

# Leveraging aspect phrase embeddings for cross-domain review rating prediction

Aiqi Jiang<sup>Corresp., 1</sup>, Arkaitz Zubiaga<sup>2</sup>

<sup>1</sup> University of Warwick, Coventry, United Kingdom

<sup>2</sup> Queen Mary University of London, London, United Kingdom

Corresponding Author: Aiqi Jiang

Email address: A.Jiang.2@warwick.ac.uk

Online review platforms are a popular way for users to post reviews by expressing their opinions towards a product or service, as well as they are valuable for other users and companies to find out the overall opinions of customers. These reviews tend to be accompanied by a rating, where the star rating has become the most common approach for users to give their feedback in a quantitative way, generally as a likert scale of 1-5 stars. In other social media platforms like Facebook or Twitter, an automated review rating prediction system can be useful to determine the rating that a user would have given to the product or service. Existing work on review rating prediction focuses on specific domains, such as restaurants or hotels. This, however, ignores the fact that some review domains which are less frequently rated, such as dentists, lack sufficient data to build a reliable prediction model. In this paper, we experiment on 12 datasets pertaining to 12 different review domains of varying level of popularity to assess the performance of predictions across different domains. We introduce a model that leverages aspect phrase embeddings extracted from the reviews, which enables the development of both in-domain and cross-domain review rating prediction systems. Our experiments show that both of our review rating prediction systems outperform all other baselines. The cross-domain review rating prediction system is particularly significant for the least popular review domains, where leveraging training data from other domains leads to remarkable improvements in performance. The in-domain review rating prediction system is instead more suitable for popular review domains, provided that a model built from training data pertaining to the target domain is more suitable when this data is abundant.

# Leveraging Aspect Phrase Embeddings for Cross-Domain Review Rating Prediction

Aiqi Jiang<sup>1</sup> and Arkaitz Zubiaga<sup>2</sup>

<sup>1</sup>University of Warwick, Coventry, United Kingdom

<sup>2</sup>Queen Mary University of London, London, United Kingdom

Corresponding author:

Aiqi Jiang<sup>1</sup>

Email address: A.Jiang.2@warwick.ac.uk

## ABSTRACT

Online review platforms are a popular way for users to post reviews by expressing their opinions towards a product or service, and they are valuable for other users and companies to find out the overall opinions of customers. These reviews tend to be accompanied by a rating, generally as a Likert scale of 1-5 stars. Existing work on review rating prediction focuses on specific domains, such as restaurants or hotels. This, however, ignores the fact that some review domains which are less frequently rated, such as dentists, lack sufficient data to build a reliable prediction model. In this paper, we experiment on 12 datasets pertaining to 12 different review domains of varying level of popularity to assess the performance of predictions across different domains. We introduce a model that leverages aspect phrase embeddings extracted from the reviews, which enables the development of both in-domain and cross-domain review rating prediction systems. Our experiments show that both of our review rating prediction systems outperform all other baselines. The cross-domain review rating prediction system is particularly significant for the least popular review domains, where leveraging training data from other domains leads to remarkable improvements in performance. The in-domain review rating prediction system is instead more suitable for popular review domains, provided that a model built from training data pertaining to the target domain is more suitable when this data is abundant.

## INTRODUCTION

In recent years, the advent and the prevalent popularisation of social media has led to a change in users' habits of surfing the Internet (Kaplan and Haenlein, 2010; Quan-Haase and Young, 2010; Perrin, 2017; Goss, 2017). Since the emergence of social media platforms, Internet users are no longer limited to browsing online contents as mere readers, but they also can also contribute by expressing and sharing their opinions (Salehan et al., 2017). Users can freely post comments and share experiences on target entities such as products, businesses or events on online review platforms like Yelp.com or Amazon.com (Salehan and Kim, 2016; Xing et al., 2019). These reviews present the subjective opinions of people on products or businesses, which are invaluable for consumers and companies (Sparks and Browning, 2011). Given the volume of these reviews and the fact that they are spread on different sites across the Internet, it becomes more challenging and costly to aggregate all the reviews on a particular product or business (Zhang and Qu, 2013). To alleviate the cost of this task, there is a need to explore the development of automated review rating prediction systems.

