

8 Modeling the Linguistic Brain with LLMs

I like nonsense. It wakes up the brain cells.

—Dr. Seuss

What went wrong in the 1980s was that we all discovered that artificial intelligence technologies were actually useful. And that led to a kind of commercial distraction from which the field has never recovered.

—Patrick H. Winston, comment during panel discussion at MIT, 2011

Brain Models for Vision Can Teach Us Important Lessons

The metaphoric term *neural networks* naturally gave way to concrete implementations of real neural mechanisms. This was pioneered by visual scientists (as has often been the case), and hence vision is a good place to start; it might teach us something important that can then be applied to the modeling of language in the brain.

In 2007, Tomaso Poggio’s computational neuroscience group at MIT successfully modeled early stages of primate vision with a feedforward network (that is, a network with no recurrence).¹ Predating Deep Learning, this work is based on a long tradition in the neurobiology of vision, and represents a new level of sophistication and detail, manifest in a complex model that features a “tentative mapping between the ventral stream in the primate visual system and the functional primitives of the feedforward model.”

In this model (figure 8.1), the number of units within each layer varies. Its creators first showed how its early recognition of visually displayed object compares to that of human observers. Next, they measured the activation levels at the model’s various layers, as evoked by objects displayed




Figure 8.1

A flow chart of monkey early vision, 2007: the functional components in the monkey (bottom left), some anatomical pathways (top left), the structure of the network (middle left), and how they apply to an image (bottom middle). Reproduced by permission from T. Serre, A. Oliva, and T. Poggio, "A Feedforward Architecture Accounts for Rapid Categorization," *Proceedings of the National Academy of Sciences of the United States of America* 104 (2007): 6424–6429.

to the model. This was followed by similar experiments with primates: they measured neurophysiological results in locations of the monkey's brain that corresponded to the model's parts. Finally, they compared the model's network to the primate brain's signal. This complex sequence led the authors to conclude that their model "is closer to the anatomy and the physiology of the visual cortex in terms of quantitative parameter values" than any previously proposed ones.

The authors aptly note that the measures taken in the stations of the model fit “basic facts about the cortical mechanisms of recognition,” attributing specific roles to each station and successfully simulating the functions of each of these stations. This remarkable success, coupled with other properties of the model, qualify it as a theory of the anatomicophysiological properties of the brain that produce early vision. Thus, in addition to technological advantages that such networks have proven to have, the level of empirical detail that this work provides speaks for the biological viability of its model, even if it is restricted to the earliest stages of vision. This is no small matter.

This successful construction of a network model for aspects of the visual domain, along with its validation with highly specific sets of brain data, was followed by many more, through the use of more advanced models.² A deep neural network-based scientific theory of visual perception and its neural bases seems to be emerging, and it seems to be composed of relatively small neuronal chunks that carry out the computation.³

Constitutive Principles for Neural Modeling of Perception and Cognition

Investigations of the visual system in the animal kingdom appear to be in a leading position in both experimental and theoretical neuroscience. It is their standards that of modeling we neurolinguists should seek to meet. Indeed, it seems that some principles they formulate might help us in our journey to model the language system in the brain. Several recent programmatic papers seek to articulate and implement a “neuroconnectionist” agenda.⁴ They call for “explicit mapping between [artificial neural networks] and biology” for the explanation of behavior with neural networks. One of them, by Stanford philosopher/neuroscientist Rosa Cao and neuroscientist Dan Yamins, posits two general requirements for “successful explanatory models in cognitive and systems neuroscience: (a) the variables in the model correspond to components, activities, properties, and organizational features of the target mechanism . . . and (b) the dependencies posited among these variables in the model correspond to causal relations among the components of the target mechanism.”

It is easy to agree with these requirements. It is equally easy to see how the feedforward network model for vision described above meets them and even goes beyond them. Its variables correspond to components of the mechanism, and there is a causal relation between these and the function;

moreover, it identifies anatomical similarity between the model's structure and direction of information flow and what is known about the brain's relevant cortical tissue, nuclei, and pathways. This feedforward network, then, *is* a scientific theory of early vision in the primate's brain.

With this in mind, let's move on to human language, where matters quickly become more complex, if not problematic.

Can LLMs Model the Linguistic Brain, and Do They?

