

# Semantic-aware transformation of short texts using word embeddings: An application in the Food Computing domain

**Andrea Morales-Garzón**

University of Granada, Spain  
andreamgm@correo.ugr.es

**Juan Gómez-Romero**

University of Granada, Spain  
jgomez@decsai.ugr.es

**Maria J. Martín-Bautista**

University of Granada, Spain  
mbautis@decsai.ugr.es

## Abstract

Most works in food computing focus on generating new recipes from scratch. However, there is a large number of new online recipes generated daily with a large number of users reviews, with recommendations to improve the recipe flavor and ideas to modify them. This fact encourages the use of these data for obtaining improved and customized versions. In this thesis, we propose an adaptation engine based on fine-tuning a word embedding model. We will capture, in an unsupervised way, the semantic meaning of the recipe ingredients. We will use their word embedding representations to align them to external databases, thus enriching their data. The adaptation engine will use this food data to modify a recipe into another fitting specific user preferences (e.g., decrease caloric intake or make a recipe). We plan to explore different types of recipe adaptations while preserving recipe essential features such as cuisine style and essence simultaneously. We will also modify the rest of the recipe to the new changes to be reproducible.

## 1 Introduction

Our dietary habits have a huge impact on health and, thus, in quality of life. In the last decades, the amount of nutritional data available has notably increased. This fact, together with the ubiquity of smartphones, has encouraged the use of machine learning techniques for automatizing some tedious and repetitive tasks as diet generation. In this context, the food computing concept refers to the use of food data to improve the quality of life as well as understanding human behavior (Min et al., 2019).

Recipes and their composition have been largely studied in food computing, especially in the food recommendation systems field (Teng et al., 2012). These systems mainly perform recipe-based nutrition assessment, looking for suitable combinations to user preferences. The use of predictive

algorithms to understand relations between recipes has emerged in the last years (Sajadmanesh et al., 2017). Recently, authors have taken advantage of these tools to generate synthetic food data. Recipe generation is a current area of research, and the latest works in the area have put their interest in the creation of synthetic recipes. However, these works have focused on automatized text generation from scratch instead of taking direct advantage of the already existing recipes to generate new versions.

In this thesis, we will address the problem of partially-generation of recipes. Particularly, we will put our effort into recipe adaptation and recipe completion tasks. Online cooking communities and social media generate daily a huge amount of food data, mostly cooking recipes that users want to share with the world. In these communities, many users review the shared recipes, often giving feedback, customization, and suggestions for tasty versions of a given recipe. We plan to use this information to generate new recipe versions. Particularly, we will modify recipes to fulfill the user's requests. There are many reasons to modify a recipe, e.g., a diet restriction such as vegan or vegetarian diets, a lack of ingredients at home, to make the recipe tastier or cooking a kid-friendly version.

Also, many users follow restricted diets linked to nutritionist personalized assessment. A user would require a light version of a given recipe or including high-protein ingredients, among others. We propose to automatize the process of ingredient modification in a recipe and extend this idea with a recipe completion task. In both cases, we can consider several criteria simultaneously, such as those mentioned before. Thus, we tackle twofold challenges here; we have to preserve the semantic of the recipe and its essence while combining heterogeneous sources to incorporate nutrition and user knowledge during the adaptation.

Here, a specific-domain language model can able us to tackle both purposes. We propose to use a fine-tuned word embedding model as the base of our contribution. We will use it to model the recipe ingredients to incorporate useful information from external sources (i.e., complete the ingredient data with nutrition information, user tips, and cuisine styles). Then, we will use the merged data in an adaptation function to find the most suitable foods to adapt a recipe to given restrictions. The semantic information combined with the external data will be the base of the adaptation engine. But adapting recipes do not only consist of dealing with ingredients. Likewise, we will use this model for a synthetic adaptation of title, recipe steps, and extra recipe data affected in the process.

## 2 Related work

Cooking recipes have been largely explored in food computing (Min et al., 2019). Last recipe-based works in food computing have surrounding agreed with the advantages of data mining techniques to understand how people cook. Regarding the use of natural language processing approaches to resolve food computing tasks, they have been mainly focused on the analysis of cuisines and ingredient relations (Min et al., 2019). From a wide perspective, these relations have been addressed by using the textual description of foods and flavor networks. The latter has been widely studied with statistical natural language processing methods (Takahashi et al., 2012; Chen, 2017; Chang et al., 2018). Our proposal is particularly related to the following topics.

