
J a p a n e s e

S e n t e n c e

P r o c e s s i n g

Edited by

**R E I K O M A Z U K A
& N O R I K O N A G A I**

Japanese Sentence Processing

Japanese Sentence Processing

Edited by

Reiko Mazuka
Noriko Nagai
Duke University



1995

LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS
Hillsdale, New Jersey

Hove, UK

Copyright © 1995 by Lawrence Erlbaum Associates, Inc.

All rights reserved. No part of the book may be reproduced in any form, by photostat, microform, retrieval system, or any other means, without the prior written permission of the publisher.

Lawrence Erlbaum Associates, Inc., Publishers
365 Broadway
Hillsdale, New Jersey 07642

Library of Congress Cataloging-in-Publication Data

Japanese sentence processing / edited by Reiko Mazuka, Noriko Nagai
(Duke University).

p. cm.

Papers of the International Symposium on Japanese Sentence
Processing, held at Duke University in 1991.

Includes bibliographical references and indexes.

ISBN 0-8058-1125-7

1. Japanese language—Psychological aspects—Congresses.
2. Japanese language—Data processing—Congresses. 3. Parsing
(Computer grammar)—Congresses. I. Mazuka, Reiko. II. Nagai,
Noriko. III. Duke University. IV. International Symposium on
Japanese Sentence Processing (1991: Duke University)

PL513.J37 1994

495.6'5—dc20

94-2032

CIP

Books published by Lawrence Erlbaum Associates are printed on acid-free
paper, and their bindings are chosen for strength and durability.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Contents

Contributors	vii
Preface	ix
1 Japanese Sentence Processing: An Interdisciplinary Approach <i>Reiko Mazuka and Noriko Nagai</i>	1
2 Information-Paced Parsing of Japanese <i>Atsu Inoue and Janet Dean Fodor</i>	9
3 Syntactic Representation and Interpretive Preference <i>Bradley L. Pritchett and John B. Whitman</i>	65
4 Constraints on Topics and Their Gaps: From a Parsing Perspective <i>Noriko Nagai</i>	77
5 An Overview of JPSG: A Constraint-Based Grammar for Japanese <i>Takao Gunji</i>	105

6 A Constraint-Based View of Language: A Unified Theory of Competence and Performance <i>Kôiti Hasida</i>	135
7 The Syntax and Semantics of the Japanese Language Engine <i>Megumi Kameyama</i>	153
8 Madama Butterfly Redux: Parsing English and Japanese with a Principles and Parameters Approach <i>Robert C. Berwick and Sandiway Fong</i>	177
9 Null Elements in Parallel Structures in Japanese <i>Susumu Kuno</i>	209
10 Licensing Constraints and the Theory of Language Processing <i>Amy Weinberg</i>	235
11 Scrambling and Probe Recognition <i>Mineharu Nakayama</i>	257
12 Transparency Between Parser and Grammar: On the Processing of Empty Subjects in Japanese <i>Tsutomu Sakamoto</i>	275
13 Can Japanese Speakers Be Led Down the Garden Path? <i>Reiko Mazuka and Kenji Itoh</i>	295
14 Japanese Trees and the Garden Path <i>Paul Gorrell</i>	331
Author Index	351
Subject Index	355

Contributors

Robert C. Berwick

Artificial Intelligence Laboratory, MIT
545 Technology Square, Rm. 838, Cambridge, MA 02139

Janet Dean Fodor

Linguistics Program, CUNY Graduate Center
33 West 42nd Street, New York, NY 10036-8099

Sandiway Fong

NEC Research Institute
4 Independence Way, Princeton, NJ 08540

Paul Gorrell

Department of Linguistics, University of Maryland at College Park
College Park, MD 20742

Takao Gunji

Graduate School of Language and Culture, Osaka University
1-1 Machikaneyama, Toyonaka, Osaka 560, Japan

Kôiti Hasida

Electrotechnical Laboratory
1-1-4 Umezono, Tsukuba 304, Japan

Atsu Inoue

Department of Economics, Kantoogakuin University
Mutsuura-cho 4834, Kanazawa-ku, Yokohama 236, Japan

Kenji Itoh

The Research Institute of Logopedics and Phoniatics,
Faculty of Medicine, University of Tokyo
7-3-1 Hongo, Bunkyo-ku Tokyo 113, Japan

Megumi Kameyama

SRI International AI Center
333 Ravenswood Avenue, Menlo Park, CA 94025

Susumu Kuno

Department of Linguistics, Harvard University
Cambridge, MA 02138

Reiko Mazuka

Department of Psychology—Social & Health Sciences, Duke University
Durham, NC 27706

Noriko Nagai

Department of Asian and African Languages and Literature
Duke University
Durham, NC 27706

Mineharu Nakayama

Department of East Asian Languages and Literatures
The Ohio State University
204 Cunz Hall, 1841 Milikin Road, Columbus, OH 43210

Bradley Pritchett

Department of Psychology, Carnegie-Mellon University
Pittsburgh, PA 15213

Tsutomu Sakamoto

Department of Applied Linguistics, Faculty of Literature
Kyushu University
6-19-1 Hakozaki, Higashi-ku, Fukuoka 812, Japan

Amy Weinberg

Department of Linguistics, University of Maryland at College Park
College Park, MD 20742

John Whitman

Department of Modern Languages and Linguistics
Morrill Hall, Cornell University
Ithaca, NY 14853

Preface

This volume came out of the International Symposium on Japanese Sentence Processing held at Duke University in October 1991. The symposium consisted of 17 papers presented in four sessions and an overall commentary by Janet Fodor. This volume includes 13 of the papers that were presented at the conference and an additional chapter by Kôiti Hasida, who was unable to attend the symposium.

Organizing an international conference and editing a book have been profoundly educational for us, and we received invaluable assistance from numerous people within Duke and elsewhere. We thank Nan Lin, the director of the Asian/Pacific Studies Institute at Duke University, whose initial support made this conference possible. Janet Fodor provided us with vital assistance in the initial planning of the conference, and generously agreed to be the overall discussant for the conference. Yasunari Harada was instrumental in organizing the Japanese side of the conference and its participants. We thank Mitch Marcus, Donald Hindle, John Hawkins, Syun Tutiya, Yasuhiro Katagiri, and Hiroyuki Suzuki for their contributions at the conference.

A special thanks goes to Mavis Mayer, who was instrumental in managing all of the local arrangements of the conference, and other staff members of the Asian/Pacific Studies Institute at Duke University for their administrative and organizational assistance. We are grateful to Yukio Otsu, Yasushi Yoshimoto, Stanley Dovinsky, and Yoko Collier-Sanuki for their assistance in running the conference. Philip Malatin and Elizabeth Abrams helped us edit and prepare the manuscript.

The conference and the editing of the volume were funded by a grant from the Japan Foundation and additional resources from the Asian/Pacific Studies Institute, the Center for International Studies, and the Office of the Deans of Arts and Sciences at Duke University.

Reiko Mazuka
Noriko Nagai

Japanese Sentence Processing: An Interdisciplinary Approach

Reiko Mazuka
Noriko Nagai
Duke University

INTRODUCTION

This volume is the result of the International Symposium on Japanese Sentence Processing held in October 1991 at Duke University. The symposium provided the first opportunity for researchers in three disciplines, from both Japan and the United States, to participate in a conference where they could discuss issues concerning Japanese syntactic processing. The goals of the symposium were three-fold:

1. To illuminate the mechanisms of Japanese sentence processing from the viewpoints of linguistics, psycholinguistics, and computer science.
2. To synthesize findings about the mechanisms of Japanese sentence processing by researchers in these three fields both in Japan and in the United States.
3. To lay foundations for future interdisciplinary research in Japanese sentence processing, as well as international collaborations between researchers in Japan and in the United States.

For the last two decades, mechanisms of natural language processing have been the focus of extensive investigation in psychology, linguistics, and computer science. In the field of psychology, the investigation of natural language processing is of crucial significance to our endeavors to understand general human cognitive mechanisms. For linguistics, in our attempt to investigate universal properties of

grammatical knowledge, it is essential to investigate how grammatical knowledge interacts with performance; that is, how humans utilize their grammatical knowledge to comprehend and generate sentences. In computer science, models for natural language processing are mathematically laid out and implemented, and their validity is tested by actual simulation.