The program for theoretical cognitive neuroscience that we just saw might help to distinguish a technological advancement from a neurobiological theory. It would therefore be wise to use it in neurolinguistics. What would the requirements spelled out by Cao and Yamins mean in the context of language? I imagine that we would need to propose neural mechanisms from which linguistic functions emanate and back up these proposals by experimental evidence of some sort. At a minimum, a theory is needed, one that would show how our structured language behavior is caused by the activity of a collection of neural mechanisms. The "pieces of language in pieces of brain" dictum would be cashed out with detailed mechanisms that causally explain linguistic functioning via an appeal to neural networks. I believe wholeheartedly in this program and am hopeful about its success; but I am not sure that we are that close to realizing it. Nevertheless, the success of LLMs has led to attempts to apply them to neurolinguistics, with the expressed hope of converging on a mechanistic neural model for language. The previous chapter showed that the path towards such a model is long and tortuous, but the excitement about LLMs has motivated some to take the risk. Here, I focus on two ambitious projects that have recently applied LLMs to the study of language mechanisms in the human brain, in an attempt to see whether they qualify as scientific theories of language or at least the beginnings thereof.

Two recent series of works on LLM-brain-language relations have received much attention, and I will briefly describe and then critique them. The first, by Princeton's Ariel Goldstein and Uri Hasson, compares aspects of the language performance of various LLMs to human performance and plainly argues that "the human brain processes incoming speech similarly to an autoregressive [deep language model]"⁵ and that Goldstein and Hasson's results provide evidence that LLMs such as GPT-2 "capture subtle statistical dependencies reflecting syntactic, semantic, and pragmatic relationships

among words” and hence seem to “support this shift from a rule-based symbolic framework to a vector-based neural code for processing natural language in the human brain.”⁶

The second series, led by Evelina Fedorenko at MIT, makes somewhat similar comparisons and claims to have discovered “the language network,” which differs from previous language-to-brain mappings.⁷ Fedorenko’s group has been studying the properties of this rather large brain volume, in quest of a neural theory of linguistic cognition. As we have seen before, looking at experiments through a magnifying glass is sometimes critical to their evaluation, so please bear with me, as I am quite convinced that taking a minute and look more closely at these two projects would pay off.

LLM–Language–Brain Relations Seem to Converge on Next Word Prediction

At the center of the Princeton project are measures taken from GPT-2, from behaving humans, and from the brains of epileptic patients undergoing a complex procedure similar to that undergone by neurosurgeon Edward Chang’s patients (who helped to discover the neural substrate for abstract phonetic features; see chapter 7): surface electrodes are placed directly in clinically predetermined locations on patients’ cortices, and as part of the clinical battery, their brains’ electrical activity is recorded while they are performing a next word prediction task. The collected data enables a comparison between human behavioral responses, as well as their brains’ and the LLM’s. Similarity in performance would result in a high correlation between all three measures, and indicate that they predict the next word to the same degree.

A similar program, but with different test materials, was adopted by Fedorenko and her colleagues. Instead of an unconstrained recorded text, they devised a “localizer” task (inspired by an earlier fMRI test battery that has been said to quickly localize face recognition in each individual participant).⁸ Like its face predecessor, this program aims to capture brain areas that support receptive language mechanisms and to set them apart from areas from which nonlinguistic perception and cognition emanate. Fedorenko’s language task is said to be efficient and general, identifying in each participant a function that resides in roughly the same anatomical region. Participants hear or read meaningful sentences, and also controls, that are typically meaningless sequences of words or word-like syllable strings. This

enables Fedorenko to zero in on a similar set of activated regions in a matter of minutes. The difference between test and control conditions, so goes the reasoning, isolates “sentence processing” in the brain. A second step studies the properties of these areas in detail, with the hope of testing specific models against the measurements that the “localizer” task has obtained.

Indeed, in a series of publications, Fedorenko and her team describe a test battery that they presented to healthy participants (and in some instances, to neurosurgical patients with electrodes installed on their cortex as I already described). Like the Princeton group, they expect to find correlations between the responses of GPT-2, the brain, and behaving humans.

