

# I-Language

An Introduction to Linguistics as Cognitive Science

**OXFORD CORE LINGUISTICS  
GENERAL EDITOR**

David Adger, University of London

**PUBLISHED**

*Core Syntax: A Minimalist Approach*

David Adger

*I-Language: An Introduction to Linguistics as Cognitive Science*

Daniela Isac and Charles Reiss

**IN PREPARATION**

*Core Semantics*

Gillian Ramchand

*Introduction to Theoretical Linguistics*

Peter Svenonius

# I-Language

## An Introduction to Linguistics as Cognitive Science

Daniela Isac  
and Charles Reiss

**OXFORD**  
UNIVERSITY PRESS

**OXFORD**

UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

Oxford University Press is a department of the University of Oxford.  
It furthers the University's objective of excellence in research, scholarship,  
and education by publishing worldwide in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi  
Kuala Lumpur Madrid Melbourne Mexico City Nairobi  
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece  
Guatemala Hungary Italy Japan Poland Portugal Singapore  
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trademark of Oxford University Press  
in the UK and in certain other countries

Published in the United States  
by Oxford University Press Inc., New York

© Daniela Isac and Charles Reiss 2008

The moral rights of the author have been asserted  
Database right Oxford University Press (maker)

First published 2008

All rights reserved. No part of this publication may be reproduced,  
stored in a retrieval system, or transmitted, in any form or by any means,  
without the prior permission in writing of Oxford University Press,  
or as expressly permitted by law, or under terms agreed with the appropriate  
reprographics rights organization. Enquiries concerning reproduction  
outside the scope of the above should be sent to the Rights Department,  
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover  
and you must impose the same condition on any acquirer

British Library Cataloguing in Publication Data  
Data available

Library of Congress Cataloging in Publication Data  
Data available

Typeset by SPI Publisher Services, Pondicherry, India  
Printed in Spain by Cayfosa Quebecor

ISBN 978-0-19-953419-7 (Hbk.)  
ISBN 978-0-19-953420-3 (Pbk.)

1 3 5 7 9 10 8 6 4 2

# Contents

<i>Preface</i>	ix
<i>Acknowledgements</i>	xi
<i>List of Figures</i>	xiii

## **PART I The Object of Inquiry**

<b>1 What is I-language?</b>	<b>3</b>
1.1 Jumping in	5
1.2 Equivalence classes	8
1.3 Partial reduplication in Samoan	9
1.4 Mentalism	12
1.5 I-language	13
1.6 Some implications of mentalism	14
1.7 Summing up	15
1.8 Exercises	17
<b>2 I-everything: Triangles, streams, words</b>	<b>20</b>
2.1 A triangle built by the mind	20
2.2 More visual construction	25
2.3 Auditory scene analysis	27
2.4 Words are built by the mind	30
2.5 Summing up	33
2.6 Exercises	33
<b>3 Approaches to the study of language</b>	<b>36</b>
3.1 Commonsense views of “language”	37
3.2 I-language	39
3.3 The kind of stuff we look at	43
3.4 Methodological dualism	48
3.5 Bilingualistics	51
3.6 And so?	52
3.7 Exercises	53

## vi CONTENTS

<b>4 I-/E-/P-Language</b>	<b>55</b>
4.1 Computation in phonology	55
4.2 Extensional equivalence	61
4.3 Non-internalist approaches	67
4.4 How is communication possible?	72
4.5 Exercises	75

**PART II Linguistic Representation and Computation**

<b>5 A syntactic theory that won't work</b>	<b>79</b>
5.1 General requirements on grammars	79
5.2 Finite state languages	87
5.3 Discussion	94
5.4 Power of grammars	97
5.5 Exercises	101
<b>6 Abstract representations</b>	<b>103</b>
6.1 Abstractness	103
6.2 Abstractness of sentence structure	104
6.3 Allophony	109
6.4 Turkish vowel harmony	114
6.5 Words are not derived from words	125
6.6 Think negative	127
6.7 Summing up	131
6.8 Exercises	132
<b>7 Some details of sentence structure</b>	<b>136</b>
7.1 Basic syntactic categories	136
7.2 Syntactic constituents	139
7.3 Labels and phrasal categories	144
7.4 Predicting syntactic patterns	154
7.5 Using trees to predict reaction times	164
7.6 To sum up	167
7.7 Exercises	168
<b>8 Binding</b>	<b>170</b>
8.1 Preliminaries	171
8.2 Anaphors	173

## CONTENTS vii

8.3	Pronouns and “referential expressions”	181
8.4	Some implications	181
8.5	Binding and <i>wh</i> -movement	182
8.6	Non-structural factors in interpretation	187
8.7	Exercises	189
<b>9</b>	<b>Ergativity</b>	<b>192</b>
9.1	Preliminaries	194
9.2	A nominative-accusative system	197
9.3	An ergative-absolutive system	198
9.4	A tense-split system	201
9.5	A nominal-verbal mismatch	202
9.6	A NP-split system	203
9.7	Language, thought and culture	206
9.8	Exercises	207
<b>PART III Universal Grammar</b>		
<b>10</b>	<b>Approaches to UG: Empirical evidence</b>	<b>215</b>
10.1	On the plausibility of innate knowledge	216
10.2	More negative thoughts	221
10.3	Exercises	233
<b>11</b>	<b>Approaches to UG: Logic</b>	<b>235</b>
11.1	Let’s play cards	238
11.2	Where does this leave us?	245
11.3	Building blocks in other domains	247
11.4	Exercises	248
<b>PART IV Implications and Conclusions</b>		
<b>12</b>	<b>Social implications</b>	<b>253</b>
12.1	Prescriptive vs. descriptive grammar	253
12.2	Negation	254
12.3	Change is constant	256
12.4	Exercises	262

## viii CONTENTS

---

<b>13</b>	<b>Some philosophy</b>	<b>265</b>
13.1	Rationalism and empiricism	265
13.2	Competence and performance	271
13.3	Reference	277
13.4	Essentialism	282
13.5	Mind and body	286
13.6	A view from neuroscience	291
13.7	Exercises	298
<b>14</b>	<b>Open questions and closing remarks</b>	<b>300</b>
14.1	You and your grammar	300
14.2	Retracing the links among key <i>-isms</i>	304
14.3	Bearing on philosophical questions	306
	<i>References</i>	311
	<i>Index</i>	315

## Preface

Our original goal was to write a popular book that would lead the reader through some fairly technical analyses of linguistic data. Instead of just reporting on the claims and results of modern linguistics, we wanted to show the reader how to think like a linguist. In the end, we realized that a textbook format was more suitable, given the depth and breadth we are aiming at. We foresee the book serving as an introduction to linguistics for students planning to continue in the field, as well as for those with interests in other branches of cognitive science. Throughout the book, linguistic issues are related to topics in vision, philosophy, ethology, and so on. We hope that we can inspire our readers to continue the search for unifying themes among these fields.

All the material in this book has been presented to undergraduate students in large classes (often over one hundred students). Much of it has also been presented to middle school students, prison inmates, and non-academic audiences. In developing and teaching the materials, we have had the advantage of being a team of a syntactician and a phonologist, but we hope that any enthusiastic teacher will be able to understand the material and help motivated students work through it. We think that the ideas are important, but, in fact, not very difficult when broken down. Additional exercises and links to material related to the text can be found on the book's companion website.

As an introduction to linguistics the book is very narrow. There are no chapters on sociolinguistics or historical linguistics, for example. And yet, we strongly believe that the best results in these fields can be attained by incorporating the approach to the study of language we develop, basically the framework of generative grammar developed by Noam Chomsky and collaborators since the 1950s. In some sense the book is an exegesis of the Chomskyan program or, rather, our understanding of the program.

In the course of writing we often found each other to be frustratingly thick-headed about various topics. The resulting heated arguments have helped us to achieve deeper understanding and greater intellectual humility

X PREFACE

---

and, we hope, to produce a better book. We expect that even sympathetic readers will find much to object to in our presentation, but we think that we have succeeded in laying out a coherent position, sometimes by openly attacking other positions, that can at least serve as the basis for fruitful debate. If any of our claims or arguments manage to get students' "blood pressure up to an appropriately high level" (to quote Morris Halle) where they seek to challenge our point of view, we will consider this to be a successful textbook.

