
Book Reviews

VIGGO STOLTENBERG-HANSEN, INGRID LINDSTRÖM AND EDWARD R. GRIFFOR
Mathematical Theory of Domains. Cambridge University Press. 1994. ISBN 0-521-38344-7 £27.50. 349 pp. hardbound.

'What is the set of values of a given datatype?' Innocent though this question may look, it turns out to be a misleading one if you want to give a mathematical account of datatypes and, subsequently, of the meaning of programs. For a proper mathematical treatment you must know not only what the values are, but how they relate to each other: which are finite, which are infinite, which are just partial approximations and what to. The mathematical structures embodying this extra information are called 'domains', so the opening question is then rephrased as, 'What is the *domain* of values of a given datatype?'

The deep nature of domains is somewhat murky, both technically (there are a number of rather different sorts with different properties) and foundationally (for instance, domains can be described either as partially ordered sets or as topological spaces). Nonetheless, there is a more or less stable mathematical account which can already be found in various books such as Winskell's *Formal Semantics of Programming Languages* or Gunter's *Semantics of Programming Languages* (or, indeed, in Plotkin's classic lecture notes, which are now readily available in electronic form).

The new book, *Mathematical Theory of Domains*, is unlikely to supplant the older ones as an introduction to the basic theory. On the whole it limits itself to Scott domains. Continuous domains (essential for domains involving real numbers) are not described at all and the SFP domains that are vital for some purposes are given only a short account towards the end. This means that the view of domains that is presented is a slightly narrow one. Moreover, though a whole chapter is devoted to giving an introduction to topology, the topological aspect of domains is presented rather shallowly (for instance, the elegant topological account of power-domains is not mentioned at all).

However, the true strength of the book lies in mathematical logic. Virtually a third of the book is devoted to issues of computability and effectiveness, including a whole chapter on basic recursion theory and a full treatment of effective domains and computable elements (including the Myhill-Sheperdson and Rice-Shapiro Theorems, deep results that are not commonly included in books).

In conclusion, then, the book is more specialist than might appear from the title: computer scientists who wish to learn something of the mathematical theory of domains would be better off using one of the texts

oriented towards semantics. However, it can be thoroughly recommended for anyone interested in computability.

STEVEN VICKERS
Imperial College

DOV M. GABBAY, C. J. HOGGER AND J. A. ROBINSON (eds)
Handbook of Logic in Artificial Intelligence and Logic Programming. Oxford University Press. 1993. Hardbound. Vol. 1, ISBN 0-19-853745-X, £60.00, 518 pp. Vol. 2, ISBN 0-19-853746-8, £55.00, 511 pp. Vol. 3, ISBN 0-19-853747-6, £60.00, 529 pp.

The first three volumes of the *Handbook of Logic in Artificial Intelligence and Logic Programming* represent a detailed and comprehensive exposition of the theoretical and computational features of a wide variety of classical and non-classical logics. The sequence covers considerable ground, beginning with introductory chapters which explain some of the fundamental properties of first-order systems and concluding with a series of highly technical, learned and, indeed, inspiring chapters on nonmonotonic and uncertain reasoning. The intended audience is stated to be graduate students and researchers, but much of the material presented is more specialist than this target implies. Given the intended audience the sequence seems to contain a rather odd mixture of material, some foundational and well within the grasp of the average graduate student, and some highly specialist, requiring considerable proficiency in philosophical and mathematical aspects of logic. The foundational chapters do not supply the necessary background for the specialist ones, so the specialist chapters seem bound to be beyond the reach of most graduate students and researchers with other specialisms. Whilst this, to some extent, undermines the description of the sequence as a *handbook*, which one expects to be accessible to the majority of AI practitioners who see a role for logic in their work, it does not prevent one from enjoying, and benefitting from, the effort involved in studying the volumes.

The objective of the first volume is to provide a foundation in classical logics and their automation. It begins by motivating the use of logic in AI and introducing first-order logic, then goes on to discuss computational aspects of deduction. In-depth surveys of the automation of various deduction systems are followed by an excellent introduction to, and analysis of, modal logics in their many different forms and interpretations. Despite a gentle introduction to the role of logic in AI, containing some quite basic material about the construction and analysis of arguments, the

