AI Inspired Recipes: Designing Computationally Creative Food Combos

Azzurra Pini
Fjord, The Dock
7 Hanover Quay, Dublin 2, Ireland
azzurra.pini@fjordnet.com

Jer Hayes
Accenture Labs, The Dock
7 Hanover Quay, Dublin 2, Ireland
jeremiah.hayes@accenture.com

Connor Upton
Fjord, The Dock
7 Hanover Quay, Dublin 2, Ireland
connor.upton@fjordnet.com

Medb Corcoran
Accenture Labs, The Dock
7 Hanover Quay, Dublin 2, Ireland
medb.corcoran@accenture.com

ABSTRACT

If chocolate and broccoli sound a strange pairing, can you imagine a broccoli chocolate bar that combines them? As a matter of fact, the two ingredients share the highest number of flavour molecules, so their combination might not be as weird as it sounds.

We applied computational creativity, that is AI systems to enhance human creativity, to the food domain, with the main goal of feeding the mind of the creative professional in the food business with new unexpected combinations.

CHI'19 Extended Abstracts, May 4–9, 2019, Glasgow, Scotland UK

ACM ISBN 978-1-4503-5971-9/19/05.
https://doi.org/10.1145/3290607.3312948

KEYWORDS

Computational creativity; food innovation; knowledge graph; user interface design; information visualisation.

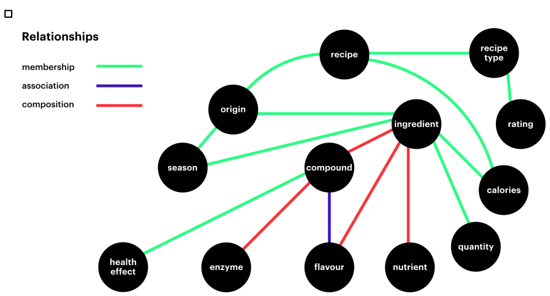

Figure 1: Knowledge graph structure. The graph shows some high level connections of the data points integrated into the KG.

INTRODUCTION

Computational creativity is a "multidisciplinary endeavor that is located at the intersection of the fields of artificial intelligence, cognitive psychology, philosophy, and the arts" [1, 2]. AI is already being used to assist people in creative processes, e.g. music, design, and food [3]. Similar systems though can be also used to inform and aid new product development based on cost pressures, trends, market segments and new insights aggregated with internal research.

The creative process is often hard and time consuming, even just for generating novel combinations among existing elements, whether they are groceries, chemical compounds, sounds or colours. Artificial intelligence can provide systems capable of creating those connections in a fast programmatic way, to support innovation and creativity by saving time and improving efficiency.

Starting from a business request and building on an existing research area at The Dock Accenture Labs, we applied computational creativity to the food domain, starting in particular with the confectionary industry, to design a system able to generate new interesting combinations.

The key aspect is that the tool is not designed to provide solutions, but ideas — possibly unexpected to be re-elaborated by humans, eg. food technicians, chefs or any other creative professional interested in experimenting new combinations.

A KNOWLEDGE GRAPH APPROACH

In order to serve this goal the project started by searching and selecting a large number of external open data sources in the food domain, for finding the connections among ingredients. Platforms such as FooDB[1], CulinaryDB[2], Yummly[3], FlavorNet[4], USDA[5], provided the datasets for ingredients, compounds, health effects, recipes, flavours and many other food-related entities, as shown in Fig. 1. The selected data sources have been cleaned, refined, organised and embedded into a knowledge graph structure.

A knowledge graph is a large network of entities, their properties, semantic types linked by different types of relationships [7]. We convert the knowledge graph into graph embeddings, i.e., a latent space that contains entities and relations as vectors. Such graph embeddings can amplify hidden information and weak signals. By this we mean it can help to find information — explicit connections among two entities in a single dataset — and at the same time create new information, that is building connections among entities from different sources, such as ingredients compounds and recipes.

