Food Recipe Recommendation

CS608 Group 8 Project Final Presentation

Felicia Anto Christy
Gordy Adiprasetyo
Ian Carter
Nikitha Banda
Ulysses Chong

July 2021

Preview

Table of Contents

1. Project Objectives

Project Objectives

To build an application that recommends recipes based on user's preference

Objective

Build a recommendation application that can recommend top recipes based on user's available ingredients and taste preference.

Background

The perfect recipe doesn't exists, or if it does, you may not have the ingredients on hand to cook it. We believe there is a demand for a recipe recommendation application that works with the constraints of the average home chef to materialise choice recipe dishes based on their recipe/food interactions and available inventory.

Project Objectives

Interaction data and multimodal supplementary data are explored across a sweep of models to identity top performing model to tune and deploy

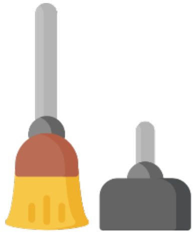

Data Preparation

Prepare interaction data between users and recipes. Identify and prepare textual data to be used as multimodal supplementary data.

Model Sweep & Selection

Explore prepared data across a sweep of models from explicit models, implicit models to multimodal models. MF, NMF, SVD, BPR, WMF, CTR, CDL.

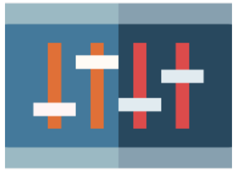

Hyperparameter Tuning

Identify top performing models and performing hyperparameter tuning of top models. Informed hyperparameter search; Hyperopt.

Train Final Model

Select top performing tuned models and train the final model and validate its generalisability.

2. Dataset

Datasets

Identified food recipe interactions data with metadata and raw text data. Raw text is to be explored and prepared for appropriate NLP treatment.

Food.com Recipe & Review Data

Description

These datasets contain recipe details and reviews from Food.com (formerly GeniusKitchen). Data includes cooking recipes and review texts.

Basic statistics

	Food.com
Number of recipes:	231,637
Number of users:	226,570
Number of reviews:	1,132,367

Metadata

- Ratings and Reviews
- Recipe Name, Description, Ingredients, and Directions
- Recipe Categories (Tags)
- Recipe Nutrition Information

Data Tasks (2) Code (22) Discussion (14) Activity Metadata Download (851 MB) New Notebook

Content

This dataset contains three sets of data from Food.com:

Interaction splits

- `interactions_test.csv`
- `interactions_validation.csv`
- `interactions_train.csv`

Preprocessed data for result reproduction

In this format, the recipe text metadata is tokenized via the GPT subword tokenizer with start-of-step, etc. tokens.

- `PP_recipes.csv`
- `PP_users.csv`

Data Source

https://cseweb.ucsd.edu/~jmcauley/datasets.html#foodcom

https://www.kaggle.com/shuyangli94/food-com-recipes-and-user-interactions?select=RAW_recipes.csv

Datasets

Sample data

Example

Recipe:

```
name:                beer mac n cheese soup
id:                  499490
minutes:             45
contributor_id:      560491
submitted:           2013-04-27
tags:                60-minutes-or-less
                     time-to-make
                     preparation
nutrition:           678.8
                     70.0
                     20.0
                     46.0
                     61.0
                     134.0
                     11.0
n_steps:             7
steps:               cook the bacon in a pan over medium heat and set
aside on paper towels to drain , reserving 2 tablespoons of the
grease in the pan
                     add the onion , carrot , celery and jalapeno and
cook until tender , about 10-15 minutes
                     add the garlic and cook until fragrant , about a
minute
                     mix in the flour and let it cook for 2-3 minutes
                     add the broth , beer , nutmeg , bacon and
macaroni and let cook until the macaroni is al-dente , about 7-8
minutes
                     add the cream , mustard , worcestershire sauce
and cheese and cook until the cheese has melted without bringing it
back to a boil
```

```
                     season with cayenne , salt and pepper to taste
description:         all of the flavors of mac n' cheese in the form
of a hot bowl of soup!  submitted by kevin lynch
ingredients:         bacon
                     onion
                     carrots
                     celery
                     jalapeno pepper
                     garlic cloves
                     flour
                     chicken broth
                     beer
                     nutmeg
                     elbow macaroni
                     heavy cream
                     dijon mustard
                     worcestershire sauce
                     cheddar cheese
                     cayenne
                     salt and pepper
n_ingredients:       17
```

