

Extraction of Atypical Aspects from Customer Reviews: Datasets and Experiments with Language Models

Smita Nannaware¹, Erfan Al-Hossami¹ and Razvan Bunescu^{1,*}

¹Department of Computer Science, University of North Carolina at Charlotte, Charlotte, NC 28223, USA

Abstract

A restaurant dinner may become a memorable experience due to an unexpected aspect enjoyed by the customer, such as an origami-making station in the waiting area. If aspects that are atypical for a restaurant experience were known in advance, they could be leveraged to make recommendations that have the potential to engender serendipitous experiences, further increasing user satisfaction. Although relatively rare, whenever encountered, atypical aspects often end up being mentioned in reviews due to their memorable quality. Correspondingly, in this paper we introduce the task of detecting atypical aspects in customer reviews. To facilitate the development of extraction models, we manually annotate benchmark datasets of reviews in three domains – restaurants, hotels, and hair salons, which we use to evaluate a number of language models, ranging from fine-tuning the instruction-based text-to-text transformer Flan-T5 to zero-shot and few-shot prompting of GPT-3.5.

Keywords

surprise, serendipity, customer reviews, language models

1. Introduction

When looking for a restaurant or a hotel, people are often faced with an overwhelming number of options matching their search constraints. Even when ranked by their average review scores, there may be numerous high quality choices that satisfy the basic search criteria, especially in a metropolitan area. This may lead to choice overload, or overchoice [1, 2], where an individual is presented with a large number of choices that are too difficult to compare, particularly under time constraints [3]. Making a decision in the presence of overchoice becomes mentally exhausting and can lead to subsequent impaired self-regulation [4], decision paralysis, and anxiety [5]. The level of satisfaction that people experience when faced with an increasing number of choices has been observed to follow the well-known Wundt curve [6], an inverted U-shape curve originally relating stimulus intensity with its pleasantness. According to this functional dependency, as the number of choices goes up, satisfaction initially increases and then decreases [7, 8]. In this context, choice overload can be alleviated by reducing the number of consumer choices [9] or by making one option stand out and appear better than the others [10]. To this end, we propose that recommender systems em-

phasize options that possess aspects with the potential to surprise the user in a positive way, i.e., serendipity. Figure 1 illustrates an example where a user Jane is looking for a ramen restaurant in her locality. The system knows that she has been passionate about creating crafts from paper since childhood. Upon being asked for recommendations, the system finds a number of highly rated ramen restaurants, of which Nikita Ramen stands out because it has an origami making station, an atypical aspect for a restaurant, in its waiting area. The system recommends this restaurant to Jane, importantly *without telling her about the origami station*. Upon entering the restaurant, she is very pleasantly surprised to see the origami making station in the waiting area, which brings feelings of nostalgia and happy memories from childhood. She takes some time making various origami figures, before being seated at her table. This serendipitous experience was facilitated by the fact that an origami making station is an atypical aspect for a restaurant, hence it would be experienced as surprising. After the dinner event, the system further confounds her expectations by asking her if she enjoyed the origami making station, which surprises her because she did not expect that the system was responsible for the initial serendipity.

To enable such recommendations with potential for serendipity, the input would ideally consist of three parts:

1. The user's query, be it a standalone request or a turn in a longer conversation. This would be used to find the initial, often large set of items that satisfy the user's information need.
2. The item's data, including not only information about the typical aspects of items from the same category, but also atypical aspects that may generate surprise.

Fifth Knowledge-aware and Conversational Recommender Systems (KaRS) Workshop @ RecSys 2023, September 18–22 2023, Singapore.

*Corresponding author.

✉ snannawa@charlotte.edu (S. Nannaware);

ealhossa@charlotte.edu (E. Al-Hossami);

razvan.bunescu@charlotte.edu (R. Bunescu)

🌐 <https://webpages.charlotte.edu/rbunescu> (R. Bunescu)

🆔 0009-0002-7185-7186 (S. Nannaware); 0000-0002-8436-8974

(E. Al-Hossami); 0000-0003-2919-3566 (R. Bunescu)

© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License

Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

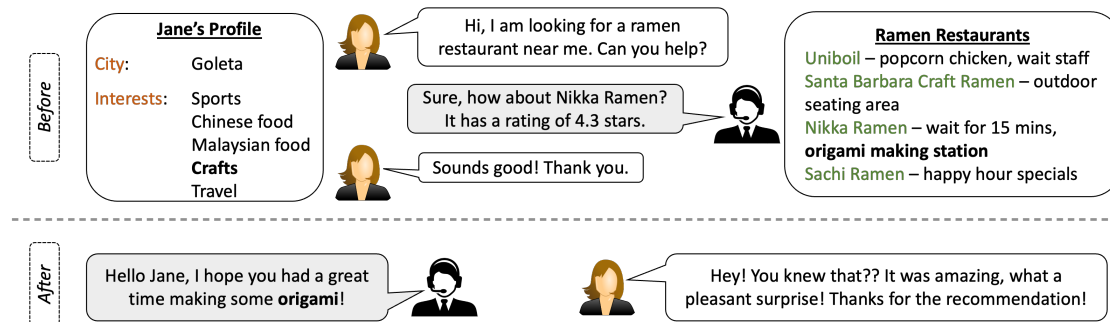


Figure 1: Example of an interaction with the recommender system showing the potential for serendipity of atypical aspects. The user profile on the left shows a user interest (crafts) that is relevant for the atypical aspect (origami) shown on the right.

3. The user’s data, especially in terms of their interests, their likes and dislikes. This would be useful for determining if an atypical aspect would be enjoyed by the user, i.e. serendipitous.

To increase the chance of serendipity, the system would need to (i) have knowledge about the user preferences, and (ii) also ensure that the user notices / takes advantage of the atypical aspect, e.g. estimating that there will be some wait involved in the origami example.

In this paper, we introduce a more focused task where we assume that the category of items requested in the user’s query is known, e.g. restaurants, and the task is confined to using the item’s data to extract aspects that are *atypical* for its category, e.g. origami station for restaurants. Because users’ expectations are shaped by what they think is common for the category of interest in their query, atypical aspects are likely to confound expectations, and hence be perceived as surprising. Henceforth, we will use the term *surprising aspects* to refer solely to *atypical aspects*. Whenever atypical aspects are observed for an item, they tend to be more noticeable and often lead to more memorable experiences. As such, they are likely to be mentioned in customer reviews of that item. Therefore, we use customer reviews as the source of an item’s data. At this time, no user data is used as input, which means that the atypical aspects that are extracted, while surprising for the user searching for that particular category of items, cannot be said for sure to lead to serendipity due to unknown user preferences. In short, the task is that of extracting atypical aspects from customer reviews, where atypical is defined to be relative to a predefined item category. To the best of our knowledge, no prior work has looked into extracting atypical aspects from reviews or any other types of item data.

