



Information Research – Vol. 31 No. 2 (2026)

Large language models

Review of Raaijmakers, S. (2025) *Large language models*. Cambridge, MA: The MIT Press. ix. 292 pp. ISBN 978-0-262-55269-1

DOI: <https://doi.org/10.47989/ir31265442>

It is only four years since ChatGPT introduced generative AI and large language models (LLMs) to the world and became the fastest growing app on the Web. Few at the time had any understanding of LLMs and how they were used and one would have to work through many Web documents to find the answers. Now Stephan Raaijmakers, Senior Scientist at TNO and Professor of Communicative AI at Leiden University, has produced a compact and accessible volume that fills the gap.

An effective feature of the book is its treatment of language not as a philosophical abstraction or literary construct, but as a computational phenomenon. This shift in perspective underpins the rest of the book: language is treated as a system of patterns and statistical regularities that can be modelled and exploited by machine learning systems. Such an approach demystifies the operation of LLMs, presenting them not as entities with human understanding, but as systems that generate plausible sequences of text based on probabilistic prediction.

The structure of the book follows a broadly logical progression. Initial chapters establish the conceptual framework before moving on to the architectures that underpin contemporary systems, including neural networks and transformer models. Subsequent sections deal with training processes and, importantly, the question of alignment. Here, the author distinguishes clearly between the large-scale absorption of textual data in pre-training and the subsequent shaping of system behaviour through techniques such as reinforcement learning from human feedback. He also points out the limitations of such approaches: human evaluators are inconsistent, biases are inevitably introduced, and systems may perform well under training conditions while behaving unpredictably subsequently.

The discussion of applications covers familiar territory: translation, summarisation, code generation and other forms of text production. These are balanced by an examination of well-known shortcomings, including the generation of erroneous or fabricated information, sensitivity to prompt formulation, and a general lack of robustness. The question of creativity is addressed with some care. While the outputs of LLMs may appear novel, the author emphasises their derivation from patterns present in training data, raising legitimate questions about whether such outputs can be said to constitute creativity in any meaningful sense.

In its later chapters, the book turns to broader societal issues, including the potential for disinformation, the impact on employment, and the concentration of technological capability within a relatively small number of corporate organisations. These topics are necessarily treated in a summary fashion, but they are sufficient to indicate the scale and nature of the challenges involved. As with earlier sections, the author avoids both alarmism and complacency, offering instead a measured account that recognises both the capabilities and the limitations of the technology.

This is a clear, well-structured and well-argued introduction to a complex and rapidly evolving field. It will be of value to students and practitioners in computing, information science and communication studies, and would be a useful addition to library collections concerned with developments in artificial intelligence.

Professor T.D. Wilson

April 2026

© [CC-BY-NC 4.0](#) The Author(s). For more information, see our [Open Access Policy](#).