hoikkaido University, Japan) made a connection between inductive hypothesis formation and proof procedures for consequence finding. The proof procedure is a special case of SOL-resolution [15], extending SLD-resolution by allowing to skip the proof of selected subgoals, adding them to an abductive explanation instead. This is then related to Muggleton's inverse entailment operator [16].

The papers in this session essentially concentrated on abduction as a tool in the learning process. Immediate questions are then when this is needed or appropriate, and how it can be done. The ability to extend imperfect background knowledge abductively was generally found to be useful, as domain theories are often incomplete in practice. One can go one step further and argue that also the learned hypothesis will in general be incomplete (i.e. nonmonotonic) to some extent, in which case we are not just using but *learning* abductive theories. This obviates the need for learning not only classification rules but also integrity constraints, which are needed in abductive logic programming to constrain the possible explanations. Presumably having an abductive coverage relation also influences the generality structure of the hypothesis space, and may have a profound impact on the learning algorithm. Finally, there should be a reasonable trade-off between treating a wrongly covered negative example as an exception vs. revising the hypothesis.

[1]The slides for David Poole's talk can be found at `http://www.cs.ubc.ca/spider/poole/talks/ind-ab.pdf`

4 Discussion panel 2: Abduction and Induction – their relation and integration

Raymond Mooney (University of Texas at Austin, USA) [8] presented an overview of work on the integration of abduction and induction in machine learning systems that his group has been doing over the last years. He argued that each inference can strengthen the other and presented practical applications to support this. In particular, he showed how abductive reasoning can be useful in inductively revising existing knowledge bases by finding appropriate places of "repair" of the knowledge base. Also he showed how inductive learning can be used to form theories for abductive reasoning.

Akinori Abe (NTT Communication Science Laboratories, Japan) [1] observed that abduction and induction can be seen as dual to each other and can be linked together when we are in a situation of similar observations. Abductive explanations for similar observations can form suitable data for inductive generalization. He presented a framework for Analogical Abductive Reasoning and showed how when this is applied to similar observations its generated hypotheses can form good examples for generalization.

Pinar Öztürk (Norwegian University of Science and Technology) was unfortunately unable to attend the workshop and present her paper [9].

The second discussion complemented in some sense the first one by considering ways in which induction could be useful in abduction. One interesting possibility would be to abductively explain several observations seperately, and then to inductively generalise them into a single explanation. In this way we are extending the capabilities of an abduction system by generating explanations that are non-ground rules. From the induction perspective we are learning rules for non-observed predicates, which points at a link with *descriptive* induction. An important issue that came up was whether the fact that induction typically requires many observations whereas abduction proceeds from a single or few observations is at all relevant. A number of people supported the intuition that abduction deals with incomplete knowledge regarding a *single* situation, whereas induction extends incomplete knowledge about a *class* of situations. Others objected by pointing at the difficulty of defining the notion of a situation in a non-syntactic way.

5 Discussion panel 3: Unifying foundations of abduction and induction

Geert-Jan Kruijff (Charles University, Czech Republic) [7] pointed out the importance of novelty in abductive reasoning and argued that this forms an important criterion for comparing the two forms of reasoning. Furthermore, he distinguished the two schools of "unification" and "cooperation" for the relation of abduction and induction and argued for the latter. In particular, he argued that induction can lend further credibility to abductive hypotheses.

Peter Grünwald (CWI, the Netherlands) [5] argued that probability can provide a conceptual unification for the two inferences. The difference lies in the ontology of the output of the inference: abduction generates data while induction generates hypotheses (in the statistical sense). He showed how the Minimum Description Length Principle, which is frequently used as a heuristic in inductive learning, is also appropriate for abductive reasoning, thus supporting the claim that the two forms of reasoning can be unified under probability.

Pei Wang (Indiana University, USA) was unfortunately unable to attend the workshop and present his paper [11].

One issue that came up in the third discussion was that induction seems to do a better job in assigning credibility to hypotheses. This is not just because induction is typically based on many observations, but also because inductive hypotheses allow us to make predictions about unseen cases and verify them by cross-validation. Some participants argued against this by pointing out that also abductive hypotheses can have additional consequences that can be verified. Related discussion topics were whether and how we can use induction to increase the credibility of abductive hypotheses, and whether the selection criteria really differ for abductive and inductive hypotheses.

6 Concluding remarks

While the preceding ECAI'96 workshop was successful in bringing people from different disciplines together and identifying some of the main general issues, this IJCAI'97 workshop approached the issue of integrating abduction and induction from a more practical AI perspective. The main conclusion to be drawn from these two workshops is that whether one perceives abduction and induction as two of a kind or as fundamentally different reasoning forms depends strongly on the domain of application and the particular AI approach employed. Hence, an appropriate question to ask is not "What is the relation between abduction and induction", but rather "What are good reasons for perceiving them as fundamentally different or fundamentally similar?" The successor workshop to be organised at ECAI'98 will take this more relativistic perspective as its starting point.