Review:

```
user_id:        8937
recipe_id:      44394
date:           2002-12-01
rating:         4
review:         This worked very well and is EASY.  I used not quite a
whole package (10oz) of white chips.  Great!
```

Dataset

Data Preprocessing - removed users with less than 5 ratings to improve sparsity. Experienced a 7.5x improvement to sparsity.

<u>Pre-split Data by author</u>
# of interactions: 1,132,367
# of users: 226,570
# of recipes: 231,637
Sparsity: 99.9978%

Issues noted:

- Extremely high sparsity

```
TEST:
...
    | NCRR@50 | NDCG@50 | Recall@50 |   Train (s) | Test (s)
--- + ------- + ------- + --------- + ----------- + --------
MF  |  0.0000 |  0.0000 |    0.0000 |      0.6461 | 104.7110
NMF |  0.0000 |  0.0000 |    0.0000 |      3.4930 |  98.5881
WMF |  0.0000 |  0.0000 |    0.0000 | 22571.0294 | 167.1049
SVD |  0.0000 |  0.0000 |    0.0000 |      0.5959 | 120.9478
BPR |  0.0000 |  0.0000 |    0.0000 |     19.9052 | 127.8410
```

<u>Processed Data</u>
# of interactions: 872,021
# of users: 23,086
# of recipes: 231,637
Sparsity: 99.9837% (7.5x improvement in sparsity)

Steps taken to process:

- Used the complete raw data
- Removed users who has rated 5 recipes or less

```
TEST:
...
    |    MAE |   RMSE |    AUC |    MAP | NDCG@-1 | Precision@10 | Recall@10 | Train (s) | Test (s)
--- + ------ + ------ + ------ + ------ + ------- + ------------ + --------- + --------- + ---------
BPR | 3.7741 | 3.8812 | 0.8347 | 0.0103 |  0.1448 |       0.0078 |    0.0199 |   18.2711 | 1995.8848
```

Dataset

Data Preprocessing – split the data into train and test sets for consistent model selection on hold out test set; 80-20 split.

78.5%

Train Data
# of interactions: 684,907
# of users: 23,086
# of recipes: 189,673

	user_id	item_id	rating
0	2046	4523	2
1	2046	4684	5
2	2046	3431	5
3	2046	13307	5
4	2312	51964	5
5	2312	1232	4
6	2312	4397	5
7	2625	471	3
8	2369	100	3

21.5%

Test Data
# of interactions: 187,114
# of users: 23,086
# of recipes: 90,614

Note: Test dataset has at least one of every user and each user does not have more than 20% of their number of interactions from the original dataset in the test dataset.

3. Experiment

Performance Metrics chosen evaluate the recall and ordering of recommendations made by model

Recall@10

To check that the top recommendations made by the models are within what the user would rate as well.

We set @10 to simulate the context of the recommendation app where people may view top 10 recipes for each session as opposed to a larger number like 50 (which may be more applicable to topics like e-commerce)

NDCG@10

NCRR@10

We included some metrics to rate the ordering performance of the models to make sure the order is similar to what the user would order their interactions in real life.