## Acknowledgements

There is probably nothing original in this book, and thus we are beholden to the community of linguists and other scholars from whom we have liberally borrowed. In some cases we have acknowledged specific debts, but in general we have not. Most obviously, the book is inspired by and draws heavily from the work of Noam Chomsky. We excuse our general failure at careful attribution by adopting Chomsky's own attitude that full attribution is not only impossible but also fairly unimportant. Our common goal as a community is to understand the object of study—the language faculty and the human mind, in general.

That being said, we will point out that several authors have been most inspiring in helping us to achieve our understanding of the place of linguistics in cognitive science. We mention these to encourage you to consult them on your own. We include their work in the readings for the course that this book grew out of, *Language and Mind: The Chomskyan Program* at Concordia University. Specific works by these cognitive scientists are listed in the reading suggestions at the end of each chapter: Albert Bregman, Morris Halle, Donald Hoffman, Ray Jackendoff, Zenon Pylyshyn. The course that this book grew out of was originally built around Jackendoff's *Patterns in the Mind*, and that book was so useful in developing our own understanding of the place of linguistics in cognitive science that it was actually a challenge to us as authors to move away from its excellent structure around three fundamental arguments.

Many of the articles in the four-volume *Invitation to Cognitive Science*, edited by Daniel Osherson, have also been instrumental as teaching resources, and we recommend them to students interested in making connections among the various branches of cognitive science.

We are also most grateful to our reviewers, Sam Epstein, Virginia Hill, and Ur Schlonsky, and the Core Linguistics Series Editor, David Adger, for useful feedback and for pushing us to not take the easy way out when dealing with technical linguistic topics. Their input has vastly improved the manuscript. It has been a pleasure to work with John Davey, our

## xii ACKNOWLEDGEMENTS

Consultant Editor at Oxford, whose encouragement and flexibility are greatly appreciated.

Two non-linguists read early drafts of the book. Harold Golubock provided encouraging feedback—we originally wanted to write a popular book for the intelligent lay reader and Harold was the perfect guinea pig. Lesly Reiss managed to proofread the entire first draft, and she was fascinated . . . fascinated that anyone would find this material interesting, a sentiment she repeatedly shared with us. We are grateful to Chris Eldridge for particularly helpful insight into the mind–body problem.

The hundreds of Concordia undergraduate students who took our course and helped us develop the materials that have become this book must be acknowledged. Of all our students, Hisako Noguchi deserves special mention. She not only took the class but she has served as a teaching assistant too many times to count. Her input has been crucial to the success of the course and the development of teaching materials. Michael Gagnon and Alexis Wellwood provided excellent comments and Sabina Matyiku was very helpful with the graphics. Francis Murchison and Kevin Brousseau contributed exercises on Kuna and Iyinu (Cree) based on their own research. Michael Barkey's work to develop the Concordia Linguistics Outreach Project (CLOUT) was instrumental in getting us to think about how to introduce difficult material to diverse audiences. These audiences are also to be thanked for their patience and feedback, especially the inmates and teaching staff at Clinton County Correctional Facility in Dannemora, New York, where CLOUT presented several workshops.

Our friend and colleague Alan Bale has taught the Language and Mind course twice, and the book owes a lot to the influence of the materials he developed and his own spin on topics we discuss.

It is impossible to say which examples, arguments or discussions contained herein were taken directly from Mark Hale—we have discussed every issue in this book with him at some point over a very long period of time. His influence on our thinking as reflected in these pages cannot be overestimated.

Finally, we are grateful to our respective spouses who, despite the sometimes cantankerous nature of much of our interaction, managed to deepen their relationship with each other while we were engaged in this writing process.

We would like to dedicate the book to our parents and our children.

## List of Figures

1.1 The equivalence class of nouns is itself an abstraction from equivalence classes abstracted from sets of tokens of individual nouns.	10
2.1 Triangle constructed by visual system.	21
2.2 An illusory triskaidecagon.	23
2.3 Unnamed form constructed by visual system.	23
2.4 Rectangles constructed by visual system—of humans and bees, who can be trained to treat the two figures as members of an equivalence class in terms of orientation.	24
2.5 The bees do not treat the illusory rectangle above as the same as either of these two figures.	24
2.6 How many objects on the left? How many on the right?	25
2.7 Mouths, snouts, lips, eyes, and ears of Oonagh and Baby Z?	27
2.8 Spectrogram of a complex wave consisting of music and speech.	28
2.9 The music (left) and the speech (right).	29
2.10 Waveform of a sentence.	30
2.11 Waveform of <i>The spotted cat skidded by</i> .	31
3.1 Which two of these look most alike?	48
6.1 The equivalence class of the phoneme /t/ is itself an abstraction from equivalence classes of allophones abstracted from individual utterances.	113
6.2 A native speaker pronouncing the eight Turkish vowels.	117
10.1 Oonagh and Baby Z are raised in a very similar environment.	217
10.2 Instinct (Innate Knowledge) in goslings	219
11.1 A non-“linguistic” card contrast	244
13.1 Where does knowledge come from?	266



# PART I

## The Object of Inquiry



# 1

## What is I-language?

- |  |  |
|--|--|
| <b>1.1 Jumping in 5</b>                          | <b>1.5 I-language 13</b>                         |
| <b>1.2 Equivalence classes 8</b>                 | <b>1.6 Some implications of<br/>mentalism 14</b> |
| <b>1.3 Partial reduplication in<br/>Samoan 9</b> | <b>1.7 Summing up 15</b>                         |
| <b>1.4 Mentalism 12</b>                          | <b>1.8 Exercises 17</b>                          |

In the summer of 1991 Charles lay in an Istanbul hotel room burning with fever, 15 percent lighter than his normal weight. In the other bed lay his friend Paul, who had just golfed his way to an MBA, also hot with fever, the inside of his mouth covered with blisters.<sup>1</sup> Paul had paid for the room on his credit card, so it was several steps above the dives they had been staying in. He had gotten the name of a doctor in Istanbul from his mother back in Kansas and was now on the phone with the hotel receptionist, who, given the price of the establishment, spoke excellent English. In vain, Paul was asking her to find the number of Dr. Ozel—“That’s right, it’s o-z-e-l, Ozel.” It wasn’t happening.

From the depths of his delirium and intestinal distress, Charles finally found the strength to call out in a hoarse voice, “Tell her to try o with two dots,” referring to the Turkish letter ö, so *Özel*. Much to Paul’s surprise, she found the number immediately. “Reiss, that’s amazing—how did you know that?” Paul asked, fully aware that Charles did not speak Turkish, and also

<sup>1</sup> Charles had recommended that he rinse his mouth in the alkaline waters of Lake Van, but that hadn’t helped at all.

## 4 WHAT IS I-LANGUAGE?

annoyed with himself for having spoken to him civilly, since they were at one of the points in the trip when they wanted to strangle each other. “If you had listened to me on the bus ride from Bucharest to Istanbul, instead of obsessing about what pork products we would sample on the passage through Bulgaria, you would know,” Charles replied brightly, suddenly energized by his ability to gloat.

So, what had Charles tried to explain on that bus ride, in the thirty seconds before Paul’s eyes glazed over? How *did* he know? The answer lies in Charles’s understanding of vowel patterns in Turkish, an example of a most wonderful linguistic phenomenon called vowel harmony. Understanding of Turkish vowel harmony happened to have a practical value in this situation, something neither of us has ever again experienced, but its real beauty lies in the fact that it reflects some of the deepest workings of the human mind.

Our goal in this book is to get you to accept this rather grandiose claim about the vowel patterns in Turkish words. We will introduce many new ideas, some of which will initially strike you as ridiculous. However, we will try to convince you with logical arguments, data-based arguments from both familiar and less familiar languages, and also appeal to general scientific methodology.

Building on linguistic phenomena, our discussion will touch on some of the most longstanding and difficult issues in philosophy including the following:

### 1.1 Big philosophical issues we will address

- The Nature–Nurture debate: How much of what we are is innate and how much depends on our experience?
- What is knowledge? How is it acquired?
- What is reality?
- Whatever reality is, how can we get access to it?
- Is there a principled distinction between mind and body?
- How can our study of these issues bear on social questions and educational practice?

Given both the incomplete nature of all scientific inquiry and the limited space we have, we will not propose complete and final solutions to all these problems, but we do hope to offer a real intellectual challenge in a fascinating domain. This should lead you to experience frustration . . . confusion . . . annoyance . . . and ultimately (we hope) . . . understanding and insight and pleasure.