Research in natural language processing in the three fields just mentioned has only recently gained a forum in interdisciplinary journals and at conferences where researchers in the various fields can present their findings to an interdisciplinary audience. The success of such journals and conferences has shown that such interdisciplinary exchange is highly productive for progress in the field.

To date, the research in natural language processing has been mostly limited to English. Obviously, however, as the ultimate goal of the research lies in the understanding of human cognitive mechanisms, it is critical that cross-linguistic perspectives are represented. Japanese is an ideal language to investigate in contrast to English because it involves reverse forms of syntactic structures in exactly those dimensions that seem likely to be critical for sentence processing (i.e., head direction, presence of extensive empty categories, and presence of scrambling).

In theoretical linguistics, which recognizes the distinction between competence and performance, it is crucial to clarify the relationship between the domain of the two components. In the past, important works by Dryer (1980), J. D. Fodor (1978), and Kuno (1973) showed that performance influences competence. What has not been so clear is *how* performance influences competence. Recent works, such as that by Hawkins (1990), have just begun to address this issue. Another important issue that concerns the exact relation between competence and performance rests on the exact grammatical status of sentences. When a sentence is judged ungrammatical, we must correctly determine whether this judgment arises because the sentence violates grammatical conditions or if it is due to parsing difficulties. Unless we know exactly which grammatical phenomena fall into the domain of competence and which into the domain of performance, we can hardly begin to discuss the relation between the two components in any meaningful way, or even study competence effectively.

In computational linguistics, over the past few decades, numerous models have been put forward as models for human sentence processing (e.g., Frazier and Fodor's Sausage Machine, 1978, Wood's Augmented Transition Network model, 1970, Marcus's Deterministic Parser, 1980). These models are assumed to characterize how humans process natural languages, and they are implicitly or explicitly assumed to be universal. Yet these models are based on detailed studies of English, leaving the study of other languages, especially non-Indo-European languages, largely neglected. When those "universal models" are applied to languages such as Japanese, which is head-final and allows extensive use of empty categories, they are found to make the unrealistic prediction that Japanese should be an extremely difficult language to process (see Mazuka &

Lust, 1988, 1989). Thus, in order to capture the fact that the human natural language processing mechanism is capable of comprehending and generating both head-initial languages, such as English, and head-final languages, such as Japanese, with equal efficiency, it is essential to investigate how Japanese speakers process Japanese sentences and then integrate the Japanese results with the English data accumulated over the years.

In psycholinguistics, recent advances in methodologies for obtaining precise measures of human behavior during on-line comprehension and production of language, such as the eye-movement monitoring technique (cf. Rayner, 1983), the cross-modal priming technique (Swinney, 1979), and the on-line grammaticality judgment task (Kurtzman, Crawford, & Nychis-Florence, 1991), have allowed researchers to test finely tuned hypotheses with regard to how humans process natural language (see Mitchell, in press, for a review). The results of these psycholinguistic experiments can provide empirical validation as to whether a particular model for natural language processing is appropriate as a model of how humans comprehend or generate language. However, experimental investigation in sentence processing has concentrated almost entirely on English and a handful of other European languages, again neglecting the non-Indo-European languages, such as Japanese. Only very recently have researchers attempted to include Japanese in their inquiries.

ISSUES IN JAPANESE SENTENCE PROCESSING

Within these three fields, numerous theories of sentence processing have been put forth. There are two major issues that have concerned researchers of Japanese sentence processing. These issues extend beyond the domain of individual disciplines, and any attempt to resolve them requires collaborative endeavors across the disciplines.

The Universality of the Human Language Comprehension Mechanism

A central concern is the issue of whether the human sentence processing mechanism is in fact universal, which has been assumed either implicitly or explicitly by virtually all researchers in the field.

Everyone agrees that the human sentence processing mechanism must be universal to a certain degree, although the grammars of languages may vary systematically from one to another. The question is how to account for differences across languages within the framework of a universal theory of language processing.

Inoue and Fodor's *information-paced parser* (this volume) provides a strong form of the universal hypothesis. According to their model, the human language

comprehension mechanism is universal, and Japanese and English can be processed by exactly the same mechanism without any parameterization. Others have proposed that while the core of the parsing mechanism is universal, a few parameters must be incorporated into the basic parsing mechanism (Berwick & Fong, this volume; Mazuka, in press; Weinberg, this volume). However, the proposals differ in how the parameterization of the grammar interacts with the specific nature of the parsing mechanism. For example, according to Weinberg's model, the parameterization of the parser is restricted to a minor component of the parser (i.e., whether IP is licensed directly by Infl [English] or indirectly licensed by a nominative marked NP [Japanese] is a parametric variation). In Mazuka (in press), parameterization concerns a more global aspect of parsing (i.e., whether clauses should be attached in relation to the topmost S on-line, or parsed to be the lowest S, until otherwise marked). Mazuka and Itoh's (this volume) *tentative attachment strategy* implies that this strategy may not apply universally.

The Relation Between Grammatical Knowledge and the Processing Mechanism

Another question that is critical in our endeavor concerns the exact nature of the relation between grammatical knowledge and the human language processing mechanism that utilizes such information. Because of its syntactic structure (as previously mentioned), Japanese provides an ideal language in which to investigate this issue.

Models differ in how they formalize the nature of our grammatical knowledge: Government and Binding theory (GB) based formalization (see, in this volume, Berwick & Fong; Gorrell, Inoue, & Fodor; Mazuka & Itoh; Nagai; Nakayama; Pritchett & Whitman; Sakamoto; Weinberg). Unification-theory-based formalization (Gunji; Hasida; Kameyama), or Functional-theory-based formalization (Kuno). Any model for human language processing must make an explicit assumption as to how such linguistic knowledge can be accessed by the parser. In this volume, the researchers' main concerns center around three basic issues.

Types of Linguistic Representation. One issue concerns the validity of the linguistic representation in relation to the parsing mechanism. Gunji argues that if linguistic knowledge is represented in the formalization of Japanese Phrase Structure Grammar, which is in the class of context-free grammars, it is "computationally tractable," thus more psychologically plausible than a GB-based grammar. Berwick and Fong, on the other hand, present a fully implemented computational model for a GB-based grammar, and show that a principles and parameters model of grammar is not computationally "intractable," and can be investigated straightforwardly in its implications for parsing methods.

Types of Processing Components. A second issue concerns the choice of architecture for the parser. Given a particular type of linguistic representation, what types of parsing architecture can best account for the phenomena of Japanese sentence processing? For example, among the researchers using GB-based representation, Inoue and Fodor's model proposes that incoming items are attached fully into a parse tree without delay in both Japanese and English. Pritchett (1992), on the other hand, has proposed a head-driven parser, where such an attachment decision is made at the grammatical head, which results in a delayed attachment decision in a head-final language such as Japanese. In Weinberg's minimal commitment parser, following Marcus, Hindle, and Fleck's (1983) D-theory parser, parsing decisions are made monotonically while individual parsing decisions are underspecified in such a way that inserting an additional node later in the parse will not require additional processing costs. Gorrell's proposal, in which a parser builds a phrase-structure tree based on "dominance" and "precedence," also assumes that parsing decisions are made monotonically.

Kameyama and Hasida each propose significantly different architectures to process Japanese, although both use Unification-theory-based representation. For example, Kameyama adopts a unification-based parser/interpreter called the Japanese Language Engine, which is a descendant of the Core Language Engine (Alshawi, Moore, & Moran, 1987), and Hasida uses a dynamic field of force to model the processing.

Competence/Performance. The dichotomy between grammatical knowledge and the parsing mechanism also has implications for generative theories that differentiate between the competence and performance. In such theories, the dichotomy between grammatical knowledge and parsing mechanisms roughly corresponds to the distinction between theories of competence and theories of performance. To clarify the relation between the two components, a number of questions must be resolved. In this volume, two specific issues relevant to this topic are discussed. The first issue addresses the exact relationship between the mental grammar and the parser. More specifically, we consider whether competence theory can predict the complexity of the human sentence mechanism (see chapter 3) and whether the parser has a transparent relation to the mental grammar (see chapters 11 and 12). Pritchett and Whitman (chapter 3) reassess the once-abandoned Derivational theory of complexity for predicting syntactic complexity. They argue that given the recent developments of linguistic theory (GB theory), in particular, trace theory and the addition of the logical form (LF) component, derivational history can correctly predict the syntactic complexity of a sentence.