The rest of the paper proceeds as follows. Section 2 introduces the task of atypical aspect extraction from customer reviews. Section 3 details the development of benchmark datasets of customer reviews that are manually annotated for atypical aspects with respect to three

categories: restaurants, hotels, and hair salons. In Section 4 we describe a number of extraction approaches that rely on language models (LM), ranging from Flan-T5 [11, 12] and ChatGPT [13] in zero-shot or few-shot setting, to fine-tuning of Flan-T5. Experimental evaluations of these models in both extractive and abstractive settings are detailed in Section 5. The paper ends with related work in Section 6 and concluding remarks.

2. Task Definition and Guidelines

Given a domain category, e.g. RESTAURANTS, and a customer review of a particular item in that category, e.g. *a restaurant*, the task is to extract aspects of that particular item that are *atypical of items in its category*. Throughout most of this paper we will use the category of restaurants as an example. All aspects that are related to the core business of a restaurant, including but not limited to food, service, price, opening hours, parking, are considered typical aspects and are not annotated. Conversely, we define and annotate an aspect as atypical if it is not related to the core business of RESTAURANTS, yet it belongs to or is a feature of *the restaurant* (RESTAURANTS refers generically [14] to the restaurant category, whereas *the restaurant* refers to a specific restaurant). Correspondingly, in Table 1 we show samples taken from 4 reviews, illustrating two types of manual annotations, *extractive* and *abstractive*, analogous to the extractive [15, 16] and abstractive [17, 18] annotation schemes widely used in summarization datasets¹. A special case is made of aspects related to the ambience or atmosphere of a restaurant: while ambience might be considered as an important part of, and thus subordinated to, the core business of restaurants, there are cases where ambience aspects stand out and become an attraction on their own. When that happens, we annotate them on a secondary, optional layer.

In the extractive annotation, only base noun phrases

¹<https://duc.nist.gov>, <https://tac.nist.gov>

► A group of work friends and I stumbled upon Upper Deck a little over a year ago and everyone from our office has turned Upper Deck into our local watering hole ever since. Their happy hour special is unbeatable, they have a good selection of draft beers, and the food is out of this world good. The stand out feature of Upper Deck is the offering of **life size beer pong** at their outside patio. This takes traditional beer pong and substitutes solo cups with garbage cans (painted to look like solo cups) and Ping Pong balls with dodgeballs They also have a **pool table** and recently added **arcade games (nfl blitz 99 beats madden 15** all day). Get some friends and bring your appetites and some quarters, you won't be disappointed.

The restaurant offers life size beer pong at their outside patio. They have a pool table. They recently added arcade games, such as nfl blitz 99 and madden 15.

► The big draw of this place is the excellent pizza, which you can have with beer on an outside deck with a view of the **parklands**. It's a nice place to hang out on a sunny afternoon. You can even go for a walk in the **Goose Creek Park** behind the restaurant afterwards to burn off the calories you just consumed. The big minus is that if you go for lunch or in early afternoon, the menu is really limited. This place also seems to attract a goodly number of families with kids at lunch times, probably because it serves pizza and there's a **playground** in the adjoining park,

The restaurant has an outside deck with a view of the parklands. Customers can go for a walk in the Goose Creek Park behind the restaurant. There's a playground in the adjoining park.

► Classic West Philly spot where you can see **local wildlife**. Everyone from **mom**s to **anarchists** to **hackers** to **organic gardeners** to **activists** hangs out there. The coffee is excellent, the baked goods are great as well, and if you're working on something, you might run into a possible collaborator there. If you're thinking of moving to West Philly, definitely check out Satellite and the farmer's market. They've replaced the cracked and chipped cups with awesome new cups, which are awesome.

In this restaurant you can see local wildlife. Everyone from moms to anarchists to hackers to organic gardeners to activists hangs out there.

► This is such a cool place! Three words was all it took to add this gem to my list of places to visit while in St. Louis - "**Good Burger Car!**" YES! They have the car from the movie Good Burger! A movie I was obsessed with as a child & have since gotten my kids to love just as much! The place is made out of **cool, colorful shipping containers** with many neat decorations, what looks like an **alien spaceship** from Toy Story with a laser on it adorned the top of the place along with a **cow**. Now, on to the food. They have many different options to choose from including create-your-own burgers & many of their own creations, sandwiches, salads, sides, kid's meals, & shakes & floats ... Such a unique place & worth a visit! They also sell **souvenir T-shirts & hats**, & my fiance had to get himself a "**HI AF**" shirt. The shirts were heavily influenced by the **movie Good Burger** & there was one in particular I had my heart set on, unfortunately they did not have it in my size.

The restaurant has the car from the movie Good Burger. They sell souvenir T-shirts and hats, and a customer got a "HI AF" shirt. The shirts are heavily influenced by the movie Good Burger.

The place is made out of cool, colorful shipping containers, with many neat decorations, what looks like an alien spaceship from Toy Story with a laser on it adorned the top of the place along with a cow.

Table 1

Examples of **extractive** and **abstractive** annotations of customer reviews in the restaurants domain. Secondary or optional annotations are shown in **blue**.

referring to atypical aspects are annotated. If multiple phrases refer to the same atypical aspects, we only annotate the grounding instance of the coreference chain. For example, the noun phrase "adjoining park" in the second review is not annotated, as it refers to the "Goose Creek Park" which is already annotated. However, if a review mentions an atypical category, such as "arcade games" in the first review or "local wildlife" in the third review, any category instance that is mentioned will also be annotated as atypical, such as "nfl blitz 99" in the first review or "anarchists" in the third review, respectively. If the atypical aspect is a more complex noun phrase, we only annotate the base noun phrase that expresses the semantic core (often the syntactic head), as in "the *Goose Creek Park* behind the restaurant". The extractive annotation is meant to be used together with the original review in downstream applications, which makes it acceptable

to annotate only the most important part of the phrase.