**Recipe generation and completion** Creative cooking is the food computing area focused on the automatic generation of new recipes. Here, there is a distinction based on the approach. Synthetic recipes are created in two main ways. One is recipe completion, able to generate synthetic partial recipes from already existing ones. Completing recipes has also been studied in the frame of food recommendation systems. In (Cueto et al., 2019), the authors tackle the problem of completing partial recipes by using context-based recommendation. Recipe generation tasks have also considered the cuisine style for adapting recipes to other cultures (Kazama et al., 2018). In this case, they propose a neural network method to change ingredients for their equivalents in other cuisines. Regarding recipe generation, cooking

recipes have been generated with natural text generation tasks (Aljbawi, 2020). Due to the repetitive results that are usually obtained with this approach, the authors in (Bosselut et al., 2018) proposed a synthetic recipe generation model that considered a reward to get more coherent and less repetitive texts.

**Word embedding in food computing** Word embedding models in food computing have been mainly focused on ingredient analysis. One of the more relevant works in this area is food2vec, where the author used a word embedding model trained with lists of ingredients to understand relations between ingredients and cuisines of the world (Altossar, 2015). Recipe2vec is another model trained in food data, in this case, for recipe retrieval purposes (BuzzFeed and Tasty, 2017). It has been mentioned the many advantages of embedding models referring to fusion heterogeneous food data for multiple purposes, where nutritional and social media textual data are integrated (Salvador et al., 2017) more specialized in resolving image recognition tasks rather than language processing. In (Chen et al., 2019), the authors used a word embedding model to detect ingredient relations to create pseudo-recipes. They used a model trained on a list of recipes to detect which ingredients appear together in recipes. They created a pseudo-recipe object based on this idea.

**Transfer learning** The state-of-the-art in NLP tasks is based in transfer learning models. It is very useful for specific-domains where data are limited since general-purpose models will perform poorly. This approach allows to train models with a bigger capacity but capturing the subtle essence of the problem addressed. The most well-known models using fine-tuning for specific tasks are BERT (Devlin et al., 2019) and GPT-2 (Budzianowski and Vulić, 2019) with excellent results. Transfer learning has been used in different specific areas, e.g., in biomedicine (Lee et al., 2019). To the best of our knowledge, transfer learning has not been proposed to extract semantic information from food item descriptions to combine heterogeneous sources.

**Conditional text generation** Controllable text generation is the area where sentences' attributes can be controlled by factors such as age, gender, or style (Prabhumoye et al., 2020). In this problem, we have a sequence output that is conditioned by the sequence input. Text generation language

models have to assess the need for controlling specific parts of the task for resolving a specific problem (Keskar et al., 2019). In this line, recent approaches have put interest in style transfer techniques. Text style transfer has allowed adapting a synthetic text to different situations such as audiences, complexity, and other contextual circumstances (Li et al., 2020). Recent style transfer algorithms employ parallel data in supervised learning approaches and non-parallel data in seq2seq architectures for unsupervised approaches. Also, Variational Auto-Encoders have been applied for this aim by separating content and style in the latent space for better adjustment of the style (Fu et al., 2018).

### 3 Proposed methodology

We have divided our approach into two tasks explained in the following subsections.

#### 3.1 Heterogeneous data-handling

The first problem that appears when modifying a recipe is obtaining enough food knowledge to be able to generate recipes that fulfill user preferences. One of the main challenges to address in food computing is the inherent difficulty in using food features from many different nutritional sources. Consequently, food items need previous processing to handle them jointly. According to this idea, we can use the item textual description to identify equivalent items between databases, allowing the joint use of these databases as a unique data collection (Morales-Garzón et al., 2020). Notice that ontology-based methods could perform well in this problem. But these models have problems when applied to ingredient-based tasks. They do not represent high detailed ingredients, and also have difficulty generalizing to online recipes. Furthermore, knowledge extraction has to be hand-crafted. To overcome this, we propose to model ingredient descriptions with a word embedding model. This unsupervised model can deal with arbitrary-sized text and capture the semantic of cooking.