Nakayama and Sakamoto (chapters 11 and 12) examine the psychological reality of linguistic theory via psycholinguistic experiments. Sakamoto, for example, considers the degree of transparency in the relationship between the mental grammar and performance theory, claiming that the most recent filler (MRF) strategy proposed by Frazier, Clifton, and Randall (1983) does not recognize the

empty category as a possible antecedent (the lexical fillers only [LFO] hypothesis). Frazier et al. assume a nontransparent relationship between the parser and the mental grammar, the latter of which recognizes the empty category as a possible antecedent of PRO. This assumption is contrary to recent psycholinguistic research, which assumes a transparent relationship between the mental grammar and the parser. Sakamoto tests Japanese sentences against the MRF strategy and the LFO hypothesis and argues on the basis of the experimental results that either the MRF strategy or the LFO hypothesis is false for Japanese. Nakayama examines whether traces left by scrambling can be detected by probe recognition, as is possible with other types of traces (e.g., NP trace) and finds that traces left by scrambling cannot be detected using this technique.

The second issue addresses whether certain grammatical phenomena can be accounted for by an analysis based on performance. Nagai clarifies the grammatical status of “topics” which involve relative clauses. She argues that topicalization out of relative clauses, which has been analyzed in terms of linguistic theory (cf. Hasegawa, 1985) is more accurately analyzed from a performance perspective. She claims that all topicalization out of relative clauses is grammatical (i.e., licensed by grammatical principles). However, some topics result in unacceptable sentences due to parsing difficulties.

For theories such as Functional theory (cf. Kuno, 1987), the domain of linguistic investigation extends beyond the sentence level to include discourse. Such theories do not demarcate competence from performance. However, in this framework as well, the relationship between linguistic knowledge and the processing component is a central concern. Kuno’s chapter examines the appearance of null elements in parallel-sentence discourse and proposes a number of functional conditions. He suggests that these conditions are difficult to implement computationally.

Hasida extends the domain of this research even further and proposes that human language processing should be seen as a part of a general cognitive system, which he proposes to model using a dynamic field of force. His model is motivated by the fact that most practical problems humans face, including language processing, must be solved when only part of the relevant information is available. He argues that the partiality of information in description (one never has every piece of relevant information: i.e., inadequacy in competence) and partiality in processing (one can never process all information exhaustively: i.e., inadequacy in performance) are in fact the same problem, and that the distinction between competence and performance is not necessary.

CONCLUSION

The chapters in this volume have been written from the points of view of three different disciplines, with various immediate objectives—from building usable speech-understanding systems to investigating the nature of competence grammars for natural languages. However, as just discussed, all of the chapters share

the long-term goal of understanding the nature of the human language processing mechanism. In particular, the chapters are concerned with two central issues, namely, the universality of the language processing mechanism, and the nature of the relationship between the components of linguistic knowledge and language processing.

The types of evidence that are used to model Japanese sentence processing differ in the three disciplines. The primary data for a linguistic model are intuitive judgments by native speakers as to whether a particular sentence causes conscious processing difficulty, or whether it is grammatically acceptable. In psycholinguistics, predictions are tested more precisely in controlled experiments, so that differences that may not be available to conscious judgment can also be evaluated. Computational models can be evaluated by implementation and simulation. However, as the chapters in this volume show, the same issue can be investigated via different methodologies, and the converging evidence from different types of data contributes in a most productive manner to our understanding of human language processing. It is our hope that this volume shows that such interdisciplinary research can be fruitful, and that the volume will provide the groundwork for further research in Japanese sentence processing.

REFERENCES

- Alshawi, H., Moore, R. C., & Moran, D. B. (1987). *Research programme in natural language processing*. Annual Report SRI Project 2989, SRI International, Cambridge Computer Science Research Centre, Cambridge, England.
- Dryer, M. (1980). The positional tendencies of sentential noun phrases in universal grammar. *The Canadian Journal of Linguistics*, 25, 123–195.
- Fodor, J. D. (1978). Parsing strategies and constraints on transformations. *Linguistic Inquiry*, 9, 427–473.
- Frazier, L., & Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition*, 6, 291–325.
- Frazier, L., Clifton, C., & Randall, J. (1983). Filling gaps: Decision principles and structure in sentence comprehension. *Cognition*, 13, 187–222.
- Hasegawa, N. (1985). On the so-called ‘zero pronouns’ in Japanese. *The Linguistic Review*, 4, 289–341.
- Hawkins, J. A. (1990). A parsing theory of word order universals. *Linguistic Inquiry*, 21, 223–261.
- Kuno, S. (1973). Constraints on internal clauses and sentential subjects. *Linguistic Inquiry*, 4, 363–385.
- Kuno, S. (1987). *Functional syntax*. Chicago: University of Chicago Press.
- Kurtzman, H., Crawford, L., & Nychis-Florence, G. (1991). Locating WH-traces. In R. Berwick, S. Abney, & C. Tenny (Eds.), *Principle-based parsing: Computation and psycholinguistics* (pp. 347–382). Dordrecht: Kluwer.
- Marcus, M. (1980). *A theory of syntactic recognition for natural language*. Cambridge, MA: MIT Press.
- Marcus, M., Hindle, D., & Fleck, M. (1983). D-theory: Talking about talking about trees. *Proceedings of the 21st Meeting of the Association of Computational Linguistics*, 129–136.
- Mazuka, R. (in press). *The development of sentence processing strategies: A cross-linguistic study between Japanese and English*. Hillsdale, NJ: Lawrence Erlbaum Associates.

- Mazuka, R., & Lust, B. (1988). Why is Japanese not difficult to process?: A proposal to integrate parameter setting in universal grammar and parsing. In J. Blevins & J. Carter (Eds.), *Proceedings of the North Eastern Linguistic Society, 18*, 333–356.
- Mazuka, R., & Lust, B. (1989). On parameter setting and parsing: Predictions from cross-linguistic differences in adult and child processing. In L. Frazier & J. de Villiers (Eds.), *Language processing and language acquisition* (pp. 163–205). Dordrecht: Kluwer.
- Mitchell, D. (in press). Sentence parsing. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics*. New York: Academic Press.
- Pritchett, B. (1992). *Grammatical competence and parsing performance*. Chicago: University of Chicago Press.
- Rayner, K. (Ed.). (1983). *Eye movements in reading: Perceptual and language processes*. New York: Academic Press.
- Swinney, D. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior, 18*, 645–659.
- Woods, W. (1970). Transition network grammars for natural language analysis. *Communications of the Association for Computing Machinery, 13*, 591–606.

Information-Paced Parsing of Japanese

Atsu Inoue
Kantoogakuin University

Janet Dean Fodor
City University of New York

A UNIVERSAL PARSER?

It is easier to argue that parsing Japanese is impossible than to explain how it is done. The special problems posed by the Japanese language result from three central facts about its grammar:

1. Japanese is head-final: Verbs follow their arguments and adjuncts; nouns follow their modifiers; complementizers, subordinating adverbs, and conjunctions follow the clauses they govern.
2. Scrambling: Argument and adjunct phrases are freely ordered within a clause, and (though less commonly) across a clause boundary.
3. Null pronouns: Any argument can be phonologically unrealized if the discourse context is sufficient to identify its reference.

As a result, every sentence of Japanese is infinitely ambiguous up to its final word, and often multiply ambiguous even when complete. The example in (1) is not atypical.

- (1) John-ni Mary-ga shookaishita tomodachi-o odorokaseta to itta.
 Dat Nom introduced friend-Acc surprised Comp said

The sequence of categories in (1) is shown in (2). Some of its many possible structural analyses, with their meanings, are shown in (3). A temporarily possible but ultimately unsuccessful analysis is shown in (4).

(2) NP-ni NP-ga V N-o V Comp V

(3) [NP-ni NP-ga [[e [[e V] N]-o V] Comp] V]

‘To John, Mary said that she had surprised the friend she had introduced to someone.’

[e NP-ni [[e [[NP-ga e V] N]-o V] Comp] V]

‘Someone said to John that he had surprised the friend to whom Mary had introduced him.’

[e [[e [[NP-ni NP-ga V] N]-o V] Comp] V]

‘Someone said that someone had surprised the friend that Mary introduced to John.’

