Other cited examples are more of the same.

[3]E.g., the crucial lemma 5.5.5 (p. 191) is described as having a proof that is "not difficult, but it is very long and tedious, and is omitted here. One may find it in Wadsworth [41]." One may find it there only if one is at Syracuse University, or was in 1974. The full reference to [41] is "C. Wadsworth. SIS830 Lecture Notes. A course at Syracuse University, 1974." We are not being picky. If the proof is not difficult, why does not Davis describe it as, e.g., by induction on the length of the environment? Elsewhere (Lemma 3.2.1, p. 54), Davis refers to [27] "for sketch of proof". We feel sure she means not [27], "D. McDermott and J. Doyle. 'Non-Monotonic Logic I'. *Artificial Intelligence*, 13: 41–72, 1980," but rather Mendelson 1964, who proposes a lemma identical to Davis's Lemma 3.2.1. Mendelson's so-called "sketch of proof" is "(Hint: induction on length of t)". Our worry is that Davis has no clear thoughts on what should be proven and what not, nor on what counts as a legitimate proof summary.

References

Barendregt, H. P. (1981), *The Lambda Calculus: Its Syntax and Semantics*, Amsterdam, North-Holland.

Davis, Ruth E. (1985), "Logic Programming and Prolog: A Tutorial," *IEEE Software* 2.5, 53–62.

Mendelson, Elliott (1964), *Introduction to Mathematical Logic*, New York: Van Nostrand.

Mendelson, Elliott (1979), *Introduction to Mathematical Logic*, 2nd ed., New York, Van Nostrand.

Stoy, Joseph E. (1977), *Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory*, Cambridge, MA: MIT Press.

Department of Computer Science, MERRIE BERGMANN
Smith College,
Northampton, MA 01063, USA

Department of Philosophy, THOMAS TYMOCZKO
Smith College,
Northampton, MA 01063, USA

Ernest Davis. *Representations of Commonsense Knowledge*, Series in Representation and Reasoning, San Mateo, CA: Morgan Kaufmann Publishers, 1990, xxiii + 515 pp., $42.95 (cloth), ISBN 1-55860-033-7.

This book is a compendium of alternative formal representations of diverse fragments of commonsense knowledge. It is centered largely on formal representations drawn from first-order logic, and thus lies in the tradition of Kenneth Forbus (1984, 1985), Patrick Hayes (1985ab), et al. (see Hobbs and Moore 1985, Weld and de Kleer 1989). The focus of the work is narrowed further by an interest in the *dynamics* of common sense, and more precisely still by an interest in the *patterns of inference* employed in commonsensical contexts, and in the automation of such inferences.

Two aims determine the nature of the results have summarized, in a way that

represents an important ambiguity in the notion of commonsense reasoning itself, as between:

1. formally rigorous and precise reasoning (reasoning on the theoretical level) about the world of common sense, and
2. reasoning as actually practiced by humans in their everyday, non-theoretical moods.

It will become clear in what follows that the manifest tension between these two aims is far from being successfully resolved in the projects with which Davis deals.

The basic approach common to all this work is that of taking a sampling of deductive inferences from a given domain and of attempting to develop formal languages in which the relevant knowledge can be axiomatically expressed and the relevant inference procedures formally captured. An introductory logic tutorial is included, together also with extensive and useful indications of the more recent literature on formal theories of the commonsense world. The range of the work is impressive, encompassing recent developments in the theory of plausible reasoning, temporal logic, qualitative physics, folk psychology, the theory of plans and goals, and the theory of communicative action. The work thus represents the first comprehensive textbook in an area that is of quite pressing importance not only for the AI community but also, in principle, for philosophers, developmental psychologists, linguists, anthropologists, and others interested in the field of common sense.

This having been said, however, it has to be stressed that the work has a number of flaws. These derive first of all from the underdeveloped state of the field in question (both as regards concrete theoretical and practical achievements and also as regards basic philosophical and methodological presuppositions). But they derive also from the author's resolute concentration on the immediate needs and expectations of those working in AI.