In the abstractive annotation, one or more sentences are generated that enumerate the atypical aspects mentioned in the review. The formulation is kept as close as possible to the original text while maintaining naturalness. The generated sentences are intended to be concise, usually maintaining details that are expressed in the same sentence in the review, however keeping out unimportant information about the atypical aspect that is mentioned in other sentences, or details that are vague or uncertain. Sentiment words are maintained only if it helps keep the text natural and faithful to the original. The abstractive annotation is meant to be standalone and used without the original review in downstream tasks, as such it may require some minimal rewriting of the original review formulation, e.g. adding the phrase "the restaurant", or removing opinion words such as "beats" in

the first example. Sometimes reviewers use metaphors to refer to atypical aspects, in which case it is important that the abstractive version preserves the metaphorical meaning. This is the case for "local wildlife" in the third review, which refers metaphorically to types of customers that are seen relatively less often in that context.

In many aspect-based sentiment analysis approaches [19], identifying *typical* aspects that are mentioned in a review is done explicitly as *aspect term extraction*. However, the task of extracting *atypical* aspects, as introduced above, cannot be solved simply by first (a) identifying all typical aspects of restaurants that are mentioned in a review, followed by (b) extracting all other noun phrases, i.e. phrases that do not refer to a typical aspect of a restaurant. In their reviews, people often mention entities or events that are not associated with the restaurant, such as "our office" in the first review, or "the farmer's market" in the third review, and these phrases should not be extracted either. Thus, it is important that the noun phrase refers to an aspect that is *associated with the reviewed restaurant*, and that at the same time is atypical of restaurants (or unexpected for a restaurant). Finally, while Table 1 may induce the perception that atypical aspects are common, the opposite is actually true. As will be detailed in Section 3 below, it takes going through at least 50 reviews in order to find one review that mentions an atypical aspect. The difficulty of manually finding this "needle in a haystack" further motivates the development of automated approaches for surprising aspect extraction.

3. Manually Annotated Datasets

We used the Yelp dataset [20] as a source of reviews for the 3 target categories: Restaurant (~5M reviews), Hotel (~190K reviews), and Hair Salon (~115K reviews). Because most aspects are expressed as nouns and less frequently as verbs, we use spaCy [21] to collect lemmas of all nouns and verbs and compute their frequencies for each domain. We rank words in ascending order based on their counts and filter out words that appear with very low frequency, e.g., less than 10 times for the Restaurant domain, as these tend to be spelling mistakes or interjections that are purposely misspelled for extra emphasis, e.g., "amaazing". We then consider the remaining rare words in ascending order of their frequency as candidate atypical words, extract the reviews that mention them, and read these reviews to determine which occurrences truly refer to an atypical aspect. When reading a review, all atypical aspects are annotated, not only the ones corresponding to the search word. Notwithstanding the heuristic selection of reviews based on the occurrence of rare words, overall this was still a very time-consuming process, because rare words very often appear in a review without necessarily referring to atypical aspects. For ex-

ample, out of the 43 restaurant reviews that contain the lemma "poncho", in only 1 review the word "ponchos" was deemed to refer to an atypical aspect (the restaurant was selling them). The other 42 reviews contained references to ponchos that were not associated with the restaurant itself, e.g. staff helping customers put their ponchos on a rainy day, or customers describing their arrival at the restaurant on a rainy day. As we went down the list of rare words, their frequency increased, resulting in a larger number of reviews to skim through for each rare word. Overall, for the Restaurant dataset, we used as search words the rare words that appeared with a frequency of up to 187. Upon semi-automatically sifting through the ~97K reviews found to contain these words, we were able to collect 114 reviews that contained atypical aspects. On average, one hour of following this process led to finding between 2 and 3 reviews containing atypical aspects for the restaurant category, whereas for the hair salon category it took on average two hours to find 1 atypical reviews. Henceforth the term *atypical review* will be used to refer to a review that contains one or more atypical aspects; analogously, the term *typical review* will be used to refer to reviews that do not contain any atypical aspect.

As illustrated in the examples from Section 2, we organize annotations of atypical aspects on two layers:

- A *primary* layer that contains atypical aspects that are clearly not connected to any core feature of that domain.
- A *secondary* layer that contains atypical aspects that are related to a typical aspect, such as ambiance or location, but that stand out and are interesting on their own, separate from the core features of the domain.

For example, 'I was even encouraged to visit their *petting zoo in the back*' would be considered a primary atypical aspect in any of the 3 categories, whereas 'There is an *interesting giant stuffed spider that goes up and down when the door leading to the bathrooms opens and closes*' would be annotated as a secondary atypical aspect.

Table 2 shows summary statistics for the 3 datasets, one for each domain (category), split between data used for training and testing, and data used for development. Under the Primary column, we show the number of atypical reviews and atypical aspect annotations contained in them. The next column shows the same statistics for when both primary and secondary atypical aspects are considered. The total number of reviews in each dataset, shown in the third column, is about double the number of primary atypical reviews, reflecting a balanced dataset where the number of typical reviews was selected to be about the same as the number of atypical reviews.

We computed inter-annotator agreement (ITA) on both the extractive and abstractive annotations in the development sets of the Restaurant and Hair Salon domains.

Table 2

Statistics for the 3 datasets, split between Train+Test and Development (Dev). The number of atypical reviews and atypical aspects are presented separately for primary (Primary) vs. both primary and secondary atypical aspects (+ Secondary).

Domain	Dataset split	Primary		+ Secondary		Total reviews
		# reviews	# aspects	# reviews	# aspects	
Restaurant	Train+Test	100	253	107	340	200
	Dev	14	32	16	46	28
Hotel	Train+Test	69	274	85	401	150
	Dev	10	33	11	49	20
Hair Salon	Train+Test	45	147	48	181	90
	Dev	5	24	5	29	10

The ITA metrics are shown in Table 3 and are calculated by assuming one annotator provides the ground truth while the other annotator is considered as the system.

4. LM-based Approaches

This section describes our Language Model (LM) based approaches to detecting surprising aspects in customer reviews. We experimented with 2 language models: FLAN-T5 and ChatGPT. The 3 billion parameter FLAN-T5 is an encoder-decoder transformer based on the T5 model [22] that was further instruction-tuned on the FLAN dataset [11, 12]. We decided to use the FLAN-T5 model due to its exposure to narratives in the style of reviews, e.g., blog posts, during pre-training on the C4 corpus [22], and also due to its instruction-tuning on summarization and sentiment analysis tasks. We also experimented with zero-shot and few-shot prompting of the much larger ChatGPT (gpt-3.5) [13] in order to evaluate the performance of a state-of-the-art language model without any fine tuning.

