

# ADVANCES IN NATURAL LANGUAGE GENERATION

An Interdisciplinary Perspective

VOLUME 2



Communication  
in Artificial Intelligence

EDITED BY

Michael Zock and Gérard Sabah

# Advances in Natural Language Generation

An Interdisciplinary  
Perspective

Volume 2

*Edited by*

Michael Zock and Gérard Sabah

江苏工业学院图书馆  
藏书章



Pinter Publishers, London

© Michael Zock and Gérard Sabah, 1988

First published in Great Britain in 1988 by  
Pinter Publishers Limited  
25 Floral Street, London WC2E 9DS

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted by any other means without the prior written permission of the copyright holder. Please direct all enquiries to the publishers.

**British Library Cataloguing in Publication Data**

A CIP catalogue record for this book is available from the  
British Library

ISBN 0 86187 995 3

Typeset by Joshua Associates Ltd., Oxford  
Printed by Biddles of Guildford Ltd.

---

## Foreword

---

Natural language generation is a young field; probably only five or six years old. True, there have been isolated efforts to make computers speak human languages in a principled way since the very first computers were developed, but the work of these early pioneers largely fell on deaf ears. Building on other people's work was the exception rather than the rule, and there was no community to provide support. Nearly all of the first people in the field soon went on to other projects where their work was better received. The problem was that their work was before its time: there can be no community without an appreciation of a common set of demands and problems, and until the early 1980s this simply did not exist. Why this has happened has a combination of several factors, not the least of which is exemplified by this book.

One factor contributing to generation's youth has been the lack, until recently, of any need for it. The people who study generation are trying to understand the human capacity to use language—with all its subtleties of nuance, and the complexity, even arbitrariness, of its motivations. They develop large and elaborate theories, with equally large programs to instantiate and test them, as befits the subject they are studying. Computers, on the other hand, do not think very subtle thoughts: indeed, their programs, even artificial intelligence programs, inevitably leave out rationales and goals behind the instructions for their behaviour, and certainly none of them (with the exception of a few programs written specifically to drive generation systems) has any emotional or even rhetorical attitudes towards the people who are using them. Without the richness of information, perspective and intentions that humans bring to what they say, computers have no basis for making the decisions that go into natural utterances—it makes no sense to include a natural language generator in one's computer system if there is no way to drive it.

But while computers have not had the knowledge that it takes properly to motivate the content and style of what they might say, they nevertheless have had strong demands placed on them to communicate—for the AI

program to explain the reasoning behind its decisions or the machine tutor to interact with its students—and this invariably means using natural language. For some of these tasks, artfully combined ‘canned’ texts—strings of words in print statements—have been sufficient. In the past they were certainly a better technique than the cumbersome and grammatically brittle generators that were available. But those days are rapidly coming to an end, and with that comes the realization that a true generation capacity will soon be a requirement, not a luxury.

A modern generator permits computer application programs to distance themselves from language in ways that print statements cannot permit. This distance makes it possible to increase dramatically the quality of the texts that are produced by allowing independent influences to enter into the process and their constraints and contributions combined. There are of course technicalities: the generator takes care of making the number of subject and verb agree and gets the right case on the pronouns. The more important contribution of a generator, however, is that it provides a framework that makes the process of assembling and couching the message a text conveys exceptionally flexible.

An applications program, on its own, may select the individual references and propositions that should be composed to form the complex sentences of, say, an explanation to a specific, specialized audience. (Though note that much of the organization in a text has to come from its position in the ongoing discourse and must accommodate the syntactic peculiarities of the wording that is used—all of which requires linguistic knowledge that will reside in the generator.) A grammatically sophisticated generator will annotate these conceptual units according to the different ways they could be combined and realized, and will establish a set of decision criteria—rhetorical, semantic and pragmatic—which govern what actually happens. The process invariably then goes through a series of stages and representational levels, each with its own capabilities and specializations, and with its own opportunities for independent influences to have their effect. This stratification and differentiation sharpens and clarifies our sense of what it is that makes an utterance effective, and makes research on generation tractable and additive.