Preliminary Critique: Methodological Requirements Are Not Met

Before getting into the actual results, let’s pause to reflect on possible outcomes of these projects and what they might mean. Suppose that the language tasks probe all or at least most aspects of the broad range of our linguistic functions. Suppose further the best possible results are obtained and all the participants—the behaving humans, their active brains, and GPT-2—predict the next word perfectly, with perfect between-group correlations. Would such results lead to the conclusion that the machine is a scientific theory of language mechanisms in the human brain? Both groups of researchers seem to answer this question in the affirmative. And yet, even with such ideal (and imagined) results, Cao and Yamins’s requirements for a theory are far from being met: excellent predictions and perfect correlations would nonetheless give no evidence that GPT-2 “correspond[s] to components, activities, properties, and organizational features of the target mechanism[s]” in the brain, nor would “the dependencies posited among these variables in the model correspond to causal relations among the components of the target mechanism.” That is, even if all stars were to be properly aligned, we would have no analysis of the brain’s or the model’s *components* of language, although as we have already seen (in chapter 7), multiple results suggest that the linguistic brain is divided into grammatically structured pieces.

I believe that my final conclusion can already be anticipated: namely, that the way both groups probe the linguistic functions of their participants (healthy humans and neurosurgical patients, their brains, GPT-2) stops short of providing a realistic picture of their language capacities. Still, these studies are serious, they deserve careful scrutiny; moreover, I haven’t told you what their findings are. This is where we go next.

The Princeton Studies

A pivotal element in these studies has obviously been the task that participants were carrying out. It was a recorded story, presented to the three groups (via headphones to the neurosurgical patients and in written form to the healthy humans and to GPT-2). The patients listened to the story passively, and the implanted electrodes recorded their brain activity as they were listening; the healthy participants viewed the same story on a screen, through a moving window, and performed a next word prediction task, with a fixed length moving window (that is, participants saw a fixed number of words from the text and were asked to predict the next one; then, they were shown a new word, and the earliest word they had seen vanished; and so on). GPT-2 was given a transcript of the story, likewise presented through a moving window; it returned predictions for each word. The researchers' analysis focused on the predictions, as measured at a time point just prior to a word's appearance.

Human participants made correct predictions at a 28% rate; GPT-2 performed better, and it naturally improved when given access to longer word strings before making its prediction, which lowered the entropy of incoming words. In about 10% of the electrodes placed on patients' brains, the levels of electrophysiological activity recorded during and immediately prior to the appearance of a word were correlated, as shown in figure 8.2. This correlation was interpreted as evidence for next word prediction. All three measures—brain, behavior, and GPT-2—were correlated.




Figure 8.2

The position of 160 electrodes on the brain's surface (pooled from all 9 patients) that displayed a signal increase for the correct next word (or rather, its word embedding). The correlation goes from yellow (low) to red (high). Reproduced by permission from A. Goldstein, Z. Zada, E. Buchnik, M. Schain, A. Price, B. Aubrey, S. A. Nastase, et al., "Shared Computational Principles for Language Processing in Humans and Deep Language Models," *Nature Neuroscience* 25, no. 3 (2022): 369–380.

(continued)

The second study used the same data set to focus on three patients and provided a more detailed description of the neural code in which these words were expressed, as extracted from electrodes placed in Broca's area. It showed similarities between this code and the code produced by GPT-2 for these words. It also showed that signals extracted from adjacent brain areas are much less similar to GPT-2's and that relations between words that are expressed by a "symbolic" code are also less similar to what the brain produced. The study concluded that "deep language-model-induced representations of linguistic information are more aligned with brain embeddings sampled from [the inferior frontal gyrus] than symbolic representation."

In sum, these highly sophisticated studies found correlations between human behavior, machine, and brain activity, when responses were evoked by a text. If valid, these claims should lead to the rejection of grammatical representations as the correct description of the brain's linguistic code. Before scrutinizing these claims more carefully, let me move to the MIT set of studies whose authors draw similar conclusions.

Like the Princeton group, these MIT researchers sought a connection between the output of various LLMs, the activation pattern recorded from the collective functional region at issue, and associated behavioral patterns.⁹ The correlation between the scores of GPT-2, brain activity, and behavior was strikingly high and, moreover, higher than that of any other LLM of the many that they tested. On the face of it, this is a remarkable set of results, which led them to claim that LLMs "are beginning to approximate the brain's mechanisms for processing language."

