

Knowledge Graphs in the age of Large Language Models

Aidan Hogan*

* Universidad de Chile, Santiago
<https://aidanhogan.com/>

Résumé. The public release of ChatGPT in November 2022 has led to the widespread popularisation and discussion of Large Language Models. Their fame (or infamy) is in large part due to the fact that not only do they encode a vast volume of information, they can also mimic our human language in unprecedented ways. These models thus call into question a fundamental tenet of Data Management: that in order to address non-trivial information needs, the first step is to explicitly structure data in order to lift them from the ambiguous swamp of our human language. This talk will address a pressing question that arises in this context: do we still need Knowledge Graphs in the age of Large Language Models? We will look at the current capabilities and limitations of Large Language Models, and contrast them with those of Knowledge Graphs. Ultimately we will propose that in many ways, Large Language Models and Knowledge Graphs can be considered complementary, with interleaving strengths and weaknesses. More specifically, we will argue that Large Language Models can help bring the power of Knowledge Graphs to the masses, and the power of the masses to Knowledge Graphs.

Summary

Aidan is a Professor and Vice-Director of the Department of Computer Science, University of Chile, and Vice-Director of the Millennium Institute for Foundational Research on Data (IMFD). His research interests primarily relate to the Semantic Web, Graph Databases, Information Extraction and Reasoning; he has published over one hundred peer-reviewed works on these topics. He is an Editor-in-Chief and Co-Founder of the upcoming journal Transactions on Graph Data & Knowledge (TGDK). He is the (co)author of three books, namely « Reasoning on the Web of Data », « The Web of Data », and « Knowledge Graphs ».