Experiment

Multiple combinations of modalities were explored to arrive at the best model

S/N	Modality	Algorithm	NCRR@10	NDCG@10	Recall@10	Harmonic Mean
1	-	MF	0.0073	0.0081	0.0095	0.0082
2	-	SVD	0.0073	0.0081	0.0095	0.0082
3	-	BPR	0.0136	0.0154	0.0190	0.0157
4	Steps	CTR	0.0155	0.0173	0.0206	0.0176
5	Steps	CDL	0.0010	0.0011	0.0008	0.0010
6	Description	CTR	0.0121	0.0142	0.0187	0.0145
7	Description	HTF	0.0006	0.0005	0.0003	0.0004
8	Ingredients	CTR	0.0133	0.0152	0.0194	0.0156
9	Ingredients	HTF	0.0004	0.0004	0.0004	0.0004
10	Nutritional Values (as text content)	CTR	0.0143	0.0162	0.0198	0.0164
11	Nutritional Values (as context)	LibFM	RMSE: 1.0465			-

Experiment

Obtain base models on interaction data only to use as benchmark for future multi-modal models.

Experiment Objective

Obtained tuned base models performance to compare further multi-modal experiments against.

Tune hyperparameters

```
# models to sweep
mf = MF(k=int(70.0),
        learning_rate=0.00010320592158418093,
        lambda_reg=0.035816774214327635,
        early_stop = True,
        use_bias=True,
        verbose=False,
        seed = seed,
        name="MF")

svd = SVD(k=int(152.0),
          learning_rate=0.00010259837855631452,
          lambda_reg=0.001690022699357155,
          early_stop = True,
          verbose=False,
          seed = seed,
          name="SVD")

bpr = BPR(k=int(371.0),
          learning_rate=0.010321497562406403,
          lambda_reg=0.00022716178962653O6,
          verbose=False,
          seed = seed,
          name="BPR")
```

Model Performance and results

	NCRR@10	NDCG@10	Recall@10
MF	0.0073	0.0081	0.0095
SVD	0.0073	0.0081	0.0095
BPR	0.0136	0.0154	0.0190

Base model performances will be use as a benchmark to evaluate further experiments on other multi-modal models.

Experiment

Recipe description added little value to the recommendation system

Experiment Objective

Explore **Collaborative Topic Regression (CTR)** and **Hidden Factors and Hidden Topics (HFT)** models using the **description** feature.

Preprocessing steps

BaseTokenizer and removal of stopwords, supported by cornac's TextModality

Tuning – Hyperparameters tuned and search space

k=20

a=1.0

b=0.01

lambda_u=0.01

lambda_v=0.01

Model Performance and results

```
TEST:
...
    | NCRR@10 | NDCG@10 | Recall@10 | Train (s) | Test (s)
--- + ------- + ------- + --------- + --------- + --------
CTR |  0.0121 |  0.0142 |    0.0187 | 6182.5077 | 550.6342
HFT |  0.0006 |  0.0005 |    0.0003 | 7046.2544 | 400.5194
```

CTR decisively outperformed HFT on the description feature. However, it did not seem to perform better than the base BPR model.

Experiment

Using ingredients as text modality improves recommendations, though not significantly

Experiment Objective

Explore **Hidden Factors and Hidden Topics (HFT) and Collaborative Topic Regression (CTR)** model using **ingredients** as feature

Preprocessing steps

BaseTokenizer and removal of stopwords, supported by cornac's TextModality

Tuning – Hyperparameters tuned and search space

k=16

max_iter=25

a=1.0

b=0.01

lambda_u=0.00015749647292245656

lambda_v=0.00013408921352280157

Model Performance and results

	NCRR@10	NDCG@10	Recall@10	Train (s)	Test (s)
HFT	0.0004	0.0004	0.0004	30465.4790	702.5297
CTR	0.0121	0.0141	0.0185	4606.9860	623.2838

CTR significantly outperformed HFT. However, it did not outperform the base BPR model.

	NCRR@10	NDCG@10	Recall@10	Train (s)	Test (s)
CTR	0.0133	0.0152	0.0194	1862.4895	648.8440

After parameter tuning, the CTR model using ingredients was able to outperform the base models. However, its performance still falls short compared the tuned CTR model using steps as additional text modality.