In what follows, I shall indicate a number of these flaws as they appear from the perspective of a philosopher concerned theoretically with the nature of common sense and with the structures of commonsense reality. As Davis himself points out:

> The most important external influence on AI theories of commonsense reasoning has been twentieth-century analytic philosophy. Most of our basic analytic tools, particularly formal logics, and much of our analysis of specific domains, particularly time, action, and mind, were developed by philosophers and mathematical logicians (p. 23).

Davis is not referring here to the so-called "commonsense philosophy" of G. E. Moore (1959) and others. For this, he holds (not quite correctly – see the works of Avrum Stroll (1988) and Lynd Forguson (1989)), was concerned not with the analysis of specific domains of commonsensical concepts but rather with the defence of common sense as such: In other words, it was concerned to establish that common sense is true. Accordingly, Davis maintains, the commonsense philosophers contributed "little if anything relevant to our enterprise" (p. 26).

The philosophy that is most directly relevant to Davis's enterprise is the so-called "mathematical philosophy" that was initiated by Alfred North Whitehead in a ground-breaking paper (not mentioned in this book) entitled "On Mathematical Concepts of the Material World" (1906). This mathematical philosophy was then pursued in some early writings of Bertrand Russell, but it was developed most systematically by Stanisław Leśniewski and his disciples (also not mentioned by Davis), who constructed a range of precise and rigorous formal-ontological theories of those general concepts that lie at the heart of common sense – concepts such as time, space, part, whole, and so on – in ways precisely in keeping with the first central aim of the projects here summarized. (Leśniewski's writings are collected in Surma et al. 1992; for an overview, see Simons 1987.)

If, then, one is serious in the attempt to build up a formal theory of the structures of reality as these are captured by our commonsense concepts, then the achievements of Leśniewski et al. will have to be taken account of. These include the formal theories of part and whole developed by Leśniewski himself (extensional mereology), the theory of dependent and independent parts developed by Edmund Husserl in his *Logical Investigations* (1900-1901), the formal theories of temporal and biological concepts developed by Joseph H. Woodger (1937), and the various systems of realistic formal ontology that have been developed on a Leśniewskian or Husserlian basis in subsequent decades (see Smith 1982). Such theories have escaped the attention of Davis and his fellows, one may presume, because they do not fall squarely within the Frege-Russell-inspired logical tradition that the AI community takes as its standard. Most importantly, they differ from the Frege-Russell tradition in adopting as the basis of their formal-ontological theories not the abstract and mathematically problematic theory of sets but rather the simpler and more commonsensical theory of parts and wholes or "mereology".

The drive toward realistic formal ontology, toward precise and rigorous theories of the concepts at the heart of common sense, is not unknown to Davis. As he points out, certain otherwise attractive primitives have to be rejected as "not really quite kosher" because:

> They do not correspond to anything much in the real world; they are arbitrary distinctions made by us, as theory builders, for the purpose of making axioms cleaner and shorter. As a result, our representation becomes less a description of the relations in the world and more a matter of logic programming (p. 206).

On the other hand, however, and in conflict with this realistic drive, is the desire of AI research on common sense to achieve faithfulness to commonsense *reasoning* via the development of theories that would themselves employ inference-patterns mimicking those of common sense. For the latter is clearly not precise and rigorous, and it is not clear that a sophisticated theory of the commonsense world (or indeed of commonsense reasoning) can be produced that is at the same time faithful to those crude processes of reasoning that serve our

everyday human purposes. This problem is compounded still further when account is taken of the fact that such commonsense reasoning seems not to follow standard patterns at all, much less the deductive patterns captured by extensional first-order logic and by those of its close cousins exploited here. (On this, see McDermott 1990.)

The drive toward realistic ontology suffers also from the fact that, in his actual practice, Davis is all too often willing to substitute artefacts of his chosen logical machinery for the treatment of commonsense concepts themselves in strict and realistic fashion. Thus, he takes it for granted that the appropriate way to analyze 'Calvin is in the living room' lies via the shamefacedly counter-commonsensical set-theoretic translation into: 'the set of spatial points making up the region occupied by Calvin is a subset of the set of points making up the living room' (p. 248). Similarly, he suggests that in order to express a sentence pertaining to family relations, for example, to the effect that Tom bears the same relation to Dick as Bruno bears to Fritz, it is necessary to conceive such a sentence as amounting to an assertion to the effect that ⟨Tom, Dick⟩ and ⟨Bruno, Fritz⟩ are both members of a certain set of ordered pairs (p. 8). Such translations are an artifice of logic, and, unfortunately, they are as far removed from commonsense ontology as they are from the representation of our actual commonsense reasoning.

