

Erik Sandewall

A Review of the Handbook of Knowledge Representation

This series of memos contains research reports and tutorial texts from the project on the Knowledge Representation Framework (KRF).

The present memo can persistently be accessed as follows:

Memo persistent URL: <http://piex.publ.kth.se/reports/krf/00n/>

Date of manuscript: 2008-10-01

Related information can be obtained through the following www sites:

KRF website: <http://piex.publ.kth.se/krf/>

The author: <http://www.ida.liu.se/~erisa/>

The newly appeared *Handbook of Knowledge Representation* is an impressive piece of work. Its three editors and its fortyfive contributors have produced twentyfive concise, textbook-style chapters that introduce most of the major aspects of the science of knowledge representation. Reading this book is a very positive experience: it demonstrates the breadth, the depth and the coherence that our field has achieved by now.

The Handbook is divided into three major sections that are intended for, in turn, general methods, domain-specific methods, and application areas. General methods include for example SAT solvers, description logics and nonmonotonic logics whereas domain-specific methods include the representation of time and change, qualitative modelling, reasoning about actions, and related topics. Application areas is understood as question-answering systems, cognitive robotics, semantic web and a few other, similar titles.

One important and difficult challenge for a scholarly handbook is to achieve coherence on several levels: uniformity of notation, consistency and non-redundancy of definitions and terms, and not too much duplication of contents – although some overlap of contents between the chapters may in fact be desirable. The Handbook of Knowledge Representation makes a good job in these respects. One can always find specific points where it could be further improved, but it appears to do considerably better than other, comparable handbooks. The logicist framework that is introduced in Chapter 1 is a de facto standard for most of the succeeding chapters, and there are frequent cross-references between the chapters. This bears witness of a concerted effort by the editors and the authors, but it is also a sign of the maturity of the field.

Although the Handbook is quite comprehensive, it does omit some important areas, including some classical topics in Knowledge representation. Nonmonotonic multiple inheritance is not covered. Reasoning under uncertainty is only addressed in the chapter on Bayesian networks. Learning is not addressed at all – it is an interesting question whether Learning in A.I. is entirely disjoint from Knowledge representation, and if so, why. The analysis of natural language is only addressed in the introductory part of the chapter on “Knowledge representation and question-answering”, and not in the systematic manner that one would have expected. These are not major defaults, however, in the sense that the chapters that are present in the Handbook do not suffer from the absence of these specific topics.

I see a bigger problem with the chapter on Planning. This chapter is introduced in the Preface as a survey of advances in automated planning. It is none of the sort, however; the chapter proposes a general framework for methods for planning, at least one level more abstract than methods for planning as usually understood. Many readers would presumably be interested in a chapter that begins with basic concepts, such as progressive vs regressive planning, hierarchical planning, and so forth, and that continues with the designs of major actual planning systems. However, such a chapter is not found in the Handbook. This is unfortunate since this Handbook would have been an excellent place for clarifying the relationships between Reasoning about actions, Planning, and Cognitive robotics.

In principle, later chapters of the Handbook rely on earlier chapters for definitions, conceptual framework, and methods. This works quite well in general, as already remarked, but one can find some points where additional connections would have been in place. For example, a section on abduction

in some of the early chapters could have been a useful basis for the treatment of diagnostic reasoning and of planning, in later chapters. (Chapter 1 only mentions abduction very briefly). If the Handbook is used as a textbook then the instructor should be prepared to point out such additional aspects and connections between the topics.

Another important design decision for a scholarly handbook concerns how to represent multiple schools of thought that address the same topic, but without very much interaction between them. Should one allow one chapter for each of these schools, so that the presentation is well in line with the literature within each school, or should one try to obtain a synthesis? The latter approach may help to advance the research in the area by introducing a framework that can be shared by all those concerned, but it may also alienate all of them. As it happens, the Handbook contains examples of both approaches. The most obvious example of multiple-school presentation concerns the topic of Reasoning about actions and change, which obtains four chapters: for the Situation Calculus, Event calculus, Time and action logics, and Nonmonotonic causal logic. This was probably unavoidable, but it does lead to redundancy as the various aspects of the frame problem are introduced repeatedly by the different authors.

While the Handbook is an excellent presentation of the achievements of Knowledge representation research, it illustrates at the same time its peculiarities and shortcomings. I would like to use this occasion for discussing one of these, namely, the role that actual knowledge bases play in our field, or more precisely, the weakness of their role. We all realize that the body of human knowledge is enormous, and that even the body of common-sense knowledge and common knowledge is enormous, so we must look for methods that scale up to the required size. It is also clear that linguistic knowledge, commonsense knowledge of facts, and commonsense ontologies are strongly connected and require a common representation scheme.

It is mentioned in a few places, in the Handbook, that knowledgebases of these kinds actually exist and are being used. Cyc is mentioned in one place, Wordnet in another place. It is mentioned that a large number of ontologies, expressed in OWL, are available on the web. Other important knowledgebases, such as Dbpedia (derived from the Wikipedia) and Freebase are not mentioned. But in any case, all of these knowledgebases appear to be entirely marginal to the topic of the Handbook.

If you think about it, this is very strange. From an apriori point of view, it would seem natural to *begin* a handbook such as this one with an account of existing knowledgebases: how are they organized, what notation (formalism, markup language) do they use, what can be said about their information quality, and so on. It would also be natural to provide some excerpts of nontrivial size from some of these knowledgebases for use in later chapters or, at least, to direct the reader to websites where these knowledgebases can be browsed conveniently, using the same notation as is used in the Handbook.

After such an introduction, subsequent chapters of the Handbook should be able to refer to the introductory knowledgebase excerpts in order to motivate concepts and methods of each chapter, and also in order to show nontrivial examples of how the methods can be used.

As it stands, the Handbook is fairly short of examples. The chapter on Time

and action logic contains a somewhat nontrivial example of reasoning about a sequence of events. The chapter on question-answering also develops two examples, and there is one or two more in other chapters. However, these examples are small and local, and none of them is anchored in a larger knowledgebase.

In fact, the Handbook does not even contain a chapter on ontologies (actual ontologies, not methods for the design and the alignment of ontologies). It would have been useful to add a chapter in section I for an overview of major ontologies that are of interest for Knowledge representation, and it should have been possible for later chapters to build on it.

A possible answer to these observations may be that the structure and contents of knowledgebases are still under development and therefore have no place in the Handbook. However, this immediately begs the question why even their most basic parts must still be in flux, after several decades of research in our field.

These observations therefore apply to our entire field and are not specific to the Handbook. In fact, the Handbook is entirely consistent with the style of research articles on Knowledge representation in our journals and conferences, and this is as it should be. I propose however that our field is missing a very important part of its mission when we focus so much on logic, algorithms and other *methods* for the representation and the computational use of knowledge, and when we disregard so regularly the actual bodies of knowledge that our methods are supposed to operate on.