## 1.1 Jumping in

Not only the average person but also experts in fields like psychology, engineering, neuroscience, philosophy, and anthropology are willing to make proclamations, sometimes in the pages of respected scholarly publications, about language—its evolution, its acquisition by children and adults, its relationship to thought, and so on. But there is a question that is prior to all of these issues, namely *What is language?* We aim in this book to provide you with a deeper appreciation of the nature of language than that of the average academic in the fields listed above.

This book is not a catalogue of cool facts about language, nor is it a report on the exciting findings of modern linguistics over the past fifty years—there are several excellent books on the market for those purposes. Instead, our strategy is to get you to think about language the way linguists do. With this in mind, we'll jump right in with some data (not Turkish—we'll come back to that later), before we even explain the somewhat obscure title of the book. We won't even tell you what “I-language” means yet. By the end of the chapter, we hope you will have an appreciation of the term that is much deeper than you would have if we just handed you a definition.

Let's begin with a simple example, the relationship between singular and plural nouns in Warlpiri, an Australian Aboriginal language.

### 1.2

#### Warlpiri plurals

SINGULAR	PLURAL	
kurdu	kurdukurdu	child/children
kamina	kaminakamina	girl/girls

In English, we form the plural of most nouns (but not all—look at *children*) by adding a suffix to the singular, as in *girl-s*. As you can see, it looks like the plural of a noun in Warlpiri is formed by repeating the singular. This is a real plural—*kurdukurdu* does not just mean “two children,” it means “children” and is used to denote two or a hundred children—any number greater than one. You can probably guess the plural form of the word *mardukuja* “woman”—it is *mardukujamardukuja*.

Processes of word formation that involve repeating material from a basic form (all or just part of the basic form) to create a derived form are called processes of reduplication. Reduplication processes are very common in the languages of the world with a variety of meanings, but are not productive in English.

## 6 WHAT IS I-LANGUAGE?

Even with this simple example, we can learn a lot about the nature of language:

### 1.3

#### Some lessons about language based on Warlpiri plurals

- a. Some aspects of language are simply memorized—it is necessary to remember certain arbitrary links between sound and meaning, for example, that *kurdu* means “child” in Warlpiri but *child* means “child” in English.
- b. Some aspects of language involve rules or patterns. Your ability to correctly guess the Warlpiri form for “women” shows that the form can be generated by a rule.
- c. If there are rules, they have to apply to some kind of input and produce some kind of output. The Warlpiri plural formation rule highlights an important aspect concerning the nature of rules of language—the units of language, the elements that rules affect, can be quite abstract. We cannot give a definite answer to the question “What sound corresponds to the plural in Warlpiri?” because the answer varies depending on context. We will illustrate this point by discussing the rule in more detail below.
- d. The rules apply to elements that are only definable in linguistic terms—for example, the Warlpiri plural rule applies to nouns, not verbs, and the noun-verb distinction is a purely linguistic one.

The first item is fairly obvious, although the arbitrary nature of the sound-meaning links of human language was only really fully appreciated about one hundred years ago by the Swiss linguist Ferdinand de Saussure, the inventor of structuralism. The point is just that one of the requirements for language is memory. A system, device, organism without memory cannot generate Warlpiri or English plural forms from singulars, since it has no way to store the singulars.

The second item will be dealt with again and again in this book. A Warlpiri speaker has to memorize that *kurdu* means “child”, but not how to say “children,” since *kurdukurdu* is generated by a rule that repeats any noun in the singular form to make a plural. Of course the rule or pattern itself must be memorized, but this is an even more abstract kind of information than that required for memorizing words.

This discussion of reduplication illustrates a property of language central to our approach: languages are *computational* systems. This term scares some people, but all we mean by it is that language can be analyzed in terms of explicit rules that apply to symbols. Given an input symbol and a rule that applies to that symbol, we can say what the output form will be. The symbols and rules are different ones than those that are familiar in

math, but the goal of a computational approach is to make them as explicit as the formulas of math or the mathematical formulas used in sciences like physics or chemistry.

To illustrate the third item, let's compare Warlpiri to English, although we will simplify greatly. In English, we can say that the rule for pluralization is something like "If a noun is of the form  $x$ , then the plural of that noun is of the form  $x$ -s" as in *girl*-s. In Warlpiri, the rule must be something like "If a noun has the form  $x$ , then the plural of the noun is of the form  $xx$ ." Both the English and the Warlpiri rules show that the rules of language must refer to VARIABLES. A variable is a kind of symbolic placeholder that can change in value each time a rule is applied. This is particularly clear for Warlpiri—the plural marker is not a constant "piece" of sound, as it apparently is in English regular forms, but rather a copy of the noun. Sometimes the variable has the value *kurdu*, sometimes *kamina*, etc.

Variables in this sense are just like the variables of math—in a function like  $y = 2x + 3$ , we can plug in different values for the variable  $x$  and derive values for the dependent variable  $y$ . If  $x$  is set equal to 4 then  $y = 2 \times 4 + 3$ , which is 11; if the variable  $x$  is set equal to 5, then  $y = 2 \times 5 + 3$ , which is 13; and so on.

In contrast to the Warlpiri rule that tells us to repeat the singular in order to generate the plural, the English rule for regular plurals takes the variable corresponding to a noun and adds a constant -s ending.<sup>2</sup>

If we really want to make the parallel to math explicit, we can think of pluralization as a function mapping members of the set of singulars (the domain of the function) to a set of plurals (the range of the function). In Warlpiri, the pluralization function is something like

$$1.4 \quad f(x) = x \hat{\ } x$$

where the variable  $x$  is drawn from the set of singular nouns and the symbol  $\hat{\ }$  denotes CONCATENATION— $a \hat{\ } b$  means "a followed by b."

In English, the function would require a variable drawn from the set of singulars and a constant corresponding to the suffix:

$$1.5 \quad f(x) = x \hat{\ } s$$

<sup>2</sup> As we said above, we are oversimplifying, but probably only those readers who have taken a linguistics course know what details we are glossing over. If you do not, you are better off, since you can concentrate on our point about variables and constants.

## 8 WHAT IS I-LANGUAGE?

Concatenation is not the same as mathematical addition or multiplication, but it may still be useful to draw a parallel in math. A function like  $f(x) = x + 3$ , where the output of the function, typically shown on the  $y$ -axis of a graph, depends on the value assigned to the variable  $x$  added to a constant, 3.

It is probably apparent that the notions of rules and variables are intimately related. By virtue of the fact that they refer to variables, rules apply to classes of entities. That is what makes the rules productive. The Warlpiri rule that says “Repeat the singular  $x$  to make the plural  $xx$ ” applies not just to *kurdu*, but to *kamina*, *mardukuja*, and in fact to all nouns.

With respect to item (1.3d.), note that nouns are just one of the categories that linguistic rules operate on, but all linguistic categories are just that—linguistic. They cannot be reduced to categories of physics, biology, psychology, or any other domain. The category noun cannot be defined as “a person, place, or thing”, despite what your English teacher told you. We’ll come back to this later.

### 1.2 Equivalence classes

Let’s elaborate on the notion of “variable” used above. The various nouns of Warlpiri have different pronunciations, and yet we are able to treat them all as members of a set or class of elements that are all subject to the same rule. In other words, any noun can stand in for the variable  $x$  in the Warlpiri rule to give the output  $x\hat{x}$ . One way of understanding this is that the rule ignores the differences among various nouns and treats them all as members of the abstract category or class “noun.”

However, there is another kind of abstraction that is necessary before we even can refer to the nouns in this class. If five Warlpiri speakers utter *kurdu*, the actual sound will be different coming from each speaker—there are differences in the shapes and masses of their vocal apparatus, so that an old man and a young child will produce tokens of *kurdu* with very different physical characteristics. And yet someone hearing all five speakers can perceive *kurdu* in each case.

Even more fundamentally, each pronunciation of *kurdu* by even a single speaker will be physically distinct with respect to the sound wave that reaches a listener, due to differences in ambient noise, the moisture in the speaker’s vocal tract, variability in muscle control of the speech organs, etc.

We will come back to these issues in several subsequent chapters, but what they illustrate is a point made about eighty years ago by the great linguist and anthropologist Edward Sapir: “No entity in human experience can be adequately defined as the mechanical sum or product of its physical properties.” In modern parlance, human perception and cognition depends upon equivalence classes—symbolic representations that may be derived from experience (tokens of a word heard) or somehow manifested in behavior (tokens of words uttered), but whose relationship with actual experience is quite complex. As Sapir noted “it is notorious how many of these physical properties are, or may be, overlooked as irrelevant” in a particular instance. In Chapter 2, we will illustrate these ideas with both linguistic examples and examples from other cognitive domains.

