

Book Review

Semantic Web for the Working Ontologist, Second dition: Effective modeling in RDFS and OWL by Dean Allemang and James Hendler, Morgan Kaufmann, 384 pp., \$55, ISBN 0-123-85965-4

The Semantic Web is one sub-discipline of Knowledge Engineering about which everyone has an opinion. Some think that it is the next big revolution in the Internet. Others have the impression that it is just a dream that will take much too long to arrive. After all, it is about the same age as the Social Web, which obviously has been the major focus of online activity for years now. Nevertheless, others think that it holds a great promise, but unfortunately it requires a good understanding of some intimidating logic details, and therefore many perceive it as only a task for pedantic mathematicians. Finally, some seem to believe that the Semantic Web is already here.

When reading the book *Semantic Web for the Working Ontologist*, I realized that the state of the art is much more advanced than I imagined. This does not mean that the discipline has become much more complicated. On the contrary, dealing with the Semantic Web has become an everyday engineering task, not unlike dealing with relational databases, for instance. The title suggests that creating and managing ontologies is now a profession that you can learn from a single book, and I think the authors came very close to fulfilling that promise.

The book is pleasingly didactic. The first two chapters guide us through the basic principles of the semantic web and semantic modeling. I really enjoyed reading this part, because aside from explaining the fundamental concepts, it uses very sophisticated rhetorics to make its case. After reading it you will be ready for battle: you will see the Semantic Web as the logical next step of advancement, but you will also be informed about the dangers surrounding it. The Semantic Web described here is more about making our computers do even more useful things than before, and less about the rather abstract notions of machine understanding and the importance of meaning representation that you will find in some other introductory works. Because of this, I would recommend this intro for Semantic Web skeptics and advocates alike as food for thought.

The following chapter explains Resource Description Framework (RDF) triples and their notations. It is probably because of the authors' experience in teaching Semantic Web that they included extra emphasis on the transformation of Structured Query Language (SQL) databases to RDF data. The RDF data are usually depicted as graphs in this book, which helps you to work through and understand their relationships with one another. A minor issue is that there is no legend for these figures, and therefore you have to figure out what the small icons in the boxes mean or why some lines are solid and others are dashed. Fortunately, these things are rather evident, and hence there is not much confusion about them.

After you have learned about RDF, you will be curious about which software and architecture you need to deal with it. In the next chapter, the authors explain the architectural elements of Semantic Web, and while they wisely do not recommend any particular software, you will have the necessary starting point to find whatever suits your needs.

In the following section, the authors explain SPARQL, the language for querying RDF stores. A good feature of this section is that you can understand almost everything in this chapter without having read the rest of the book (of course, you need to know the basics of RDF). This is helpful because otherwise an easy-to-understand but still rather complete SPARQL tutorial is not easy to come by. You will also learn how new RDF triples are constructed from old ones by using SPARQL—a way to inference on the Semantic Web. Actually, the authors decided to explain how ontologies work with SPARQL expressions throughout the whole book. I think this was a good

decision, because it allows you to exercise the language and learn ontology theory without being confronted by too much formal logic. This book is the proof that you can be accurate and go quite deep in discussing ontologies, without scaring away most of your audience with the math.

The next six chapters are about the creation of ontologies. It starts with the rather simple RDF Schema and ends with advanced OWL, the state-of-the-art ontology language. The authors guide the reader step-by-step, and thus in every chapter you learn a little, a bit more manageable. They bridged the gap between RDF Schema and OWL by inventing RDF Schema Plus, a small but useful subset of OWL. On the way, you will also learn some of the most important Semantic Web applications, such as the data sets of the US government (data.gov), the Friend of a Friend Project (FOAF), Facebook's open graph protocol (OGP), the Simple Knowledge Organization System (SKOS), the Good Relations ontology (GR), NASA's Quantities/Units/Dimensions/Types (QUDT) and some interesting biological ontologies.

I think that this book, 384 pages and 16 chapters in length, could be the ideal basis for a university class in the Semantic Web for computer scientists/software engineers. There are many ways to select the 12–14 weeks worth of material you need in full a semester. If you have less time, you can still teach some of the chapters individually, such as the one about SPARQL. The chapters contain several elaborate challenges with solutions that people can use to test their understanding of the material.

However, if you are looking for a long discussion of more theoretical aspects, you probably need to look elsewhere. While not leaving out the essentials you need to understand the current technology, the book is very concise on theory. In addition, this book is not an ideal text to demonstrate the connections between logical calculi, reasoning and the semantic web, but this clearly was not a goal of the authors.

Semantic Web for the Working Ontologist is an ideal book to start or restart learning the Semantic Web. Its didactics are streamlined for the best understanding as it explains the latest technology. In addition, it gives you a fresh picture about the industry with the many examples. If you did not follow the advancements of the field closely, I'm sure that some of these examples will surprise with their advanced state. For me, it turned out that by deliberately leaving out the usual 30–50 pages about logic that one is used to finding in Semantic Web books, and by cutting back on the abstract philosophical grounding, one is actually better able to understand the Semantic Web technology.

Mihály Héder

Department of Philosophy and History of Science
Budapest University of Technology and Economics
Budapest, Hungary;

e-mail: mihaly.heder@filozofia.bme.hu