There has been work in the scientific literature looking at review rating prediction (Li et al., 2011; Fan and Khademi, 2014; Seo et al., 2017a; Wang et al., 2018; Xing et al., 2019; Wang et al., 2019). This work has however been limited to the prediction of ratings of reviews in very popular domains, such as restaurants (Ganu et al., 2009; Zhang et al., 2018b; Laddha and Mukherjee, 2018; Xing et al., 2019), hotels (Zhao et al., 2016; Laddha and Mukherjee, 2018) or movies (Ning et al., 2018). For these domains, it is relatively easy to get a large-scale dataset from online sources for training a review rating prediction system, thanks to sites like TripAdvisor or Yelp where large collections of rated reviews are publicly accessible. The task can however become more challenging for less popular domains, where the dearth of rated reviews available online inevitably means that the scarcity of labelled data available to train a

rating prediction model will be rather limited. Moreover, the variance in vocabulary across different domains makes it more difficult to develop a prediction system that generalises to different domains; while reviewers are expected to use phrases like *well written* or *entertaining book* to express that they liked a book, a different vocabulary is expected to indicate that they liked a dentist, such as *careful treatment* or *clean office*. Our work builds on the idea that these phrases, associated with different aspects that vary across domains, can be effectively leveraged for the review rating prediction if the barriers between domains are reduced.

To the best of our knowledge, review rating prediction for non-popular domains has not been studied in previous work, and our work is the first attempt to do so. We propose to pursue review rating prediction for non-popular domains by developing a cross-domain rating prediction system, where rated reviews from popular domains can be leveraged to build a model which can then be generalised to non-popular domains. To facilitate and ensure the effectiveness of building a model that will generalise to other domains, we propose an approach for generating aspect phrase embeddings and polarised aspect phrase embeddings, where phrases that vary across domains can be brought to a common semantic space by using word embeddings.

In this work we make the following contributions:

- We introduce the first cross-domain review rating prediction system that creates semantic representations of aspect phrases using word embeddings to enable training from large-scale, popular domains, to then be applied to less popular domains where labelled data is limited.
- We perform experiments with 12 datasets pertaining to 12 different types of businesses of different levels of popularity.
- Our analysis shows that, while an in-domain review rating prediction system is better for popular domains, for the least popular domains our cross-domain review rating prediction system leads to improved performance when we make use of aspect phrase embeddings.
- Different from an in-domain prediction system, our cross-domain system can be effectively applied to a wide range of product domains found on the Internet that do not necessarily have proper review rating systems to collect labelled data from.
- Our classifier outperforms the results of A<sup>3</sup>NCF (Cheng et al., 2018), a state-of-the-art rating prediction system, in 11 out of 12 domains, showing its generalisability across domains of different levels of popularity.