Another way in which this workshop differed from its predecessor is that several participants (e.g. Poole and Grünwald) approached the issue from a probabilistic perspective. The debate between logical and symbolic approaches is an important one because, as Poole pointed out, both can profit from the other's expertise. The discussions at the workshop tentatively suggested that from the probabilistic viewpoint there is no essential difference between abduction and induction. However, one must keep in mind that logical approaches tend to concentrate on hypothesis formation, while probabilistic approaches are concerned with evaluation and selection of hypotheses. It may very well be that the evaluation process is the same for abduction and induction (and indeed for any other form of reasoning), while the formation process is different in each case.

Our experience with both of these worskhops has shown that it is premature to expect universally agreed positions on these difficult issues. Nevertheless, one generally accepted conclusion of this workshop was that, if we perceive abduction and induction as separate inferences, these can be integrated in a **cycle** of knowledge generation governed by the 'equation' $B \cup H \models O$ where H is the new knowledge generated. On one side of the cycle this new knowledge then feeds in the place of O and on the other side it feeds in the place of B. Depending on where we break this cycle we identify the separate inferences of abduction and induction: abduction generating new elements for O and induction for B. This then raises the important question of how can one inference be used to justify (or affect the selection of) the hypotheses generated by the other inference.

The workshop notes contain 12 short papers and are available on-line through the workshop's WWW-pages (`http://www.cs.bris.ac.uk/~flach/IJCAI97/`).

Acknowledgements

This workshop has been made possible by financial support from the European Network of Excellence CompulogNet. Writing of this report has been partially supported by the Esprit Long Term Research Project 20237 (Inductive Logic Programming 2).

References

[1] Akinori Abe, 'The relation between abductive hypotheses and inductive hypotheses', *Proc. IJCAI'97 Workshop on Abduction and Induction in AI*, pp. 1-6.

[2] John Bell, 'Inductive, abductive and pragmatic reasoning', *Proc. IJCAI'97 Workshop on Abduction and Induction in AI*, pp. 7-12.

[3] Philippe Codognet, 'Abductive reasoning: backward and forward', *Proc. IJCAI'97 Workshop on Abduction and Induction in AI*, pp. 13-16.

[4] Randy Goebel, 'Abduction and its relation to constrained induction', *Proc. IJCAI'97 Workshop on Abduction and Induction in AI*, pp. 17-18.

[5] Peter Grünwald, 'The Minimum Description Length principle and non-deductive inference', *Proc. IJCAI'97 Workshop on Abduction and Induction in AI*, pp. 19-23.

[6] Takashi Kanai & Susumu Kunifuji, 'Extending inductive generalisation with abduction', *Proc. IJCAI'97 Workshop on Abduction and Induction in AI*, pp. 25-30.

[7] Geert-Jan Kruijff, 'Concerning logics of abduction – on integrating abduction and induction', *Proc. IJCAI'97 Workshop on Abduction and Induction in AI*, pp. 31-36.

[8] Raymond Mooney, 'Integrating abduction and induction in Machine Learning', *Proc. IJCAI'97 Workshop on Abduction and Induction in AI*, pp. 37-42.

[9] Pinar Öztürk, 'An AI criterion for an account of inference: how to realize a task', *Proc. IJCAI'97 Workshop on Abduction and Induction in AI*, pp. 43-48.

[10] Chiaki Sakama, 'Inductive extension of abduction', *Proc. IJCAI'97 Workshop on Abduction and Induction in AI*, pp. 49-52.

[11] Pei Wang, ‘Return to term logic’, *Proc. IJCAI’97 Workshop on Abduction and Induction in AI*, pp. 53-57.

[12] Akihiro Yamamoto, ‘Representing inductive inference with SOLD-resolution’, *Proc. IJCAI’97 Workshop on Abduction and Induction in AI*, pp. 59-63.

[13] Peter Flach & Antonis Kakas, ‘Abductive and Inductive Reasoning: report of the ECAI’96 workshop’, *Logic Journal of the IGPL* 5(5):773-778, 1997.

[14] Ross Quinlan, ‘Learning logical definitions from relations’, *Machine Learning* 5(3):239-266, 1990.

[15] Katsumi Inoue, ‘Linear resolution for consequence finding’, *Artificial Intelligence* 56:301-353, 1992.

[16] Stephen Muggleton, ‘Inverse entailment and Progol’, *New Generation Computing* 13:245-286, 1995.