Scientists, when they conduct experiments and build theories, also make idealizations and consciously exclude certain observations from consideration. In describing equivalence classes, however, we are saying something different. We are claiming that the human mind and cognitive systems act as a filter on experience—they are built to collapse certain detectable differences when categorizing input.

Returning to Warlpiri, then, we see that we need to recognize that words, themselves, are just equivalence classes. The word *kurdu* is one such class, as is the word *kamina*. But then the category noun is also an equivalence class, an abstraction over the set of abstractions that correspond to words.

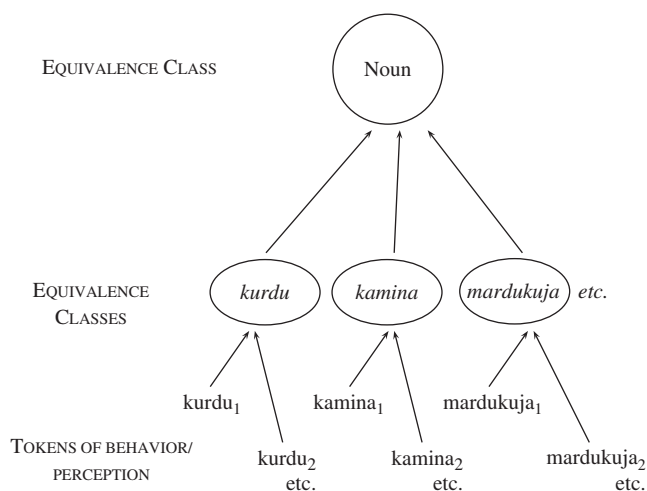
In Fig. 1.1 we see that individual nouns represent an abstraction from various tokens of words that are spoken and perceived. The category noun is itself an abstraction over the set of individual nouns. The use of symbols that represent equivalence classes is one of the most important notions for understanding language.

There is much more philosophizing to be drawn out of the Warlpiri example, but just for fun we will broaden our empirical base with another example of reduplication before returning to big picture issues.

### 1.3 Partial reduplication in Samoan

In the case of Warlpiri, the input symbol corresponded to the singular form of a noun, call it  $x$ , and the output form could be denoted  $x\hat{\ }x$ . This pattern is called “total reduplication” because the whole base form is repeated. In

10 WHAT IS I-LANGUAGE?



**Fig 1.1** The equivalence class of nouns is itself an abstraction from equivalence classes abstracted from sets of tokens of individual nouns.

the following discussion of Samoan, we will discover a pattern of partial reduplication, where only part of the base is repeated.

In Samoan, the singular of the verb “sit” is *nofo* “she sits” and the plural is *nonofo* “they sit” as shown in (1.6).

1.6	Samoan verbs: sg-pl
	nofo nonofo “sit”
	moe momoe “sleep”
	alofa alolofa “love”
	savali savavali “walk”
	maliu maliliu “die”

If you compare the singular with the plural, are you tempted to posit a rule that adds a prefix *no-* to the singular to get the plural? We can reject this by considering some more data: the singular and plural for the verb “sleep” is *moe/momoe*—clearly there is no prefix *no-* here. So, maybe the rule in Samoan involves reduplication, just as in Warlpiri, but in this case reduplication just repeats part of the base word, say the first syllable.<sup>3</sup> Well, this idea fails when we get to another pair, the forms for the verb meaning “love”: *alofa/alolofa*—the first syllable of the singular *alofa* is *a-*, and this is not repeated in the plural. Instead, the syllable *lo* is repeated in

<sup>3</sup> We will assume that you have an intuitive notion of what a syllable is—it is a technical term in linguistics.

*alolofa*. Perhaps these forms show us that the correct rule involves starting at the beginning of the word, looking for the first consonant and the vowel following that consonant, and then repeating the two of them. This would work for the three verbs considered so far, but there is more data to consider: the last two verbs in example (1.6) show the forms *savavali* and *maliliu*, which shows that the correct rule involves copying the second to last syllable of the singular to make the plural.

We thus see that the Samoan rule requires a variable that constitutes a part of a root word. We won't worry too much about how to represent this—it is an advanced topic, beyond the scope of this book, but here is one approach: suppose that we represent each verb as a sequence of numbered syllables starting from the end of the word. So a two-syllable verb would correspond to (1.7a.) and a three-syllable word to (1.7b.), where the symbol  $\sigma$  stands for a syllable.

#### 1.7 Representing syllable sequences

- a.  $\sigma_2-\sigma_1$
- b.  $\sigma_3-\sigma_2-\sigma_1$
- c.  $\sigma_n-\dots-\sigma_2-\sigma_1$

The representation in (1.7c.) corresponds to a word with an arbitrary number of syllables,  $n$ . The rules for plural formation can now be stated by referring to the variable  $\sigma_2$ :

#### 1.8 If $\sigma_n-\dots-\sigma_2-\sigma_1$ is a singular verb, then the plural is $\sigma_n-\dots-\sigma_2-\sigma_2-\sigma_1$

We will revise our view of Samoan later, but for now we have an idea of what is needed. The Samoan and Warlpiri both are instances of the same process of reduplication. What differs is the nature of the variable that gets repeated in each case: a full word in Warlpiri, a syllable in Samoan. It is exactly because we are able to abstract away from the different nature of the two variables that we can see that the two languages are in fact using the same computational process, reduplication.

Our discussion of Samoan has also illustrated a crucial aspect of linguistic analysis—we examined pieces of data and made hypotheses that we have then tested against more data, revising the hypotheses as necessary to match the full data set. This is a good example of how language data can be subjected to the scientific method. The same methodology is used in all sciences. However, as we will discuss later, there is a deeper level of analysis than just coming up with a rule that is consistent with the data.

## 12 WHAT IS I-LANGUAGE?

---

### 1.4 Mentalism

We have posited some rules to account for patterns of nouns in Warlpiri and verbs in Samoan. Let's now ask what those rules are. Well, in some sense, they are our creation, hypotheses we made to account for data sets on the page. However, unless we have some kind of mystical view of our own creative powers, and assuming the data on the page is somehow related to what Warlpiri speakers and Samoans say, it seems reasonable to think that these rules reflect something that existed prior to our analysis—in other words, we have discovered them, not invented them.

Even if the data we analyzed had never been written down, it seems that the rules describe a property of Warlpiri and Samoan speakers. In fact, the memorized singular forms needed to generate the plurals also describe a property of the speakers. Actually spoken words have a sound associated with them, but the rules and the variables they refer to do not—and, as we have seen, even the constant parts do not, since each token is different physically. The rules, the variables, and also the memorized forms of the singulars constitute properties of Warlpiri and Samoan speakers. Similarly, the information that *cat* is pronounced as it is, that it is subject to the regular plural formation rule, and that this rule adds *-s* to the end of the singular is a property of you. We will assume that these properties are a kind of information somehow encoded in the brains of the speakers, and we will refer to that kind of information as a kind of knowledge in the mind of the speakers. Linguistic analysis aims to discover what speakers know—we have discovered, for example, that Samoan speakers *know* (that is, have as one of their properties) a rule that generates plural verbs by reduplicating the second to last syllable of the singular.

The preceding discussion falls under the *mentalist* approach to linguistics. It considers the information and rules and patterns that can be used to analyze linguistic behavior to reflect mental properties, properties of the minds of individuals—the mind consists of information and rules and patterns, some of which constitute knowledge of language. We will later argue that what is mental is part of the biological world, and thus our approach is also known as *biolinguistics*.

Neuroscientists who are trying to understand how cognition arises from the physical matter of the brain need linguists to tell them what kinds of powers inhere in the brains they are studying. If they cannot come up with a model of the brain that accounts for the ability to memorize words (like

Warlpiri singulars) and also store and apply rules that contain variables (the pluralization via reduplication rule of Warlpiri and Samoan, for instance) then their work is not done.

### 1.5 I-language

You now have all the pieces that are necessary to understand what I-language is. An I-language is a computational system that is encoded in, or a property of, an individual brain. It is a system of rules (a grammar) that computes over symbols that correspond to equivalence classes derived either from experience or other symbols. The mind contains (or perhaps is composed of) many such systems, for vision, language, etc., and an I-language is the name given to that one of these systems that generates the structures associated with speaking and understanding speech.