## RELATED WORK

There has been a body of research in sentiment analysis in recent years (Liu and Zhang, 2012). This research has worked in different directions by looking into lexicon-based approaches (Hu and Liu, 2004; Taboada et al., 2011), machine learning methods (Pang et al., 2002; Ye et al., 2009; Tripathy et al., 2016) and deep learning techniques (Wang et al., 2016; Poria et al., 2017; Zhang et al., 2018a; Xing et al., 2019). Different from the sentiment analysis task, the review rating prediction task consists of determining the score –as a rating between 1 and 5– that a user would give to a product or business, having the review text as input. While this may at first seem like a fine-grained sentiment analysis task, which is a translation of a textual view to a numerical perspective and consists of choosing one of five labels rather than three labels (positive, neutral, negative), there are two key characteristics that make the review rating prediction task different. First, the sentiment of a review is not necessarily the same as the rating, as a user may express a positive attitude when sharing a low rating opinion of a business, e.g., “I very much enjoyed my friend’s birthday celebration, however the food here is below standard and we won’t be coming back”. And second, a review tends to discuss different aspects of a business, some of which may be more important than others towards the rating score, e.g., in a review saying that “the food was excellent and we loved the service, although that makes the place a bit pricey”, a user may end up giving it a relatively high score given that key aspects such as the food and the service were considered very positive. In addition, the review can discuss different aspects, and the set of aspects discussed in each review can vary, with some users not caring about certain aspects; e.g., one user focuses on food while another one focuses more on price (Ganu et al., 2009). Using the same star rating mode to express the score of specific features can be more helpful to find users’ interests. Hence, we argue that a review rating prediction system needs

to consider the opinions towards different aspects of a business. We achieve this by extracting aspect phrases mentioned in different sentences or segments of a review, which are then aggregated to determine the overall review rating. Despite the increasing volume of research in review rating prediction, which we discuss in what follows, research in cross-domain review rating prediction is still unstudied. This is particularly important for less popular review domains where labelled data is scarce and hence one may need to exploit labelled data from a different, more popular domain for training a classifier.

Review rating prediction has commonly been regarded as a classification task, but also as a regression task on a few occasions (Li et al., 2011; Pang and Lee, 2005; Mousavizadeh et al., 2015). The text of the review is the input that is consistently used by most existing works (Qu et al., 2010; Leung et al., 2006; Mousavizadeh et al., 2015), however many of them also incorporate other features extracted from the product being reviewed or from the user writing the review (Qu et al., 2010). Ganu et al. (2009) improved the review rating prediction accuracy by implementing new ad-hoc and regression-based recommendation measures, where the aspect of user reviews is considered; their study was however limited to only restaurant reviews and other domains were not studied. A novel kind of bag-of-opinions representation (Qu et al., 2010) was proposed and each opinion consists of three parts, namely a root term, a set of modifier terms from the same sentence, and one or more negation terms. This approach shows a better performance than prior state-of-the-art techniques for review rating prediction. Datasets including three review domains were used for their experiments, however they ran separate experiments for each domain and did not consider the case of domains lacking training data. Similarly, Wang et al. (2010) performed experiments predicting the ratings of hotel reviews, using a regression model that looked at different aspects discussed in a review, with the intuition that different aspects would have different levels of importance towards determining the rating.

Recent years have seen an active interest in researching approaches to review rating prediction but are still limited to popular domains and do not consider the case of domains lacking training data. In a number of cases, features extracted from products and users are being used (Jin et al., 2016; Lei et al., 2016; Seo et al., 2017b; Wang et al., 2018), which limits the ability to apply a prediction system to new domains and to unseen products and users. Tang et al. (2015) studied different ways of combining features extracted from the review text and the user posting the review. They introduced the User-Word Composition Vector Model (UWCVM), which considers the user posting a review to determine the specific use that a user makes of each word. While this is a clever approach to consider differences across users, it also requires observing reviews posted by each user beforehand, and cannot easily generalise to new, unseen users, as well as to new review domains where we lack any information about those users. An approach that predicts review ratings from the review text alone is that by Fan and Khademi (2014). They performed a study of two kinds of features: bag-of-words representations of reviews, as well as part-of-speech tags of the words in the review. They studied how the top K keywords in a dataset contributed to the performance of the rating prediction system, finding that a multinomial logistic regression classifier using the top 100 keywords performed best. Others have used topic modelling techniques like LDA or PLSA to identify the aspects that users discuss in restaurant reviews, however they did not study their effectiveness for review rating prediction (Titov and McDonald, 2008; Huang et al., 2014; Zhou et al., 2015; Vargas and Pardo, 2018).

