

Kent Academic Repository

Full text document (pdf)

Citation for published version

Orchard, Dominic A. (2014) Computational Semantics with Functional Programming, by Jan van Eijck and Christina Unger. Review of: Computational Semantics with Functional Programming by van Eijck, Jan and Unger, Christina. Journal of Functional Programming, 24 (4). pp. 524-527. ISSN 0956-7968.

DOI

<https://doi.org/10.1017/S0956796814000057>

Link to record in KAR

<http://kar.kent.ac.uk/57485/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Book review: Computational Semantics with Functional Programming, by Jan van Eijck and Christina Unger

Dominic Orchard
Computer Laboratory, University of Cambridge
(e-mail: dominic.orchard@cl.cam.ac.uk)

Context This book intersects two disciplines: functional programming and computational linguistics (the study of natural language using computational techniques). This review is from a functional programming perspective.

In computational linguistics, the approach of *Montague semantics* (named after its pioneer Richard Montague) uses mathematics and logic as a metatheory for a mathematically precise semantics (Janssen, 2012). The core principle is that there is “*no important theoretical difference between natural languages and the artificial languages of logicians*” (Montague, 1970). Central to the approach is a compositional, high-order semantics, with the λ -calculus and predicate logic as foundations (see the example in Figure 1 on the following page). It is within this framework that Eijck and Unger write *Computational Semantics with Functional Programming* (Van Eijck & Unger, 2010). Their thesis: typed functional languages can naturally express Montague semantics, providing the benefits of a rich, executable metalanguage. From this basis they provide a guidebook through the core principles of the *Montagovian* approach and subsequent research, written in the languages of predicate logic and Haskell.

Overview Whilst the book is interdisciplinary between functional programming and computational semantics, there is a slight bias towards the latter. The introductory material focusses more on logic, the λ -calculus, and Haskell than on natural language semantics. From a functional programmer’s perspective, most of the natural language concepts are however easily understood from the examples and definitions.

Its thirteen chapters span roughly 400 pages, about half of which are tutorial, explaining the mathematical and programming basis. The core aspects of the Montagovian approach are introduced, providing a survey. Chapters on set theory, the simply-typed λ -calculus, induction, polymorphism, Haskell, and parser combinators will already be familiar to most functional programmers. Small examples with a linguistic flavour appear throughout these chapters, with larger examples uniting the introductory material, such as a model checker for predicate logic and an inference engine over first-order propositions written in English.

The second half of the book applies Montagovian principles using higher-order predicate logic combinators to more complex linguistic phenomena, such as quantification, context dependence, pronoun scoping, and the effects of common knowledge on meaning.

The book does not offer any new research results. Instead it overviews various computational semantics concepts from the literature, using Haskell to give a complete account which can then be experimented with. References to key results over the last forty years are provided throughout, giving a useful starting point for further study.

The technical material is sound and the writing is mostly easy to follow. However, particularly in the second half, the book would benefit from more discussion on how chapters relate to each other and to the wider research field.

2


Book review

Fig. 1. Example Montague-style semantics for a sentence fragment (based on Chapter 7 of the book) where $\llbracket \text{every giant that laughed} \rrbracket_{\text{NP}} = \lambda Q. \forall x. (\text{Giant } x \wedge \text{Laugh } x \wedge Q x)$.

Complete source code is provided throughout and is available free online without purchasing the book¹ (although the online version lacks any documentation or comments, which are in-line in the book's text). The book's index includes function names from the code examples, which is extremely helpful given the backward-dependencies on code presented in earlier chapters.

Concepts and ideas Figure 1 above gives an example *compositional* interpretation of a sentence fragment (based on Chapter 7) comprising a *noun phrase* (NP), a *determiner* (DET), *relative clause* (RCN), *common noun* (CN), and *verb phrase* (VP). Figure 2 shows relevant fragments (from the book) of the Haskell interpretation for this example.

The key concept of this book (and the Montagovian approach) is *compositionality*, and maintaining this compositionality even for more complex language features which, on the surface, appear to be non-compositional. Haskell serves the development well in this area.

One of the most interesting approaches is the use of continuations to denote linguistic contexts, in the sense of sentences with a “hole”. This is used to describe different evaluation orders (as is done similarly in programming language semantics) for quantifiers which explains how ambiguities arise from different readings of quantifier scoping in a sentence.