The I-language approach to linguistics thus studies *individual* mental grammars, entities that are *internal* to each person. In addition to these two words beginning with the letter *I*, there is a third relevant term implicit in the notion of a grammar as a system of rules or patterns. In mathematics a set can be defined *extensionally*, by listing its members, or *intensionally*, by providing a formula or description that characterizes all and only the members of the set. For example, {2, 4, 6, 8} extensionally defines the same set as the intensional description “even numbers less than 10.” Notice that an intensional definition is more practical for large sets, and required for infinitely large ones like the set of *all* even numbers. A Warlpiri speaker need not store the set of plurals as an extensionally defined list, since the reduplication rule defines this set intensionally as a function from the set of singulars.

#### 1.9 Two characterizations of the set of Warlpiri plurals

**Extensional:** {kurdukurdu, kaminakamina, mardukujamar-  
dukuja, ... }

**Intensional:** { $x \wedge x$  such that  $x$  is a singular noun}

The intensional characterization reflects the rule-governed nature of the relationship between singulars and plurals. I-language is meant to suggest all three of these notions—internal, individual, and intensional.

The study of the shared properties of all I-languages is thus the study of what is sometimes called the human language faculty. This study is

## 14 WHAT IS I-LANGUAGE?

sometimes called Universal Grammar, the goal of which is to discover the core of properties common to all I-languages.<sup>4</sup> We will address implications of the I-language approach and also contrast it with other approaches throughout the book.

### 1.6 Some implications of mentalism

This mentalistic, I-language approach to language has several implications. First of all, we need to recognize the difference between our conscious knowledge of Warlpiri and Samoan reduplication that we developed as a scientific analysis, and the unconscious knowledge that the speakers have. Samoans, for example, may have no idea what a syllable is, and thus could not tell us how the singular and plural verb forms they produce are related. They acquired these rules as pre-literate children without any direct instruction from their parents—they were not given organized data sets as you were.

Furthermore, if all speakers of Warlpiri were to die tomorrow, then nobody in the world would have the kind of knowledge that they have, and the language would cease to exist. We might have some writings that describe our analysis of aspects of their language, but that is all. A language, for linguists, is a system of representations and rules in the mind of a person. If the person ceases to exist, that particular person's language ceases to exist. In other words, we have been talking about the Samoan language and the Warlpiri language, but we have been doing so informally. From a linguistic perspective, each Warlpiri speaker and each Samoan speaker has his or her own set of symbols and rules, what we call his or her own *mental grammar*, his or her own I-language.

If this is so, that each Warlpiri speaker actually has his or her own individual mental grammar, then how can Warlpiri speakers communicate with each other? Why do they seem to have the same grammar? The answer is simple—they have mental grammars that are quite similar because they are all humans and they were exposed to similar linguistic experiences when they were acquiring their language as children.

<sup>4</sup> Just as the terms *physics* and *history* refer both to objects of study (the physical world or the events of history) and the study itself (as in “He failed physics”), the term Universal Grammar is also used sometimes to refer to the common core of the human language faculty.

Everything we have just said about Warlpiri and Samoan holds as well for English. If we take the mentalistic approach seriously, then we have to admit that there is no entity in the world that we can characterize as “English.” There is just a (large) bunch of people with fairly similar mental grammars that they can use to communicate in a way that is typically more efficient than between what we call Japanese and English speakers, because the so-called English mental grammars are more similar to each other. We will continue to use terms like “the English language,” “Warlpiri plurals,” and “Samoan verbs,” but bear in mind that each name is a just practical label for a set of individual mental grammars that are identical with respect to a given phenomenon under analysis.

### 1.7 Summing up

So, at this point, we hope you have an idea of the I-language approach. The ultimate goal is an understanding of the human language faculty, which is instantiated in individual minds/brains, in the same way that we talk of a human visual faculty. Each individual person, based on their particular experience of language acquisition, ends up with a language faculty that is in a particular state.

We told you earlier that we would not review the major findings of modern linguistics, but we have changed our mind—here they are:

#### 1.10 The fruits of linguistic research

- Every language is different AND
- Every language is the same.

Believe it or not, both of these claims have elicited virulent criticism. Obviously, we have stated the claims like this for rhetorical effect, but we have suggested that they can both, in fact, be true in some non-trivial way. The two claims are complementary rather than contradictory.

We have illustrated the sense in which linguists say that each language is different: each language corresponds to information in a particular mind. Since each person has at least slightly different experiences of language acquisition, it is not surprising that each ends up with different grammars, different bodies of information. When we say that two people speak the same language, it is rather like saying that they are “near” each other. This is a useful expression whose definition depends on numerous

## 16 WHAT IS I-LANGUAGE?

factors—Montreal is near Kingston, only three hours away; we work near Mary, only three blocks away; we are sitting near Mary, only three feet away; Paul’s liver is near where his gall bladder used to be, only three inches away (N.B. We know nothing about anatomy). What does *near* mean? There is no formal definition of the everyday word *near*, and there is no formal definition for the everyday term “English.” Linguistically, there are no dependable criteria for defining a speaker of English—some dialects share properties with Hungarian that others dialects do not share, for example.

The situation becomes even clearer if we look at other languages (using the term in the everyday sense). Spanish and Italian are called different languages, but speakers of the standards feel like they can communicate with each other quite well. On the other hand, the various Italian dialects are often mutually incomprehensible—they are called dialects of the same language because they are spoken within the political boundaries of Italy, not for any linguistic reasons.

The second claim is just the hypothesis of Universal Grammar, an idea we have already hinted at. We will try to show in later chapters that Universal Grammar is more of a logical necessity than a hypothesis. However, in order to understand the claims, and to decide whether to accept or reject them, we propose to continue developing an understanding of what language is.

As promised, we have already argued for one apparently ridiculous notion, the non-existence of English! As with any scientific endeavor, it is to be expected that our results will surprise us from time to time, and that they will be at odds with our everyday intuitions and common sense. In the same way that modern science departs from our common sense, which tells us that light should behave as either a particle or a wave, not both, or that our bodies and our cars must be made of fundamentally different substances, we expect the scientific study of language to overturn some of our most dearly held intuitions. This commitment to science and its ability to surprise us is expressed well in the following quotation from Zenon Pylyshyn, a psychologist and computer scientist whose work inspired much of what you will find in the following pages:

[I]f you believe  $P$ , and you believe that  $P$  entails  $Q$ , then even if  $Q$  seems more than a little odd, you have some intellectual obligation to take seriously the possibility that  $Q$  may be true, nonetheless. [Zenon Pylyshyn (1984), *Computation and Cognition*: xxii]

Throughout the book, we intend to mind our *Ps* and *Qs* in accordance with Pylyshyn’s dictum.

1.8 Exercises

**Exercise 1.8.1. Ethnologue:** Throughout the book we refer to languages in the everyday sense of English, ~~Warpii~~, Spanish, and so on. Find information about where languages are spoken, how many speakers they have and what family they belong to by consulting the *Ethnologue* at <http://www.ethnologue.com>. Go to the website and write up a description of the language that immediately follows your family name alphabetically and the language that immediately follows your given name. (If your name is James Jameson, or something else that gives the same language twice, use the language that *precedes* your family name alphabetically.)

**Exercise 1.8.2.** How do you express the meaning *very* in Pocomchi? Fill in the blanks.

adjective		<i>very</i> + adjective		
saq	white	saqsaq	very white	} all three times
ras	green	rasras	very green	
q’eq	black	q’eqq’eq	very black	
q’an	ripe		very ripe, rotten	
nim	big		very big	
kaq	red		very red	

**Exercise 1.8.3.** Can you see how to generate the set of definite nouns (like *the bird*) from the set of bare nouns (like *bird*) in Lyle? Note that vowels in Lyle can bear one of three tones: a = mid tone; á = high tone; à = low tone. These tonal differences are distinctive—they can differentiate meaning.

kúmí	bird	kúmíí	the bird
yálá	millet	yáláá	the millet
kùlí	dog		the dog

## 18 WHAT IS I-LANGUAGE?

Things may be a bit more complex than you thought:

nà	foot	nàá	the foot
yiji	church	yijií	the church
ya	market	yaá	the market
cèlé	parrot	cèléé	the parrot

To make the definite form (*the* + N) repeat \_\_\_\_\_ but always use a \_\_\_\_\_ tone.

What equivalence classes are relevant to a discussion of these Lyele noun forms?

**Exercise 1.8.4. Is it English?** Here are some sentences rendered in Standard orthography that we have heard spoken in various places that are referred to as English-speaking places. Identify differences from your own variety of English, if you can figure out the intended translation into your own dialect. Are these sentences all English? How does the I-language approach bear on the issue?