[1] An open biochemical nutrition database supported by The Metabolomics Innovation Centre. http://foodb.ca/
[2] A repository of structured data of recipes and ingredients across over 22 world regions. https://cosylab.iiitd.edu.in/culinarydb/
[3] A dataset on ingredients, recipes and cuisines hosted by Kaggle. https://www.kaggle.com/c/whats-cooking
[4] A biochemical compilation of aroma compounds found in human odor space. http://www.flavornet.org/
[5] The USDA Food Composition Databases. https://ndb.nal.usda.gov/ndb/search/list

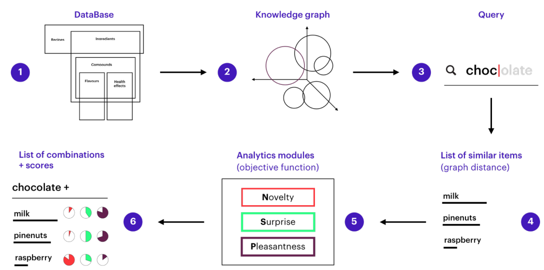

Figure 2: Pipeline of the tool

By populating the database with various sources, it becomes possible to build a knowledge base that is at the same time specific with respect to the knowledge domain, but also flexible and open to other collateral fields, such as health (through chemical compounds and health effects) and culture (through recipes and their origin).

The key idea of the tool is to support human creativity with surprising insights, something they would have never (or difficultly) thought of.

COMPUTATIONAL SURPRISE

Previous work on computational creatity in the food domain has used objective measures such as "novelty", "pleasantness" [4] and "surprise" [6]. With the idea of designing a tool to support human creativity, and not replacing it, our challenge was to relate these concepts to quantitative values and to extend previous research in this area. Pleasantness, although a subjective experience, was modelled using human ratings and associated chemical flavour compounds. Given a compound, or ingredient(s), which are mapped to compounds, then a pleasantness score can be predicted via a deep learning model.

Following the creation of the embedding space from the knowledge graph, we define three different complementary metrics for "surprise" for combination of ingredients. Specifically, the three metrics applied are:

- Pleasantness based Surprise: ingredients are mapped to flavour compounds. An aggregated KL-divergence[6] is calculated for the pleasantness distribution of each recipe.
- Graph based Surprise: Uses link prediction from graph embeddings and KL-divergence between sequence probabilities to score recipes.
- Novelty based Surprise (Bayesian approach): KL-divergence of probability distributions of ingredient sequences.

The similarity of entities can also be calculated by the graph distance obtained from the embedding space, which determines how close two items are based on their shared relationships and entities in the dataset. Given the structure of the graph and the derived embeddings, two ingredients will have a high similarity if they share a large number of flavour compounds, they have the same health effects or they appear frequently in the same recipe.

The three metrics — novelty, pleasantness and surprise — are then integrated as further calculations into an objective function: this includes the values calculated by our algorithms to rank each combination that can be queried from the tool.

In Fig. 2 the tool pipeline is described into six main steps:

1) the different datasets are collected and integrated into a single database;
2) a knowledge graph is created by mapping and classifying all the relationships among the entities in the database;

[6] Kullback–Leibler divergence, or information divergence. https://en.wikipedia.org/wiki/Kullback%E2%80%93Leibler_divergence [Visited: Feb. 2019]

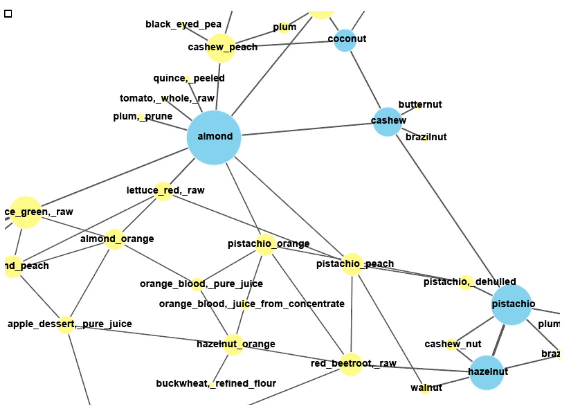

Figure 3: Snapshot of exploratory visualisation: network of co-occurrences of fruits, nuts and chocolate in our database. (In blue the original entities queried from the tool as traditionally combined with chocolate, in yellow the resulting combinations in the knowledge graph space).

