The Book Review Column [1]
by **Nicholas Tran** (ntran@scu.edu)
Department of Mathematics & Computer Science, Santa Clara University

1 Notable New Releases

Approximation Algorithms for Traveling Salesman Problems (Cambridge University Press, 2024) by Vera Traub and Jens Vygen (both of University of Bonn) presents the state of the art of approximation algorithms for the Traveling Salesman Problem (TSP), covering both classic and recent results as well as outstanding open problems.

Understanding Cryptography: From Established Symmetric and Asymmetric Ciphers to Post-Quantum Algorithms, Second Edition (Springer, 2024) by Christof Paar (Max Planck Institute), Jan Pelzl (Hamm-Lippstadt University) , Tim Güneysu (Ruhr University Bochum) focuses on modern applications of cryptography, with special emphasis on post-quantum methods.

The Elements of Computing Systems: Building a Modern Computer from First Principles, Second Edition (The MIT Press, 2021) by Noam Nisan (Hebrew University of Jerusalem) and Shimon Schocken (Interdisciplinary Center Herzliya) updates the popular textbook that inspires Nand-to-Tetris courses on how to build a general-purpose computer system from scratch. This new edition is divided into two standalone parts, Hardware and Software, and features four new appendices.

The Joy of Quantum Computing: A Concise Introduction (Princeton University Press, 2025) by Jed Brody (Emory University) is an accessible introduction to the subject requiring only calculus and no quantum mechanics background.

2 This Column

Dan Gusfield's latest book *Proven Impossible: Elementary Proofs of Profound Impossibility from Arrow, Bell, Chaitin, Gödel, Turing, and More* was named an American Library Association's Choice Outstanding Academic Title in 2024. Bill Gasarch recommends it to both laypersons and experts alike.

The latest edition of John Vince's *Foundation Mathematics for Computer Science: A Visual Approach* covers a wide range of the mathematics used in computer science, both continuous and discrete. David J. Littleboy finds it a useful reference for those needing a quick preview of the material or a review of concepts learned in the past.

Terrence J. Sejnowski's *ChatGPT and the Future of AI: The Deep Language Revolution* is an engaging demonstration of the capabilities and limitations of large language models such as

ChatGPT to the general audience with the help of ChatGPT itself. I took the cue from the book and asked ChatGPT to help correct my review's spelling and grammar.

3 How to Contribute

Don't let AI have all the fun; consider contributing a human-made book review to SIGACT News. Either choose from the books listed below, or propose your own. In either case, the publisher will send you a free copy of the book. Guidelines and a LaTeX template can be found at `https://algoplexity.com/~ntran`.

BOOKS THAT NEED REVIEWERS FOR THE SIGACT NEWS COLUMN

Algorithms, Complexity, & Computability

1. Vaze, R. (2023). *Online Algorithms.* Cambridge University Press.
2. Ferragina. P. (2023). *Pearls of Algorithm Engineering.* Cambridge University Press.
3. Downey, R. (2024). *Computability and Complexity: Foundations and Tools for Pursuing Scientific Applications.* Springer.
4. Traub, V., & Vygen, J. (2024). *Approximation Algorithms for Traveling Salesman Problems.* Cambridge University Press.
5. Brody, J. (2025). *The Joy of Quantum Computing: A Concise Introduction.* Princeton University Press.

Miscellaneous Computer Science & Mathematics

1. Grechuk, B. (2019) *Theorems of the 21st Century.* Springer.
2. Nisan, N., & Schocken, S. (2021). *The Elements of Computing Systems: Building a Modern Computer from First Principles, 2nd ed.* The MIT Press.
3. Chayka, K. (2024). *Filterworld: How Algorithms Flattened Culture.* Doubleday.
4. Valiant, L. (2024). *The Importance of Being Educable: A New Theory of Human Uniqueness.* Princeton University Press.

Discrete Mathematics and Computing

1. Ross, S., & Peköz, E. (2023). *A Second Course in Probability.* Cambridge University Press.

Cryptography and Security

1. Chen, K., & Yang, Q. (2023). *Privacy-Preserving Computing for Big Data Analytics and AI.* Cambridge University Press.

