

in which the original motivation for affix support is undone. This tight interweaving of narrow syntax and PF-mapping contrasts starkly with a recurring speculation in Chomsky's recent work (see e.g. Berwick & Chomsky 2016). Chomsky's thesis, motivated in part by evolutionary considerations, is that the mechanisms concerned with converting hierarchical structure into an externalizable form are 'ancillary' to grammar proper, and as such are the locus of much of the superficial complexity of language, including crosslinguistic variation. With regard to the latter aspect, R's work is entirely consistent with Chomsky's suggestion, deriving as it does important aspects of syntactic variation from variable phonological properties. But if R is right, the phonology is more than an ancillary mapping relating the internal computational system to articulation and perception: it is an 'active player' in syntactic computation.

*Contiguity theory* is a fascinating and thought-provoking attempt at making sense of the seemingly arbitrary distribution of displacement phenomena across languages in terms of a modest inventory of universal mechanisms. One can only be impressed by the way in which R tackles this Herculean task: with untiring optimism fueled by a firm belief in theory.

#### REFERENCES

- BERWICK, ROBERT C., and NOAM CHOMSKY. 2016. *Why only us*. Cambridge, MA: MIT Press.
- BIBERAUER, THERESA; ANDERS HOLMBERG; and IAN ROBERTS. 2014. A syntactic universal and its consequences. *Linguistic Inquiry* 45(2).169–225. DOI: 10.1162/LING\_a\_00153.
- CHOMSKY, NOAM; ÁNGEL J. GALLEGÓ; and DENNIS OTT. 2017. Generative grammar and the faculty of language: Insights, questions, and challenges. *Catalan Journal of Linguistics*, to appear.
- RICHARDS, NORVIN. 2010. *Uttering trees*. Cambridge, MA: MIT Press.
- RICHARDS, NORVIN. 2017. Deriving contiguity. Cambridge, MA: MIT, ms. Online: <http://ling.auf.net/lingbuzz/003289>.

Department of Linguistics  
University of Ottawa  
70 Laurier Avenue East  
Ottawa, ON, K1N 6N5, Canada  
[dennis.ott@post.harvard.edu]

**Computational models of referring:** A study in cognitive science. By KEES VAN DEEMTER. Cambridge, MA: MIT Press, 2016. Pp. 339. ISBN 9780262034555. \$34 (Hb).

Reviewed by WILLIAM S. HORTON, *Northwestern University*

In this book, Kees van Deemter provides a wide-ranging synthesis of work from computational linguistics on the natural language generation of referring expressions. vD has been a leading member of a highly productive research community focused on empirical and computational approaches to REFERRING EXPRESSION GENERATION (REG). Here, he places these efforts into a comprehensive theoretical and historical context, outlining a range of issues and questions important to the development of algorithms for reference generation and providing a guide for where such work is likely to go in the future. In the preface, vD characterizes reference as the 'fruit fly' of the study of language, given all the ways that researchers from many disciplines have used reference to explore various phenomena in language and communication. And indeed, this book makes a compelling case for linguistic reference as a central intellectual topic within the cognitive sciences, while at the same time illustrating how computational models, in particular, can be useful tools for exploring an assortment of thorny theoretical issues.

The volume is organized into four major parts. Part I orients the reader to some of the most relevant issues and questions related to the generation of referring expressions, and introduces so-called 'classic' approaches to modeling reference generation. 'Reference' is a large and diverse

research area, so vD takes the opportunity in Ch. 1 to be explicit about his treatment of this topic. For example, the focus is mostly (but not exclusively) on reference production, and especially on algorithms intended to capture how human speakers refer in particular contexts (versus applications of natural language generation that prioritize utility for listeners). There is also a focus on CONTENT DETERMINATION, or the conceptualization of referential expressions at the level of logical form (i.e. semantic specifications like {*car*, *blue*}). Other aspects of production, like LINGUISTIC REALIZATION, are not addressed because they involve tasks (such as specifying word order) not specific to REG. Finally, most of the models presented are concerned with ‘one-shot’ reference, which excludes forms of reference that rely on the discourse context, like demonstrative NPs or pronouns. vD points out that these contextual phenomena often involve language-dependent mechanisms that lie beyond the scope of content determination. While acknowledging the dangers of restricting attention to instances of so-called ‘literary’ reference, in which speakers generate referring expressions intended to uniquely identify referents for attentive addressees, vD argues (and successfully demonstrates throughout the book) that there are still a sufficient number of open issues that justify this focus.