Let me pause to reflect on the broader significance of these claims: if true, they go way beyond the Turing test. These researchers claim not only that the model they have simulates human language behavior, as Turing had modestly hoped, but also that it provide the *true* neurological model of how the human brain brings about this simulation. That's a very tall order, which deserves careful scrutiny.

The MIT Studies

Fedorenko and colleagues contrasted meaningful sentences with lists of scrambled words and meaningless syllables. The brain areas that responded more strongly to the former were marked in each individual participant, then pooled together. They dubbed the resulting contrast map “the language network,” seen in figure 8.3.¹⁰




Figure 8.3

Scores of an fMRI experiment with healthy adult participants, who were given a task that involved reading meaningful sentences, and contrasted with the scores of their reading of meaningless sequences of pronounceable nonwords. Depicted is (A) a map that combined the scores of the whole group of participants (hence the fuzzy borders) and (B) the map of an individual participant. Reproduced by permission from I. Blank, N. Kanwisher, and E. Fedorenko, “A Functional Dissociation Between Language and Multiple-Demand Systems Revealed in Patterns of BOLD Signal Fluctuations,” *Journal of Neurophysiology* 112, no. 5 (2014): 1105–1118.

The group map (A) shows participants whose scores converged on a very large area of the left hemisphere (and displays just 68% of the participants, who were at the maximal value, meaning that this group map excludes about one third of the participants that had been tested). The boundaries of this area are murky, because it is drawn on the basis of a large group of participants. Right-hemispheric parts of the alleged “language network” are not shown; an individual map (B) obviously has much sharper borders.

Next, Fedorenko and colleagues went looking for various cognitive properties of the marked area. For example, they attributed “formal linguistic competence” to their “language network,” claiming that it “is sensitive to linguistic regularities” and arguing that this pattern of activation contrasts it with other brain networks serving other cognitive functions.¹¹

The Main Critique: LLMs Are Not a Theory of the Linguistic Brain

Both groups have expressed optimism and excitement about LLMs as brain models for language. I wish I could be as positive. I contend that, contrary to their claims, the works discussed above provide no evidence that any current LLM models the brain's language processing device. In addition to the preliminary critique above, the main critique consists of several points. First, the neurolinguistic benchmarks by which these groups seek to sample performance are impoverished—the simple tests they deploy are plainly inadequate; second, LLMs may be good technologies, but they stop short of delivering a scientific theory of human language function; third, the studies in question display very poor anatomical precision, despite the existence of a rich neurolinguistic record, which points to a need to break language down into its component parts. This record is left unaccounted for. These are potentially devastating criticisms, so I'd better be a bit more detailed.

A. Neurolinguistic Benchmarks Are Far from Adequate

Both groups have been using language tests to probe human, machine, and brain activities. But what exactly are these tests, and what can we learn from the fact that they gave rise to the next word prediction behavior? We have already seen the critical importance of language test structure for the probing of human behavior and brain activity as well as machine behavior. Throughout this book we have seen both the necessity and usefulness of linguistically structured tests, leading to the conclusion that they are the ultimate evaluation tool for linguistic ability. In the context of the brain, we saw it through the remarkable success of Chang's neurosurgical group in identifying neuronal clusters that respond to abstract phonetic features. These researchers presented their participants with a highly structured corpus, which consisted of *the complete phonetic inventory of American English*. Without such a corpus, their conclusions would be unwarranted. As we have seen in previous chapters, scoping the broad range of the brain's sentence processing abilities is a truly formidable task. An informative "localizer" must include good representatives of this whole range of sentence probes (and their controls). Thus, testing the totality of neural processes that underlie sentence processing is a complex project that involves, at a bare minimum, a wide range of sentence types, chosen in a manner that is both theoretically and empirically reasoned: structural choices must be

made; the relative proportion of each sentence type within the test needs to be properly evaluated; and test contrasts have to be constructed by control conditions, resulting in a well-structured, comprehensive test that isolates the cognitive component at issue and at the same time is short. We must therefore ask: did the above-discussed work on sentence processing in the brain employ the same high standard as the work on phonetics? Unfortunately, the answer is negative. The scope of the language materials is very narrow, and their structure stops short of delivering a proper test.