Experiment

Using recipe steps as text modality improves recommendations provided to users

Experiment Objective

Explore **Collaborative Topic Regression (CTR)** and **Collaborative Deep Learning (CDL)** models using the **steps** feature.

Preprocessing steps

BaseTokenizer and removal of stopwords, supported by cornac's TextModality

Tuning – Hyperparameters tuned and search space

k = 30
max_iter = 110
a=2.5108
b=0.2558
lambda_u=0.01
lambda_v=0.01

Model Performance and results

```
TEST:
...
    | NCRR@10 | NDCG@10 | Recall@10 |   Train (s) | Test (s)
--- + ------- + ------- + --------- + ---------- + --------
CTR |  0.0155 |  0.0173 |    0.0206 | 11291.9330 | 807.8510
```

```
TEST:
...
    | NCRR@10 | NDCG@10 | Recall@10 | Train (s) | Test (s)
--- + ------- + ------- + --------- + --------- + --------
CDL |  0.0010 |  0.0011 |    0.0008 | 5806.4249 | 966.0430
```

CTR model trained using recipe steps as additional text modality performed the best compared to other modalities

Experiment

Using Nutrition as Text Modality

Experiment Objective

Explore **Collaborative Topic Regression (CTR)** for the topic of **Nutrition**

Tuning – Hyperparameters tuned and search space

K= 20

a= 5

b= 0.2

lambda_u= 0.001

lambda_v= 0.001

Model Performance and results

```
TEST:
...
    | NCRR@10 | NDCG@10 | Recall@10 | Train (s) | Test (s)
--- + ------- + ------- + --------- + --------- + --------
CTR |  0.0143 |  0.0162 |    0.0198 | 1886.0074 | 830.7849
```

Pre-processing steps:

calories (#), total fat (PDV), sugar (PDV) , sodium (PDV) , protein (PDV) , saturated fat, carbs

nutrition	nutrition_type
[51.5, 0.0, 13.0, 0.0, 2.0, 0.0, 4.0]	High Sugar
[173.4, 18.0, 0.0, 17.0, 22.0, 35.0, 1.0]	High Saturated Fat
[368.1, 17.0, 10.0, 2.0, 14.0, 8.0, 20.0]	High Carbs
[160.2, 10.0, 55.0, 3.0, 9.0, 20.0, 7.0]	High Sugar
[380.7, 53.0, 7.0, 24.0, 6.0, 24.0, 6.0]	High Total Fat

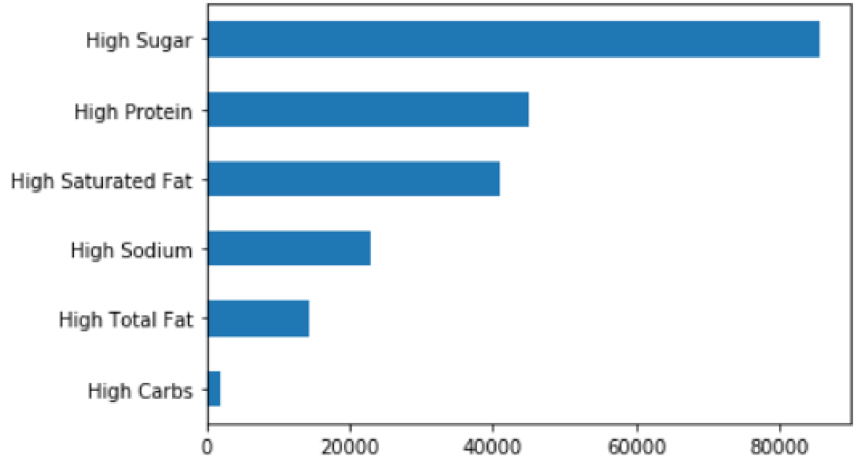

4. Results & Findings

Most model recommendations overlap (id=2046), but some information is lost by using only one text modality