2. Paar, C., Pelzl, J., & Güneysu, T. (2024). *Understanding Cryptography: From Established Symmetric and Asymmetric Ciphers to Post-Quantum Algorithms, 2nd ed.* Springer.

3. Garfinkel, S. (2025). *Differential Privacy.* The MIT Press.

Combinatorics and Graph Theory

1. Landman, B., Luca, F., Nathanson, M., Nešetřil, J., & Robertson, A. (Eds.). (2022). *Number Theory and Combinatorics: A Collection in Honor of the Mathematics of Ronald Graham.* De Gruyter.

Review of [1]

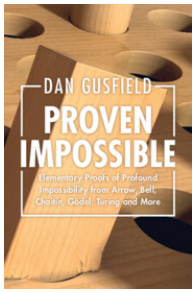

Proven Impossible:
Elementary Proofs of Profound Impossibility
from Arrow, Bell, Chaitin, Gödel, Turing, and More
by **Dan Gusfield**

Cambridge University Press, 2024
$79.99 Hardcover, $34.99 Paperback, eBook, 270 pages

Review by **William Gasarch** (gasarch@umd.edu)

1 That's Impossible!

Some people outside of math think that statements like *You cannot write a program for HALT* are defeatist and pessimistic. They do not realize that these are rigorous theorems, and that it is good to know what you can't do, so you can modify your goals.

So clearly the layperson needs a book that gives coherent explanations of problems that are impossible to solve. The current literature seems to be in three overlapping categories:

1. Books for the layperson that are too fluffy and don't really get to the point.
2. Books for the layperson that oversell, for example, claiming *Gödel's Incompleteness Theorem proves that humans are creative!* or *Quantum computing will solve world hunger!*
3. Technical articles for experts that are not helpful to the layperson, even if they give (allegedly) simpler proofs of known theorems.

So this is a book for the layperson. How would it be for readers of this column? I suspect that $\frac{2}{3}$ of the people reading this review will enjoy $\frac{2}{3}$ of the book.

The chapters of the book do not quite correspond to theorems on impossibility, since some such theorems have two chapters about them or relate to other chapters. Hence I review the book not chapter by chapter but impossibility topic by impossibility topic.

2 Bell's Theorem

When quantum mechanics was first studied, the question *Can we model this using classical physics?* arose. This question would seem hard to formalize. Nevertheless, Bell's Theorem does just that: classical physics is formalized, and it is shown that quantum mechanics cannot be so described.

The book's explanation of Bell's Theorem is excellent. Amazingly it does not require knowing any quantum mechanics. The layperson will benefit; however, I suspect that my readers who don't live and breath quantum mechanics (that is, most of them) will also benefit from this chapter.

[1]

I note for future reference that Bell's Theorem is a profound statement about how our universe works. I will later comment on the profundity of the other chapters; however, none will top Bell's Theorem.

There are two chapters on Bell's Theorem.

3 Arrow's (and Friends) Theorem

As the readers of this column probably know, Arrow's Theorem states that, assuming reasonable assumptions about how an election can be run, the only system that satisfies them is a dictatorship. The *and Friends* in the title of the chapter refers to Gibbard-Satterthwaite Theorem, a similar result which uses different (more natural, according to Gusfield) assumptions and has a simpler proof.

Since voting is a man-made phenomenon, the results here cannot be as profound as Bell's Theorem. Also, there could be more discussion of what real voting systems look like and how they seem to not lead to dictatorships.

That said, this chapter gives a clean exposition of both Arrow's Theorem and Gibbard-Satterthwaite Theorem. The results are interesting. There is some discussion of why the real world seems to manage with voting despite these theorems (the consistency assumption is suspect, though it looks fine to me).

There is one chapter on voting theorems.

4 Clustering

Clustering is a grouping of data in the form of a set of points in the plane (or in d-dimensions) that will (hopefully) lead to a meaningful statement (e.g., *Basketball players tend to be tall.*) or a useful statement (e.g., *If we have to split up a club into smaller clubs, here is a way to do it so that close friends are in the same club.*).

