Abstract:

Recipes are the universal language of cooking. Through a recipe you quickly understand how a certain dish can be made.
The goal of this master's thesis is to explore the potential for computational creativity in gastronomy by developing an ingredient-driven recipe generation system. Having an ingredient centered generator was the most attractive way of solving this task. We focused on the ingredient semantics and also on the modelling interactions between the ingredients. The task was broken down into: learning continuous representations of ingredients that capture their semantics and building a way of generating event chains based on the learnt representations of ingredients.
In this thesis several new approaches to automatic recipe generation are proposed: firstly, a map from ingredients into cooking instructions events was learned. Also, by posing the model as a multi translation tasks the model proved to be more flexible in the generated results and can be ported easily; secondly we overcome similar approaches by handling the amounts of ingredients and also by forcing the generation to include the events learned and allowing to generate an instruction that include a given ingredient multiple times. Thirdly we experiment with computational creativity by making use of the two distributional spaces created - one for ingredients and one for the events of cooking.
The methods used via distributional semantics were evaluated both intrinsically and extrinsically. Because our generator system was learned by posing the task as a machine translation one, the metrics to evaluate the generator were metrics from the domain of machine translation and language modelling. Our model was found to be better than the baseline, a sequence to sequence model that encoded ingredients and decoded the instructional cooking steps.