The generators actually available right now of course have their limitations and rough spots—a great many of them. This motivates research into alternative theories of grammar, alternative control flows, experiments with particular decompositions of the process into modules, and especially the study of language itself: how goals are actually achieved, nuances expressed, the intricacies of grammatical constructions and references dealt with, and so on. Such research is the subject of this volume.

Research does not flourish without a community to nourish it. This

brings us to the other principal factor—the key to the emergence of generation as a field—the international workshops on natural language generation. The first workshop was in Germany in 1983, riding on the coat-tails of an international AI meeting; the second at Stanford the next year; and the third in Holland in 1986. The most recent meeting is the subject of this book: the European workshop in January of 1987 at the Abbey of Royaumont in France.

These workshops have established generation as a field: they have set its (rather wide) boundaries, promulgated a common sense of the state of the art, identified the central questions. Even more important, they have provided the means for the individuals in the field to come to know each other, to become more than just anonymous names in papers, to meet new people and learn about the range of work that they are doing. This process has created a sense of community: people who share the same concerns, puzzle over the same problems, react to the same odd things that they hear while talking over dinner. The pioneers have gained colleagues; the newcomers have been inducted into a strong, if thinly populated, tradition.

Part of this tradition is a diversity of views and backgrounds. Unlike some schools within the field of natural language understanding (easily twenty years older), generation research embraces not just computational linguistics and artificial intelligence—its core discipline—but also ‘conventional’ linguistics and psycholinguistics. Perhaps youth brings with it a lack of bias; on the other hand, simply making the effort to bring these diverse groups together to a single place for several days allows the common elements and goals within their work to be recognized. This interdisciplinary group is well represented in this European workshop, as it was in the one before, continuing the cross-connection and exposition that was established there.

Language generation is not a simple process. It is not language understanding in reverse, and it challenges many of the established ideas within linguistic theory. It has its own subtleties and depths, with surprises waiting for those who explore them. Its study has created a new field with enormous vigour and infectious enthusiasm. More so than any of its complementary fields, generation will provide a real ‘window into the mind’, since it alone forces one to synthesize actual mental models—the conceptual sources of what the generators say—before any other work is possible, models that can be tested just by looking at the utterances they lead to and seeing how reasonable they are. It is research that is fun to do, and challenging, and it promises to change the way we think about language and ourselves.

*David D. McDonald*  
*Northampton, Massachusetts*  
*November 1987*

---

## Introduction \*

---

Text production is decision-making under specific social, pragmatic and psychological constraints. It consists of determining, organizing and expressing thoughts in order to achieve some goal.

Despite its youth, the modelling of text production language generation has become a rapidly growing field. In the last decade, several international workshops have been held, special sessions have been scheduled at major AI conferences, and more than a dozen books have appeared (see bibliography). The chapters in this volume are the result of one of those workshops. More precisely they are the revised version of chapters presented at the First European Workshop on Language Generation at Royaumont Abbey in France. These chapters reflect the kind of problems researchers are currently struggling with.

We start Volume 1 of this book with a chapter by McKeown and William Swartout, two prominent American scientists who survey the literature on text generation. This review is useful in various respects: it defines the problems of text generation, it provides a general framework within which the different contributions can be located, and it describes the techniques developed by researchers. As this chapter is primarily concerned with work done in American institutions, it provides a good setting for the rest of the chapters presented in this book, which comes mainly from European researchers.

The work presented here is grouped into four sections: grammatical formalisms, psychological issues, implementations and educational applications.

### **Grammatical formalisms**

Russell Block analyses the suitability of LFG as a formalism for parsing and generation. On linguistic grounds, he comes to the conclusion that

\* We are greatly indebted to ARI Communication and the German French Program of the CNRS who sponsored this workshop.

separate formalisms are necessary. While a non-transformational theory like LFG might be perfectly adapted for parsing, it seems to be inadequate for generation. Standard Theory Transformational Grammar seems to be a better candidate.

Harry Bunt is also concerned with the problem of designing a single grammar formalism for analysis and synthesis. His formalism, a semantically and pragmatically augmented phrase structure grammar, is capable of generating discontinuous constituent structures, a notorious problem in language generation.