LMs are known to be sensitive to word choices [23]. In the case of zero-shot/few-shot experiments, we tried between 10 and 20 different prompts for each type of annotation and selected the ones that performed the best on the development data. We noticed that, given the description of the task along with examples of typical aspects in the prompt helped model to identify the multiple atypical aspects present in a review. We also observed that the more examples of typical aspects were given in the prompt the easier it was for the model to identify atypical aspects. We also had to try multiple prompt formulations to instruct the LM to align with an expected output format. For 5-shot experiments, various sets of 5 examples from the development set were first selected at random, then manually checked for diversity in terms of difficulty, number, and types of atypical aspects. Of these, we selected the 5-shot examples that yielded the highest performance on the development set.

With the exception of abstractive generation for Hotels,

which did not benefit from a prompt, the FLAN-T5 fine-tuning experiments employed the following prompts:

► **Fine-tuning FLAN-T5 Extractive Prompt:** *question: Based on the following restaurant review, list aspects that are atypical for a restaurant. Separate them using commas. context: {{Review}}*

► **Fine-tuning FLAN-T5 Abstractive Prompt:** *question: Based on the following restaurant review, what are the atypical aspects for a restaurant? context: {{Review}}*

In the 0-shot setup for ChatGPT, we include an instruction to either extract lists of atypical aspects (extractive) or to generate naturally sounding text about the atypical aspects in the review (abstractive):

► **0-shot ChatGPT Extractive Prompt:** *Given the following restaurant review, can you list atypical aspects for a restaurant? Atypical aspects are not related to service, food, drinks, location, price, menu, discounts, policies, staff, customer satisfaction, or other items commonly associated with a restaurant. Please be precise in your response; it should contain only atypical aspects associated with the restaurant that is reviewed. Extract base noun phrases in the output format: 'Atypical aspects: aspect 1, aspect 2, aspect 3.' Output <None> if there are no atypical aspects. Please follow the output format strictly. Passage: {{Review}}*

► **0-shot ChatGPT Abstractive Prompt:** *Which aspects mentioned in the review are atypical for a restaurant? Unlike common aspects such as service, food, drinks, location, price, menu, discounts, policies, staff, or customer satisfaction, atypical aspects are not commonly associated with a restaurant. In the output, formulate each aspect as sentences, e.g., "Atypical aspects: – The restaurant has <aspect 1>. – The restaurant has <aspect 2>. – The restaurant has <aspect 3>." If there are no atypical aspects, output "None". Passage: {{Review}}*

In the few-shot setup for ChatGPT, we include in the prompt both the instruction and 5 worked-out examples:

► **Few-shot ChatGPT Extractive and Abstractive**

Prompt: *Given the following restaurant review, can you list atypical aspects for a restaurant? Atypical aspects are not related to service, food, drinks, location, price, menu, discounts, policies, staff, customer satisfaction or other types of items that are commonly associated with a restaurant. Please be precise in your response, which should contain only atypical aspects that are associated with the restaurant that is reviewed. Output <None> if there are no atypical aspects.*

Example 1: {{Example Review 1}} Atypical aspects: {{comma-separated extractive annotations OR bullet-listed abstractive sentences}} ... Examples 2, 3, 4, 5

Can you try for the review below? {{Review}}

We use the Hugging Face Transformers package [24] for fine-tuning Flan-T5 with the following hyper-parameters: an effective batch size of 32, a number of epochs of 30, a learning rate of 3e-5 for Restaurants and 5e-5 for Hotels and Hair Salons, a weight decay of 0.001, and a generation max length set to 512. Those hyper-parameter values were found through tuning on the development portion of each dataset. We perform the fine-tuning experiments on a high-performance computing cluster using 8 CPU cores, 128 GB RAM, and 2 A100 80 GB GPUs, for around 96 hours. We use the OpenAI API Python package [25] for ChatGPT, where we do greedy decoding by setting the temperature parameter to 0.

5. Experimental Evaluations

The LM-based approaches are evaluated in a 10-fold scenario where the Train+Test reviews dataset is partitioned into 10 folds, 9 folds are used for training and 1 fold is used for testing. This process is repeated 10 times until each fold in the dataset is used as a test fold. The metrics computed across the 10 folds are then micro-averaged yielding the final evaluation metric. For the extractive evaluation, we report the precision, recall, and F_1 scores for the **exact** and **partial** matches of the extracted base noun phrase (BNP) with the ground truth phrase.

In the exact match method, an extracted phrase is considered correct if and only if it matches exactly a ground truth (gold) phrase. Precision (P) and Recall (R) are then computed as follows:

$$P = \#correct\ extracted\ BNPs / \#extracted\ BNPs$$

$$R = \#correct\ extracted\ BNPs / \#gold\ BNPs$$

In the partial match method, we use the greedy method shown in Algorithm 1 to compute a bipartite matching between gold phrases gp and extracted phrases ep that aimed at maximizing their word overlap $|ep \cap gp|$ in total. The overlaps between extracted and gold phrases are then used to compute precision, recall, and F_1 .

Algorithm 1: PartialMatchMetrics(GP, EP)

Input : Gold Phrases GP , Extracted Phrases EP

Output: Precision (P), Recall (R), and F_1

```

1 // TPe = # True Positives w.r.t EP
2 // TPg = # True Positives w.r.t GP
3 // FP = # False Positives
4 // FN = # False Negatives
5 for gp in GP do
6   if EP is empty then
7     FN ← FN + |gp|/|gp|
8   else
9     Find ep ∈ EP that has maximum Jaccard
      similarity with gp
10    TPe ← TPe + |ep ∩ gp|/|ep|
11    TPg ← TPg + |ep ∩ gp|/|gp|
12    FP ← FP + |ep - gp|/|ep|
13    FN ← FN + |gp - ep|/|gp|
14    Remove ep from the set EP
15 for ep in EP do
16   FP ← FP + |ep|/|ep|
17 P ← TPe/(TPe + FP), R ← TPg/(TPg + FN)
18 return P, R, F1 ← 2PR/(P + R)

```

For the abstractive evaluation, we follow prior work in summarization [26, 27, 28] and compare the generated output with the ground truth using BERT F1 Score [29] instantiated with DeBERTa [30], Rouge-1, Rouge-2, and Rouge-L-Sum [31]. In the case of typical reviews, the prediction, i.e. the generated text, should be empty, reflecting no atypical aspect. When using RougeLsum and BERTScore to evaluate the model output on a typical review, F_1 is calculated as 1.0 if the model generates an empty text, and 0.0 if it generates a non-empty text.