Kleinberg gave three axioms about what a good clustering algorithm should achieve, and then showed that no clustering algorithm could achieve all three. This is reminiscent of Arrow's Theorem. In fact, once again one of the troubling axioms is about consistency.

Since clustering is a man-made phenomenon, the results here cannot be as profound as Bell's Theorem. However, in this case (unlike the chapter on voting) there is a discussion of why this theorem does not seem to be a problem in the real world of clustering.

Unlike Arrow's Theorem, this result is likely new to the reader (at least it was new to me). The presentation is clear and has the advantage of the proof being simplified since Kleinberg.

There is one chapter on clustering theorems.

5 Gödel's and Chaitin's Incompleteness Theorems

The readers of this review likely know Gödel's Incompleteness Theorems. Even so, it is good to revisit it from time to time. In addition, Chaitin's Theorem (which we discuss below) is related.

The first chapter on Gödel's Theorems contains some simple Gödel-ish statements and a proof of a simpler variant of the first incompleteness theorem. The second chapter on Gödel's Theorems contains the following more technical versions; indeed, this is the most technical chapter in the book.

1. *Gödel's First Incompleteness Theorem.* I paraphrase Theorem 9.5.1 (p. 207):

   Suppose a formal system Π *contains the language of arithmetic* $\mathcal{L}_A$*, and that* Π *is sound. If a particular subset of Gödel numbers is expressible in* $\mathcal{L}_A$*, then there is a true sentence that can be written in* $\mathcal{L}_A$ *that cannot be derived in* Π*, so* Π *is incomplete.*

   I prefer the following simpler version: *In all axiomatic systems commonly used in mathematics there are statements in math that are true but not provable in that system.* However, the formulation above serves to remind us that we need to state things carefully.

2. *Gödel's Second Incompleteness Theorem.* I quote the book exactly (p. 215):

   Any rich enough formal system Π *that can derive (inside* Π*) a statement implying that it is consistent is in fact inconsistent.*

   I *really* prefer the following simpler version: *If a system can prove its own consistency, then that system is inconsistent.* However, the formulation above serves to remind us that we need to state things carefully.

Gödel's Incompleteness Theorems have not affected how non-logicians go about their business. Why? Because the statements that Gödel proved were true but not provable were not statements of interest. They were contrived for the sole point of being true but not provable. One could argue that the Second Incompleteness Theorem is less contrived: one would want to prove that (say) ZFC is consistent. Even so, that is a concern of logicians.

Chaitin's Theorem gives a possibly less contrived example. We state it informally.

Let $\mathcal{L}$ *be a programming language. There is a constant* $u_{\mathcal{L}}$ *such that for every string* x *the statement "The shortest program in* L *that prints* x *is of length* $\geq u_{\mathcal{L}}$*." is not provable in any consistent formal system.*

Informally Chaitin's Theorem says that some strings (actually an infinite number of strings) do not have short descriptions, *and* this cannot be proven in any formal systems. In other words, it says that proving that a string is complicated is complicated.

These theorems are profound statements about mathematics.

There are two chapters on Gödel's Incompleteness Theorems and one on Chaitin's Theorem.

6 Turing Undecidability

The readers of this review likely know that HALT is undecidable. Even so, it is good to revisit it from time to time. I'll point out two things in this chapter that are of interest both to the layperson and the readers of this review.

1. Its not just HALT. Rice's Theorem states that *all nontrivial* properties of programs are undecidable.

2. The undecidability of HALT can be used to prove Gödel's First Incompleteness Theorem.

The chapter does a good job of leisurely explaining what HALT is, showing that it's undecidable, and exploring some consequences of this. The chapter does not say how this affects actual programmers and what they do about it.

That HALT (and other problems) are undecidable is a profound statement about computation.

There is one chapter on HALT.

7 Opinion

This is a great book both for the layperson and for people who already know some of the contents (the readers of this review are likely in the second category). So perhaps buy one for yourself and one for your math-inclined great-niece.

That said, here are some comments that are... not quite negative, but need to be said.