(4) [NP-ni NP-ga [[e e V] N]-o . . .

‘To John, Mary (surprised) the friend she had introduced to someone . . .’

There are three basic designs for a sentence processing system that has to cope with ambiguity: parallel processing; serial processing with revision in case of garden paths; and delay of decisions until disambiguation. But *none* of these is a good design for a language as massively ambiguous as Japanese. A parallel parser is safe; because it computes all analyses, it never misses the correct one. But when there are so many analyses, parallel parsing would consume vast processing resources. A strict cut-off might be imposed on how many analyses to pursue in parallel at each choice point, but reliability would suffer, and since there are many choice points the total burden would still be enormous. A delay model is also safe, because it does nothing until all relevant information is in hand. But the fact that unstructured material is a strain on memory (Miller, 1956) implies that for a realistic parser, delays ought to be few and quite short; yet to parse Japanese reliably, delays would need to be many and long. A serial parser lives dangerously, because it makes guesses on the basis of insufficient information. It wastes neither computation space nor memory, because it structures the input immediately and computes only one analysis. Its costs are all in the *re*-parsing process, when it must revise analyses that prove incompatible with later items in the input. For Japanese it would have to revise at many points, and possibly many times at each point, until it finally finds the one analysis among many that is successful throughout.

Thus, Japanese is not congenial to any familiar kind of parser (or even to mixtures of them). Yet people do parse Japanese successfully—at least, the everyday evidence is that this is so. Some of the success might be due to heavy reliance on semantic/pragmatic information, as is sometimes suggested. Some may

be due to intonation patterns, or to the punctuation that encodes them in writing (see Marcus & Hindle, 1990). There is clearly some truth to both of these claims. It seems to be more difficult to parse Japanese than English when nonsense words are substituted for the “content” words of a text, suggesting that there may be a greater reliance on word meaning. And certainly there are characteristic intonation patterns, and rules for inserting commas in writing, that can help to reduce ambiguity in some constructions. But we submit that semantics/pragmatics and intonation/punctuation are not sufficiently informative to resolve all, or even most, of the on-line ambiguities in Japanese sentences.¹ Throughout this chapter, our working assumptions are that there *is* a heavy syntactic component to sentence comprehension in Japanese, and that the syntactic processor does operate tolerably successfully. Our goal is to understand how this could be so, and to identify what kind of device this processor is.

The first basic cut to be made is between *committal* and *non-committal* parsers. The difference is whether or not the parser adopts a hypothesis about the analysis of an ambiguous input prior to disambiguation. A serial parser is committal; both a (pure) parallel and a (pure) delay parser are non-committal; various mixed systems may be committal with respect to some decisions but not others (see Weinberg, this volume). Pre-disambiguation hypotheses adopted by the parser could reveal themselves as beliefs about what the input means, before it is complete; as expectations about what category of word will be next in the input; and as garden-path experiences when the hypothesized structure turns out to be wrong. Any one of these (early interpretation; expectations; garden paths) could be investigated in psycholinguistic experiments on Japanese, as they have been for English (see Mazuka, 1991; Sakamoto, this volume; Yamashita, Stowe, & Nakayama, 1993). We have no experimental results to report here. All of our data are drawn from the intuitions of native speakers (originally those of the first author, then checked with six other speakers) about parsing difficulty, about on-line semantic interpretations, and about expectations concerning later input. Our reason for relying on intuitions is that psycholinguistic experiments are most informative when there is one specific question to be answered about the truth of some hypothesis. For Japanese, with so many questions unanswered, and so few hypotheses that seem antecedently plausible at all, it seems more useful at this stage of research to spread the net wide and shallow, to get a sense of how parsing proceeds by considering and comparing as many different constructions

¹One grammar book gives the rule: place a comma after each non-final adjective in a string of adjectives; and also the rule: place a comma after an adjective if it does not modify the first noun that follows it. Therefore, the comma in *omoshiroi, shinsetsuna tomodachi-no okaasan* (funny, kind friend's mother) fails to disambiguate which noun is modified by the first adjective. In *shinsetsuna, tomodachi-no okaasan-no oishasan* (kind, friend's mother's doctor) the comma also does not fully disambiguate. It signals that the adjective doesn't modify the first noun, but there are two other nouns to choose between. A comma (its presence versus absence) is a binary signal, but Japanese constructions are usually more than two-ways ambiguous.

and types of ambiguity as possible. Once a feasible model has been formulated, it could be tested experimentally.²

Our claim is that the parser for Japanese is exactly like that for English; only the mental grammar that provides the parser with facts about the language differs between English and Japanese. We believe that the parser for English is a serial device that formulates and revises hypotheses on-line,³ and this is exactly the conclusion we have been led to in the case of Japanese as well. Because wishful thinking was an obvious danger, we did our best to study Japanese parsing impartially, in its own right, setting aside opinions about English. Indeed, we admit here to an initial prejudice that there must be significant *delays* of decision in the case of Japanese. It was a surprise, as we examined more constructions in more detail, to find the evidence building up that Japanese parsing is overwhelmingly serial and immediate. In every case we have looked at, it appears that perceivers are semantically interpreting sentences on-line, and have expectations about how they will continue, indicating that they have formed a hypothesis about sentence structure. Even more surprising, in view of this, is that when an on-line expectation is disconfirmed by the subsequent input, there is typically little or no disruption, no conscious garden-path experience. How to reconcile these two facts is a puzzle. The solution we propose is what we call a “ranked flagged serial parser,” which permits many (though not all) parsing errors to be corrected rapidly and almost painlessly. This model also works well for English. So, as far as the evidence from English and Japanese is concerned, we see no reason to doubt that the human sentence parsing mechanism is fully universal. A universal parser is antecedently the most attractive prospect, as is agreed on all sides. We would expect the performance mechanisms for human language to be universal because we would expect them to be innate: Whatever survival value may be gained by having innate linguistic knowledge would be wasted if extensive learning of performance systems were necessary before that knowledge could be put to practical use in talking and understanding.

To summarize, despite the manifest structural differences between English and Japanese which suggest that they could not be parsed in similar fashion, we argue that exactly the same parsing mechanism is at work in both. More specifically, we argue that:

1. Japanese parsing is serial, without parallelism or delay; in particular, it does *not* postpone decisions about the structure of a phrase until the head of the phrase is processed.

²We are planning experiments using tasks such as auditory sentence completion, speech shadowing, and reading aloud (in which the syntactic structure anticipated by the reader may be revealed by the prosodic contour), to look for confirmation of the on-line expectations and subtle garden paths that our model predicts.

³This is a widely shared, though not universally accepted, view of English processing. For a recent review of the relative merits of this and other models, see Gibson (1991).

2. As entailed by 1, many false hypotheses are adopted by the parser in processing Japanese, but typically their failure is not accessible to consciousness, and their revision is (almost) cost-free.
3. Every known principle of English parsing that could apply to Japanese does apply to Japanese; there is no need for parameterization of the parser.

In the next section we discuss some Japanese examples. In the following section we present our account of the universal parser, including both first-pass processing and reanalysis. Finally, we compare this ranked flagged serial parser with other claims that have been made in the literature about Japanese parsing.

EVIDENCE FOR NON-DELAY SERIAL PARSING

In discussing the examples that follow, we draw a distinction between severe and mild garden paths. Severe garden paths are misanalyses that are consciously detectable and difficult to recover from; mild garden paths are unnoticed during normal language use, non-disruptive, and easily recovered from.⁴ We do not need to claim that there is a sharp dividing line between the two, but only that some garden paths are more troublesome than others. In a later section we give examples of severe garden paths in Japanese. Familiar examples from English include those in (5):

- (5) a. The horse raced past the barn fell.
- b. The cotton clothing is made from grows in Mississippi.
- c. The daughter of the king's son admires himself.
- d. They told the boy that the girl met the story.

In (6) are some English examples that we would claim (in company with Frazier, 1978, and others) exhibit mild garden paths.

- (6) a. They knew the girl at the bakeshop was hungry.
- b. He put the book on the table into his pocket.
- c. We saw the woman with the revolver.
- d. Which book did you read to the children from?