The overall experimental results are shown in Table 3. For Restaurants, we show results on extracting primary atypical aspects as well as results on extracting both primary and secondary atypical aspects. Since fine-tuned Flan-T5 and ChatGPT (5-shot) obtained the best results on Restaurants, they were selected to be evaluated on the other two domains, using solely primary atypical aspects. Fine-tuning FLAN-T5 yields the best performance in the extractive task across all domains. While we observe a big performance gap between ChatGPT and fine-tuned FLAN-T5 in the extractive setting, that gap shrinks considerably in the abstractive setting for Hotels and Hair Salons, where ChatGPT (5-shot) occasionally outperforms FLAN-T5 on some of the metrics. For both LMs, the Hotel domain appears to be more challenging. Compared to the other domains, atypical aspects are more common and more diverse in hotels, likely because hotels try to differentiate themselves from other hotels more than restaurants or hair salons do. Table 2 shows that indeed there are more primary and secondary atypical aspects

Table 3

Extractive (exact vs. partial match) and abstractive results (%), on primary (default) and primary + secondary atypical aspects across the 3 domains. Precision (P), Recall (R), and F1 are reported for the extractive setting. F1 for Rouge-1 (R-1), Rouge-2 (R-2), rougeLsum (RLS) and BERTScore is reported for the abstractive setting. Best results in each domain are in **bold**.

Language Model & Experimental Setup	Exact Match			Partial Match			Abstractive			
	P	R	F1	P	R	F1	R-1	R-2	RLS	BERT
Restaurants										
ChatGPT (0-shot)	22.0	37.7	27.7	27.9	56.8	37.4	35.0	24.0	52.0	66.0
+ secondary	25.1	35.2	29.3	32.5	56.0	41.2	42.0	29.0	53.0	68.0
ChatGPT (5-shot)	33.0	43.6	37.6	39.3	61.7	48.0	42.0	31.0	59.0	71.0
+ secondary	26.8	38.9	31.7	33.5	60.7	43.2	46.0	36.0	57.0	70.0
FLAN-T5 (0-shot)	28.9	23.5	25.9	32.4	28.7	30.5	32.0	24.0	44.0	54.0
+ secondary	27.9	18.5	22.3	32.3	23.9	27.4	30.0	22.0	44.0	56.0
FLAN-T5 (fine-tuned)	67.5	60.2	63.4	72.9	65.3	68.6	58.0	50.0	73.0	79.0
+ secondary	57.7	58.6	56.4	64.2	64.4	62.3	56.0	48.0	68.0	75.0
Human ITA (est.)	80.0	78.26	79.12	85.56	82.61	84.06	65.0	51.0	75.0	80.0
+ secondary	75.4	79.3	77.3	82.4	86.8	84.5	71.0	54.0	76.0	86.0
Hotels										
ChatGPT (5-shot)	30.5	35.2	32.7	34.8	46.9	39.9	34.0	26.0	56.0	65.0
FLAN-T5 (fine-tuned)	60.2	54.9	55.9	63.8	57.6	59.0	34.0	26.0	59.0	63.0
Hair Salons										
ChatGPT (5-shot)	34.7	43.2	38.5	42.4	61.1	50.1	48.0	37.0	57.0	69.0
FLAN-T5 (fine-tuned)	66.6	62.9	63.9	75.1	69.1	71.1	38.0	33.0	61.0	65.0
Human ITA (est.)	92.3	82.8	87.3	94.2	85.1	89.4	59.0	41.0	78.0	91.0
+ secondary	74.2	67.6	70.8	79.0	71.8	75.3	65.0	47.0	77.0	91.0

per review in the hotel domain.

To determine how well Flan-T5 generalizes to unseen atypical aspects in the Restaurant domain, we manually created groupings of atypical aspects where semantically similar atypical aspects, e.g. greeting cards and anniversary gifts, are grouped together, such that aspects in different groups are semantically very different. We then partition the set of groups into 10 folds of groups, which ensures that the atypical aspects that the language model sees in the test fold have not been seen during training (either literally or semantically similar). Upon fine-tuning and evaluating FLAN-T5 on this dataset, we observe a similar precision as reported in Table 3, however, there is a significant drop in recall from 60.2 to 46.1 for primary atypical aspects and from 58.6 to 49.3 when extracting both primary and secondary atypical aspects. Improving generalization to semantically novel atypical aspects is therefore an interesting avenue for future work.

To get a sense of the real performance in the abstractive setting for primary aspects, we also performed a manual evaluation of the fine-tuned Flan-T5 and 5-shot ChatGPT outputs. To enable calculation of precision and recall, we manually label and count the following: (A) Atypical aspects from the model output that are semantically similar to gold aspects; (B) Gold aspects that are missing from the model output; (C) Typical aspects or other types of entities that should not have been extracted as atypical; (D) Extra details about atypical aspects; (E) Secondary

atypical aspects; (F) Typical reviews for which the model correctly generates an empty string. The first 4 types are illustrated on the example below, where A = 2, B = 1, C = 1, D = 1, and E = 1:

- Gold = ["On the weekends kids can experience 'Mark the Balloon Guy' at the restaurant. They have stuffed animal/puppets for sale at the front.^B", "The restaurant has therapeutic sketching at every table."]
- Output = ["The restaurant has a balloon artist on weekends.^A The restaurant has traditional wooden baseball stadium seats for waiting.^E", "The restaurant offers therapeutic sketching at every table.^A If your sketch is good, it will be on the wall.^D The restaurant has limited menu options before 4:30 pm.^C"]

These counts are then used to compute a lenient precision as $P = (A + F)/(A + F + C)$ and a strict precision as $P = (A + F)/(A + F + C + D + E)$, whereas recall is computed as $R = (A + F)/(A + F + B)$. Correspondingly, the fine-tuned Flan-T5 obtains a lenient and strict F_1 of 77.5% and 76.5%, respectively, whereas ChatGPT obtains a lenient and strict F_1 of 71.3% and 65.9%, respectively. Looking at Table 3, the strict F_1 is in between RLS and BERT, further supporting the use of these automated metrics. Overall, manual evaluation shows that models perform better in the abstractive vs. extractive setting.