volume is not really aimed at the uninitiated. Davis' swift introduction to first-order logic is too compressed to be helpful to someone without background and, if the objective of the chapter is to provide a working knowledge of logic, would perhaps have been better substituted for some good references to basic logic texts. Israel's chapter introduces various themes that are not pursued in the volume, e.g. the important distinction between the applications of logic at the object and meta-levels of an AI system, which is mentioned there, is not sustained as the volume progresses. The meta-level use of logic, i.e. its use as a formal description language to specify the intended behaviour of AI systems, is implicit in subsequent volumes, in some of the discussions concerning various modalities and their uses in the axiomatization of intelligent systems. The surveys which follow Davis' chapter concentrate on the importance of logic at the object level, and it is assumed that the reader knows how logic can be employed and is looking for help in choosing the best approach to the construction of an efficient proof procedure. A great deal of useful information is found in these chapters, including interesting insights into the role and value of heuristics in general logic and in theories, but the important role of logic as a formal description language is largely ignored within the volume.

Fitting's chapter on modal logic serves both as a thorough and clearly presented introduction to the logical properties of the first-order modal logics and as an advanced text, supplying considerable technical detail on the construction and interpretation of non-standard modal logics. It is an enjoyable read as well as a valuable source of technical information, addressing a large number of the issues that arise in the design of a modal logic. The interpretations of these logics through their accessibility relations are discussed and the consequences of varying domains are considered through the discussion of rigid and non-rigid designators. The decision to allow the objects that populate alternative possible worlds to vary between these worlds results in important philosophical consequences which are reflected in the syntactic descriptions of the resulting modal systems. For example, it is interesting to note that the Barcan formula, which allows the transposition of necessity and the universal quantifier, does not hold of these systems. Multimodal logics are considered and a variety of proof systems are explored. Although it is stressed that modal logics can be custom-built, no help is offered to decide how best to use modal operators to design a custom-built modal logic, or to assist in the construction of logics suited to particular kinds of AI applications. Although modal logics are surely of considerable value at the meta-level of AI systems, there is no discussion of its special role in the formal description of non-deterministic systems. These omissions do not seriously detract from the excellence of a chapter which might be considered the high point of the volume.

The volume is of considerable potential value to the AI

practitioner who is already convinced of the importance and relevance of logic to AI. The chapters that deal with efficient automation of deduction might win a few others over, those whose only objection to logic is that theorem proving is hard. Although the volume is likely to be the most accessible of the three, and in some senses the most handbook-like, it is not really aimed at, and probably will not appeal to, those who need convincing that logic has a role to play, or those that lack adequate formal background to really make use of it.

The second volume is somewhat oddly entitled 'Deduction Methodologies', a title which is difficult to interpret and difficult to fit with the included material. The first three chapters are about aspects of deduction methodology and therefore seem appropriate, but the remaining three do not discuss these issues directly. The first two of these are dedicated to the presentation of higher-order logics and meta-languages, their philosophical implications, and their role in the representation and use of knowledge and beliefs in intelligent agents. The final chapter relates the non-classical logics, required to achieve appropriate representation within these systems, to the more stable classical logics. The volume begins with Wos and Veroff's chapter on automated reasoning. The authors consider the use of heuristic strategy to direct reasoning, together with subsumption and simplification techniques used to remove redundancy from a database of derivations. The subsumption techniques described have relevance beyond theorem proving, to the organized maintenance of automatically acquired information in any AI learning context. The authors provide a useful section in which the state of the art in theorem proving is discussed and open problems are presented, including such problems as how to automatically generate counter examples and how to restrict the application of inference rules in contexts where they are applicable but unlikely to contribute to the construction of efficient proofs. This chapter is followed by a thorough and highly technical presentation of unification, one of the foundational mechanisms supporting automatic theorem proving, and a discussion of approaches to the automatic construction of inductive proofs.

If genuine intelligence is the goal of AI there must be techniques for modelling self-awareness, as the ability to reflect on the consequences of one's own beliefs is fundamental to intelligent reasoning. Perlis and Subrahmanian discuss some of the theoretical consequences of reification, necessary to allow an agent's beliefs to be the subject of its own reasoning, and its role in nonmonotonic and other reasoning systems in which the introduction of knowledge or belief modalities makes the ontological status of propositions ambiguous. In particular, they consider the importance of maintaining consistency in different approaches to reified reasoning. Attention is drawn to the consequences of *situated* reasoning on the need for self-awareness in the reasoner, a discussion of particular relevance to applied areas of AI

such as planning and robotics. Leivant's presentation of higher-order logic appears to be aimed at the logician, rather than at the AI practitioner. Its relevance to AI is less clear than is the case for Perlis and Subrahmanian's discussion of reification and reflection. It takes a mathematical and philosophical view of higher-order languages in order to evaluate their claims to logical status and their expressive power, particularly with respect to the finiteness and infiniteness of domains. The relationship between higher-order languages and polymorphic type systems is explored, and computational features of higher-order deduction are examined.