One of the most recent methods by Cheng et al. (2018) combines topic modelling with a neural network classifier. The approach, called A<sup>3</sup>NCF, extracts users' preferences and items' characteristics on different aspects from reviews. An LDA topic modelling layer is used to then generate topic-specific representations of users and items. Subsequently, the method uses an attention network which combines these user and item features extracted from reviews. This method is shown to outperform an earlier method by Catherine and Cohen (2017), called TransNet, which uses two convolutional neural networks to learn latent representations of users and items. We will use the approach by Cheng et al. (2018), namely A<sup>3</sup>NCF, as a state-of-the-art baseline classifier.

In this paper, we propose to perform a textual analysis of reviews that goes beyond the sole exploration of the whole text as a unit. Our method also looks at the aspect phrases mentioned in the text, which contain the core opinionated comments. This is the case of *good food* for a restaurant or *comfortable bed* for a hotel. While these differ substantially across review domains, we propose to leverage aspect phrase embeddings to achieve representations of aspect phrases to enable generalisation. This allows us to perform cross-domain review rating experiments, where one can build a rating prediction model out of larger training data thanks to leveraging labelled data from other domains. To the best of our knowledge,

ours is the first work to perform rating predictions for review domains with scarce training data, such as dentists, as well as the first to propose and test a cross-domain review rating prediction system.

## DATASETS

To enable our analysis of review rating prediction over different domains, we make use of 12 different datasets, each associated with a different domain. We use three different data sources to generate our datasets: (1) a publicly available collection of 6 million reviews from Yelp<sup>1</sup>, (2) a collection of more than 142 million reviews from Amazon provided by McAuley et al. (2015); He and McAuley (2016), and (3) a collection of more than 24 million reviews retrieved from businesses listed in TripAdvisor's top 500 cities. We filter 12 categories of reviews from these 3 datasets, which gives us 12 different datasets that enable us to experiment with review rating prediction over 12 different types of businesses and products. All of the datasets include both the review text and the review rating in a scale from 1 to 5. The use of a standardised 1-5 rating scale facilitates experimentation of rating prediction systems across domains.

Table 1 shows the list of 12 datasets we use for our experimentation. Our datasets comprise more than 58 million reviews overall, distributed across different types of businesses, where some businesses are far more popular than others. The number of reviews per type of business ranges from 22 million reviews for books to 36,000 reviews for dentists, showing a significant imbalance in the size and popularity of different domains. The variability of dataset sizes and popularity of review domains enables our analysis looking at the effect of the availability of large amounts of in-domain data for training.

Business/Product	Source	# Reviews
Books	Amazon	22,507,155
Restaurants	TripAdvisor	14,542,460
Attractions	TripAdvisor	6,358,253
Clothing	Amazon	5,748,920
Homeware	Amazon	4,253,926
Hotels	TripAdvisor	3,598,292
Nightlife	Yelp	877,352
Events	Yelp	387,087
Casinos	Yelp	115,703
Hair salons	Yelp	99,600
Resorts	Yelp	57,678
Dentists	Yelp	36,600
TOTAL	–	58,583,026


**Table 1.** Details of the 12 datasets, sorted by number of reviews. The number of reviews that each type of business/product receives varies drastically.

In Figure 1 we show a breakdown of the 1-5 star ratings for each dataset. We observe an overall tendency of assigning high ratings across most categories, except in the cases of casinos and resorts, where the ratings are more evenly distributed. Most categories show an upwards tendency with higher number of reviews for higher ratings, as is the case with attractions, books or restaurants. Interestingly, in the case of dentists and hair salons, the ratings that prevail are 1's and 5's, showing that users tend to either love or hate these services.