1. We are allowed running here. (Montreal)
2. We are allowed to run here. (Brooklyn)
3. I did nothing today. (Brooklyn)
4. I didn't do nothing today. (Brooklyn)
5. The government has decided to raise taxes. (Montreal)
6. The government have decided to raise taxes. (London)
7. I'm going to the dep to get some cigarettes and beer. (Montreal)
8. That's all the faster I can run. (Michigan)
9. That's as fast as I can run. (Brooklyn)
10. I might could go. (Alabama)
11. I might be able to go. (Brooklyn)
12. He been try make me mad. (Cajun English, Louisiana)
13. I ate a egg. (Ypsilanti)
14. I ate an egg. (Brooklyn)

### Further Readings

- Chapters 1 and 2 of *Patterns in the Mind* by Ray Jackendoff (1994). This is an excellent book that inspired much of this book—we actually recommend reading it all.

- Recapturing the Mohawk Language by Marianne Mithun and Wallace Chafe, in Timothy Shopen (1979) (ed.) *Languages and their Status*, (3–33). We have our students read this partly because Mohawk is spoken in the vicinity of Montreal where we teach, and partly because it gives interesting illustrations of productive grammar in a language that is very different from English. There are aspects of the article we disagree with, but this can lead to useful discussion.

# 2


## I-everything: Triangles, streams, words

<b>2.1 A triangle built by the mind</b>	<b>20</b>	<b>2.4 Words are built by the mind</b>	<b>30</b>
<b>2.2 More visual construction</b>	<b>25</b>	<b>2.5 Summing up</b>	<b>33</b>
<b>2.3 Auditory scene analysis</b>	<b>27</b>	<b>2.6 Exercises</b>	<b>33</b>

In the last chapter we introduced two important notions related to I-language: computation and equivalence classes. As we suggested, these ideas have quite broad relevance for an understanding of the human mind, and in this chapter we will provide demonstrations from various domains in addition to linguistic ones. Abstracting away from physical properties and providing analyses in terms of equivalence classes is something that all scientists do, including linguists and other cognitive scientists. In the case of cognitive science, this process of forming equivalence classes actually constitutes the object of study. The human mind/brain automatically filters incoming stimuli in such a way as to collapse even grossly distinct signals and treat them identically. This kind of information processing is what cognitive science studies.

### 2.1 A triangle built by the mind

Look at Fig. 2.1. If you are a normal human being you will see a white triangle with its vertices at the center of the three Pac-Man figures. You can see the edges of the triangle and trace them with your finger, but if you cover up the Pac-Men the edges seem to disappear. The area of the triangle is exactly the same shade of white as the background of the page, so it is not



**Fig 2.1** Triangle constructed by visual system.

surprising that no edges remain visible. But why do you see a triangle in the first place—are you hallucinating? If so, why does every other human who looks at such a page also see the triangle?

From the point of view of physics, which can measure things like the light reflecting off the page, there is no distinction between the area inside the triangle, its edges and the background. So is the triangle not *real*? We could decide to say that the triangle is not part of the real world and thus adopt a pure physicalist definition that accepts as real only that which can be defined using the categories of physics, like mass, wavelength, velocity, etc. But that is not very satisfying—it leaves us with a big mystery: Why does everyone who looks at the page see the triangle? Isn't *that* a real fact about real humans?

So, is there really a triangle on the page? The solution offered by cognitive science to the triangle mystery is this. The human visual system interprets certain stimuli in such a way as to construct a representation of a triangle. In other words, the triangle is not a physical property of the page but a result of how you process physical stimuli like this page under certain circumstances—for example, when your head is oriented correctly, your eyes are open, and it is not pitch dark in the room. In other words, your mind imposes the triangle interpretation on the page. Now, one could just declare that there is no triangle since its edges cannot be physically detected. One could decide that the only things that are real are those that can be described in physical terms, and the edge of the triangle has no mass, or charge, or luminance, and so it is not a physical entity and thus not real. If the edges aren't real, the triangle itself cannot be real.

As we said, one can arbitrarily decide to use the term *real* in this way, but this puts us in an uncomfortable situation. On the one hand, we have to accept that every single person who looks at the triangle figure sees the same thing, a triangle, and so do certain animals, as can be shown by experiment,

## 22 I-EVERYTHING: TRIANGLES, STREAMS, WORDS

despite the fact that the thing they see is not real. Are we all deluded? How come we are all deluded in exactly the same way, then? On the other hand, if we want to study the human visual system scientifically, we have to accept the idea that science can study what is not real. Rather than arbitrarily defining the *real* to include only that which has mass, charge, luminance, location, and so on, we can recognize that physics contains certain categories and vision science others, but it is not the case that the categories of one are necessarily more real than those of the other. In fact, the categories of modern physics are so remote from our everyday experience of what we call the physical world, that they too must be considered abstractions. We'll elaborate on this later on.

People sometimes think that the fact that we see a triangle on the page has to do with the fact that we have the word *triangle* that we can apply to certain experiences. There are at least two problems with this view. The first problem is that if we couldn't recognize triangles in the first place, we would not know what to apply the word to—it just doesn't make sense to say that the word allows us to perceive the object.


The second problem is that our visual system constructs percepts of edges and corners that compose objects even when we have no name for them. Fig. 2.2 contains an illusory regular polygon with thirteen sides. You experience the illusion even if you do not know that such a figure is called a *triskaidecagon* by mathematicians.

We even see illusory shapes that nobody has a name for, as in the blob of Fig. 2.3.

We see the contours of a blob because of the way our visual system processes the information received by our eyes. The triangle or blob we perceive is not part of the physical input to our eyes but is rather an information structure, or representation, constructed by the visual system based on the input it receives and its own rules.

Note that we make no effort to see the triangle or the blob, and in fact we can't help but see the edges of these figures, even when it is pointed out that there is no difference in luminance between the figure and the background. Our visual system works the way it does despite contradictory conscious knowledge.


We mentioned that certain animals will also see shapes when presented with a display with illusory contours like Fig. 2.1. Nieder (2002) reviews the evidence for animal perception of such shapes: for example, bees have been trained in a Y-shaped tunnel to choose the correct branch to a sugar




**Fig 2.2** An illusory triskaidecagon.

solution when that route is marked with stripes oriented in a particular direction, say rising towards the right. The bees are then tested with the rightward rising pattern replaced by various displays. If the display contains a solid rectangle or one with illusory edges as in Fig. 2.4, the bees treat it like stripes with the same orientation.

However, if the display contains a solid triangle with the wrong orientation, or crucially with the Pac-Men oriented in a way that does not produce illusory contours (even for us humans), the bees treat the display as different from the rightward rising stripes. The two sides of Fig. 2.4 are




**Fig 2.3** Unnamed form constructed by visual system.




**Fig 2.4** Rectangles constructed by visual system—of humans and bees, who can be trained to treat the two figures as members of an equivalence class in terms of orientation.

processed as members of an equivalence class that excludes the two sides of Fig. 2.5. Clearly, from a purely physical perspective, one could argue that each side of Fig. 2.4 is more like one of the members of Fig. 2.5 than like the other member of Fig. 2.4. However, the well-defined contours of the solid rectangle and the illusory contours of the other figure can be treated as equivalent by a bee (and a human).

However, we can't just say that the rectangle or triangle is "out in the world." If they were out in the world, then a computer vision system with a robotic eye that is way more sensitive than a human eye should be able to detect these shapes at least as easily as a person can. However, it is, in fact, very difficult to get an artificial system to recognize these displays as a rectangle or triangle. It is only a rectangle or triangle to a system that processes information about incoming patterns of light in such a way as to construct the representation of such shapes. Nature has given us such a system, but we haven't yet figured out how to endow computers with such a system. The rectangle or triangle is a symbolic representation, a member of an equivalence class, that is internal to the entity (bee, cat, human,



**Fig 2.5** The bees do not treat the illusory rectangle above as the same as either of these two figures.



**Fig 2.6** How many objects on the left? How many on the right?

whatever) that is constructing it. Since we assume that bees do not have words for various shapes, we now have a third argument against relating our perception of the triangle to our linguistic experience.