Helmut Horacek disagrees with Block's conclusion. The research group he works in uses LFG both for parsing and generation. The author shows in his chapter how unification grammars can be adapted to overcome some of the shortcomings of LFG, in particular backtracking and handling of free word order.

Starting from the premise that linguistic well-formedness is a necessary but not sufficient condition for communicative adequacy, Simon Dik suggests linking the linguistic component to a knowledge base and to a reasoning component. The task of the latter is to draw conclusions on the basis of existing or incoming information. As we have three components we are faced with the problem of how to represent knowledge in each one of them so as to allow for communication. An optimal solution would be the use of a unique representation language. According to the author, functional grammars are particularly well suited to solve this problem.

Robin Fawcett describes the theoretical underpinnings of COMMUNAL, a project he is currently conducting. The acronym of the system stands for Convivial Man-Machine Understanding through Natural Language. Its goal is to enable ordinary people to interact naturally with knowledge-based systems. Discussing some of the basic decisions which must be made before starting to build such a system, he argues specifically for the use of meaning-based or systemic functional grammars.

### **Implementational issues: sentence-, discourse- and text generation**

Giovanni Adorni discusses two methods of converting a deep structure into a surface string. The first one starts from conceptual dependency networks, using a transformational grammar. The second one uses logical predicates as input. The task is divided into a lexicalization and a linearization procedure. The former produces an unordered set of words composing the final sentence; the latter integrates these elements into a syntactic structure. The author claims a quasi-universal status for the underlying meaning representation, while the linearization procedure is language-specific.



An interesting feature of the first system is the fact that it only knows about syntax. What conceptual primitives or relationships are necessary to build the surface structure, or how those conceptual primitives are mapped on lexical items is learned by the system.

Chris Mellish's goal is to organize large amounts of data in order to generate smooth readable text. Rejecting the 'hill climbing' method as too expensive he presents another technique of generating text from plans. Critically analysing the shortcomings of his method, he suggests a series of improvements.

Nathalie Simonin's chapter deals with the task of structuring an unordered set of information and converting this set of messages into coherent text. Starting from a knowledge base, her system produces short summaries of the agricultural, economic and political situation of various countries. The process is decomposed into different tasks: extraction of relevant information, dynamic building of the text plan by grouping the information according to some point of view (for example, topic), elimination of redundancies, determination of linear order of the messages, generation of inter-sentence links, sentence generation. It should be noted that her system is bi-directional (top-down and bottom-up). There is constant interaction between choices concerning content and the choice of textual structure. Another interesting feature of this system is the fact that the various levels all use the same representation mechanism, namely functional descriptions.

George Houghton and Mark Pearson are concerned with the production of dialogues. The communication setting is composed of two 'actors' (robots) which try to solve practical problems in a microworld. The user of the program determines the actors' goals as well as the physical arrangement of the elements composing this microworld. The actors attempt to make plans to achieve their goals. Whenever they lack knowledge to form or execute a plan, they engage in dialogue, querying the other participant for help. The output from this system is synthesized speech.

Domenico Parisi and Alessandra Giorgi present a lexically driven word ordering component. This approach is interesting in that it contrasts with the traditional philosophy of syntax-, i.e. structure-driven processing. As far as the role of the lexicon is concerned, there seems to be a convergence between psychological and various recent linguistic theories. All of them emphasize the role of the lexicon as an important aspect of structure building.

Robert Dale is concerned with the generation of referring expressions in structured discourse. As successful reference implies adequate description of a referent in order to allow its discrimination from a set of alternatives, the question arises of how to reduce the set of possible candidates (search

space). Examining Grosz's and Sidner's approach, he concludes that their model is inadequate in several respects; however, a revised version could be good candidate.

A striking feature of Hans Joachim Novak's system is the fact that it takes visual data as input. His system generates coherent text describing an image sequence of a traffic situation. As the hearer cannot see the scene, the system has to describe it in terms that allow the hearer to imagine it, i.e. build a corresponding visual representation of the scene. In order to achieve this goal, particular attention must be paid to the description of the objects and to specification of their relative position in the scene.