## METHODOLOGY

One of the key challenges of dealing with reviews pertaining to different domains is that the vocabulary can vary significantly. We can expect that people will express that they like or dislike a product in different ways for different domains. For example, one may make a reference to *good food* when reviewing a restaurant, a *comfortable bed* for a hotel, an *inspiring novel* for a book and a *fun party* for an event. All of these could be deemed similarly positive for each domain, however without a proper method to capture these semantics, predictions made across domains may not be accurate enough. To tackle this problem, we propose to use aspect phrase embeddings.

<sup>1</sup><https://www.yelp.com/dataset>



**Figure 1.** Distributions of star ratings in the 12 datasets.

### Aspect Phrase Extraction and Representation

Different review domains tend to have different aspect categories associated with them. While one may care about *food quality* in restaurants, the focus is instead on the *engagement* when reviewing a book. Even if other aspect categories such as *price* are more widely generalisable across domains, most aspect categories and associated vocabulary vary across domains. In the first instance, we are interested in capturing those phrases associated with aspect categories, which we refer to as aspect phrases. We define aspect phrases as tuples providing opinionated judgement of a particular aspect of a product, e.g., *excellent food* for restaurants or *interesting reading* for books. Once we have these tuples, in a second step we use word embeddings to achieve a generalisable semantic representation of the aspect phrases.

**Aspect Phrase Extraction** To extract the tuples corresponding to aspect phrases from the reviews, we rely on the assumption that these opinionated tuples will be made of (1) a sentiment word that judges the object or action concerned and (2) a non-sentiment word that refers to the object or action being judged. To restrict the context in which a tuple can be observed, we consider segments of the reviews, i.e., sentences or parts of sentences that are independent from each other. We perform the extraction of aspect phrases by following the next steps:

- 1. POS tagging:** We extract the part-of-speech (POS) tags of all words in the reviews by using NLTK's POS tagger (Bird and Loper, 2004), hence labelling each word as a noun, verb, adjective, adverb, etc.
- 2. Identification of sentiment words:** We use the sentiment lexicon generated by Hu and Liu (2004), which provides a list of over 6,800 words associated with positive or negative sentiment. With this list, we tag matching keywords from reviews as being positive or negative.
- 3. Segmentation of reviews:** We break down the reviews into segments. To identify the boundaries of segments, we rely on punctuation signs ( , . ; : ) and coordinating conjunctions (for, and, nor, but, or, yet, so) as tokens indicating the end of a segment. Text at either side of these tokens are separated into different segments.
- 4. Extraction of aspect phrases:** At this stage, we only act within the boundaries of each segment. Within each segment, we identify pairs of words made of (1) a sentiment word labelled as positive or negative and (2) a word labelled as noun or verb by the POS tagger and identified as a non-sentiment word, i.e., not matching any of the keywords in the sentiment lexicon. Each pair of words matching these criteria within a segment is extracted as a tuple pertaining to an aspect phrase.

215 Through the process above, we extract aspect phrases for all review domains. Table 2 shows the most  
 216 salient aspect phrases for each of the review domains. We observe that aspect phrases vary substantially  
 217 across domains, with phrases referring to the *ease of use* for homeware, *comfort* for clothing, *food quality*  
 218 for restaurants or *happy hour* for nightlife, among others.

Category	Top aspect phrases
books	(great, book) (good, book) (like, book) (recommend, book) (well, written)
restaurants	(good, food) (good, was) (delicious, was) (great, food) (good, service)
attractions	(worth, visit) (beautiful, is) (worth, is) (free, is) (attraction, is)
clothing	(comfortable, are) (well, made) (worn, have) (perfectly, fit) (love, shoes)
homeware	(easy, clean) (easy, use) (great, works) (well, works) (stainless, steel)
hotels	(nice, hotel) (good, hotel) (great, hotel) (recommend, hotel) (clean, was)
nightlife	(happy, hour) (good, was) (pretty, was) (good, were) (good, food)
events	(nice, was) (clean, was) (pretty, was) (clean, room) (nice, room)
casinos	(nice, was) (nice, room) (like, casino) (like, room) (clean, room)
hair salons	(great, hair) (like, hair) (amazing, hair) (best, hair) (recommend, salon)
resorts	(grand, mgm) (lazy, river) (nice, was) (nice, room) (like, room)
dentists	(best, dentist) (wisdom, teeth) (work, done) (clean, office) (recommend, office)