Error analysis reveals that fine-tuned Flan-T5 is more succinct in its answers, leading it to sometimes ignore

atypical aspects in its response. Conversely, ChatGPT tends to be more verbose, often generating unnecessary details about the atypical aspects that it extracts, or mistaking typical for atypical aspects. For instance, ChatGPT extracted a food-related aspect in the second sentence in *"The restaurant has an annual Valentine's Date Night with table service, a caricature artist, and a piano player. The restaurant offers in-house made chocolate-covered strawberries."* ChatGPT also tends to extract one-time customer experiences which are not atypical aspects of a restaurant, as shown in the second sentence in *"The restaurant has a sister restaurant next door with a lively fun band. The server, Peter, gave very helpful suggestions."* On the other hand, fine-tuned FLAN-T5 mistakenly considers *"farmer's market"* to be an atypical aspect in the third review from Table 1. Also, when the sole atypical aspect of a business is in its proximity, the fine-tuned FLAN-T5 tends to not extract it, as in the model failing to generate the ground truth abstractive sentences *"They are close to a dock area where customers can board paddle cruises."*, *"Customers can catch an airboat ride down the road."*

6. Related Work

Addressing overchoice is a core focus of recommender systems and it is typically addressed by recommending between 5 to 20 attractive and diverse items [32], based on user preferences [33, 34], user ratings, item attributes, or user reviews [35, 36, 37]. Recommending items with potential for serendipity is one way of diversifying an item set. In [38], unexpectedness is defined as the distance of an item from a set of obvious items for that user, relative to the user's preferred level of unexpectedness. Li et al. [39] recommended unexpected items by modeling user interests as clusters of historical data in a latent space and calculating the weighted distance between a new item and the clusters of interests. Kotkov et al. [40] crowd-sourced serendipity labels for a movie dataset using multiple definitions of serendipity.

User reviews have been used to learn latent features of users [37], extract sentiment [36], derive user preferences [34], or to perform aspect-based sentiment analysis in order to recommend better quality products with aspects relevant to the user [35]. Conversational recommender systems use reviews to provide explanations [41], to maintain fluency in conversation [42], or to understand the user's requirements by asking questions about aspects mentioned in reviews [43].

Aspect-based sentiment analysis (ABSA) is a technique that aims to extract topic-specific aspects mentioned in a review, together with any associated sentiment. There are different approaches to extracting aspect terms, with some focusing on using term frequency in the corpus [44, 45] or using semantic clustering of prominent

aspects [46]. Qiu et al. [47] only extract aspects that have an expressed opinion, using expanded opinion lexicons containing adjectives. As detailed at the end of Section 2, extracting atypical aspects cannot be addressed simply as the logical complement of ABSA.

Prior work investigated using Language Models (LMs) in recommender systems. Zhang et al. [48] investigate using LMs as a recommender system by formulating a movie recommendation task as a multi-token cloze task and find that LMs underperform traditional recommender systems such as GRU4Rec in both zero-shot and fine-tuned settings. LMs have also been used to elicit user preferences in natural language given historical interactions and prior selections for more personalized recommendations [49]. Such a model can benefit a complete implementation of the system illustrated in Figure 1, where it would automatically infer Jane's interests. Furthermore, LMs have been used for explaining why a specific recommendation was made to the user [50].

Atypical aspects are valuable due to their potential to create serendipity, i.e. experiences that are both unexpected (surprising) and relevant (positive) for the user. Unlike Kotkov et al. [40], we do not consider novelty to be a required dimension of serendipity. Finally, surprise can appear from other sources, be it the experience of a typical aspect that stands out (such as a unique dish offered by a restaurant), or the accidental discovery of a restaurant as shown in the first sentence in Table 1 and studied in [51]. All of these different sources of surprise have great potential for setting up serendipity.

7. Conclusion and Future Work

We introduced the new task of extracting atypical aspects from customer reviews. To enable training and evaluation of atypical aspect extraction models, we manually annotated two layers of atypical aspects in customer reviews from three domains. While experimental evaluations using few-shot prompting of ChatGPT and fine-tuning of Flan-T5 show promising results, there is still a substantial gap relative to human performance, as shown by the higher ITA. Future work includes enhancing reproducibility by using open-source LMs [52], prototyping a recommender system that leverages atypical aspects, and a user study verifying their utility.

To facilitate reproducibility and future progress, we make the code and the datasets publicly available at <https://github.com/smitanannaware/XtrAtA>.

Acknowledgments

This research was partly supported by the United States Air Force (USAF) under Contract No. FA8750-21-C-0075.