The next two chapters provide important background concerning relevant philosophical, linguistic, and psycholinguistic issues within the study of reference. Ch. 2 provides a useful introduction to fundamental distinctions, such as denotation versus connotation, extensional and intensional contexts, and referential versus attributive descriptions. Generally, the focus is on concepts that have had particular implications for computational models, especially in their application to more complex cases. Then, Ch. 3 reviews key psycholinguistic work, including debates on the role of shared knowledge in the production and comprehension of referring expressions. Particular attention is devoted to evidence concerning possible Gricean constraints on the types of information that speakers mention when referring to objects. Crucially, vD contrasts the notion of DISCRIMINATORY POWER, which emphasizes properties that successfully distinguish referents from other objects in the domain, with INTRINSIC PREFERENCE, which captures how some object properties, such as color, appear to have psychological priority. In general, vD emphasizes the fact that human speakers are not always strictly rational (in the Gricean sense) and that computational models interested in capturing ‘humanlikeness’ must be explicit about the constraints that determine how REG algorithms consider particular object attributes.

Part II is the heart of the book, focusing on classic computational approaches to REG. Ch. 4 begins by describing early artificial-intelligence work on natural language generation, which often emphasized problems that were difficult to explore systematically. In this context, vD highlights Dale and Reiter’s (1995) influential INCREMENTAL ALGORITHM as a crucial advance, focusing as it did on the ‘smaller’ challenge of referent IDENTIFICATION. This narrower focus set the stage for more rigorous empirical testing through the classic REG task, which involves reference to a single entity within a domain of equally salient objects, and no context-dependent, vague, or complex properties or relations. Under the incremental algorithm, referent identification is accomplished by monotonically considering relevant object attributes one at a time, in a particular PREFERENCE ORDER, with the goal of being ‘logically complete’—that is, successful at producing a distinguishing description in every situation in which such a description exists. Ch. 5 describes how, following Dale and Reiter’s groundbreaking work, computational work on REG took an empirical turn, most commonly in the form of organized ‘evaluation challenges’ in which research teams take part in open tests of alternative algorithms designed to address a particular research problem, using common data sets and agreed-upon methods of evaluation. The bulk of this chapter details the logic of this type of empirical evaluation. Using data from the so-called TUNA reference corpus (van Deemter et al. 2006), vD demonstrates how different assumptions about attribute preference order across versions of the incremental algorithm can provide better-or-worse matches to human data. vD gives a balanced appraisal of the utility of these evaluation challenges, noting that while they have provided researchers with important common ground, one may reasonably wonder how their conclusions generalize to other domains or alternative evaluation metrics.

Following this introduction to the classic approach, Ch. 6 briefly considers several alternative approaches that (unlike the incremental algorithm) view REG nondeterministically, often as a

probability distribution over a set of possible referring expressions. Among other work, vD highlights Frank & Goodman 2012, which considers the choices of both speakers and listeners within a rational, Bayesian framework, and also presents some of his own work on PROBABILISTIC REFERENTIAL OVERSPECIFICATION (van Deemter et al. 2012), designed to capture within-speaker variation in patterns of overspecification. This chapter also highlights other computational approaches that have emphasized communicative goals other than simple referent identification per se (such as summarization; Jordan 2000).

In Part III, vD considers various ways that recent REG models have attempted to move beyond classic reference expression (RE) types to account for more than just single referents or one-place predicates. Ch. 7 presents a number of ideas about the challenges posed by reference via proper names, while Ch. 8 discusses reference to sets of objects (e.g. identifying ‘the white poodles’ among dogs of various kinds and colors) as a generalization of the original incremental algorithm. This chapter draws heavily on formal set theory to consider issues such as collective properties, negation when referring to exceptions to a set (‘the one that is not a poodle’), and disjunctive expressions (‘the horses and cows’). Ch. 9 considers the challenge for REG algorithms posed by gradable properties, such as ‘tall’ or ‘huge’, that invoke a value along a continuum. In this context, vD describes some of his own influential work on the generation of ‘vague’ descriptions (van Deemter 2006). This chapter also considers issues of multidimensionality and salience, suggesting that salience itself may best be regarded as a multidimensional, gradable property (i.e. an RE like ‘the brown dog’ may be understood as picking out the most salient brown dog).

In Ch. 10, vD argues that computational approaches to REG would benefit further by incorporating insights from recent developments in formal approaches to KNOWLEDGE REPRESENTATION and DESCRIPTION LOGIC. On his account, incorporating formalisms from these areas would allow REG algorithms to successfully deal with more complex referential situations, including diverse forms of quantification and negation such as ‘the woman who feeds two cats’ or ‘is not loved by all dogs’, which involve not just relations between objects but relations between sets. Then, Ch. 11 returns to the issue of logical completeness and considers whether the expressive power of the algorithms discussed in the previous chapters, which deal with a range of increasingly complex referential situations, is such that these algorithms are able to generate a distinguishing description whenever one exists. After providing a proof of what it means for an entity to be ‘referable’ via first-order predicate logic, vD suggests that there may be limits to what such models may be expected to accomplish, given that even human speakers struggle to produce adequate descriptions in complex circumstances.