**Table 2.** List of most salient aspect phrases for the 12 review categories.

219 **Aspect Phrase Representation** Despite the ability of our method above to capture aspect phrases  
 220 independent of the domain, these still vary in terms of vocabulary. To achieve semantic representations of  
 221 aspect phrases extracted for different domains, we make use of Word2Vec word embeddings (Mikolov  
 222 et al., 2013). We train a word embedding model using the 58 million reviews in our datasets. This model  
 223 is then used to achieve semantic representations of the aspect phrases with the following two variants:

- 224 • **Aspect Phrase Embeddings (APE):** we aggregate all the aspect phrases extracted for a review.  
 225 We generate the embedding vector for each review by adding up the word embeddings for all words  
 226 included in those phrases.
- 227 • **Polarised Aspect Phrase Embeddings (PAPE):** we aggregate the aspect phrases for a review in  
 228 two separate groups, one containing positive phrases and the other containing negative phrases.  
 229 Following the same method as for aspect phrase embeddings, here we generate a separate embedding  
 230 representation for each group, which leads to an embedding representation for positive aspect  
 231 phrases and another embedding representation for negative aspect phrases. We then concatenate  
 232 both embeddings to get the combined vector.

## 233 Experiment Settings

234 **Cross-validation** We perform different sets of experiments for comparing the performance of our  
 235 rating prediction systems in in-domain and cross-domain settings for different domains. As we are  
 236 interested in performing 10-fold cross-validation experiments for both settings, we randomly split the  
 237 data available for each domain  $d \in \{1..12\}$  into 10 equally sized folds,  $f \in \{1..10\}$ . Each time one  
 238 of these folds is considered as the test data,  $Test_{ij}$ , while the training data depends on the setting, i.e.,  
 239 (1)  $\sum_{f \in \{1..10\}, f \neq j, d=i} Train_{df}$  for in-domain experiments, and (2)  $\sum_{d \in \{1..12\}, d \neq i, f=j} Train_{df}$  for cross-  
 240 domain experiments. We ultimately average the performance scores across all 10 folds in each setting.

241 **Classifiers** In setting up the experiments for our analysis, we tested all the classifiers proposed by  
 242 Fan and Khademi (2014) and some more, including a multinomial Logistic Regression, a Gaussian  
 243 Naive Bayes classifier, Support Vector Machines and Random Forests. Our experiments showed that the  
 244 multinomial Logistic Regression was clearly and consistently the best classifier, and hence for the sake of  
 245 clarity and brevity we report results using this classifier in the paper.

246 **Features and Baselines** We implement four different sets of features: two baselines and two methods  
 247 that we propose:

- 248 • **Baseline (A<sup>3</sup>NCF):** Introduced by Cheng et al. (2018), A<sup>3</sup>NCF is a model that extracts users' preferences and items' characteristics on different aspects from reviews. They introduce an attention network which utilises these user and item features extracted from reviews, which aims to capture the attention that a user pays to each aspect of an item in a review. In the absence of cross-domain rating prediction systems in previous work, we use A<sup>3</sup>NCF as the state-of-the-art rating prediction system to beat.
- 254 • **Baseline using Word Embeddings (w2v):** we implement a second baseline consisting of a word embedding representation of the entire text of a review by averaging all the embeddings. We use the word embedding model we trained from our review datasets, which serves as a comparable baseline to measure the impact of aspect phrase embeddings on the performance of the rating prediction system.
- 259 • **Aspect Phrase Embeddings (APE):** we concatenate the word embedding vectors obtained for the entire texts of reviews with the aspect phrase embeddings, i.e., word embedding representations of the aspect phrases extracted from a review.
- 262 • **Polarised Aspect Phrase Embeddings (PAPE):** we concatenate the word embedding vectors obtained for the entire texts of reviews with the polarised aspect phrase embeddings, i.e., a concatenation of the word embedding representations of positive aspect phrases and the word embedding representation of negative aspect phrases.