References

- [1] A. Toffler, *Future shock*, Random House, New York, 1970.
- [2] A. Chernev, U. Böckenholt, J. Goodman, Choice overload: A conceptual review and meta-analysis, *Journal of Consumer Psychology* 25 (2015) 333–358. URL: <https://www.sciencedirect.com/science/article/pii/S1057740814000916>. doi:<https://doi.org/10.1016/j.jcps.2014.08.002>.
- [3] Y. Inbar, S. Botti, K. Hanko, Decision speed and choice regret: When haste feels like waste, *Journal of Experimental Social Psychology* 47 (2011) 533–540. URL: <https://www.sciencedirect.com/science/article/pii/S002210311100028X>. doi:<https://doi.org/10.1016/j.jesp.2011.01.011>.
- [4] K. D. Vohs, R. F. Baumeister, B. J. Schmeichel, J. M. Twenge, N. M. Nelson, D. M. Tice, Making choices impairs subsequent self-control: a limited-resource account of decision making, self-regulation, and active initiative, *Journal of Personality and Social Psychology* 94 (2008) 883–898. doi:10.1037/0022-3514.94.5.883.
- [5] S. S. Iyengar, M. R. Lepper, When choice is demotivating: Can one desire too much of a good thing?, *Journal of Personality and Social Psychology* 79 (2000) 995–1006. doi:10.1037/0022-3514.79.6.995.
- [6] D. E. Berlyne, Aesthetics and psychobiology, *Journal of Aesthetics and Art Criticism* 31 (1973).
- [7] A. M. Shah, G. L. Wolford, Buying behavior as a function of parametric variation of number of choices, *Psychological Science* 18 (2007) 369 – 370.
- [8] D. Kaimann, N. Stroh-Maraun, J. Cox, Variety in the video game industry: An empirical study of the wundt curve, *Managerial and Decision Economics* 39 (2018) 354–362. doi:<https://doi.org/10.1002/mde.2909>.
- [9] B. Schwartz, *The paradox of choice: Why more is less*, Harper Perennial, New York, 2004.
- [10] B. Scheibehenne, R. Greifeneder, P. M. Todd, Can There Ever Be Too Many Options? A Meta-Analytic Review of Choice Overload, *Journal of Consumer Research* 37 (2010) 409–425. URL: <https://doi.org/10.1086/651235>. doi:10.1086/651235, eprint: <https://academic.oup.com/jcr/article-pdf/37/3/409/5173186/37-3-409.pdf>.
- [11] J. Wei, M. Bosma, V. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, Q. V. Le, Finetuned language models are zero-shot learners, in: *International Conference on Learning Representations*, 2022. URL: <https://openreview.net/forum?id=gEZrGCozdqR>.
- [12] H. W. Chung, L. Hou, S. Longpre, B. Zoph, Y. Tay, W. Fedus, Y. Li, X. Wang, M. Dehghani, S. Brahma, A. Webson, S. S. Gu, Z. Dai, M. Suzgun, X. Chen, A. Chowdhery, A. Castro-Ros, M. Pellat, K. Robinson, D. Valter, S. Narang, G. Mishra, A. Yu, V. Zhao, Y. Huang, A. Dai, H. Yu, S. Petrov, E. H. Chi, J. Dean, J. Devlin, A. Roberts, D. Zhou, Q. V. Le, J. Wei, Scaling instruction-finetuned language models, 2022. arXiv:2210.11416.
- [13] OpenAI, *Introducing chatgpt*, 2022. URL: <https://openai.com/blog/chatgpt>.
- [14] G. Carlson, F. Pelletier, *The Generic Book*, University of Chicago Press, 1995.
- [15] H. Jing, K. R. McKeown, The decomposition of human-written summary sentences, in: *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '99*, Association for Computing Machinery, New York, NY, USA, 1999, p. 129–136. URL: <https://doi.org/10.1145/312624.312666>. doi:10.1145/312624.312666.
- [16] K. Knight, D. Marcu, Summarization beyond sentence extraction: A probabilistic approach to sentence compression, *Artificial Intelligence* 139 (2002) 91–107. URL: <https://www.sciencedirect.com/science/article/pii/S0004370202002229>. doi:[https://doi.org/10.1016/S0004-3702\(02\)00222-9](https://doi.org/10.1016/S0004-3702(02)00222-9).
- [17] H. Jing, Using hidden Markov modeling to decompose human-written summaries, *Computational Linguistics* 28 (2002) 527–543. URL: <https://aclanthology.org/J02-4006>. doi:10.1162/089120102762671972.
- [18] A. M. Rush, S. Chopra, J. Weston, A neural attention model for abstractive sentence summarization, in: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Lisbon, Portugal, 2015, pp. 379–389. URL: <https://aclanthology.org/D15-1044>. doi:10.18653/v1/D15-1044.
- [19] SemEval, Task description: Aspect based sentiment analysis (absa) < semeval-2014 task 4, 2014. URL: <https://alt.qcri.org/semeval2014/task4/>.
- [20] Yelp, *The yelp dataset*, 2023. URL: <https://www.yelp.com/dataset>.
- [21] spaCy, *Industrial-strength natural language processing*, 2023. URL: <https://spacy.io/>.
- [22] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, *The Journal of Machine Learning Research* 21 (2020) 5485–5551.
- [23] A. Webson, E. Pavlick, Do prompt-based models really understand the meaning of their prompts?, in: *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Association for Computational Linguistics,

- tics, Seattle, United States, 2022, pp. 2300–2344. URL: <https://aclanthology.org/2022.naacl-main.167>. doi:10.18653/v1/2022.naacl-main.167.
- [24] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, et al., Transformers: State-of-the-art natural language processing, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, 2020, pp. 38–45.
- [25] OpenAI, Openai python library, 2022. URL: <https://github.com/openai/openai-python>.
- [26] R. Tangsali, A. J. Vyawahare, A. V. Mandke, O. R. Litake, D. D. Kadam, Abstractive approaches to multidocument summarization of medical literature reviews, in: Proceedings of the Third Workshop on Scholarly Document Processing, Association for Computational Linguistics, Gyeongju, Republic of Korea, 2022, pp. 199–203. URL: <https://aclanthology.org/2022.sdp-1.24>.
- [27] O. Ahuja, J. Xu, A. Gupta, K. Horecka, G. Durrett, ASPECTNEWS: Aspect-oriented summarization of news documents, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 6494–6506. URL: <https://aclanthology.org/2022.acl-long.449>. doi:10.18653/v1/2022.acl-long.449.
- [28] K. Mrini, C. Liu, M. Dreyer, Rewards with negative examples for reinforced topic-focused abstractive summarization, in: Proceedings of the Third Workshop on New Frontiers in Summarization, Association for Computational Linguistics, Online and in Dominican Republic, 2021, pp. 33–38. URL: <https://aclanthology.org/2021.newsum-1.4>. doi:10.18653/v1/2021.newsum-1.4.
- [29] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, Y. Artzi, Bertscore: Evaluating text generation with bert, in: International Conference on Learning Representations, 2020. URL: <https://openreview.net/forum?id=SkeHuCVFDr>.
- [30] P. He, X. Liu, J. Gao, W. Chen, Deberta: Decoding-enhanced bert with disentangled attention, arXiv preprint arXiv:2006.03654 (2020).
- [31] C.-Y. Lin, ROUGE: A package for automatic evaluation of summaries, in: Text Summarization Branches Out, Association for Computational Linguistics, Barcelona, Spain, 2004, pp. 74–81. URL: <https://aclanthology.org/W04-1013>.
- [32] D. Bollen, B. P. Knijnenburg, M. C. Willemsen, M. Graus, Understanding choice overload in recommender systems, in: Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys ’10, Association for Computing Machinery, New York, NY, USA, 2010, p. 63–70. URL: <https://doi.org/10.1145/1864708.1864724>. doi:10.1145/1864708.1864724.
- [33] N. Jakob, S. H. Weber, M. C. Müller, I. Gurevych, Beyond the stars: Exploiting free-text user reviews to improve the accuracy of movie recommendations, in: Proceedings of the 1st International CIKM Workshop on Topic-Sentiment Analysis for Mass Opinion, TSA ’09, Association for Computing Machinery, New York, NY, USA, 2009, p. 57–64. URL: <https://doi.org/10.1145/1651461.1651473>. doi:10.1145/1651461.1651473.
- [34] Y. Zhang, G. Lai, M. Zhang, Y. Zhang, Y. Liu, S. Ma, Explicit factor models for explainable recommendation based on phrase-level sentiment analysis, in: Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR ’14, Association for Computing Machinery, New York, NY, USA, 2014, p. 83–92. URL: <https://doi.org/10.1145/2600428.2609579>. doi:10.1145/2600428.2609579.
- [35] C. Musto, M. de Gemmis, G. Semeraro, P. Lops, A multi-criteria recommender system exploiting aspect-based sentiment analysis of users’ reviews, in: Proceedings of the Eleventh ACM Conference on Recommender Systems, RecSys ’17, Association for Computing Machinery, New York, NY, USA, 2017, p. 321–325. URL: <https://doi.org/10.1145/3109859.3109905>. doi:10.1145/3109859.3109905.
- [36] S. Selmene, Z. Kodia, Recommender system based on user’s tweets sentiment analysis, in: Proceedings of the 4th International Conference on E-Commerce, E-Business and E-Government, ICEEG ’20, Association for Computing Machinery, New York, NY, USA, 2020, p. 96–102. URL: <https://doi.org/10.1145/3409929.3414744>. doi:10.1145/3409929.3414744.
- [37] P. Li, A. Tuzhilin, Learning latent multi-criteria ratings from user reviews for recommendations, IEEE Transactions on Knowledge and Data Engineering 34 (2022) 3854–3866. doi:10.1109/TKDE.2020.3030623.
- [38] A. Panagiotis, T. Alexander, On unexpectedness in recommender systems: Or how to expect the unexpected, in: Workshop on Novelty and Diversity in Recommender Systems, volume 2011, 2011. URL: <https://ceur-ws.org/Vol-816/paper2.pdf>.
- [39] P. Li, M. Que, Z. Jiang, Y. HU, A. Tuzhilin, Purs: Personalized unexpected recommender system for improving user satisfaction, in: Proceedings of the 14th ACM Conference on Recommender Systems, RecSys ’20, Association for Computing Machinery, New York, NY, USA, 2020, p. 279–288. URL: <https://doi.org/10.1145/3383313.3412238>. doi:10.1145/3383313.3412238.