Severe garden paths are unmistakable, but mild garden paths can be difficult to distinguish empirically from the absence of garden paths. In cases like (6) there

⁴The term *garden path* is sometimes reserved for severely disruptive, consciously detectable parsing failures (e.g., Marcus, 1980; Mazuka & Itoh, this volume; Pritchett, 1991), but we are using it to cover all cases in which an incorrect analysis turns out to be incompatible with subsequent input.

can be genuine disagreements about the facts. Yet we must try to settle these facts, because they are just what would distinguish one type of parsing model from another. A garden path is a sign of prior commitment on the part of the parser. So, if the right verdict on disputed cases like (6) is that they exhibit mild garden paths, that would support a committal model; if there are no garden paths in these examples, that implies a parser which commits itself rarely (only in cases like (5)).

Our committal (flagged serial) model is formulated on the basis of observations about Japanese processing which indicate that there are a great many minor, fleeting garden paths comparable to (or even weaker than) those in (6). This contrasts with other models, which entail that where there is no severe garden path there is no garden path at all. For instance, the decisions of a deterministic parser are “indelible”; at most, errors are revised by special patch-up processes more like cognitive problem solving than like “normal” (first-pass) parsing. It is usually held that invoking the patch-up procedures is what causes a severe garden-path experience. Most of the time, in order to avoid such disasters, a deterministic parser does not guess ahead of the input, so it makes no errors. What cannot occur, if parsing is deterministic, is a *mild* garden path, that is, an error from which recovery is easy.

It is widely agreed that though there are some Japanese garden paths quite as severe as the most notorious English examples, these are rare in Japanese (except perhaps where the ambiguity is lexical; see Mazuka & Itoh, this volume). But it cannot be concluded from this that *garden paths* are rare in Japanese. Only a very careful look at the data can distinguish absence of garden paths from presence of easily corrected garden paths. And though serial parsing of ambiguous sentences is incompatible with the former, it is the only model compatible with the latter.

Verbs and Their Arguments

Predicate–argument structure within clauses offers plenty of evidence for on-line anticipation of one analysis of an ambiguous input string. For instance, if we ask native speakers to complete the initial sentence fragment in (7), they will more likely do so as in (8a) than as in (8b).

(7) Bob-ga Mary-ni ringo-o . . .
 -Nom -Dat apple-Acc

(8) a. Bob-ga Mary-ni ringo-o ageta.
 -Nom -Dat apple-Acc gave
 ‘Bob gave Mary the apple.’

b. Bob-ga Mary-ni [_Sringo-o tabeta] inu-o ageta.
 -Nom -Dat apple-Acc ate dog-Acc gave
 ‘Bob gave Mary the dog that ate the apple.’

et al., 1993). So it is interesting that there is no sense of surprise when a sentence continues beyond the verb that ends a clause. In (9), the initial sequence *Bob-ga tegami-o kaita* could be a one-clause sentence.

- (9) Bob-ga tegami-o kaita tomodachi-ga okotta.
 -Nom letter-Acc wrote friend -Nom became angry
 ‘The friend to whom Bob wrote a letter became angry.’

When the noun *tomodachi* is subsequently encountered, it is clear that the initial clause must be a relative modifying that noun, and hence that there must be a clause above the first one. It must take some amount of work to attach *tomodachi* into the structure, but there is no sense of surprise or disruption when *tomodachi* is encountered.⁷ This is important because it indicates that parsing operations are not *inherently* costly. It must be concluded that measurable difficulty occurs only when some specific problem arises.

We present one more example of on-line anticipation of structure within the clause. It appears that the parser tends to assume that an incoming noun phrase will be an argument of the verb, rather than an adjunct (modifier) of a following noun.⁸ This is revealed in expectations about what case marking a noun will have.⁹ For example, if we ask for completion of a sentence-initial fragment such as (10), native speakers will more often add a noun with an argument case marking, as in (11a), than a noun with the genitive marker *-no* that is characteristic of pronominal modifiers, as in (11b). (Note that we have included a determiner and adjective to stretch out the time between the beginning of the NP and the arrival of the case marker. Expectations strong enough to create surprise at their failure seem to take a little while to build up; but this is not peculiar to Japanese.)

- (10) John-ga kono totemo urusai . . .
 this very noisy

- (11) a. John-ga kono totemo urusai inu-ni hone-o ageta.
 this very noisy dog-Dat bone-Acc gave
 ‘John gave this very noisy dog a bone.’

⁷Judgments in relative clause examples depend in part on the case marking of the head noun and of the prior NPs, and on what thematic role the head noun is bound to within the relative, which is fully ambiguous in Japanese (e.g., our translation of (9) with the head noun associated with the **goal** role in the relative clause is natural but not necessary). See Inoue (1991, chap. 3) for more detailed discussion. We return later to why it is not significantly more disruptive for the head noun to displace a phrase that had previously been included in the relative clause.

⁸See Frazier (1987) for discussion of an argument versus adjunct preference for PPs in Dutch, though the sentence structures are not the same as here.

⁹We are not drawing a principled distinction in this chapter between case inflections and postpositions in Japanese. It has been argued by Saito (1985) that *-ga*, *-o*, and *-no* are case inflections whereas *-ni*, and *-de*, among others, are postpositions. But for expository convenience we refer to all such postnominal elements as case markers.

- b. John-ga kono totemo urusai inu-no kainushi-ni keikokushita.
 this very noisy dog-Gen owner-Dat warned
 ‘John warned the owner of this very noisy dog.’
 (or: ‘John warned this very noisy owner . . .’ etc.)

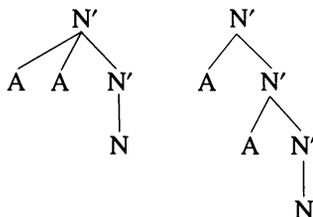
This preference for argument over adjunct¹⁰ is confirmed by judgments on complete sentences. Speakers will sometimes report a fleeting sense of surprise at the point of the *-no* morpheme in (11b), but nobody seems surprised to encounter *-ni* in (11a).

As noted previously, we are not primarily concerned in this chapter with identifying the specific principles underlying the parser’s preferences. But it is worth noting that, as for the earlier examples, the preference in (10)/(11) is consistent with structural choice principles familiar from English, such as Minimal Attachment and/or Local Association. The source of the preference cannot be that NP-*no* is much less frequent than argument-marked NP. NP-*no* is extremely frequent in Japanese, even a sequence of two adjacent NP-*no* is quite common (and three or more in a row can occur). The preference might be due to thematic expectations (see Gibson, 1991; Pritchett, 1987, 1988), but once again it should be noted that the verb (or other predicate) that assigns thematic roles has not been received by the parser at the point of interest in these examples.

Nouns and Their Modifiers

Within the noun phrase, also, it seems that structural decisions are made on-line before the input string is disambiguated. Consider first the unambiguous example (12). (It is unimportant where this noun phrase occurs in a sentence; for convenience, it can be assumed to be sentence-initial. Also, for simplicity, the distinction between N’ and NP nodes will be suppressed in the tree diagrams that follow.)

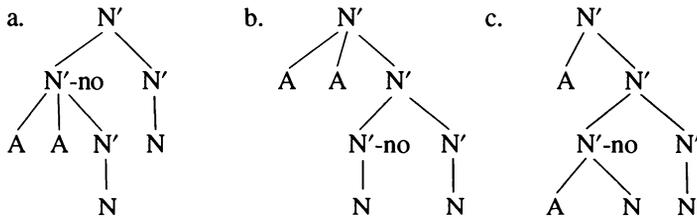
- (12) toshitotta takai shuurinin
 old (animate) expensive repairer
 old expensive repairer (of things)



¹⁰One might worry that these examples show only that people—or their processing mechanisms—are lazy and give the shortest sentence completions they can get away with. However, not all argument continuations are shorter than all *-no* continuations. Furthermore, the results make it clear that brevity is not the driving force in the completion task; people often add more words than is necessary. For instance, the NP-*o* in (11a) is not needed; a shorter response would have been *John-ga kono totemo urusai inu-o butta* (*butta* = hit).

This phrase refers to a repairer (of things unspecified) who is old and charges a lot for his or her services. Either of the structures shown is possible, depending on whether multiple modifiers of the same noun must be nested or can be structural sisters. This distinction is not important to our point (though for some other adjective combinations a semantic scope contrast can hang on it), so for convenience we will adopt the flat structure here. Now consider example (13), which is ambiguous. It has two nouns, either of which might, at least in principle, be modified by two, one, or none of the adjectives.