We now have a nice parallel to the discussion of the non-existence of languages from the previous chapter. There is no such thing as Warlpiri or *the* Warlpiri word for “child” or *the* Warlpiri reduplication rule; there are just a bunch of humans whose minds contain similar kinds of rules and symbols that we informally group together as Warlpiri. Similarly, there is no triangle or rectangle on these pages, but humans (as well as members of some other species), who all share the same kind of visual system, all construct the same percept upon exposure to this page. Our nervous systems just process information in this way. (As an aside, note that there are, in fact, *no* triangles in the physical world—triangles are geometric figures with sides consisting of perfectly straight line segments meeting at vertices whose angles add up to exactly  $180^\circ$ . Perfectly straight line segments, for example, do not exist in the physical world.)

## 2.2 More visual construction

Our discussion of vision has already led us to some surprises, and further consideration will, as you suspect, only show us greater complexity. Let’s assume a computational approach to vision that parallels in many ways the approach we introduced for language in Chapter 1. On the topic of representation in perception Bregman (1990:3) makes the following point:

In using the word ‘representations’, we are implying the existence of a two-part system: one part forms the representations and another uses them to do such things as calculate . . .

Let’s now apply Bregman’s observation to Fig. 2.6.

## 26 I-EVERYTHING: TRIANGLES, STREAMS, WORDS

On the one hand, our visual system must clearly detect and represent shading, textures, and edges. On the other hand, it must perform the calculations or inferences that lead us to see the left side of the figure as representing an ellipse partly occluded by a rectangle, to group the two gray regions together. Note that our visual inference system cannot help but see things this way, and it does not matter that there is no right way to experience the image—it may be a picture of a rectangle occluding an ellipse, or it may be a picture of three distinct objects, as suggested by the right-hand side of the figure. In fact, it is just a pattern of ink on the page: we can specify its physical properties; and we can tell you what numbers we entered in the graphics program that we used to design it. But none of this matters—we, as humans, cannot help but perform the computations that lead to the perception of one object occluding another on the left-hand side. Note that the only difference between the two sides is the black perimeter of the rectangle on the left. The fill of the rectangle and the empty space perceived on the right-hand side are both just regions of the page without any ink.

The output of the visual system, representation of objects with colors, textures, shapes, and sizes feeds into other systems that also appear to involve computations and constructions.

Consider Figure 2.7. On the one hand, we see a nose, a snout, some ears, eyes, and lips, but, on the other hand, we see a picture of Oonagh and Baby Z. Is there any reason to even make such a part/whole distinction, or are we just being pedantic? Well, consider the following description of the condition *prosopagnosia* from the Preface of Hoffman's (1998) *Visual Intelligence*:

After his stroke, Mr. P still had outstanding memory and intelligence. He could read and talk, and mixed well with the other patients on his ward. His vision was in most respects normal—with one notable exception: he couldn't recognize the faces of people or animals. As he put it himself, "I can see the eyes, nose and mouth quite clearly, but they just don't add up. They all seem chalked in, like on a blackboard... I have to tell by the clothes or by the voice whether it is a man or a woman... The hair may help a lot, or if there is a moustache..." Even his own face, seen in a mirror, looked to him strange and unfamiliar. Mr. P had lost a critical aspect of his visual intelligence.

So, Mr. P appears to see normally in the sense of seeing objects like ears and noses and lips, but further computation by the face recognition system, involving the output of the visual system, is somehow impaired. We




**Fig 2.7** Mouths, snouts, lips, eyes, and ears of Oonagh and Baby Z?

typically think of faces as objects in the world, but this case suggests that face perception requires construction of a complex symbolic representation from the objects that themselves are constructed by the visual system. These processes of construction occur inside individual minds/brains according to rules and principles (we might say grammars) of vision and face recognition.

### 2.3 Auditory scene analysis

Just as our mind actively constructs the objects of visual perception and face recognition, it also constructs the objects of auditory perception, what we hear. Imagine you are listening to the hum of an air conditioner and then hear the footsteps of someone walking down the hall to your office. The hum is continuous, but the footsteps are a sequence of short sounds. From a physical point of view, each step is a separate event, yet you perceive the sound of footsteps as a single auditory “object.” Your mind integrates the sequence of steps into what is called a single auditory stream. Notice that the continuous hum of the air conditioner constitutes another stream. Although this may seem obvious, in fact there is a tremendously




**Fig 2.8** Spectrogram of a complex wave consisting of music and speech.

complicated issue to explain. Every time a footstep occurs, the sound originating from the step combines with the sound of the hum, and the vibrations that reach your ears are a composite of these two sources and any others that may be present, such as a person talking on the phone at the next desk. Yet your mind is somehow able to segregate the complex sound wave into two or more separate streams.

Auditory scene analysis is a framework for studying auditory perception developed by Albert Bregman and his collaborators. Auditory scene analysis can be broken down into two main components. One problem, given the fact that sound waves from various sources are combined into a single wave that reaches the eardrum, is that of simultaneous spectral integration and segregation. The auditory system integrates into a single representation parts of the sound spectrum reaching the ear within a temporal window that “go together.” Of course, the decision that spectral regions “go together” is determined by properties of the auditory system, and in the case of an illusion, the decision may lead to a non-veridical percept. An example of spectral integration is the perception of a played musical note and the overtones that give the instrument its unique timbre as emanating from the same source. The process of assigning parts of the spectrum to different perceptual sources is called spectral segregation: attending to speech while a fan provides a high-frequency hum in the background requires spectral segregation.

The other main component of auditory scene analysis is sequential integration—acoustic events occurring separated in time may be integrated into a single auditory stream. Examples of streams include a sequence of footsteps or the continuous sound of falling rain. Individual sounds of a foot striking the ground are separated by silence or other sounds, yet the steps are integrated into a single perceptual object, a stream.

The complexity of the task of auditory scene analysis can be appreciated by considering the spectrogram in Fig. 2.8. This is the spectrogram of a



**Fig 2.9** The music (left) and the speech (right).

wave created by mixing a sample of recorded speech and some music. The spectrograms of the music and speech separately are shown in Fig. 2.9. In this example, we were able to display the music and speech separately because we had the separate recordings. The mind has to extract such information from a complex stimulus, like the mixed signal, to construct distinct streams from a single physical signal.


The following quotation expresses the extent to which we construct our auditory experience—just as the edges of the triangle above are constructed by our minds, so are the edges of auditory events:

The perceptual world is one of events with defined beginnings and endings... An event becomes defined by its temporal boundary. But this impression is not due to the structure of the acoustic wave; the beginning and ending often are not physically marked by actual silent intervals. [Handel 1989]

This quotation suggests that our minds impose the structure of our auditory perception, just as with our visual perception, and it is pretty easy to find parallels between the two domains. Suppose we remove the border of the rectangle on the left side of Fig. 2.6, giving the right side. It is less likely that you perceive the two curved regions as belonging to a single elliptical object, since they appear separated by “empty space.” The presence of the border on the white region on the left lets us perceive it as belonging to a white object which can mask the non-visible part of a (continuous) ellipse. An exact parallel can be designed for audition.

If we take a tone and replace a portion of it with silence, we’ll hear the resulting sound as having a gap in the tone. However, if we replace the silence with broad-frequency white noise of a loudness that would be sufficient to mask the tone, then we actually perceive the tone as continuing behind the noise. Interestingly, we will be able to fill in a gap in a sound behind a mask even if the surrounding portions are not constant. For

## 30 I-EVERYTHING: TRIANGLES, STREAMS, WORDS



**Fig 2.10** Waveform of a sentence.

example, a gap in a tone that rises in frequency can be restored by our perceptual system if it is masked by noise.<sup>5</sup>

In vision and in audition the mind plays an active role in constructing our experience. In the next section we will discover that the perception of boundaries applies even to sound corresponding to speech—even word boundaries are constructed by our minds.


## 2.4 Words are built by the mind

So, what is this discussion of vision and audition doing in a book on linguistics? The point is that just as our visual and auditory percepts are due to active mental processes, our linguistic cognition is also a function of processing by our mental grammar. This is most easily demonstrated by our perception of speech sounds and their organization into words.

Before we proceed, note that speech perception depends upon prior general auditory processing, since the words we hear are sounds. This relationship between audition and speech perception is somewhat like that between object perception and face recognition discussed above: the output of one system is fed into another.

The display in Fig. 2.10 shows a waveform of a recorded utterance, *The spotted cat skidded by*. The horizontal axis shows time and the vertical axis is basically a record of the loudness of the signal at each point. When the display reaches up and down from the horizontal axis, the speaker's voice was most loud, and where the waveform is basically just a horizontal line, the speaker was silent. (Because of background noise there is never perfect

<sup>5</sup> A demonstration and further discussion of this phenomenon can be found <http://ego.psych.mcgill.ca/labs/auditory/Demo29.html>, which is accessible from the webpage of Al Bregman whose work inspired much of this discussion.