Dietmar Roesner sketches the evolution of his system from a machine translation system (SEMSYN) to a full-fledged top-down driven text generator (SEMTEX). The former translates titles of scientific papers on computer science from Japanese into German, whereas the latter produces newspaper stories about job market development. GEOTEXT, a more recent application, produces descriptive texts from a set of commands given by a student, expressing what this latter is doing while solving a geometry problem.

Kwee Tjoe Liong discusses the linguistic and computational problems he encountered when trying to extend the grammar of a very well-known system, TEXT. He also raises the problem of conflicting interests between the goals of computer scientists and linguists. While the former are concerned with optimization, concrete results and performance, the latter are more concerned with long-term issues and possible generalizations. Every system designer is faced with the problem of where to draw the border line, while trying to reconcile these two conflicting goals.

## Psychological issues

Trevor Harley suggests that normal speech production uses two types of mechanisms. One is fast, most likely parallel, and not prone to interference from the contents of working memory: this mechanism is called AUTOMATIC. EXECUTIVE PROCESSING, on the other hand, is slow, serial and prone to interference from working memory. The advantage of this slower mechanism is that it provides the speech production system with a flexibility that it would otherwise not have. The author provides empirical data to substantiate his claims.

H. Schriefers and Thomas Pechmann argue that planning and expressing content are done in parallel. They even claim partial parallelism to be a prerequisite to account for the fluency and speed of human performance. While the idea of incremental processing is appealing, at least on intuitive

grounds, little empirical evidence has been gathered up to now. These authors back up their claims by providing such evidence.

Vanda Zammuner presents an empirically motivated model of discourse planning and production. She analyses the different knowledge sources and shows how these components interrelate. Viewing discourse production as a special case of problem-solving, she is particularly concerned with strategic planning, i.e. how to be efficient given a set of constraints.

Annie Piolat and Fernand Farioli present a system which assists students in text planning (composition). Their system is also a research tool. Keeping a trace of various aspects of human performance—for example, time taken for a specific operation: text planning, wording and so on—the psychologist may draw on these facts in order to elaborate his theory. The authors' goal is to provide empirical evidence of the interrelationship of the planning, composing and revising processes.

## Application

Michael Zock starts by analysing why natural languages are so hard to explain, to learn, or to use. Coming to the conclusion that the difficulties are caused by the interaction of the various constraints, he presents an outline of a system whose task is to disentangle these intricate dependencies. The purpose of his system is to assist a student in expressing the content. Since the user can test the outcome of his choices, he should learn what the choices are, what they depend upon, and how they interrelate. The author puts particular emphasis on this last point. Being concerned with the temporal aspects of the process, he tries to provide a means to visualize the interaction of conceptual, textual and linguistic choices.

D. Bakker, Bieke van der Korst and Gerjan van Schaaijk's system teaches students the basics of a specific linguistic theory, namely functional grammar. The system helps the students to specify what they want to say, translating this message into linguistic form. While the user expresses the content, the system checks for syntactic and semantic consistencies and provides feedback in case of error. The major processing steps may be traced, thus showing what grammatical component is activated at any given time. This facility is useful in that it may give the user some insight into the overall organization of the grammar at work.

## Conclusion

Language generation is a complex task; consequently it does not pertain to one sole domain. The meeting at Royaumont was an attempt to bring

together people from various backgrounds in order to have them look at the same object from different perspectives.

Text generation is also a real challenge both because of its complexities and its potential social, cultural and scientific impact. It is a fascinating and young discipline. In the past a great deal of effort has been devoted to investigating the structure of sentences and more recently to the structure of text. Surprisingly little is known about the processes used to produce these structures. This is about to change, because studying text structures amounts to studying products of thought, while studying dynamic text production amounts to studying the process of thought, that is to say, thought in action.