**Models** Since the food domain is very-specific, general-purpose word embedding models will perform poorly. This issue can be solved by using pre-trained models and perform transfer learning. Deep models will be trained in large unlabeled text databases and, then, fine-tuned to the cooking domain. This approach will be able to capture automatically the semantic of cooking without human

supervision. First, we will do a transfer learning task with a BERT language model (Devlin et al., 2019). Using BERT will able us to deal with one of the more compound facts when cooking: a same ingredient can be used in different forms and meals (e.g., a user could use flour for a cake, but also frying fish). In a sentence-based model, we will be able to represent the current context in which an ingredient is used. This fact will able us to find better food alternatives for each ingredient.

We plan to test the performance of our model replicating the process with GPT-2 (Budzianowski and Vulić, 2019). The main difference between BERT and GPT-2 is while BERT looks at the context of the word, GPT-2 only looks backward. In this thesis, we will explore both and discern the advantages of each one for the cooking domain.

**Distance metrics** We understand ingredient mapping as the search for an equivalent food in an external source. This similarity can be obtained by calculating the distance between ingredient descriptions. We consider an ingredient description as a short description text (e.g. “almonds toasted”). We plan to use food representations obtained with the embedding vectors to find food equivalences within databases. We plan to test the model performances with different metrics including word mover’s distance as a baseline metric. We also plan to use a distance metric proposed in (Morales-Garzón et al., 2020), which has demonstrated to work remarkably well with food data descriptions.

**Dataset** We plan to use a pre-trained word embedding model trained on Wikipedia and Book Corpus datasets<sup>1</sup>. We will re-train the model in a food-based textual corpus. To do this, we will use a large recipe dataset available in archive.org<sup>2</sup>. The dataset contains more than 200,000 recipes with their preparation step texts. These texts contain meaningful information about the science of cooking such as ingredient combinations and cooking processes.

#### 3.2 Adaptation engine

Deciding the most profitable version for a recipe is a very subjective process. Consequently, following human adaptation rules is difficult and very tedious. Our approach consists in using word embedding vectors to represent an existing cooking

<sup>1</sup><https://huggingface.co/models>

<sup>2</sup><https://archive.org/details/recipes-en-201706>




Figure 1: An overview of the procedure

recipe. For that, we will extract the ingredients from a recipe, and we will obtain their embedded representations with the transfer learning model. Once we have a representation of the ingredients, we proceed with adapting them to fit the user requirements. We will take advantage of the captured information in the model to adapt the ingredients (e.g., similarity relations between foods), with the aim of preserving the recipe essence. In this way, semantic relations between ingredients can influence the decision when changing an ingredient for other that fits in the recipe. Besides, not only changing the ingredients will result in a finished recipe. We will also generate automatic text from the ingredient list to make coherent cooking instructions. Thus, the process consists of three steps: (1) obtain a semantic representation of the ingredients, (2) adapt the recipe by changing the ingredients to other foods that fit, (3) modify the rest of the recipe accordingly, i.e., recipe preparation steps, nutrition information, and title if need. Since title and nutrition data can be easily obtained from the final ingredients, the challenge resides in altering the preparation text. Conditional generation and style transfer techniques will be used in this last step. At the end of this process, the user will have the full recipe with the list of ingredients and the cooking procedure, being able to reproduce it at home. See Figure 1 for a better understanding of this process.

**Recipe modeling** First, we will model a recipe with the transfer learning model. The ingredient in-

formation contained in an online recipe is short and may not be sufficient for making a quality adaptation. As introduced, we plan to combine the ingredients with food features such as cuisine style, nutrition information, packaging information, cooking tips, and potential ingredient relations. Unfortunately, this information has to be obtained from external heterogeneous sources. We will join this information in one object, merging the ingredient data with food knowledge from these external databases. Subsection 3.1 describes this procedure.

**Ingredient adaptation** One part of the recipe is properly adapting the ingredients. There are two main ways of adapting a recipe. In the first case, some ingredients of a recipe are replaced following a criterion, e.g., converting a given recipe into a vegan version, and, in the second case, it consists of suggestions to add new ingredients. In both cases, we can consider several criteria simultaneously. For example, the users would like to do a recipe but with fewer calories or more proteins. Here, we will design the proper adaptation function according to a multiobjective optimization problem with restrictions, e.g., maximizing the use of sweet ingredients while minimizing the calories. This has to be subjected to maintaining the coherence of the recipe.

Notice that only similarity-based functions will be suitable for maintaining the coherence of the recipe but they do not take into account other factors like calories. Thus, the ingredient adaptation task will consider the joint ingredient data obtained from the combination of the ingredient with exter-

nal sources. Thus, we will be able to add adaptation knowledge to this step.