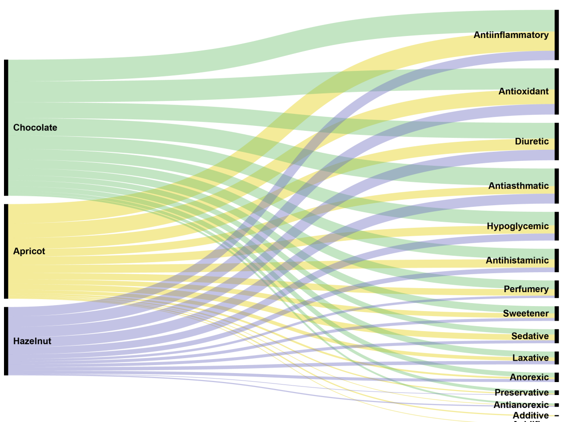

Figure 4: Exploratory visualisation: Alluvial diagram showing the distribution of health effects among central ingredients in the knowledge graph.

3) the user can query the tool for a specific base ingredient;
4) the system returns a list of ingredients ranked by similarity to the queried one;
5) the objective function is computed on the output ingredients list;
6) for each ingredient of the list the system returns the values for novelty, pleasantness and surprise.

VISUALISING SURPRISING COMBINATIONS

Exploratory visualisations

In order to make the complexity of the knowledge graph accessible a first exploration was made through TensorBoard, which allows us to visualise the entirety of the dataset and the different types of data sources that have been colour coded for better exploration.

The knowledge graph has also been explored through partial network graphs and different visualisations to explore the graph structure, to identify patterns and clusters of elments (e.g. most popular and central ingredients). As shown in Fig. 3, some common combinations, mostly taken from chocolate industry popular recipes (e.g. chocolate + fruit & nuts), have been queried from the knowledge graph to find similar ones and the results have been visualised through a network. While the tool is currently accessible though a Jupyter notebook, the single exploratory visualisations rather than functioning as final outputs, have been extremely useful to explore patterns in data, detect errors and provide insights for further analyses (Fig. 4).

Combination explorer

As a next step, the design of a user interface was required to allow potential stakeholders to access the tool in a simple and user friendly way. The interface design had to consider at the same time the representation of the similarity among entities (described by the distance in the embedding space) and the scores related to novelty, pleasantness and surprise as derived from the objective function. Fig. 5, 6 display the two different access points that have been designed to explore the tool results: the first one by querying for combinations of ingredients through attributes, the second to query for combinations based on the computed values of novelty, pleasantness and surprise.

Similarly to the previously described pipeline, the UI allows the user to explore novel combinations by typing a base ingredient, returning firstly a list of possible combinations sorted by similarity, represented by the frequency of that combination into the database. As a further action the user can add different constraints to refine the list of results. The constraints can be based on some of the entities included in the knowledge graph — e.g. flavour profile or health effect — allowing to return combinations that satisfy real world conditions, trying different combination of attributes, exploring and discussing different lists of results. The same list of results can be sorted by similarity/ frequency to highlight the combinations that are more or less frequent, therefore to stress the surprise effect. A second access point displayied in Fig. 6 allows to query combinations starting from particular values of plesantness, novelty and surprise.

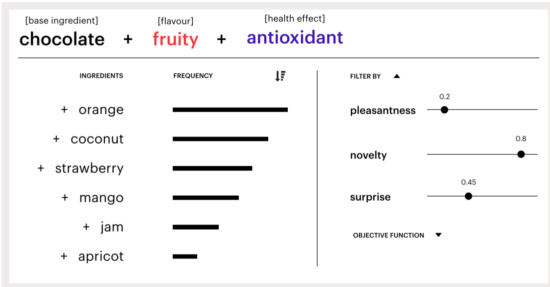

Figure 5: Combination explorer (1)