*Michael Zock and Gérard Sabah*  
*LIMSI/Orsay*

## Bibliography

- Aitchinson, J. (1976), *The Articulate Mammal*, London, Hutchinson & Co.
- Appelt, D. (1985), *Planning English Sentences*, Cambridge, Cambridge University Press.
- de Beaugrande, R. (1984), *Text Production: Towards a Science of Composition*, New Jersey, Ablex.
- Butterworth, B. (ed.) (1980), *Language Production*, vols 1 and 2, London, Academic Press.
- Danlos, L. (1987), *The Linguistic Basis of Text Generation*, Cambridge, Cambridge University Press.
- Davey, A. (1979), *Discourse Production*, Edinburgh, Edinburgh University Press.
- Deese, J. (1984), *Thought into Speech*, Englewood Cliffs, NJ, Prentice Hall.
- Kempen, G. (ed.) (1987), *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics*, Dordrecht, Martinus Nijhoff Publisher.
- Leont'ev, A. (1975), *Psycholinguistische Einheiten und die Erzeugung sprachlicher Äusserungen*, Munich, Hueber.
- Levelt, W. (1988), *Speaking: From Intention to Articulation*, Cambridge, Mass., MIT Press.
- McKeown, K. (1985), *Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural Language Text*, Cambridge, Cambridge University Press.
- Meehan, J. (1980), *The Metanovel: Writing Stories by Computer*, Library of Congress Cataloging in Publication Data.
- Patten, T. (1988), *Text Generation*, Cambridge, Cambridge University Press.
- Popov, E. (1982), *Talking with Computers in Natural Language*, Berlin, Springer Verlag.
- Rosenberg, S. (ed.) (1978), *Sentence Production: Developments in Research and Theory*, Hillsdale, NJ, Lawrence Erlbaum Associates.
- Schlesinger, I. (1977), *Production and Comprehension of Utterances*, Hillsdale, NJ, Lawrence Erlbaum Associates.
- Zamuner, V. (1981), *Speech Production*, Hamburg, Buske Verlag.

---

# Contents

---

## Volume 1

Foreword	
<i>D. McDonald</i>	ix
Introduction	
<i>Michael Zock and Gérard Sabah</i>	xiii
<b>Part I State of the art</b>	<b>1</b>
1 Language generation and explanation	
<i>K. R. McKeown and W. R. Swartout</i>	1
<b>Part II Linguistic approaches: in defence of a particular theory-formalism</b>	
2 Can a 'parsing grammar' be used for natural language generation? The negative example of LFG	
<i>R. Block</i>	53
3 The application of unification for syntactic generation in German	
<i>H. Horacek</i>	63
4 Concerning the logical component of a natural language generator	
<i>S. C. Dik</i>	73
<b>Part III Implementational issues</b>	
5 Two approaches to natural language generation	
<i>G. Adorni</i>	93
6 The production of spoken dialogue	
<i>G. Houghton and M. Pearson</i>	112

7	Natural language generation from plans <i>C. Mellish</i>	131
8	An approach for creating structured text <i>N. Simonin</i>	146
<b>Part IV Psychological issues</b>		
9	Automatic and executive processing in semantic and syntactic planning: a dual process model of speech production <i>T. A. Harley</i>	161
10	Incremental production of referential noun-phrases by human speakers <i>H. Schriefers and T. Pechmann</i>	172
<b>Part V Educational applications</b>		
11	Natural languages are flexible tools: that's what makes them hard to explain, to learn and to use <i>M. Zock</i>	181

Index to Volume I

Volume 2

Foreword <i>D. McDonald</i>		ix
Introduction <i>Michael Zock and Gérard Sabah</i>		xiii
<b>Part I Linguistic Approaches: in defence of a particular theory-formalism</b>		
1	Discontinuous Phrase-Structure Grammar and its use in sentence generation <i>H. Bunt</i>	1
2	Language generation as choice in social interaction <i>R. P. Fawcett</i>	27
<b>Part II Implementational issues</b>		
3	A lexically distributed word ordering component <i>D. Parisi and A. Giorgi</i>	51

