

Tutorial at AJCAI 2022 - A Practical Guide to Knowledge Graph Construction from Technical Short Text



Dr Michael Stewart

Research Fellow

Theme 1
2022-12-05

Hyatt Regency Perth

Tutorial at AJCAI 2022 ("A Practical Guide to Knowledge Graph Construction from Technical Short Text"). Abstract below:

Have you ever wondered how to harness the significant volume of knowledge buried within unstructured text? Approximately 80% of all data in organisations is unstructured, a large portion of which exists in the form of technical language such as doctor's notes, maintenance work orders, and traffic reports. Natural Language Processing (NLP) provides the means to construct knowledge graphs from unstructured short text, enabling the querying of knowledge held within the text. Knowledge graphs are employed by a wide range of top companies – eBay, Walmart and Volvo to name a few. But what exactly is a knowledge graph? Why are leading companies actively building knowledge graphs, and how is one created?

This tutorial provides a practical guide to knowledge graphs. We will begin by providing an overview of graph databases, highlighting their unique advantages when compared to structured data models such as relational tables. We will then detail the underlying natural language processing techniques involved in knowledge graph construction from text, namely named entity recognition (NER) and relation extraction (RE). We will motivate the need for knowledge graphs via a simple, practical example in the maintenance domain. This Python notebook-based example will demonstrate how noisy, unstructured text such as maintenance work orders can be transformed into a knowledge graph to visualise and query unstructured data and allow domain experts to make informed business decisions.

<https://nlp-tlp.github.io/ajcai-tutorial/>

Presentation Link