A link is made between seemingly non-compositional linguistic phenomena and side-effecting computations, as developed by Shan using monads for natural language semantics (Shan, 2005). Whilst the composition of impure computations using monads is well known in semantics and functional programming, this idea is not developed further. The book provides a good starting point for considering whether further programming language semantics results can be used for natural language semantics. For example, since *comonads* are used to structure context-dependent computations, *e.g.* (Uustalu & Vene, 2008), can they be used to abstract various forms of context in natural language semantics?

Discovering new abstractions and common idioms is a common activity in functional programming (both development and research). The advanced type systems found in many functional languages, particularly Haskell, help to organise and describe such abstractions. Whilst the book makes use of well-known abstractions from the Montagovian approach (*e.g.*, using parametric polymorphic definitions of relations) it does not extend these with the additional power provided by modern Haskell extensions. All of the examples are written using a simple subset of Haskell features: functions, pattern matching, algebraic data types, and parametric polymorphism. Type classes are used occasionally for defining custom *Show* instances. This is appropriate for the natural semantics audience, but there is plenty of room for the advanced features of Haskell to be leveraged, particularly as a way of exposing new abstractions.

For example, the authors define mutually recursive type-level and term-level functions for *intensionalizing* a semantics (lifting all entities and propositions to take an additional world parameter)

¹ <http://www.computational-semantic.eu/>

Book review: *Computational Semantics with Functional Programming*, by Jan van Eijck and Christina Ungert

```

intNP :: NP -> (Entity -> Bool) -> Bool
intNP (NP2 det rcn) = (intDET det) (intRCN rcn)
...

intDET :: DET -> (Entity -> Bool) -> (Entity -> Bool) -> Bool
intDET Every p q = all q (filter p entities)
...

intRCN :: RCN -> Entity -> Bool
intRCN (RCN1 cn _ vp) = \e -> ((intCN cn e) && (intVP vp e))
...
intVP Laughed = \x -> laugh x
...
intCN Giant = \x -> giant x

```

Fig. 2. Fragments of the Haskell implementation for the Montague semantics of Figure 1.

and inversely *extensionalizing* (distributing a single world parameter to all subterms). The transformations are described on types and terms, and a proof is given that these transformations form an isomorphism. The chapter then specialises these operations to different types in the semantics, writing out the code for each instance. Instead, this could be straightforwardly implemented generically, once for all types, using type families (Chakravarty *et al.*, 2005) (for the type functions) and type classes (for the overloaded value functions) in Haskell.

Conclusion Using a programming language as a metalanguage is a standard approach in programming language research. It is encouraging to see this book applying the same principle outside of the field: using a programming language as a metalanguage for natural language semantics, reaping the benefits of the type checker and runtime. As the authors put it, “*your programming efforts will give you immediate feedback on your linguistic theories*” (Van Eijck & Unger, 2010)[p.11]. This book is therefore highly valuable for those natural language semanticists who ascribe to compositional, Montagovian approaches to semantics. It should be noted that compositional approaches are not universally accepted by linguists (see, for example, the brief survey of arguments against compositionality in (Szabó, 2013)).

For functional programming researchers, this book is a helpful starting point for anyone looking to get involved in natural language semantics, or at least looking for information on the Montagovian approach. There are many tantalising threads to follow and, as described briefly above, many areas of the book for which advanced, modern functional programming techniques apply. Additionally, for those involved in language design, this book may provide some food-for-thought on how to add more natural language-like features to an artificial language.

Acknowledgments Thanks to Alan Mycroft, Tomas Petricek, and Martin Szummer for their comments on this review.

References

- Chakravarty, Manuel MT, Keller, Gabriele, Jones, Simon Peyton, & Marlow, Simon. (2005). Associated Types with Class. *Pages 1–13 of: In POPL’05: Proceedings of the 32nd ACM SIGPLAN-SIGACT symposium on Principles of Programming Languages.*
- Janssen, Theo M. V. (2012). Montague Semantics. Zalta, Edward N. (ed), *The Stanford Encyclopedia of Philosophy*, Winter 2012 edn. Stanford University.
- Montague, Richard. (1970). Universal grammar. *Theoria*, **36**(3), 373–398.
- Shan, Chung-chieh. (2005). *Linguistic side effects*. PhD dissertation, Harvard University.

Szabó, Zoltán Gendler. (2013). Compositionality. Zalta, Edward N. (ed), *The Stanford Encyclopedia of Philosophy*, Fall 2013 edn. Stanford University.

Uustalu, Tarmo, & Vene, Varmo. (2008). Comonadic notions of computation. *Electronic Notes in Theoretical Computer Science*, **203**(5), 263–284.

Van Eijck, Jan, & Unger, Christina. (2010). *Computational Semantics with Functional Programming*. First edn. Cambridge University Press.