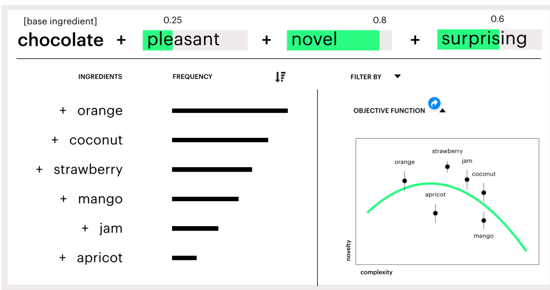

Figure 6: Combination explorer (2)

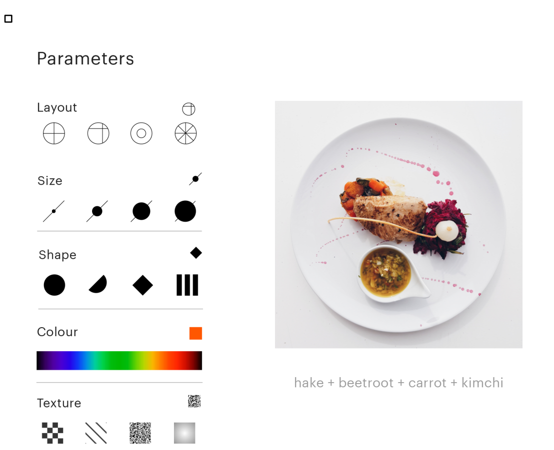

Figure 7: Plating suggestions. Visual browser for food layouting.

As the knowledge graph is designed to be extended with other data sources the database can exponentially increase in size, both in the number of entities and connections. A filtering option is therefore required to select the results by specific levels of novelty, pleasantness and surprise. The tool also allows to visualise the position of each combination on an inverted U-curve of novelty against surprise. The results are therefore visualised on the curve to help the user understand how "wild" the returned combinations are, and how likely they are to be appreciated.

Eye candy: suggesting visually trendy combinations

Although the tool had been originally designed to match the needs of the food industry, our early explorations demonstrated the complexity of such industry, which is subject to a number of constraints from the market, involving aspects that are difficult to consider into a small experimental environment. At the same time most of the cutting edge research on food innovation seemed to happen at the intersection between high-end cusine and food technology, which are lately finding several spaces of collaboration. A small design research project has been therefore conducted interviewing potential stakeholders that are more traditionally involved into the creative process, such as chefs and food technology experts. In particular, by interviewing chefs from fine-dining contexts we wanted to identify other opportunities to expand the tool to match the expectations and the needs from a different type of catchment area.

During this research, food has been stated by several interviewees to be analogous with the fashion industry, with drivers coming from visual culture in general as well as from the various trends circulating through social media and online platforms. Plating layouts, colours, proportions and textures appear nowadays even more relevant than the taste and the actual quality of ingredients and their combinations. Some famous chefs are even known to draw their own recipes before cooking them.

Our challenge was therefore to align the data perspective characterising an AI to provide chefs and food creatives with insights covering also collateral aspects such as its appearance and alignment with online trends. For this reason we imagined two different interfaces to suggest combinations based on plating compositions on the one end and social media trends on the other. (Fig. 7, 8). Fig. 7 in particular proposes a visual layouter for food, which allows to apply a set of visual parameters to filter possible combinations and suggesting matching examples found online. By using as filters the same visual variables traditionally belonging to the data visualisation theory and practice [5], such as shape color, texture and size, we aimed at stressing the idea of food parametrisation from the visual perspective, treating a plate as a canvas. Fig. 8 presents a possible interface for scanning online services and social media platforms to provide suggestions for combinations, techniques and layouts based on the latest trends.

FUTURE WORK

Further steps from a technical perspective will be the refinement of the analytics modules and the knowledge graph structure, as well as the full implementation of the user interface that is currently still at prototype level.