- (13) *toshitotta takai kutsu-no shuurinin*
 old (animate) expensive shoe-Gen repairer
 a. 'repairer of *old expensive shoes'
 b. 'old expensive repairer of shoes'
 c. 'old repairer of expensive shoes'



There are three potential structures, as shown (as well as nested-adjective versions of (a) and (b) which, as noted, we will not consider here). Both adjectives modify *kutsu* on analysis (a); both modify *shuurinin* on analysis (b); the first modifies *shuurinin* and the second modifies *kutsu* on analysis (c).¹¹ In fact, analysis (a) is unacceptable for (13), because the adjective *toshitotta* is selectionally restricted to human (or at least animate) entities; it could be predicated of a shoe only if

¹¹English has head-final adjective-noun order like Japanese so it has the potential for similar ambiguities, but it has much more limited head-final nominal modification structures. In the (slightly edited) English translation of (13), *an expensive old shoe repairer*, the sequence *shoe repairer* is a noun-noun compound. This entails that *shoe* cannot be modified by adjectives. The initially favored (a) analysis [*an expensive old shoe*] therefore has to give way at *repairer* to the right-branching (b) analysis. (Or, just possibly, in this particular case the last three words could be read as a triple compound, so that the structure is like (c), but with different node labels.) English does have productive head-final nominal modification by possessives (cf. (14)), as in *foreign butchers' knives* which can have either structure. The initial left-branching structure (the butchers are foreign) may flip at *knives* to the right-branching structure (the knives are foreign). But this shift occurs only when the two nouns form a natural compound; otherwise, the preference for the left-branching structure persists, as it does in Japanese (e.g., *foreign butchers' eyeglasses* where the butchers, not the eyeglasses, are foreign). It appears that a lexical compound analysis tends to be favored where it exists, perhaps because it minimizes the structural complexity for which the syntactic processor is responsible.

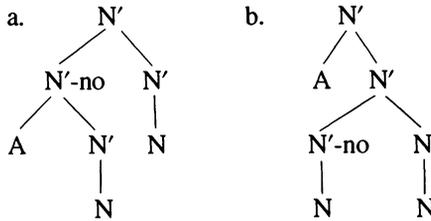
the shoe were presented as animate (e.g., in a children's story). The adjective *takai*, on the other hand, can apply either to inanimate objects, such as shoes, or to people who sell their services, such as shoe-repairers; so structures (b) and (c), in which *toshitotta* modifies *shuurinin*, both give acceptable meanings. Post-sentence judgments overwhelmingly associate *toshitotta* with the later noun *shuurinin*, as in (b) or (c), rather than with the first noun *kutsu*. But on-line, there is a rather strong preference for local association of the two adjectives with the adjacent noun, as at the beginning of structure (13a) (and as in (12)).¹² This preference is discernible in the fact that perceivers experience some surprise on encountering the inanimate noun *kutsu*; what is expected is an animate noun compatible with both adjectives. This suggests that before the head noun is encountered, the parser anticipates a tree in which the two adjectives are structured together with each other and with the adjacent noun.¹³

We can explore this intuition further by systematically changing or adding words in the example. For instance, if we replace *toshitotta* with the word *furubita*, which also means "old" but is restricted to inanimates, then there is no longer any surprise at *kutsu*. The sentence completion task leads to a similar conclusion. If we give native speakers the first two words of (13) and ask for a completion, the noun supplied will usually be animate (mostly human); if we omit the animate-restricted adjective the proportion of animate-noun completions declines; if we substitute *furubita*, the completions are almost all inanimate nouns. Again, this implies that the processor opts for local attachment of one modifier to a preceding one, and for local attachment of a noun to a preceding modifier, as in (a) rather than (b) or (c). In the case of example (13) this preferred structure is anomalous and must be abandoned, and a mild sense of garden path is experienced. Other examples are fully ambiguous, and then the attachment decision shows up simply as a preference for one reading over the other. In (14), for instance, though readings (a) and (b) are both fully acceptable, there is a preference for reading (a).

¹²Punctuation, or prosody in spoken sentences, can influence this preference. If (13) is pronounced with a pause and/or intonation rise at the end of the first adjective, the (a) analysis is less common. (The fact that intonational cues may be used to resolve syntactic ambiguities does not affect our account of how ambiguous input is processed in the absence of disambiguating information. Prosodic cues need not be present; they are typically optional. Also, they are often not fully disambiguating; see footnote 1.)

¹³This is an interesting example from the point of view of head-driven parsing because the ambiguity is about which word is the head. A natural strategy for a head-driven parser (HDP) would be to opt for the earliest possible head, since that would minimize attachment delays. However, it would not account for the observation that the *first* noun in (13) violates an expectation, which indicates that a structural hypothesis was formed *before either* noun occurred. More to the point would be a strategy favoring maximal incorporation of as many previously encountered elements as possible when structure is built on encountering a head (see Pereira, 1985; Shieber, 1983), though this isn't entailed by HDP per se. It would result in a preference for structure (13a) at *kutsu*, and so might explain why there is surprise on encountering *kutsu*—surprise that the first adjective is selectionally odd if encompassed in the phrase that *kutsu* heads. However, in (25) and (26) we show that the surprise effect can arise before any potential head has been encountered at all.

- (14) *shinsetsuna gakusei-no imotoo*
 kind student-Gen (younger) sister
 a. [kind student]'s sister = sister of kind student
 b. kind [student's sister] = kind sister of student



Examples such as this are quite common and natural in Japanese conversation and writing.¹⁴

It can be shown that this preference for local association of modifiers and noun exists even after the parser has encountered the genitive case-marker *-no* on the first noun, which flags it as a prenominal modifier.¹⁵ This informs the processor that another noun—which the adjective would be able to modify—will appear later. So the preference for modification of the immediately following noun is not a simple consequence of a bird-in-the-hand principle: If you need a noun and you find a noun, look no further. Unlike the unmarked first noun in English *shoe repairer* or *student loan*, the *-no* marking in Japanese is a guarantee that there will be a bird in the bush as well.¹⁶ The fact that the parser is indifferent to this information suggests that it is following some processing principle such

¹⁴We acknowledge that there are examples not unlike (14), in which the *non*-local association of adjective and first noun may seem to be preferred once processing is completed, even though local association was preferred along the way. This is potentially an important observation, because a change of preference would violate the principle of Revision As Last Resort (RALR; see Fodor & Frazier, 1980). However, there are reasons to think that the later-preferred reading in those cases is *not* an instance of non-local association as in (b), but represents a third structural analysis in which the two nouns form a noun-noun compound, as discussed for English in footnote 11. An example is *shinsetsuna jukensei-no haha* (preferred reading: kind mother of student), where *jukensei-no haha* can be a compound denoting the mother of a student who is preparing for exams. In this case the shift to a right-branching analysis would not be an RALR violation within the syntax, but would be a response by the syntax to a new proposal from the lexicon.

¹⁵The modifier reading is not the only one possible. An NP-*no* could be the subject of a relative clause, as in *furubita takai kutsu-no kazarareteiru mise* (the shop in which the old expensive shoes are exhibited), or the subject of a *koto* (fact) clause, as in *furubita takai kutsu-no nusumareta koto* (the fact that the old expensive shoes were stolen). Both these constructions with *-no* are quite natural, but they are less frequent than NP-*no* as a prenominal modifier, and in any case they are both guaranteed to supply another noun which the adjective could modify (even if less probably in the case of *koto*).

¹⁶The *-no* marking on a prenominal modifier is optional in Japanese, but only where the modifier and the noun constitute a lexical compound (as determined by the modification test; see footnote 11).

as Minimal Attachment or Local Association, rather than just generally trying to guess how the sentence might continue. The parser's indifference to the subordinate status of the first noun cannot easily be illustrated with example (13), since processing of *kutsu-no*, which is pragmatically incompatible with *toshitotta*, would initiate revision of the local attachment analysis.¹⁷ But the point can be demonstrated with a lexically ambiguous animate/inanimate noun such as *kutsuya* in place of *kutsu*. The word *kutsuya* can mean either "shoe-shop" or (approximately) "shoe-shop-keeper." In isolation, the preferred interpretation for most speakers seems to be the inanimate reading. But this changes to a preference for the animate reading if we put *kutsuya* after an adjective that can modify only animates—even if *kutsuya* is marked with *-no*, signaling that another noun will follow. For instance, if we ask for sentence completions of the fragment in (15), we typically find completions indicating that the adjective is associated with the immediately following noun (where it is possible to tell). If speakers are offered a choice of completions, they will generally prefer (16) over (17).