**Feeding the adaptation procedure** We can make use of user interactions with recipes to obtain information about how users react to some recipes. We will use online user interaction data with recipes to be considered in the adaptation function. We can exploit this data to measure which ingredient combinations are more appealing for the users. We will analyze this data to extract knowledge to feed the adaptation function.

**Adaptation of the rest of the recipe** Adapting a recipe does not just consist of changing the ingredients for another suitable option. We also need to adapt the preparation step to fit with the new ingredients. This part is compound because it needs to remain the coherence of the original recipe when possible. We plan to explore the use of word embedding approaches to partially-generate synthetic text using keywords. We will part from the original recipe, detecting those steps that must be modified. Notice that some recipe objects also contain nutrition tags for a serving. In this case, we will adapt this information using the ingredient data if allowed.

**Dataset** We plan to study recipes in specific cuisines. For that, we will use recipes extracted from Yummly. One of the tags stored in Yummly recipe data is the geographical origin of the recipe. There are several Yummly datasets online that we can use, with ingredients, preparation texts, and cuisine type<sup>3</sup>. Additionally, the Yummly website provide users' reviews, with their suggestions for altering the recipe, and recommendations of ingredients substitutions (and additions) to improve the taste of the dish.

Regarding nutritional food data, there are open-source nutrition dataset available for obtaining food data from the most common foods and dishes. One example is the USDA database, maintained by the Department of Agriculture in the United States (Gebhardt et al., 2008). There are also market product sources for access to typical food in specific zones of the world. Open Food Facts<sup>4</sup> is an open-source project with the aim of make worldwide food products accessible.

<sup>3</sup>[http://123.57.42.89/FoodComputing\\_Dataset.html](http://123.57.42.89/FoodComputing_Dataset.html)

<sup>4</sup><https://es.openfoodfacts.org>

There are available resources about how users interact with recipes. The Food.com dataset<sup>5</sup> available in Kaggle provide this info for more than 200,000 recipes from the popular cooking site Food.com<sup>6</sup>.

## 4 Evaluation

Validating recipe adaptations is a subjective procedure. Depending on the cultural factor, the type of meal, the flavors, and other intrinsic combinations, what could be an excellent recipe for a user, could result to be untasted for another different one. This variability makes it difficult to measure the adequacy of an adapted recipe. To tackle this variability, we plan to evaluate the proposed method with an online survey on both regular and expert users. For this, we will generate adapted recipes for different circumstances. Each recipe will receive a score, where the lowest value represents that the adapted recipe is disgusting and the highest is a very succulent recipe. Also, we plan to obtain adaptation suggestions in this step to use them as feedback for future improvement.

## 5 Strengths

With the arising of technology and, consequently, the large amount of recipes shared on the internet, food computing has played an undeniable role in recipe retrieval systems. These systems allow access to online recipes to speed up the recipe searching whenever a user wants to prepare a dish. We believe that the integration of our approach in the cited software could meet user needs when looking for cooking inspiration. Additionally, it is worth noting that a recipe-based word embedding model could be able to participate in multiple problems of food computing. One of its applications is using them for detecting recipe similarity to ensure variety in nutrition assessment systems.

We believe that food computing is not the only application of our approach. Personalized beauty treatment is another area in which our proposal could be useful. Commonly, there can be found on the internet many natural beauty care recipes consisting of a list of ingredients and instructions to create beauty remedies for different purposes. Among other many factors, this kind of treatment handles user expectations, allergies, and the cos-

<sup>5</sup><https://www.kaggle.com/shuyangli94/food-com-recipes-and-user-interactions>

<sup>6</sup><https://www.food.com>

metic composition of the treatment. A transfer learning model in this area could be applied to adapt these kinds of treatments to the user's needs.

## 6 Summary

Our proposal consists of using a transfer learning model in the food domain to adapt recipes to fulfill user needs. The challenge remains in using the model for two different tasks. First, we plan to use the model to complete ingredients information with data from external sources, such as nutritional data or cuisine traditions. Thus, we will employ this joined data for adapting a recipe to fulfill a need. Then, we will use the language model to adequate the rest of the recipe to be consistent with the adapted ingredients.

