Knowledge Graphs and Creative Applications *

Jer Hayes, Nicholas McCarthy, and Xu Zheng

Accenture Labs
The Dock, 7 Hanover Quay, Dublin
jeremiah.hayes@accenture.com
https://labs.accenture.com/

Abstract. We applied computational creativity to the food domain to create a system that uses a domain specific knowledge graph and graph embeddings to aid a user in finding novel and interesting flavour combinations. Metrics like novelty and surprise are calculated using graph embeddings to retrieve the most useful combinations.

Keywords: Computational Creativity · Knowledge Graphs · Graph Embeddings

1 Background

Computational creativity is the application of AI technologies to emulate and stimulate human creativity. Fast Moving Consumer Goods food (FMCG) companies often find it difficult to innovate fast enough to satisfy the market in a timely manner, *e.g.*, a market trend showing people would like a healthier snack bar might drive an FMCG to ask what ingredients could be pair with chocolate that also had a health benefit? We created a domain specific knowledge graph and graph embeddings to aid a user in finding novel and interesting flavour combinations with the main goal of 'feeding' the mind of creative professionals in the food business with new unexpected combinations.

2 Our approach

A knowledge graph of triples in the format <*subject, predicate, object*> that covered ingredients, recipes, nutrients, compounds, flavour molecules was compiled from multiple public data sources. For instance, we extracted triples like <*recipe 1, contains, ingredient 1*>, <*recipe 1, comes_from, location 1*>, etc. Ampligraph [2] was used with the ComplEx scoring function [3] to input triples into a number representing a probability that a given triple is correct. The resulting embeddings are used as the input data for several metrics. This allows us to compute a probability distribution $d_{I_1,...,I_n}$ for a set of ingredients $\{I_1, ..., I_n\}$ in a space of known recipes. Then the likelihood $p(I_1, ..., I_n)$ is computed as a mean of

* Copyright © 2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR-WS.org/Vol-2721/paper602.pdf

that distribution and the novelty is taken to be $1 - p(I_1, ..., I_n)$. Finally, the one measure of surprise of $\{I_1, ..., I_n\}$ is the Kullback–Leibler divergence between its distribution $d_{I_1,...,I_n}$ and likelihood $p(I_1, ..., I_n)$.

Other metrics in the system are expectedness - what ingredient is likely to come next given the recipe data in the knowledge graph, and pleasantness which is based on the chemical structure information of the flavour compounds. However, this system ultimately relies on the human user to make the final selection and as such is an ideation tool that assists the human expert[1].

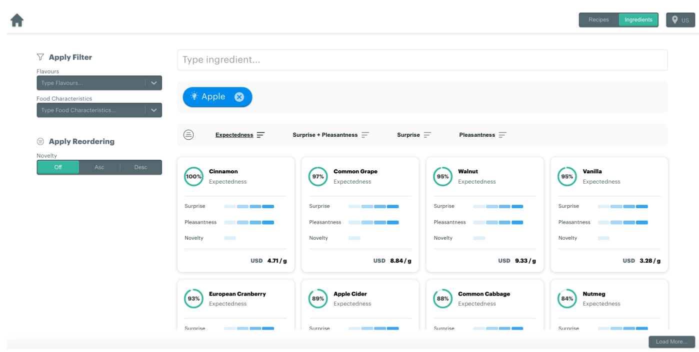

Fig. 1. An ingredient based view from the UI which shows potential additional ingredients to pair. These can be sorted by the metrics and filtered by flavour labels (*e.g.*, 'nutty') and other characteristics.

3 Conclusion

We built a system that shows how knowledge graphs and graph embeddings can be used in the new food product ideation process. The metrics for novelty and surprise are calculated from the graph embeddings and help give product developers fast, curated inspiration and companies a quicker time to market for products targeting momentary markets. Future work includes a focus on the formulations of FMCG products and the creation of new formulations.

References

1. Enhancing culinary creativity with ai, https://www.youtube.com/watch?v=NjT2OjCdGp4
2. Costabello, L., Pai, S., Van, C.L., McGrath, R., McCarthy, N., Tabacof, P.: AmpliGraph: a Library for Representation Learning on Knowledge Graphs (Mar 2019). https://doi.org/10.5281/zenodo.2595043, https://doi.org/10.5281/zenodo.2595043
3. Trouillon, T., Welbl, J., Riedel, S., Gaussier, E., Bouchard, G.: Complex embeddings for simple link prediction (2016)