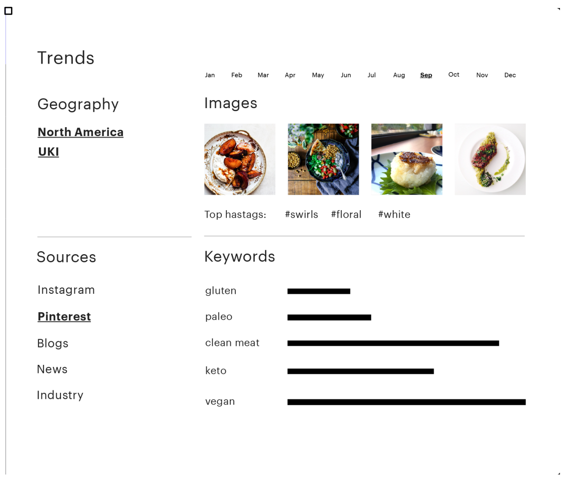

Figure 8: Trend scouting. Visual browser for exploration of food related online trends.

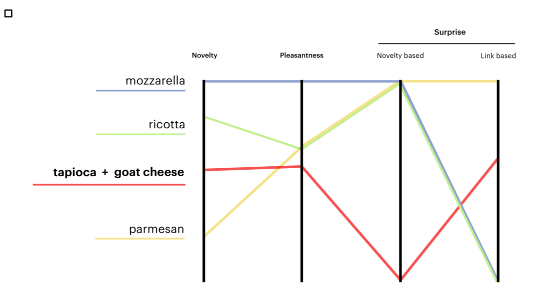

Figure 9: Parset of ingredient substitution. Comparison of different scores for cheese variants in combination with tapioca.

We are also looking at extending the database with other data points that emerged to be particularly relevant during the research, such as cooking processes and access to online trends, both for the visual aspect and the popularity of ingredients or techniques.

Ingredient substitutions have also been lately reported to be quite disruptive, both in the dining business and for food technology in the industry. The reasons are numerous: allergies and intolerances (such as to lactose or gluten), increasingly common special diets (eg. veganism) and access to certain ingredients, based on seasonality or simple unavailability. Some first explorations in that sense have been already performed, by visually ranking the values for novelty, pleasantness and surprise in combining a base ingredient with a set of other similar ones as shown in Fig. 9.

Other future areas of research that are currently being explored are in the field of formulations for the chemical industry, with a greater focus on compounds and molecular structures.

CONCLUSIONS

In summary, the presented research is still at a very experimental level and we are currently in the process of refining both the tool from a computational point of view and the design of the user interface. Nevertheless these explorations opened a number of different research directions in the field of computational creativity for formulation and combination suggestion as well as on the relation between human and AI. In a time when AI systems are criticised for trying to approach problems in a very dry and inhuman way, instead of trying to build a tool to imitate the human mind we wanted to stress its quantitative aspect. The tool doesn't in fact provide any objective or unique solution to a given problem, but rather a set of possible suggestions.

While the aid from the AI side is explicitly responding to quantified parameters which are, like surprise and novelty, not univocally quantifiable, the user is left with the freedom to filter, select, complement and discard the results following his/her own judgement, experience or feeling.

REFERENCES

[1] Amílcar Cardoso, Tony Veale, and Geraint A. Wiggins. 2009. Converging on the divergent: The history (and future) of the international joint workshops in computational creativity. *AI Magazine*, 30(3), 15–22.
[2] Association for Computational Creativity, https://computationalcreativity.net/home/about/computational-creativity/ [Visited: Feb. 2019]
[3] Bilow, R. 2014. How IBM's Chef Watson actually works. Retrieved from http://www.bonappetit.com/entertainingstyle/trends-news/article/how-ibm-chef-watson-works [Visited: Feb. 2019]
[4] Florian Pinel and Lav R. Varshney. 2014. Computational creativity for culinary recipes. In CHI '14 Extended Abstracts on Human Factors in Computing Systems (CHI EA '14). ACM, New York, NY, USA, 439-442.
[5] Jacques Bertin. 1967. Sémiologie Graphique. Les diagrammes, les réseaux, les cartes. With Marc Barbut [et al.]. Paris : Gauthier-Villars
[6] Kazjon Grace and Mary Lou Maher. 2016. Surprise-triggered reformulation of design goals. In Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16). AAAI Press 3726-3732.
[7] Markus Kroetsch and Gerhard Weikum. 2015. Journal of Web Semantics: Special Issue on Knowledge Graphs. http://www.websemanticsjournal.org/index.php/ps/announcement/view/19 [Visited: Feb. 2019]