Feature used	Top 3 topics	Top 5 recipes
Steps	pan, salt, chicken, cup, preheat, dough, bowl, small, dry, pour stirring, garlic, stir, cheese, tablespoons, vinegar, temperature, longer, beef, x drain, cover, add, stir, lightly, heat, bowl, sheet, mixture	39087 creamy cajun chicken pasta 27208 to die for crock pot roast 63689 my family s favorite sloppy joes pizza joes 22782 jo mama s world famous spaghetti 28148 oven fried chicken chimichangas
Description	recipes, simple, pretty, chips, onion, … chocolate, different, prep, version, loved, … fresh, adapted, butter, cooking, ago, garlic, …	27208 to die for crock pot roast 39087 creamy cajun chicken pasta 28148 oven fried chicken chimichangas 22782 jo mama s world famous spaghetti 63689 my family s favorite sloppy joes pizza joes
Ingredients	Sugar, butter, brown, egg, water, unsalted, … Oil, olive, vegetable, garlic, vinegar, wine, … Fresh, chicken, pepper, garlic, broth, cloves, …	54257 yes virginia there is a great meatloaf 27208 to die for crock pot roast 30987 creamy cajun chicken roast 10744 delicious chicken pot pie 22782 jo mama s world famous spaghetti
Nutritional Values	Not enough text to generate topics	39087 creamy cajun chicken pasta 27208 to die for crock pot roast 32204 whatever floats your boat brownies 22782 jo mama s world famous spaghetti 28148 oven fried chicken chimichangas

5. Future Improvements

Future Improvements

Ensembling the models leverages the information of each model to produce a more informed set of recommendations

ILLUSTRATION

Recipe Name	Steps	Description	Ingredients	Nutritional Values	Ensemble
to die for crock pot roast	4.835930	6.330528	6.202894	5.185534	5.6387215
creamy cajun chicken pasta	4.854353	6.004121	6.0547867	5.214962	5.532055675
jo mama s world famous spaghetti	4.354705	5.774430	5.8057384	4.647672	5.14563635
oven fried chicken chimichangas	3.965155	5.871809	5.573654	4.476286	4.971726
japanese mum s chicken	3.745926	5.501774	5.338853	4.136297	4.6807125

	CTR + Steps	Ensemble
Top 5 Recommendations	39087 creamy cajun chicken pasta 27208 to die for crock pot roast 63689 my family s favorite sloppy joes pizza joes 22782 jo mama s world famous spaghetti 28148 oven fried chicken chimichangas	27208 to die for crock pot roast 39087 creamy cajun chicken pasta 22782 jo mama s world famous spaghetti 28148 oven fried chicken chimichangas 68955 japanese mum s chicken

Future Improvements

Use of distributed computing tools to further train, tune and deploy models.

<u>Challenges encountered</u>

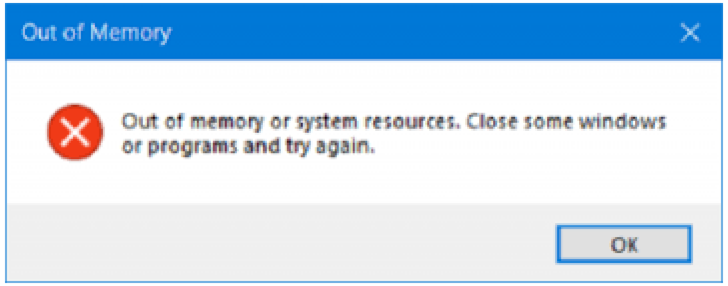

Our job training on our local machines took very long due to the training algorithms for the multi-modal recommendation systems.

We encountered many 'out-of-memory' errors due to our local machines not having the necessary RAM memory to train models with larger parameters.

<u>Solution (Future improvement)</u>

Use of distributed computing tools such as Apache Spark to spin up clusters of nodes to train and tune models in a distributed manner.

Node memory size can be configured to cater to the memory requirement of the models being trained on the node. Hyperopt can be used to managed distributed training of multiple models simultaneously.

Questions & Answers