(15) toshitotta kutsuya-no . . .
 shoe-shop(-keeper)-Gen

(16) toshitotta kutsuya-no horosukoopu
 horoscope

(17) toshitotta kutsuya-no keieisha
 proprietor

In principle, the animate-restricted *toshitotta* could modify either *kutsuya* construed as animate, or the second noun, which *kutsuya-no* modifies. The fact that the noun *horosukoopu* is judged more natural in this context than the noun *keieisha* reveals the parser's tendency to locally associate the adjective with the adjacent noun. In the preferred (16), *horosukoopu* forces *kutsuya* to be construed as animate, for semantic/pragmatic reasons: a shoe-shop-keeper, but not a shoe-shop, can have a horoscope. So in (16), *toshitotta* can and must modify *kutsuya-no*, the first noun. In the non-preferred (17), *keieisha* forces *kutsuya* to be construed as inanimate, since only a shoe-shop, not a shoe-shop-keeper, can have a proprietor. But that then forces *toshitotta*, which needs an animate noun, to modify the second noun *keieisha*.¹⁸ Thus, the preference for (16) reveals a preference for local association of the adjective and noun in (15).

¹⁷We assume that use of pragmatic information to reject a syntactic parse does not violate the modularity thesis (J. A. Fodor, 1983), but the issue of modularity in Japanese parsing deserves more attention. It is relevant to the bird-in-the-bush argument, because the *-no* marking assures the parser only that there will be a noun that the adjective can syntactically modify; it cannot guarantee that this noun will fit with the adjective's selectional requirements and meaning. So we concede that if parsing is non-modular, there might be some security to be gained from the adjective's grabbing the first candidate noun that is compatible with it.

¹⁸To check that people do not just prefer the word *horosukoopu* as a word, we can replace the animate-restricted *toshitotta* with the inanimate-restricted *furubita*:

We have argued on the basis of unambiguous examples like (13), and ambiguous examples like (14) and (15), that there is a preference for local association of modifiers and nouns, and we have attributed this to on-line decisions that the parser makes without waiting for disambiguating input. This is *prima facie* evidence for serial processing of Japanese. But some possible objections to this claim should now be considered. In particular: Could the source of the preference be something other than the processor? It could not be pragmatic plausibility, because the preference persists when plausibility is equated. The grammar is another possible source, but is not likely. The reason local attachment is preferred is not that non-local attachment is ungrammatical. In all of the preceding examples, the non-local analysis is syntactically well formed. For instance, analyses (b) and (c) of (13) are perfectly grammatical noun phrases (although the English counterpart of (c) may not be; see footnote 11). What about stylistic principles? Perhaps these constructions are stylistically impeccable on the (a) analysis but are awkward as expressions of the (b) meaning. Particularly, one might wonder whether, for the (b) meaning, the modifiers would typically occur in the opposite order. (We thank Richard Kayne for this point.) For example, (18) might be a better way to express “kind sister of a student” than (14), because some stylistic principle encourages genitives to precede adjectives when they modify the same noun.

- (18) *gakusei-no shinsetsuna imooto*
 student-Gen kind (younger) sister
 ‘student’s [kind sister]’ = ‘kind sister of student’

However, although (18) *can* express this meaning, it appears not to be a privileged form for expressing it; both forms are equally natural. (They differ subtly in semantic scope; see previous discussion.) Word order preferences do exist, but they appear to be due to a preference for long constituents before short constituents in Japanese.¹⁹ In (19) the genitive modifier of the previous examples has been replaced by a relative clause, a longer (heavier) constituent.

-
- (i) *furubita kutsuya-no horosukoopu/keieisha*
 ‘old [INANIM] shoe-shop(-keeper)’s horoscope/proprietor’

Then, as expected, the judgment flips, and speakers select *keieisha* as continuation rather than *horosukoopu* (although the latter is perfectly acceptable if *horosukoopu* denotes a piece of paper with a horoscope written on it).

¹⁹This contrasts with the English preference for short before long, but it does not indicate that the parser is parameterized (see discussion in the final section of this chapter). It would follow from Local Association, just as the mirror-image preference in English does. In a left-branching structure, the longer a constituent is, the greater the structural distance is likely to be between its first word and the last word of the preceding constituent. The best way to satisfy Local Association would be for long constituents not to follow any other constituent. This is indeed common in Japanese: Complement clauses in Japanese tend to prepose, whereas in English (and even German, which is verb-final but Comp-initial) they tend to extrapose. See also Hawkins (1990) (cf. footnote 38) for a somewhat different view of interactions between length and branching direction in parsing.

- (19) *kawaii shoojo-ga daiteita ningyoo*
 pretty girl-Nom was holding doll
 a. ‘doll that the pretty girl was holding’
 b. ‘pretty doll that the girl was holding’
- (20) *shoojo-ga daiteita kawaii ningyoo*
 a. *
 b. ‘pretty doll that the girl was holding’

In (19) as elsewhere, the (a) reading is preferred over the (b) reading. And in this case it *is* true that the reverse order of modifiers, shown in (20), is preferred where it is possible. It is possible for the (b) reading, though not for the (a) reading (because the relative clause must end at the verb, its head, and so cannot include a following adjective). But we have seen that there is no such asymmetry of preference when both modifiers are short—or indeed when both are long. Moreover, if we make the first modifier in an ambiguous sequence longer than the second, the long-before-short principle results in a slight preference in favor of the (a) word order even for the (b) meaning. Consider (21) and (22), where an adverb adds weight to the adjective.

- (21) *kyokutanni shinsetsuna gakusei-no imooto*
 extremely
 a. ‘[extremely kind student]’s sister’
 b. ‘[extremely kind] [student’s sister]’
- (22) *gakusei-no kyokutanni shinsetsuna imooto*
 a. *
 b. ‘student’s [extremely kind] sister’

For (21) the (a) reading is dominant as usual;²⁰ but in this case the preferred word order for expressing (b) is (21), not (22). The preference for local association remains constant through the waxing and waning of these other word order preferences that vary with constituent length, and it seems not to have any explanation in terms of stylistic principles.

Let us consider one last factor that might be creating a misleading appearance of local association. It could be suggested that the preference for the (a) analyses of these constructions is due to cooperative principles of conversation as proposed by Grice (1989). The (b) reading can be expressed unambiguously (e.g., (18) unambiguously expresses the (b) reading of (14)); but there is no natural way to

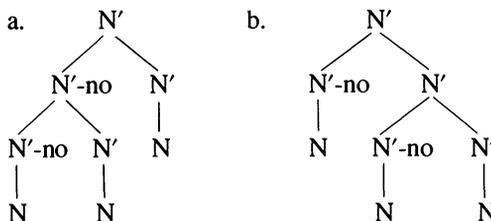
²⁰This is the majority view, but a few speakers judge that reading (b) is more natural than reading (a) for (21). We do not know why this is so; it does not follow from anything in our model at present. It is as if the phrase prefers to split into two equal parts, perhaps for prosodic reasons.

disambiguate in favor of the (a) meaning.²¹ By Gricean principles a hearer might reason that, if a speaker uses the ambiguous form, without taking advantage of the unambiguous form for the (b) meaning, that must be because (b) was not the meaning he or she intended to express; hence, the intended meaning must have been (a). In other words: If you want to express “kind sister of a student” in the way most helpful to your hearers, you should utter (18) rather than (14) because (18) is unambiguous; hence you should utter (14) only if you intend “sister of a kind student,” for which there is no more cooperative means of expression; hence (if I can trust in your cooperativeness), I can conclude that in uttering (14) you meant “sister of a kind student.”

We think there is some truth to this Gricean outlook, particularly for a language like Japanese where fully ambiguous sentences abound and could create serious misunderstandings. But even if conversational principles are at work, this does not undermine the basic claim that parsing is serial for Japanese, for several reasons. First, it does not contradict the claim that the processor selects an analysis on-line, but merely changes the basis of the selection from structural preferences to conversational principles. Second, the Gricean account exhibits the usual cause/effect uncertainty: It could be *because* the processor favors the (a) analysis of the ambiguous form, that it is helpful to use the unambiguous form to express (b). Third, and most significantly, there are examples in which the (b) reading does *not* have any unambiguous expression, and yet the ambiguous form shows the usual preference for the (a) reading; in such a case there can be no Gricean explanation for the preference. This is the case in (23), where we have replaced the adjective with a possessive NP.