1. The topics tackled are of two types:

   (a) Those that say something profound: Bell's Theorem, Gödel's Incompleteness Theorems, Chaitin's Theorem, HALT is undecidable.

   (b) Those that are about man-made phenomena and hence, by their nature, are just not that profound: Arrow's Theorem on voting, Kleinberg's Theorem on clustering.

   I found the profound chapters to be more interesting.

2. I would have preferred to see more on the question *Once it is known that X is impossible, what happens next?*

3. I note the following omissions. This is *not* a complaint, since if the author puts in everything that could be put in, he would have a 1000-page book (that is why calculus textbooks are so big).

   (a) The three problems of antiquity: constructions of trisecting the angle, doubling the cube, squaring the circle. These are all impossible. There is a great book on these for the layperson: *Tales of Impossibility: The 2000-Year Quest to Solve the Mathematical Problems of Antiquity* by David Richeson. It was reviewed in a SIGACT News Book Review column here:

   `https://mathcs.clarku.edu/~fgreen/SIGACTReviews/bookrev/53-1.pdf`

   (b) P vs. NP. There is a great book on this for the layperson: *The Golden Ticket: P, NP, and the Search for the Impossible* by Lance Fortnow. It was reviewed in a SIGACT News Book Review column here:

   `https://www.cs.umd.edu/~gasarch/bookrev/44-3.pdf`

   (c) Independence of the Continuum Hypothesis. I do not know of any account of this for the layperson. There could be a book with some history and context; however, it would be hard to do any real math. There is a book with a chapter on the history and context of this problem: *The Honor Class: Hilbert's Problems and Their Solvers* by Ben Yandell. It was reviewed in a SIGACT News Book Review column here:

   `https://www.cs.umd.edu/~gasarch/bookrev/44-4.pdf`

   (d) Unsolvability of the quintic. I know of two books for the layperson on this problem, though I have not read them: (1) *Evariste Galois 1811-1832* by Laura Rigatelli, and (2) *The Equation that Couldn't Be Solved* by Mario Livio.

I re-iterate that this is a great book both for the layperson and for people who know some of the material.

Review of [1]

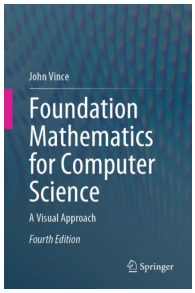

Foundation Mathematics for Computer Science
A Visual Approach
John Vince

Springer Nature Switzerland, 2024
$69.99 Paperback, $59.99 eBook, 342 pages

Review by
David J. Littleboy (djl@alum.mit.edu)
Technical translator, retired
Tokyo, Japan

1 Overview

This book is a wide-ranging overview of much of the mathematics used in computer science from the basics through, and with an emphasis on, computer graphics. The coverage of the large numbers of topics means that said coverage is shallow. Some topics (Fermat's Little Theorem, continued fractions) are just tossed in with little or no discussion. This trade-off has both pluses and minuses. Everything mentioned is something that a computer science major should be familiar with, yet insight into its nature and significance will require looking elsewhere. This book emphasizes calculation, with worked examples at the end of every chapter. However, it does not provide homework problems.

2 What's in the Book?

This is a long, heavy book. It has 26 chapters and two appendices in 638 pages. Which means that each chapter is quite short. The first nine chapters, through page 177, cover the basics, numbers from an introductory mathematical perspective and then from a computational perspective, logic, then combinatorics, probability, and statistics, and then a chapter on modular arithmetic. Starting with chapter 10, the book leaves discrete mathematics behind and moves on to the mathematics used in computer graphics.

Despite the length and number of topics, Vince does not cover the mathematics used in cryptography, AI, or quantum computing.

3 Criticisms

I found a lot to criticize here, so let me start by saying that I think there's also a lot to like here. The approach of presenting the basics, throwing out some complicated ideas without much explanation,

and grinding through examples makes this book useful to people needing a quick preview of the covered material, or as a review of material learned or seen in the past. But.

The book has an enormous flaw for an American audience: it is written in British English. Thus exponents are "indices", GCD is "highest common factor", and "revises" means reviews. The second problem is that the brevity required by the amount of material covered means that explanations are often inadequate or missing. It's an overview book, not a textbook.

