

Abstract

Traditional recommender systems have used statistical and machine learning models to predict and curate content for users based on their past behaviours and the trends of a larger population. This method involves continuous monitoring and tracking of user activity to train and refine a model. Despite the obvious limitations of needing vast amounts of data to train and test the model, this approach also poses a risk of being invasive to the user.

As an alternative, embracing Tim Berners-Lee's vision of linked data and the semantic web, knowledge graphs can be used to semantically model data that already exists on the world wide web, publicly. The semantic links with entities and their relationships can then be queried and matched with user preferences. Graph databases provide flexibility with schema for the data while providing a query interface that is able to search through semantic annotations within the data.

This dissertation provides an approach to a recommendation system that utilizes knowledge graph as the source of recommendations queried using the preferences of a user. It aims to avoid monitoring user behavior to be secure by design and utilise knowledge graphs to provide deterministic recommendations based on semantic links within the graph.