The Princeton group claims that its LLM-based method “can capture many aspects of the latent structure of human language, including syntactic trees, voice, coreferences, morphology and long-range semantic and pragmatic dependencies.” But this claim is backed by no evidence, because syntactic trees, voice, coreferences, morphology, and long-range semantic and pragmatic dependencies are not tested in any of their studies, whether directly or indirectly; nor is there an indication as to how such relations could be tested in their paradigm.

Instead of following the path I charted above, the Princeton test amounts to an auditory presentation of a famous segment of National Public Radio’s popular *This American Life* called “Monkey in the Middle,” which tells an amusing story of a monkey who pushed the button of a photographer’s field camera and took a selfie that ended up being widely distributed and resulted in (rather juicy) legal disputes. Curious as this story may be, its inventory of utterances does not comprise a representative sample of most, let alone all, aspects of language. It is, rather, an unselected text that mostly consists of active declarative sentences: *Our story begins deep in the rainforests of Indonesia on an island called Sulawesi. A few years ago, the photographer David Slater traveled there from his home in England to photograph a troop of monkeys.*

The MIT researchers claim that “participants read semantically and syntactically diverse sentences,” contrasted with meaningless sequences of words or syllables of roughly equal length. In some cases, they were given “naturalistic” texts—stories found on the web. A review of the sentence materials, however, indicates that the variety of sentences used was extremely limited, both syntactically and semantically, as it mostly consisted of simple active declarative sentences (with an occasional embedding). For example, *Senator Hubert Humphrey is obviously a man with a soul and heart*, contrasted with meaningless sequences such as *quey etoa pinning strosh chsture en freir polyll denumpments ro veep urper*. The stories used are short and simple, and the controls are nonsensical sequences of syllables. A simple design like this is an advantage, at times,

(continued)

but not here: if one seeks to study the manner by which the brain processes language in all its complexities at a level of sophistication that is similar to that of vision neuroscience, then the bar must be higher: for these standards to be met, simple and complex syntactic types must be used, in a manner already discussed.

Both groups' language materials, which aimed to evoke behavioral, neural, and machine responses, actually consisted of arbitrarily selected sentences. As such, the nature of their language localizer is not comparable to the theoretically motivated selection of materials that is customary in vision, audition, or phonetics. No wonder, then, that in the Princeton studies, a random collection of electrodes throughout the brain was modulated by the input and that, in the MIT series of studies, the brain area activated was huge and bihemispheric, localized in a set of regions whose borders converge on no known anatomical landmarks.

The Princeton studies concluded that, compared to network code, "symbolic" representations produced weak brain-machine correlations. But upon careful examination, their "symbolic" representations contain no syntactic or semantic rules but are, rather, representations of arbitrarily assigned categories, developed and used for automatic word tagging. These do not contain rules, nor are they stated in a format that is part of any known linguistic theory. They may eventually turn out to be the correct ones. However, at this point, basing an analysis on this classification method cannot lead to a conclusion regarding the advantage (or lack thereof) of one system of representation to another.

The unsystematic sentence samples modulate bihemispheric clusters that are not only too large, but also . . . too small: the test materials lack examples of many syntactic and semantic components that are essential for human communication. For example, they hardly contain any questions, and they do not properly represent imperatives, quantifiers, modal predicates, embeddings, displacements, ellipses, and a host of intrasentential dependency relations. Nor do they evince a systematic choice of phonetic and phonological materials. *This set of tests stops short of selectively localizing any recognizable cognitive or linguistic component, let alone a syntactic one.*

The near uniformity of sentence types in these tests almost amounts to multiple presentation of the same sentence, from a grammatical perspective.

Imagine that you hear the same sentence structure one million times (*John loves Mary; The man saw the woman; Elephants eat grass; etc.*). Before long, your brain would get adapted to this Subject–Verb–Object sequence, and your behavior would align its expectations to it, as would your brain, which would likely make behavioral and brain responses different than if you heard a multiplicity of sentence types. Recall that the goal of both research groups has been to investigate sentence processing (the Princeton group even claimed that its model has the capacity to capture specific complex syntactic relations). What was actually tested turns out to be an extremely narrow range of structures that provides no empirical basis for such claims. Some readers may recall the movie *The Shining*, in which the protagonist, Jack Torrance, seems to be writing a novel, which turns out to consist of repetitions of the same sentence: *All work and no play makes Jack a dull boy*. Linguistically, the tests in both the Princeton and MIT studies bear some similarity to Jack Torrance’s novel: these researchers believe that they are testing language, while repeatedly testing structurally similar sentences.