Davis's book demonstrates that the AI community, in its effort to understand the dynamics of common sense, has produced a surprisingly large amount of interesting and sophisticated theory. The book demonstrates also, however, how difficult is the task of producing anything like a total or adequate theory of common sense itself and of commonsensical reality.

References

Forbus, Kenneth (1984), 'Qualitative Process Theory', *Artificial Intelligence* **24** 85–68.

Forbus, Kenneth (1985), 'The Role of Qualitative Dynamics in Naive Physics', in Hobbs and Moore 1985, pp. 185–226.

Forguson, Lynd (1989), *Common Sense*, London: Routledge.

Hayes, Patrick J. (1985a), 'The Second Naive Physics Manifesto', in Hobbs and Moore 1985, pp. 1–36.

Hayes, Patrick J. (1985b) 'Naive Physics I: Ontology for Liquids', in Hobbs and Moore 1985, pp. 71–107.

Hobbs, Jerry R., and Moore, Robert C. (eds), (1985), *Formal Theories of the Commonsense World*, Norwood: Ablex.

Husserl, Edmund (1900-1901), *Logical Investigations*, trans. by J. N. Findlay, London: Routledge and Kegan Paul, 1970.

McDermott, Drew (1990), 'A Critique of Pure Reason', in M. A. Boden (ed.), *The Philosophy of Artificial Intelligence*, Oxford: Oxford University Press, pp. 206–230.

Moore, G. E. (1959), 'A Defence of Common Sense', in G. E. Moore, *Philosophical Papers*, London: George Allen and Unwin, pp. 60–88.

Simons, Peter M. (1987), *Parts: A Study in Ontology*, Oxford: Clarendon Press.

Smith, Barry (ed.) (1982), *Parts and Moments: Studies in Logic and Formal Ontology*, Munich: Philosophia.
Stroll, Aurum (ed.) (1988), *Surfaces*, Minneapolis: University of Minnesota Press.
Surma, Stanislaw J.; Srzednicki, Jan T.; Barnett, D. I.; and Rickey, V. Frederick (eds.) (1992), *Stanisław Leśniewski: Collected Works*, Dordrecht: Kluwer Academic Publishers.
Weld, Daniel S., and de Kleer, Johan (eds.), (1989), *Readings in Qualitative Reasoning about Physical Systems*, Los Altos, CA: Morgan Kaufmann.
Whitehead, Alfred North (1906), 'On Mathematical Concepts of the Material World', *Philosophical Transactions of the Royal Society of London*, Series A, Vol. 205, pp. 465–525.
Woodger, Joseph H. (1937), *The Axiomatic Method in Biology*, Cambridge, Eng.; Cambridge University Press.

Department of Philosophy
and Center for Cognitive Science,
State University of New York at Buffalo,
Buffalo, NY 14260, U.S.A.

BARRY SMITH

John F. Sowa (ed.), *Principles of Semantic Networks: Explorations in the Representation of Knowledge*, San Mateo, CA: Morgan Kaufmann, 1991, xi + 582 pp., $42.95 (cloth), ISBN 1-55860-088-4.

1. Introduction

Principles of Semantic Networks is based on a workshop on *Formal Aspects of Semantic Networks*, held at Catalina Island in 1989. The change of title was made to reflect the broader spread of contributions. The book is more than a conference proceedings: A strong editing process has produced a cohesive survey of current theory and applications of semantic networks. Extensive changes were made to each of the papers to make each a self-contained introduction to a subfield, leading into an original piece of current research.

John Sowa is the author of *Conceptual Structures: Information Processing in Mind and Machine* (1984), around which an international group of researchers has gathered, meeting annually at a different workshop for a number of years. By organizing the Catalina workshop, he has attempted to unify diverging fiefdoms of the semantic-network world, such as the KL-ONE and SNePS communities.

2. Organization

Principles of Semantic Networks is organized into three parts: seven chapters on issues in knowledge representation, which discuss theoretical topics independent of particular implementations; six chapters on formal analyses, which treat the