**Fig 2.11** Waveform of *The spotted cat skidded by.*

silence indicated.) Based on this information, try to figure out where each word of this sentence begins and ends.

You were probably tempted to place your word boundaries wherever the display indicates a silence. However, you will be surprised to see the transcription we have provided of some of the sounds and the word boundaries in Fig. 2.11. There are two particular aspects of the display to note. First, approximate word boundaries are indicated by the symbol #, but note that there is not necessarily silence between words. Second, note that there is sometimes silence (apart from the slight noise in the signal) inside of words—this is normal when the utterance contains sounds like those corresponding to the letters *p*, *t*, *k*. In the words *spotted* and *skidded* there is an *s* before a consonant in the same word, yet the waveform shows that there is silence between the two sounds.<sup>6</sup>

On the other hand, there is no silence between the words *the* and *spotted*. This situation is typical, and if we presented you with recorded speech from an unfamiliar language, you would not be able to find the word boundaries by either looking at the waveform or listening to the recordings. You need a mental grammar of the language to impose word boundaries on the signal.

We learn from this example that, like the triangle we saw, the words we perceive in speech are the result of information processing. The sound wave that reaches our ears does not inherently contain words. Our minds impose words on signals we hear; the words are not part of the signal.

In the case of the perceived triangle, it turns out that any normal human (or bee) who can see will see the triangle, because we all have a visual system that processes information in the same way. In the case of language, there

<sup>6</sup> This silence is there because of the way these consonants are articulated—a detail you would learn more about in a phonetics course.

## 32 I-EVERYTHING: TRIANGLES, STREAMS, WORDS

is some learning involved—we perceive word boundaries in speech in ways that depend on the languages we have learned. So linguistic information processing appears to be more plastic, more influenced by experience, than visual information processing.

We will not go into further detail analyzing waveforms here, but we will just mention that our perception of speech as consisting of discrete segments is also due to processing by our speech perception cognition—the actual signal does not consist of a well-defined sequence of segments, as you can perhaps tell by examining the waveform above. The portions of the waveform corresponding to each speech sound blend together, showing us that, like words, our perception of segments is due to the constructive information processing carried out by our minds.

The point we have just made about triangles, auditory streams, faces, and words turns out to be true of all categories of human experience—they are not definable by their actual physical properties. For example, to again cite Edward Sapir, there is no way to define in physical terms, the difference between a club and a pole. An object is called a club if we use it as a club or believe that it was intended to be used as a club by the person who fashioned it—there are no necessary and sufficient physical criteria to make something a club as opposed to a pole.

Similarly, there are no necessary and sufficient physical criteria to determine where word boundaries fall in a waveform. The perception of word boundaries depends on which mental grammar is being used to process the signal. The signal itself has no information about word boundaries, since words are not physically definable entities. We could even imagine a situation in which a given signal would be parsed into different words by speakers of different languages.

Even closer to home, we have all had the experience of misparsing, of assigning word boundaries in a fashion not intended by the speaker. Misassignment of word boundaries is one of the sources of “mishearing” that leads to *mondegreens*<sup>7</sup> like hearing Jimi Hendrix say ‘*Scuse me while I kiss this guy* instead of the intended ‘*Scuse me while I kiss the sky*. The [s] of *sky* is misparsed as belonging to the previous word. As we will see in Exercise 2.6.2, [k] after [s], is usually indistinguishable from [g].

<sup>7</sup> According to the Wikipedia, the term was coined by Sylvia Wright in *Harper’s Magazine*, November 1954, in a discussion of her understanding as a child of the poetic phrase *And laid him on the green* as *And Lady Mondegreen*.

## 2.5 Summing up

So, to reiterate, the triangles and the words we perceive are related in a very complex and indirect fashion to the physical stimuli we receive. The fact that we can imagine words and triangles in our mind's ear and eye, without any outside stimulus at all, further demonstrates that perception of these entities is due to construction by the mind.

So why is this chapter called “I-Everything”? The “I” of *I-language* is chosen to suggest *individual*, *internal*, and *intensional*. It should be obvious that the triangle you see, the auditory streams you hear, and the words you identify in an utterance are all the output of “I”-systems. For example, each of us has our own individual visual system, and this system is clearly internal to us, part of our make-up as organisms. Moreover, it is not the case that we can only perceive a limited number of objects whose images are stored in some kind of mental list. Like the productivity of our linguistic grammars, our visual computational systems are productive. They construct edges given an infinite range of stimuli. You do not see a triangle only when looking at Fig. 2.1 from a single angle under one set of lighting conditions from a particular distance—try moving the page around or visit <http://www.cut-the-knot.org/Curriculum/Geometry/EdgeIllusion.shtml> to explore the visual grammar that lets you construct the edges of a triangle.

## 2.6 Exercises

**Exercise 2.6.1. Word boundaries:** The purpose of this exercise is to give you firsthand experience with the abstractness of linguistic representation. You will see that the word and segment boundaries we perceive are typically not present in the acoustic signal, but instead are imposed by our minds. This is a linguistic example of the construction of experience that we have discussed in relation to vision and hearing.

You will work with a recorded sentence and try to find word boundaries. Using a sound editing program such as Praat ([www.praat.org](http://www.praat.org)) examine the waveform of the sound file *ilang.wav* available from the companion website. If you have trouble saving the file from your browser, then download the .zip file containing all the materials you need and unzip it. You'll end up with a folder containing the sound file.

## 34 I-EVERYTHING: TRIANGLES, STREAMS, WORDS

There is also an image of the waveform on the website, but you need a sound editing program to zoom in and play selections. You may find it useful to print this waveform, or one from within Praat, to mark your word boundaries on. In order to complete the exercise you need to be able to see the waveform (a graph of intensity vs. time) and select and play portions of the sound file. You also need to be able to find points in time in the waveform window. This is pretty easy in Praat and most other phonetics programs.

You can also get a manual from the Praat homepage or get a manual written by our former student Tom Erik Stower from the companion page. (This manual contains more detail than you will need.) Write your answers down on scrap paper as you proceed, so that you do not lose your work if the computer crashes or if your session is interrupted.

- a. Provide an orthographic transcription of the sentence—that is, just write it in normal English writing.
- b. For each word of the ten or so words in the sentence, write the ending time of the word in milliseconds. (Count contractions like *can't* as two words, if there are any.) For example:  
End word 1 “the”: 136 msec  
End word 2 “cat”: 202 msec  
and so on.
- c. Are there any cases of silence within a word? Give at least one example and say where the silence occurs—between which sounds? Example: The word “Casper” has silence between the *s* and the *p*. This can be heard and also seen because the waveform has almost no amplitude between those two sounds.
- d. Is there generally silence or a pause between words? Give an example of two adjacent words where you had difficulty deciding on where to place the boundary. Example: It was hard to decide on the boundary between “the” and “apple.”
- e. Comment on any difficulties or interesting issues you encountered in any part of this exercise. (Say something coherent and relevant—if you found nothing interesting, fake it.)

**Exercise 2.6.2.** Take either a pre-recorded sentence or record your own and mark the segment boundaries on the waveform. In other words, find the division between adjacent sounds, like the *s* and the *k* of a word like *sky*. Comment on any problems you run into. See how your results compare

to those of your classmates. Tom Erik's manual will tell you how to mark the boundaries on a waveform in Praat and print out the results.

**Exercise 2.6.3. More construction:** Visit a website with optical illusions and find examples of illusions that demonstrate your mind's role in the construction of color, motion, and shape. Here is one excellent site:

<http://www.michaelbach.de/ot/index.html>

### Further Readings

These readings are all fantastic, and we have borrowed freely from them in this book. *Visual Intelligence* is the most accessible.

- Chapters 1, 2, 7 of *Visual Intelligence* by Donald Hoffman (1998).
- Chapter 1 of *Auditory Scene Analysis* by Albert Bregman (1990).
- “The problem of reality” by Ray Jackendoff. *Noûs*, Vol. 25, No. 4. Special Issue on Cognitive Science and Artificial Intelligence (Sep., 1991), pp. 411–33. Reprinted in Jackendoff's *Languages of the Mind: Essays on Mental Representation* (1992).
- “Seeing more than meets the eye: Processing of illusory contours in animals” by A. Nieder (2002). *Journal of Comparative Physiology* **188**: 249–60.