266 Note that baseline methods from more recent works are not directly applicable to our task as they make use of user data, which is not applicable in cross-domain scenarios where users are different across domains and platforms. Because our objective is to build a review rating prediction system that would then be applicable to other sources on the Web, relying on user metadata is not realistic, given that we often do not know the author of a text on the Web or the authors are different from those in the training data. Hence, we use the state-of-the-art text-only review rating prediction system by Fan and Khademi (2014), as well as a baseline using word embeddings, which enables direct comparison with the use of aspect phrase embeddings as additional features.

## 274 Evaluation

275 The review rating prediction task can be considered as a rating task with five different stars from 1 to 5. As a rating task, we need to consider the proximity between the predicted and reference ratings. For instance, if the true star rating of a review is 4 stars, then the predicted rating of 3 stars will be better than a predicted rating of 2 stars. To account for the difference or error rate between the predicted and reference ratings, we rely on the two metrics widely used in previous rating prediction work (Li et al., 2011), which include the Root Mean Square Error (RMSE) (see Equation 1) and Mean Absolute Error (MAE) (see Equation 2).

$$276 \text{ RMSE} = \sqrt{\frac{\sum_{i=1}^n (\hat{y}_i - r_i)^2}{n}} \quad (1)$$

$$277 \text{ MAE} = \frac{\sum_{i=1}^n |y_i - r_i|}{n} \quad (2)$$

282 where:

- 283  $n$  denotes the total number of reviews in the test set.
- 284  $y_i$  is the predicted star rating for the  $i^{th}$  review.
- 285  $r_i$  is the real star rating given to the  $i^{th}$  review by the user.

286 We report both metrics for all our experiments, to facilitate the comparison for future work.



# RESULTS

The results are organised in two parts. First, we present and analyse results for in-domain review rating prediction, where the training and test data belong to the same domain. Then, we show and analyse the results for cross-domain review rating prediction, where data from domains that differ from the test set is used for training. A comparison of the performance on both settings enables us to assess the suitability of leveraging a cross-domain classifier as well as the use of aspect phrase embeddings.

MAE						
	A <sup>3</sup> NCF	w2v	ape	pape	w2v +ape	w2v +pape
books	0.821	0.542	0.664	0.663	0.530	<b>0.521</b>
restaurants	0.779	0.484	0.598	0.611	0.471	<b>0.466</b>
attractions	0.644	0.567	0.771	0.803	0.554	<b>0.549</b>
clothing	0.988	0.544	0.760	0.742	0.529	<b>0.518</b>
homeware	1.055	0.581	0.760	0.751	0.563	<b>0.543</b>
hotels	0.746	0.455	0.553	0.579	0.442	<b>0.441</b>
nightlife	1.061	0.579	0.691	0.715	0.555	<b>0.554</b>
events	1.127	0.595	0.678	0.709	<b>0.565</b>	0.568
casinos	1.127	0.769	0.771	0.803	<b>0.689</b>	0.705
hair salons	1.096	0.410	0.498	0.521	<b>0.405</b>	0.414
resorts	1.208	0.793	0.760	0.777	<b>0.689</b>	0.692
dentists	1.767	0.322	0.404	0.399	<b>0.317</b>	0.324

**Table 3.** MAE results for in-domain review rating prediction. Categories are sorted in descending order by number of reviews, with the most popular review categories at the top of the list.