---

4	The generation of subsequent referring expressions in structured discourses <i>R. Dale</i>	58
5	Generating referring phrases in a dynamic environment <i>H.-J. Novak</i>	76
6	The generation system of the SEMSYN project: towards a task-independent generator for German <i>D. Rösner</i>	86
7	Natural language generation: one individual implementer's experience <i>T. L. Kwee</i>	98
<b>Part III Psychological issues</b>		
8	Discourse planning and production: an outline of the process and some variables <i>V. L. Zammuner</i>	121
9	The effect of the macro-control of information on the temporal characteristics of text production <i>A. Piolat and F. Farioli</i>	144
<b>Part IV Educational applications</b>		
10	Building a sentence generator for teaching linguistics <i>D. Bakker, B. van der Korst and G. van Schaaijk</i>	159
	Index to Volume 2	175

---

# Part I Linguistic approaches: in defence of a particular theory—formalism

---

## 1 Discontinuous Phrase-Structure Grammar and its use in sentence generation

---

HARRY BUNT\*

### 1.1 Introduction

The task of sentence generation, as conceived in this chapter, is the systematic translation of meaning representations into strings of words. This task is not easy, since there is in general no simple relation between the elements of meaning representations and the words in sentences. Instead, we are dealing with a complex relation between substructures in meaning representations and substructures of sentences.

A specific, notorious problem arises when semantic substructures correspond to syntactic constituents of which the words should not be placed immediately after each other, i.e. when discontinuous constituents have to be generated. For many sentence generators this is a major stumbling block, for which special provisions have to be made (see e.g. Block 1988; Horacek 1988; Kempen 1987). The reason for this is that the grammar formalisms incorporated in these generators do not describe both the relation between semantic and syntactic constituents and the word order constraints on syntactic constituents in a sufficiently general way, powerful enough to deal with discontinuous constituents. This causes a problem especially if one aims at incremental sentence generation, i.e. at a process where sentences grow in an 'on-line' fashion as the elements of the meaning representation are processed in a certain order (cf. Kempen 1987).

The problems of generating discontinuous constituents parallel those of their interpretation in an on-line syntactic-semantic parser, which deals with the same complex relation between meaning representations and sentence constituents, but now in reversed direction (see Bunt 1987). A

\* *Institute for Language Technology and Artificial Intelligence Tilburg University, The Netherlands*



grammar which describes this relation in general, including the case where the syntactic constituents are discontinuous, should be a sound basis for the design of both parsing and generation algorithms which can handle the phenomenon of syntactic discontinuity in an elegant way.

In the TENDUM dialogue project (see Bunt *et al.* 1985) we have designed and implemented a grammar formalism precisely for this purpose, the syntactic side of which is called Discontinuous Phrase-Structure Grammar (DPSG). In this chapter we describe the DPSG formalism, how it is coupled to semantic and pragmatic interpretation in the TENDUM dialogue system, and how it can be used for sentence generation.

The generation of sentences in the context of natural-language dialogue is in some respects simpler than in the context of text generation, and in some respects more difficult. It is less difficult in that several problems that arise in the generation of multi-sentence text, such as the use of tense and of anaphora, hardly arises in this context. It is more difficult in that it requires not only the expression of a certain semantic content, but also that of the communicative function with which the sentence is used in the dialogue; the task is in fact more accurately described as utterance generation, rather than sentence generation.

We begin by taking a look at the phenomenon of syntactic discontinuity and its representation.

## 1.2 Syntactic discontinuity and its treatment in DPSG

### 1.2.1 *Discontinuous constituents*

In languages with rigid word order, like English, French and Dutch, but in many other languages as well, it happens quite often that the constituents of an expression are not adjacent. This phenomenon occurs in a wide variety of constituent types: noun phrases, verb phrases, adjective phrases, adverb phrases, etc. The examples (1)–(7) illustrate this: in (1) we see a discontinuous verb phrase (*talked about politics*); in (2) a discontinuous preposition phrase (*from which children*); in (3) a discontinuous adjective phrase (*better than I expected*); in (4) a split verb (*wake up*); in (5) a split negation (*ne pas*) and a discontinuous noun phrase with ‘extraposed’ relative clause (*un de tes cousins qui demeure à Bruxelles*); in (6) a discontinuous adverb phrase (*harder dan ooit tevoren*); in (7) a discontinuous noun phrase with ‘extraposed’ prepositional phrase (*ein Auto mit 6 Türe*).

(1) John talked, of course, about politics

(2) Which children did Jasper expect to get a present from?