B. LLMs Do Not Deliver a Theory of Human Language Function

My second criticism starts by questioning the suitability of GPT-2 as a model of human language, independent of brain implementation. Here, the picture is clear: based on the conceptual, methodological, and empirical arguments I gave in chapters 5 and 6, no LLM thus far created can qualify as a theory of human language—these models are unconstrained and have proven unsuccessful in correctly classifying complex cases as $\pm$ grammatical.¹² We therefore *know* that GPT-2 is not a candidate theory. Digging deeper, we see that the specialized attention heads in GPT-2 classify the grammaticality of sentences with MOVE at a mediocre level (under 80% accuracy).¹³ These models thus do not qualify as scientific theories of brain function, at least not in their current form.

C. Brain Localization of Processing Components Is Far from Precise

My final point regards brain localization. Our brain is noisy, and language processing is a highly complex activity. Listening to the linguistic brain and deciphering its signal is, moreover, difficult. Still, when one compares past work on the localization of language functions to that of the Princeton and MIT studies, the anatomical resolution of the latter is low (as seen in figures 8.2 and 8.3), certainly when compared to the sharper picture

obtained by experiments that use more focused language materials: chapter 7 displayed relatively sharp images of small brain locations for specific components of language, aligned with microscopic anatomy, thereby opening the door to a detailed functional anatomy of a mosaic of syntactic and semantic processes. Most of these results, displayed at a much higher functional as well as anatomical resolution, are not discussed by either set of studies. This is no accident: for all we know, the fine linguistic distinctions that have been found in the brain cannot be accommodated by GPT-2 or by any of its successors.

Coda: Can LLMs Model the Brain's Language Mechanisms?

Much is required for an apt characterization of receptive language, at the empirical, methodological, anatomical, and linguistic levels. Chapter 7 concluded with a divide-and-conquer (or better yet, parcellate-and-model) agenda for neurolinguistics, which would begin with the parcellation of the microanatomy on the one hand and the function on the other hand. Overall modeling, such as that done by the direct implementation of LLMs, misses the point: not only does it fail to divide language into its components (as is necessary for both theoretical and empirical reasons), but also, it provides no tools for the discovery of anatomically identifiable pieces of cortical tissue that are entrusted with specific language functions. The prospects of such a global analysis are, therefore, rather grim. More focused projects that seek to model specific parts of language in the brain are more likely to succeed, similar to those on the visual and sensorimotor systems.

Taking Stock

This chapter was the second to get under LLMs' hood. We saw:

- Studies of LLM, human, and brain performance and the relation between their results
- That in the visual domain, methodological principles are well-articulated and models are constructed in keeping with what is known from brain neurophysiology and microanatomy
- That in language studies this is not always the case
- That all reviewed neurolinguistic studies that incorporate LLMs have used tests that are only weakly informative, mostly due to poor selection of localizer tasks and test materials
- That the conclusion that the brain's specialized underlying mechanisms are LLMs is premature, to say the least
- That the claim that a novel brain location of language mechanisms has been identified is largely unwarranted
- That a more viable strategy for the discovery and modeling of the linguistic brain is (a) parcellating both the function, using linguistic tools and the structure, using anatomical tools, followed by (b) deep network modeling of the parts

This is a section of [doi:10.7551/mitpress/15932.001.0001](https://doi.org/10.7551/mitpress/15932.001.0001)

How Deeply Human Is Language?

Chomsky, the Brain, and the AI Fantasy

By: Yosef Grodzinsky

Citation:

How Deeply Human Is Language?: Chomsky, the Brain, and the AI Fantasy

By: Yosef Grodzinsky

DOI: 10.7551/mitpress/15932.001.0001

ISBN (electronic): 9780262052023

Publisher: The MIT Press

Published: 2026



The MIT Press

The MIT Press
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
mitpress.mit.edu

© 2026 Massachusetts Institute of Technology

All rights reserved. No part of this book may be used to train artificial intelligence systems or reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data is available.

ISBN: 978-0-262-05200-9

EU Authorised Representative: Easy Access System Europe, Mustamäe tee 50, 10621 Tallinn, Estonia | Email: gpsr.requests@easproject.com