I think Vince has sold himself short here. His informal introductions to and comments on technical terms/issues are often excellent, but the lack of follow-through leaves the reader with unanswered questions. I recently ran across two incorrect statements of the proof that there are an infinite number of primes: they fail to notice that one plus the product of a set of primes might be divisible by some prime not in that set. Vince not only states the proof correctly but also provides a table showing examples of one plus a product of primes being prime and of being composite. More of this sort of thing would improve the book.

4 Quibbles

I found myself wanting many sections rewritten and expanded. Again, the trade-off between brevity and coverage vs. depth forces this to be the case. There is much more that could be said, at least for the basic material, both in descriptions of the material and in explaining the use of that material.

For example, the section on number systems makes perfect sense to someone who has programmed in assembler on various architectures but needs more explanation of why doing arithmetic in various bases and complement notations is necessary.

Presumably, a reviewer who works in computer graphics would feel the same way about the latter half of the book.

Were this book divided into two and written with more depth, then specific quibbles would be worth discussing in a review, but the point of this book is to get the terms out there. Vince's intuition on what should be said in a wide-ranging overview largely agrees with mine, and my quibbles are mostly a desire for more depth.

5 Conclusions

I can't see this book being used as a textbook for an undergraduate course in the United States: the use of British terminology alone is fatal. Also, the lack of homework problems means it's not a textbook in the usual sense. But a student planning on taking a course on any of this material would find their time well spent reading this book beforehand, especially since it presents examples of calculations. It would also certainly be useful as a review of material studied years ago and not used since.

Review of [1]

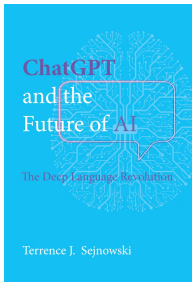

ChatGPT and the Future of AI
The Deep Language Revolution
Terrence J. Sejnowski
The MIT Press, 2024
$22.95 Paperback, $22.99 eBook, 264 pages

Review by
Nicholas Tran (ntran@scu.edu)

1 Overview

Aimed at the general audience, this book is a field report on the current state of the art in artificial intelligence, focusing on the technical development of large language models (LLMs) and their potentially profound impact on society and humanity. It presents a fascinating demonstration of the capabilities and limitations of modern chatbots such as ChatGPT through a series of conversations with them about their own nature and about the book itself (e.g., there is a ChatGPT-generated summary at the end of each chapter). An early pioneer in neuroscience, the author provides a front-row view of the development neural networks, leavened with engaging personal anecdotes as well as his unique perspective on the future of AI research and its synergy with brain research.

2 Summary of Contents

Preface introduces large language models (LLMs) as intelligent chatbots that can assist with a wide range of language tasks previously performed only by humans. Sample unedited dialogs with ChatGPT are presented to showcase its capability to gather and summarize vast amounts of data *about LLMs* in conversational form, which may contain errors. Their performance prompts the question of whether LLMs understand or are conscious and leads the author to explore the nature of intelligence and consciousness.

Part I: Living with Large Language Models

Chapter 1 explains that the arrival of LLMs marks a new stage in the Information Revolution when machines can form internal representations of data provided to them, which is then used to generate novel human-like responses to novel queries. This new technology will likely to produce scores of new jobs and industries as its predecessor, the Internet, has.

Chapter 2 provides a glimpse of how LLMs might revolutionize existing industries, such as generating and analyzing doctors' notes in healthcare, reducing teacher workload in education, drafting contracts in law, analyzing sentiments in natural language processing, generating images, music, videos, or even digital influencers in entertainment. It also notes concerns that must be addressed in such applications, such as bias, misinformation, and loss of privacy or security.

Chapter 3 presents four remarkable conversations with LLMs that expose their sophistication as well as their weaknesses. The two conversations with LaMDA by Google could be interpreted as evidence of its ability to model human social interactions and to explain and defend its sentience. The other two conversations with ChatGPT by OpenAI portray a chatbot going off-the-rails with its profession of love to the human interlocutor and its authoritative-sounding answers regarding non-existent events or impossible scenarios.