(23) Mary-no akanboo-no shashin
 -Gen baby-Gen photograph

- a. '[Mary's baby]'s photograph' = 'a photograph of Mary's baby'
 b. 'Mary's [baby's photograph]' = 'Mary's photograph of a baby'



²¹If it were for some reason important (e.g., in a legal document) to disambiguate the (a) meaning without benefit of intonation, one might say *shinsetsuna gakusei, sono hito-no imooto* (where *sono hito* = 'that person'). This appears to be an appositive construction, with a coreference relation between *gakusei* and *sono hito* (i.e., a kind student, that person's sister). This would not be a natural sentence in a normal conversational context.

The structures for (23) are exactly like those for (14) except for the category of the first modifier.²² The other lexical changes are then needed for naturalness. (Because ‘Mary’s sister of a student’ makes little sense we have turned it into ‘Mary’s photograph of a baby’ by replacing *imooto* with *shashin*, and *gakusei* with *akanboo*.) Just like (14), example (23) prefers local association of the first modifier (in this case an NP-*no*) with the noun that immediately follows it, as shown in (a). But in this case the (b) reading cannot be brought to the fore by reversing the modifiers, as in (24).

(24) *akanboo-no Mary-no shashin*

Example (24) happens to be well-formed, but its only meaning is quite different from both (a) and (b); it means ‘a baby’s photograph of Mary.’ Evidently, possessive modifiers are not freely ordered relative to other elements in a noun phrase: The first of two NP-*no* can be a possessive modifier and the second an argument of the noun, but not vice versa. Whatever the reason for this restriction, it has the consequence that there is no easy way to disambiguate the (b) reading of (23).²³ Hence, the proposed Gricean explanation for the preference for the (a) reading fails in this case.

Finally, because it is so important to theoretical concerns to determine whether heads of constituents play a special role in the scheduling of decision making, we end this section with one very clear demonstration of decision making prior to the head. The examples have a sequence of two adjectives, as before, but now the adjectives are mutually incompatible, so for pragmatic reasons they must modify *different* nouns. Consider (25) and (26).

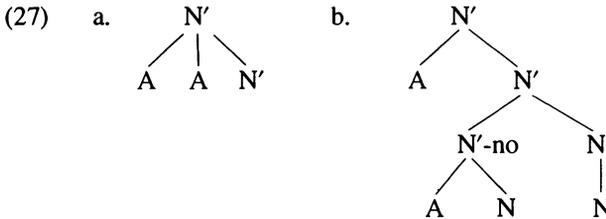
(25) *yabuketa hanashibetana resuraa-no buutsu*
 torn [INANIM] inarticulate wrestler-Gen boots
 ‘[torn [inarticulate wrestler’s] boots]’ = ‘torn boots of inarticulate wrestler’

²²This is not quite true. An NP-*no* associated with *shashin* could be its argument rather than a modifier; and in these examples it probably is, even though we diagram it, to parallel (14), as a modifier. This distinction is not important here. Incidentally, there appears to be no great preference for a compound analysis of the second and third words in (23). English has the rather prominent lexical compound *baby pictures*, which denotes pictures of someone *as* a baby, but in Japanese a comparable meaning is expressed by different means. Something like this interpretation is possible for (24), on the analysis “a photograph of baby-Mary”; but this interpretation is not possible for (23).

²³To express the (b) meaning unambiguously one could use the relative clause construction *Mary-no motteiru* (‘own/have’) *akanboo-no shashin* (‘photograph of a baby which Mary owns’). Since it is not natural to talk of owning a baby, it is the photograph which is owned. We assume that the existence of this rather cumbersome paraphrase does not save the suggested Gricean explanation. Grice’s principles do not require speakers to be unambiguous at any cost; it is only where complexity is roughly comparable that relative informativeness would come into play.

- (26) kimuzukashii semai apaato-no ooya
 hard-to-please narrow [INANIM] apartment-Gen landlord
 '[hard-to-please [narrow apartment's] landlord]' = 'hard-to-please land-
 lord of narrow apartment'

In (25), *yabuketa* is restricted to inanimate nouns, while *hanashibetana* is pragmatically limited to humans. Therefore, these two adjectives could not sensibly appear in the structure (27a), but must appear separated in a structure such as (27b).



This is clear enough by the end of processing. The (27b) structure is well-formed and perfectly natural for these examples once they have been processed, considerably more natural than their English counterparts (perhaps because Japanese lacks the right-branching alternatives that English allows). However, there is a clear preference to locally attach the two adjectives together as in (27a). This preference is revealed by a sense of puzzlement when people hear or read the second adjective. (An intonation break between the adjectives may reduce this effect but doesn't fully eliminate it; see footnote 12.) The parser apparently makes the structural assumption that both adjectives modify the same noun, even though it recognizes that semantically they couldn't make sense that way. Note that these examples rule out any account which claims that all preference effects are due to strategies (Minimal Attachment, etc.) applying to retroactive tree-building processes that take place when the head is processed. There should be no puzzlement at the point of processing the second adjective, if no decision were made until the noun.²⁴

To summarize this section, we have observed clear tendencies on-line: (a) to assume that arguments are co-arguments of the same predicate, (b) to anticipate arguments rather than adjuncts, and (c) to attach a modifier or a noun locally to a preceding modifier. We have argued that the source of these preferences is the processing routines rather than any fact about what is grammatical, or what is

²⁴Here, as before, we must ask whether any other natural explanation exists. All we can think of here is that there might be some surprise just that these two adjectives are in the same sentence. This possibility can be controlled for by presenting speakers with *yabuketa buutsu-o hanashibetana . . .*. Now *yabuketa* and *hanashibetana* are in the same sentence, but they do not modify the same noun. And now the surprise at *hanashibetana* disappears.

stylistically preferred or conversationally most cooperative. We conclude that the Japanese parser does not postpone its decisions until it finds disambiguating material, or until it finds the head of a constituent, and it does not parallel process all (or some) alternative analyses without choosing between them. Rather, a sentence is assigned a structure and an interpretation as it is received. In the next section, we take on-line decision to have been established, and consider its implications for the design of the parser.

INFORMATION-PACED PARSING

First Pass

The task is to explain how the human sentence processing mechanism can successfully parse the massively ambiguous sentences of a head-final language such as Japanese. One approach would be to claim that the processor makes few errors, because it makes few decisions about ambiguous structures, and makes them only where there is sufficient information to offer a fairly good chance of getting them right. This might reasonably be called “information-paced” parsing, and it is how we began thinking of the problem (see Inoue, 1985, 1991). Differences in the timing of decision making between languages would be attributable to language-specific differences in the ebb and flow of information over time as the input is received and the relevant aspects of the grammar are accessed. In Japanese, information is scanty, especially early in a parse, so it is predicted that in Japanese there are considerable delays of decision making. In English, information is quite plentiful, even early on, so it is predicted that many English parsing decisions are made early. Thus, an information-paced model of parsing has the merit that it could be assumed to be universal and innate, and yet simultaneously accommodate both a standard view of English parsing and the special challenge of Japanese parsing.

However, we no longer advocate a “pure” information-paced model, for two reasons. First, the evidence presented in the previous section argues for a much less cautious Japanese processor, which makes decisions at many (perhaps all) points on-line, often long before relevant information is available. Second, perfectly information-paced parsing would be wrong for English as well, because it predicts the total absence of garden paths. The only mistake a pure information-driven parser could make would be to guess wrong, at the end of a fully ambiguous sentence, about what meaning the speaker intended. Thus, for *both* languages, some early decision making must be built into the model, where we use “early” as a technical term meaning “before the decision is uniquely determined by grammar plus input.”

If the model must allow for some early decisions, the strongest move is to assume that *all* decisions are made as they arise (i.e., as the word they concern is first encountered). This would be a fully committal system, abiding by a principle of Full Attachment, which requires each incoming word to be attached, as it is encountered, into the *whole* current partial phrase marker (CPPM) for the sentence. Full Attachment amounts to permitting only the highest node of the whole CPPM