Next, Part IV consists of four brief chapters that outline a handful of other situations with the potential to complicate existing approaches to modeling reference generation—mostly because they violate one or more presuppositions of the classic REG task. When possible, vD considers how computational models might address these challenges. Ch. 12 describes the challenge of reference in large or complex domains, in which relevant properties are not always immediately apparent and listeners may have to ‘search’ to identify the necessary information. Ch. 13 considers the challenge posed by uncertainty about a listener’s knowledge, such as when the audience is unknown or open ended (e.g. in communicating online). Ch. 14 examines cases of ‘approximate’ reference, in which the aim is not to uniquely identify a specific target (e.g. by eliminating all distractors) but rather to produce an expression that is good enough given the communicative goal. Then Ch. 15 considers the challenges posed by goals other than identification, like persuasion, in which choices about which properties to mention are based not on discriminability but on factors such as ‘interestingness’ or ‘unexpectedness’. As vD points out, notions of relevance are likely to shift depending on the purpose behind reference generation.

Finally, Ch. 16 reflects on some of the strengths and limitations of computational approaches to REG to date. Although REG algorithms historically have not been taken as models of reference production, vD argues that, given their emphasis on capturing the behaviors of human speakers, these algorithms may plausibly be situated within the more general framework of COMPUTATIONAL COGNITIVE MODELING—even as he acknowledges that these models have not always

been strongly constrained by cognitive considerations and have tended to focus more on product than process. vD also revisits the Gricean maxims and considers how various algorithms presented throughout the book may be viewed through the lens of factors such as truthfulness and ‘quantity’. He also raises some important questions about the notion of a fixed preference order (central to the incremental algorithm), observing that not all values of a particular attribute may be preferred to the same degree. For example, in some cases it may be the EXTREMITY of a property (how ‘outside the norm’ it is) that makes it worth mentioning. Finally, vD enumerates a number of open problems in REG that are likely to occupy researchers in the near and longer term, such as how to generate logically complex REs and proper names, variation within and between speakers, generalization to more naturalistic domains, and the possibility of incorporating insights from neuroscience.

Overall, this book is a thought-provoking tour through the conceptual and methodological landscape surrounding computational approaches to reference expression generation. Although many chapters delve into the nitty-gritty of algorithms devoted to the complexities of reference generation, vD takes care to ensure that the concepts and logical issues remain accessible to a broader cognitive science audience, even going so far as to provide a flowchart for how readers from fields like computer science or psychology might wish to read through the volume—for example, by skipping or reordering certain chapters. Also, while more formal notations of specific algorithms are frequently presented, most are also presented in a gloss form that captures their logic in ways that facilitate understanding. Anyone interested in gaining a better handle on the intricacies of linguistic reference—whether or not they are involved in computational modeling directly—is likely to learn a great deal from reading this book. At the same time, for those engaged in modeling, vD clearly wishes to push computational work on REG forward, calling for more attention to forms of reference that have barely been addressed in the empirical literature, much less in computational models. *Computational models of referring* will undoubtedly serve as a valuable roadmap for the cognitive modeling of reference generation for years to come.

#### REFERENCES

- DALE, ROBERT, and EHUD REITER. 1995. Computational interpretations of the Gricean maxims in the generation of referring expressions. *Cognitive Science* 19(2).233–63. DOI: 10.1207/s15516709cog1902\_3.
- FRANK, MICHAEL C., and NOAH D. GOODMAN. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084).998. DOI: 10.1126/science.1218633.
- JORDAN, PAMELA W. 2000. Can nominal expressions achieve multiple goals? An empirical study. *Proceedings of the 38th annual meeting of the Association for Computational Linguistics*, 142–49.
- VAN DEEMTER, KEES. 2006. Generating referring expressions that involve gradable properties. *Computational Linguistics* 32(2).195–222. DOI: 10.1162/coli.2006.32.2.195.
- VAN DEEMTER, KEES; ALBERT GATT; ROGER P. G. VAN GOMPEL; and EMIEL KRAHMER. 2012. Toward a computational psycholinguistics of reference production. *Topics in Cognitive Science* 4(2).166–83. DOI: 10.1111/j.1756-8765.2012.01187.x.
- VAN DEEMTER, KEES; IELKA VAN DER SLUIS; and ALBERT GATT. 2006. Building a semantically transparent corpus for the generation of referring expressions. *Proceedings of the 4th International Conference on Natural Language Generation (Special session on Data Sharing and Evaluation) (INLG-06)*, 130–32.

Department of Psychology  
Northwestern University  
2029 Sheridan Rd.  
Evanston, IL 60208-2710  
[whorton@northwestern.edu]