The final chapter in the volume is Gabbay's consideration of the processes by which non-classical logics are translated into (first-order) classical logic in order to obtain the computational advantages that accrue from the well-established body of research in automated deduction methods for classical logic. Since much of the development of non-standard logics has progressed proof theoretically, without the corresponding construction of formal semantics, Gabbay introduces the labelled deductive system (LDS) as a framework for the description of systems which are to be ultimately realised as classical logic. Translation of the LDS to classical logic is straightforward and the LDS provides, in Gabbay's view, a form of semantics for the original non-classical system. Classical logic is not claimed to be universally expressive—it is proposed that non-classical formulations might be better suited to the description and formal analysis of a given non-standard reasoning system.

The objective of the volume, indicated in Gabbay's preface, is to consider the combination of logics and logical mechanisms in the application of an appropriate 'multi-logic' to the construction of flexible intelligent systems. Although the volume presents thoughtful analyses of many issues related to the use of logics it does not clearly address this notion of combination. The title of the volume does not convey, at first reading, any such objective and one might be misled by the title into expecting a volume more similar in content to volume 1. Having studied the three-volume sequence it is possible, in retrospect, to understand better what is the role of volume 2 but it does not stand alone as volumes 1 and 3 do.

The third volume begins with a motivating chapter, clearly aimed at the non-specialist, which discusses some reasoning contexts in which defeasible reasoning is deemed necessary. Some of the properties of the various approaches to this form of reasoning are then introduced, by way of an introduction to the eight highly technical and learned chapters which follow. Ginsburg's introduction is widely accessible, but the intended audience of the subsequent chapters is mathematically sophisticated, with a detailed knowledge of, and conviction in, the subject matter. It is taken for granted that the importance of nonmonotonic and other defeasible reasoning techniques is accepted and that the

reader is primarily interested in comparing the formal properties of the logics discussed.

As if to confound the criticism that nonmonotonic logics have been inadequately formally developed, there is considerable formal content. Proof theoretic and model theoretic approaches to the formalization of different systems are compared and there are detailed analyses of alternative consequence relations and inheritance strategies. An annoying oddity is that the intended meaning captured in the inferencing strategies in some of these systems seems not to coincide with one's intuitive interpretation of what should be going on. The Tweety example, which occurs often in the volume, illustrates this. It is proposed that the conclusion that Tweety flies will be contradicted when it is learnt that Tweety is an ostrich. The conclusion is obtained by *modus ponens* before it is known that Tweety is an ostrich. After this fact is known it can also be inferred that Tweety does *not* fly. The contradiction is claimed to be due to incomplete knowledge about Tweety, but in fact it is due to the premise that 'all birds fly'—a blatantly untrue claim. Knowing that, in fact, most birds fly, but some don't, and knowing nothing about Tweety, one would not conclude that Tweety flies and then be surprised to find that he does not fly. Before one knows what kind of bird Tweety is one would surely conclude only that Tweety *might be expected* to fly, and this could be expressed in a way that is not inconsistent with the subsequent discovery and inference that he does not fly. The use of statistical evidence, as discussed in Kyburg's chapter on uncertainty logics, offers a way of characterizing the knowledge that most birds fly, although the interpretation of *most* as a statistical measure is unintuitive. There is no discussion of the relative merits of alternative solutions to the Tweety problem cast in this way, such as the introduction of *sorts* into the language.

This is not to say that all of the reasoning situations that have inspired research into nonmonotonic logics have simple classical solutions, but the Tweety example is not one that convinces the sceptical reader of the necessity to abandon the security of classical logic. A sceptical reader hoping to find convincing examples of nonmonotonic reasoning in context might be disappointed to encounter Tweety and his misfortunes. Other light-hearted examples including the Yale shooting problem, the quakers and the potato in the exhaust pipe are also offered, but there are no weightier examples to really show the advantages of giving up the formal benefits of classical logic. The consequences of moving to non-standard systems, about which there is so little agreement, are so great as to require really compelling motivation. The examples which are presented create a feeling of unease—they are different in character and seem to support different types of solutions and yet they are commonly introduced to illustrate the same supposed shortcomings of classical systems. The Tweety problem, in which Tweety is discovered to be an ostrich or a penguin, seems to arise because of the misplaced use of

the word 'all' in the premise 'all birds fly'. The problem in which Tweety has his feet set in concrete is altogether different—surely this contingent fact does not necessitate revision of the inference that Tweety can fly?