- [40] D. Kotkov, J. A. Konstan, Q. Zhao, J. Veijalainen, Investigating serendipity in recommender systems based on real user feedback, in: Proceedings of the 33rd Annual ACM Symposium on Applied Computing, SAC '18, Association for Computing Machinery, New York, NY, USA, 2018, p. 1341–1350. URL: <https://doi.org/10.1145/3167132.3167276>. doi:10.1145/3167132.3167276.
- [41] C. Musto, P. Lops, M. de Gemmis, G. Semeraro, Justifying recommendations through aspect-based sentiment analysis of users reviews, in: Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization, UMAP '19, Association for Computing Machinery, New York, NY, USA, 2019, p. 4–12. URL: <https://doi.org/10.1145/3320435.3320457>. doi:10.1145/3320435.3320457.
- [42] Y. Lu, J. Bao, Y. Song, Z. Ma, S. Cui, Y. Wu, X. He, RevCore: Review-augmented conversational recommendation, in: Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021, Association for Computational Linguistics, Online, 2021, pp. 1161–1173. URL: <https://aclanthology.org/2021.findings-acl.99>. doi:10.18653/v1/2021.findings-acl.99.
- [43] Y. Zhang, X. Chen, Q. Ai, L. Yang, W. B. Croft, Towards conversational search and recommendation: System ask, user respond, in: Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM '18, Association for Computing Machinery, New York, NY, USA, 2018, p. 177–186. URL: <https://doi.org/10.1145/3269206.3271776>. doi:10.1145/3269206.3271776.
- [44] K. Bauman, B. Liu, A. Tuzhilin, Aspect based recommendations: Recommending items with the most valuable aspects based on user reviews, in: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '17, Association for Computing Machinery, New York, NY, USA, 2017, p. 717–725. URL: <https://doi.org/10.1145/3097983.3098170>. doi:10.1145/3097983.3098170.
- [45] L. S. d. Souza, M. G. Manzato, Aspect-based summarization: An approach with different levels of details to explain recommendations, in: Proceedings of the Brazilian Symposium on Multimedia and the Web, WebMedia '22, Association for Computing Machinery, New York, NY, USA, 2022, p. 202–210. URL: <https://doi.org/10.1145/3539637.3557002>. doi:10.1145/3539637.3557002.
- [46] Z. Luo, S. Huang, F. F. Xu, B. Y. Lin, H. Shi, K. Zhu, ExtRA: Extracting prominent review aspects from customer feedback, in: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 3477–3486. URL: <https://aclanthology.org/D18-1384>. doi:10.18653/v1/D18-1384.
- [47] G. Qiu, B. Liu, J. Bu, C. Chen, Opinion Word Expansion and Target Extraction through Double Propagation, Computational Linguistics 37 (2011) 9–27. URL: https://doi.org/10.1162/coli_a_00034. doi:10.1162/coli_a_00034. arXiv:https://direct.mit.edu/coli/article-pdf/37/1/9/1810309/coli_a_00034.pdf.
- [48] Y. Zhang, H. Ding, Z. Shui, Y. Ma, J. Zou, A. Deoras, H. Wang, Language models as recommender systems: Evaluations and limitations (2021).
- [49] Z. Chen, Palr: Personalization aware llms for recommendation, arXiv preprint arXiv:2305.07622 (2023).
- [50] Y. Gao, T. Sheng, Y. Xiang, Y. Xiong, H. Wang, J. Zhang, Chat-rec: Towards interactive and explainable llms-augmented recommender system, arXiv preprint arXiv:2303.14524 (2023).
- [51] Z. Fu, X. Niu, L. Yu, Wisdom of crowds and fine-grained learning for serendipity recommendations, in: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23, Association for Computing Machinery, New York, NY, USA, 2023, p. 739–748. URL: <https://doi.org/10.1145/3539618.3591787>. doi:10.1145/3539618.3591787.
- [52] A. Liesenfeld, A. Lopez, M. Dingemane, Opening up ChatGPT: Tracking openness, transparency, and accountability in instruction-tuned text generators, in: Proceedings of the 5th International Conference on Conversational User Interfaces, CUI '23, Association for Computing Machinery, New York, NY, USA, 2023. URL: <https://doi.org/10.1145/3571884.3604316>. doi:10.1145/3571884.3604316.