Chapter 4 explains the outcomes of the four conversations in Chapter 3 in terms of priming, the process of providing LLMs with context and examples (called prompts) to limit their responses to those derived from a particular subspace of their training data.

Chapter 5 explores various views of intelligence, thinking, and consciousness and tests some of them against LLMs. It speculates that the mirroring of their interlocutors by LLMs could be an important component of intelligence and suggests basing future theories of intelligence, thinking and consciousness on a mathematical theory of LLMs.

Part II: Transformers

Chapter 6 summarizes the evolution of neural network models, beginning with the perceptron modeling a single neuron, to multilayer networks using backpropagation, to deep learning networks with hundreds of layers. Transformers are deep learning neural networks for language processing that transform words to numeric vectors to capture semantic, positional and contextual information among them. They are trained to predict the next word in a sentence based on the previous words. The author explains the architecture of transformers and their training process, drawing parallels with the evolution of the human cerebral cortex.

Chapter 7 attributes the unexpected success rate of the stochastic gradient descent algorithm used for training neural networks, even when overparameterized, to different dynamical properties of high-dimensional spaces. The author suggests that developing new mathematics for high-dimensional spaces will be useful to understanding LLMs and brains alike.

Chapter 8 surveys the current rush to invest in data centers, networks, and research and development of chip technologies to support the energy-intensive training and deployment of LLMs. The author predicts a law similar to Moore's Law will reduce the energy demand of LLMs exponentially over time and points to the human brain as proof of the possibility of portable LLMs.

Chapter 9 discusses the anxiety and uncertainty among experts in artificial intelligence regarding the possibility of creating a superintelligence that could lead to human extinction. The author points out the drawbacks of *not* continuing AI research and suggests that humanity has learned to manage the risks of new technologies in the past and can do so with AI as well.

Chapter 10 reviews inchoate efforts by the European Parliament and US Congress to regulate AI, a task that is complicated by the competing interests of different stakeholders. Issues such as copyrights and authorship of LLM-generated content will need to be addressed in the future through discussions or in court.

Part III: Back to the Future

Chapter 11 reviews the ascendance of the neural network model over the symbolic model in AI research and argues that the success of the former is due to its similarity in architecture to the human brain.

Chapter 12 proposes new features for LLMs, drawing inspiration from brains, such as long-term memory, working memory, sensorimotor systems, a reward system and others.

Chapter 13 discusses the synergy between AI and brain research, suggesting that the former can aid the latter by providing new tools for analyzing brain data and by proposing new theories of brain function. The author also suggests that the brain can inspire new AI models.

Chapter 14 reviews early scientific advances obtained with or inspired by LLMs at the dawn of AI and predicts rapid progress ahead. The author calls for a scientific theory of LLMs that can explain their behavior and guide their development.

Afterword briefly mentions a recent development in the transformer architecture. **Acknowledgments** thank the author's colleagues, students, and collaborators, including ChatGPT.

3 My Opinion

This book has two aims: to provide a popular account of the development of large language models and to demonstrate how they can assist in carrying out a deeply human task, namely, writing a book. I believe the author has achieved both objectives admirably. Chapters 3 and 4 are especially useful to the general reader in understanding how to use LLMs effectively, while Chapters 6 and 7 do a commendable job in explaining in plain English the architecture of transformers and why they succeed beyond expectations. The author's personal anecdotes and reflections on the future of AI research and its synergy with brain research are engaging and thought-provoking; they are welcome additions to a growing body of retrospectives by pioneers in the field.

The book reads like a first draft at places, especially the unedited end-of-chapter summaries by ChatGPT. I think the lack of editing is intentional to highlight the technology's limitations. Additionally, I would have preferred to see the dialogs differently, e.g., **Dialog 12.2** instead of **GPT 12.2**, since some of them are with LLMs other than ChatGPT and the numbering could be confused with the OpenAI chatbot's version number.

I highly recommend this accessible book to those wishing to gain a good overview of neural networks and machine learning for further reading on the subject.