## References

- Bushra Aljbawi. 2020. Health-aware food planner: A personalized recipe generation approach based on gpt-2.
- Jaan Altossar. 2015. Food2vec-augmented-cooking-machineintelligence. *Jaan Altossar's blog, last modified Dec*, 17.
- Antoine Bosselut, Asli Celikyilmaz, Xiaodong He, Jianfeng Gao, Po-Sen Huang, and Yejin Choi. 2018. [Discourse-aware neural rewards for coherent text generation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 173–184, New Orleans, Louisiana. Association for Computational Linguistics.
- Paweł Budzianowski and Ivan Vulić. 2019. [Hello, it's GPT-2 - how can I help you? towards the use of pre-trained language models for task-oriented dialogue systems](#). In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 15–22, Hong Kong. Association for Computational Linguistics.
- BuzzFeed and Tasty. 2017. [Recipe2vec: How word2vec helped us discover related tasty recipes](#).
- Minsuk Chang, Léonore V Guillain, Hyeungshik Jung, Vivian M Hare, Juho Kim, and Maneesh Agrawala. 2018. Recipescape: An interactive tool for analyzing cooking instructions at scale. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, page 451. ACM.
- Meng Chen, Xiaoyi Jia, Elizabeth Gorbonos, Chnh T Hong, Xiaohui Yu, and Yang Liu. 2019. Eating healthier: Exploring nutrition information for healthier recipe recommendation. *Information Processing & Management*, page 102051.
- Yuzhe Chen. 2017. *A statistical machine learning approach to generating graph structures from food recipes*. Ph.D. thesis, Brandeis University.
- Paula Fermín Cueto, Meeke Roet, and Agnieszka Słowik. 2019. Completing partial recipes using item-based collaborative filtering to recommend ingredients. *arXiv preprint arXiv:1907.12380*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2018. Style transfer in text: Exploration and evaluation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Susan Gebhardt, L Lemar, D Haytowitz, P Pehrsson, M Nickle, B Showell, R Thomas, J Exler, and J Holden. 2008. Usda national nutrient database for standard reference, release 21. *United States Department of AgricultureAgricultural Research Service*.
- Masahiro Kazama, Minami Sugimoto, Chizuru Hosokawa, Keisuke Matsushima, Lav R Varshney, and Yoshiki Ishikawa. 2018. A neural network system for transformation of regional cuisine style. *Frontiers in ICT*, 5:14.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [BioBERT: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36(4):1234–1240.
- Xiangyang Li, Guo Pu, Keyu Ming, Pu Li, Jie Wang, Yuxuan Wang, and Sujian Li. 2020. Review of text style transfer based on deep learning. *arXiv preprint arXiv:2005.02914*.
- Weiqing Min, Shuqiang Jiang, Linhu Liu, Yong Rui, and Ramesh Jain. 2019. A survey on food computing. *ACM Computing Surveys (CSUR)*, 52(5):1–36.
- Andrea Morales-Garzón, Juan Gómez-Romero, and María J Martín-Bautista. 2020. A word embedding model for mapping food composition databases using fuzzy logic. In *International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems*, pages 635–647. Springer.

- Shrimai Prabhumoye, Alan W Black, and Ruslan Salakhutdinov. 2020. Exploring controllable text generation techniques. *arXiv preprint arXiv:2005.01822*.
- Sina Sajadmanesh, Sina Jafarzadeh, Seyed Ali Ossia, Hamid R Rabiee, Hamed Haddadi, Yelena Mejova, Mirco Musolesi, Emiliano De Cristofaro, and Gianluca Stringhini. 2017. Kissing cuisines: Exploring worldwide culinary habits on the web. In *Proceedings of the 26th international conference on world wide web companion*, pages 1013–1021.
- Amaia Salvador, Nicholas Hynes, Yusuf Aytar, Javier Marin, Ferda Ofli, Ingmar Weber, and Antonio Torralba. 2017. Learning cross-modal embeddings for cooking recipes and food images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3020–3028.
- Jun Takahashi, Tsuguya Ueda, Chika Nishikawa, Takayuki Ito, and Akihiko Nagai. 2012. [Implementation of automatic nutrient calculation system for cooking recipes based on text analysis](#). *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, pages 789–794.
- Chun-Yuen Teng, Yu-Ru Lin, and Lada A Adamic. 2012. Recipe recommendation using ingredient networks. In *Proceedings of the 4th Annual ACM Web Science Conference*, pages 298–307.