Kyburg's chapter on uncertainty logics notes that defeasible reasoning is only really necessary when there is real uncertainty in the reasoning environment. Under these conditions, conclusions drawn in good faith from premises expressed as precisely as is possible given the state of knowledge of the reasoning agent, turn out subsequently to be in error, based on incomplete information. Whether these conclusions represent inconsistencies requiring revision seems to depend critically on whether or not the reasoning agent has access to the real world. If an agent holds the beliefs b at time t , allowing him to infer p , and beliefs $b \setminus q \cup a$ at time $t + k$ allowing him to infer $\neg p$, and the inferences of the system are evaluated purely with respect to its beliefs, no inconsistency arises because of the passage of time. In the real world, $b \setminus q \cup a$ might have been true all along, so that p was always inconsistent with the facts. Without discussing this distinction, Ginsburg mentions the planning context as one in which inferences must be defeasible—the STRIPS assumption is, he claims, a default assumption which could give rise to erroneous inferences. However, planning systems typically do not have access to the real world, only to internal models of the world which reflect the agent's beliefs about the world. An axiomatization of planning might avoid inconsistency by introducing time stamps to differentiate between the developing belief states of the planning agent. If p is concluded at time t from a set of beliefs b , and $\neg p$ is concluded at time $t + k$ from a set of beliefs b' , $\neg p$ simply overrides p and there is no inconsistency. When planning systems are connected to physical devices which execute actions in the real world a link is made between the real world and the internal model giving rise to considerable problems concerning the maintenance of consistent internal models. The question of whether such a link is significant to an application seems to be of great significance in deciding whether non-monotonicity is relevant, but it is not discussed in these terms in the volume.

The external imposition of time stamps allows planning systems to manage changes in world states that follow the applications of actions. However, planning systems cannot reason about their own beliefs in a way that would allow them to infer new beliefs on the basis of the absence of knowledge. Konolige's chapter on autoepistemic logic shows how a reasoning agent can be given access to its own beliefs so that it can reason introspectively. It is observed that the ability to draw inferences on the basis of what is *not* known is probably equivalent to the application of default rules to obtain new conclusions from what is not inconsistent with current knowledge.

Kyburg's thorough discussion of probabalistic approaches to plausible reasoning with uncertainty

presents persuasive arguments in favour of the acquisition and use of evidential support in uncertain reasoning. This is followed by the chapter by Dubois *et al.* on possibilistic, or fuzzy, reasoning. One of the key problems of fuzzy reasoning, its semantics, is dealt with by interpreting possibilistic distributions as preference orderings on possible worlds in which the associated fuzzy proposition is true. Worlds corresponding to values closest to 1 are worlds in which the belief that the proposition is true are most strongly supported, and so these worlds are preferred, in an epistemic semantics. This interpretation removes the numerical significance of the values and hence resolves one of the traditional criticisms of fuzzy reasoning, that the assignment of arbitrary values in the range $[0, 1]$, to stand for degrees of truth of a proposition, do not have clear intuitive meaning. Under this interpretation the numerical values have only ordering significance, not quantifiable significance. The chapter presents many other formal results indicating the link between fuzzy and nonmonotonic reasoning and considers extensions of the basic possibilistic logic and the impact of these on its semantics.

Volume 3 stands alone as a specialist's text on nonmonotonic and uncertain reasoning, supported, but not encroached upon, by the previous volumes. It is an extremely enlivening presentation, highlighting many complex and unresolved philosophical problems. The volume has many strengths, as a reliable and detailed source of information on an area of research in logic and AI which is both philosophically and practically problematic, but it is not a handbook. To call it a handbook seems to diminish its intellectual contribution.

The first three volumes of the Handbook contain far more than can possibly be assimilated in a single reading. It is possible that subsequent examination will show some of the criticisms mentioned above to have been addressed, after all, in a way that was missed by the reviewer on her first reading. Even if this is not the case, the Handbook can be unreservedly recommended to AI practitioners with proficiency in logic and commitment to its role in the development of AI systems.

MARIA FOX

University College London

V. STAVRIDOU

Formal Methods in Circuit Design. Cambridge University Press. 1993. ISBN 0-521-44336-9. £27.95. 193 pp. hardbound.

The past decade has seen tremendous progress in the application of formal methods to the specification and verification of hardware designs. Right from its infancy this area of research has, by and large, been application driven. The decade has seen a number of large realistic