RMSE						
	A <sup>3</sup> NCF	w2v	ape	pape	w2v +ape	w2v +pape
books	1.068	1.057	1.223	1.225	1.038	<b>1.023</b>
restaurants	0.976	0.856	1.007	1.031	0.836	<b>0.828</b>
attractions	<b>0.823</b>	0.996	1.132	1.126	0.975	0.968
clothing	1.211	1.043	1.343	1.320	1.021	<b>1.001</b>
homeware	1.281	1.144	1.381	1.370	1.115	<b>1.081</b>
hotels	0.926	0.809	0.942	0.990	0.786	<b>0.784</b>
nightlife	1.251	1.001	1.139	1.172	0.961	<b>0.958</b>
events	1.306	1.054	1.153	1.196	<b>1.004</b>	1.006
casinos	1.304	1.245	1.218	1.252	<b>1.118</b>	1.128
hair salons	1.271	1.019	1.150	1.168	<b>0.993</b>	0.997
resorts	1.395	1.295	1.224	1.238	1.143	<b>1.136</b>
dentists	1.891	1.016	1.157	1.121	0.994	<b>0.983</b>

**Table 4.** RMSE results for in-domain review rating prediction. Categories are sorted in descending order by number of reviews, with the most popular review categories at the top of the list.

## In-Domain Review Rating Prediction

Tables 3 and 4 show results for in-domain review rating prediction, where the training and test data belong to the same domain. Experiments are performed using 10-fold cross validation within each of the 12 datasets. Note that lower scores indicate a smaller amount of error and hence better performance. These results show that both APE and PAPE combined with word embeddings (w2v+APE and w2v+PAPE) consistently outperform the sole use of word embeddings (w2v). We also observe that the sole use of phrase embeddings (APE and PAPE) does not suffice to outperform word embeddings, whereas their combination with w2v leads to clear improvements in all cases. This proves the importance of phrases,

however we also see that using the rest of the text in the review is useful to boost performance. Likewise, both of our combining methods clearly outperform the baseline method ( $A^3NCF$ ) by Cheng et al. (2018), with the exception of attractions when the performance is measured using RMSE, which is the only case where the baseline performs better; the combining methods leveraging phrase embeddings are clearly better for the rest of the domains. This emphasizes the importance of using word embeddings for capturing semantics of aspect phrases, even when the experiments are within the same domain. A comparison between our two proposed methods using phrase embeddings shows that polarised phrase embeddings outperform phrase embeddings for the most popular review categories, whereas the difference is not so clear for less popular categories.

All in all, these results show that the use of either form of phrase embeddings combined with word embeddings leads to improvements in the review rating prediction when the training and test data belong to the same domain. The main goal of our work is however to show their effectiveness on cross-domain review rating prediction, which we discuss in the next section. The in-domain results presented in this section enable us to perform a comparison with the cross-domain results presented next.

MAE							
	$A^3NCF$	BID	w2v	ape	pape	w2v +ape	w2v +pape
books	0.913	<b>0.521</b>	0.725	0.858	0.826	0.691	0.672
restaurants	0.866	<b>0.466</b>	0.514	0.624	0.657	0.502	0.501
attractions	0.736	<b>0.549</b>	0.613	0.702	0.757	0.608	0.596
clothing	0.995	<b>0.518</b>	0.712	0.886	0.834	0.684	0.656
homeware	1.061	<b>0.543</b>	0.821	0.993	0.964	0.801	0.776
hotels	0.826	<b>0.441</b>	0.494	0.588	0.654	0.483	0.490
nightlife	1.111	0.554	0.541	0.695	0.723	0.530	<b>0.524</b>
events	1.155	0.565	0.543	0.682	0.729	0.525	<b>0.521</b>
casinos	1.151	0.689	0.652	0.799	0.851	0.630	<b>0.623</b>
hair salons	1.169	0.405	0.353	0.494	0.569	<b>0.337</b>	0.341
resorts	1.106	0.689	0.633	0.762	0.819	0.606	<b>0.602</b>
dentists	1.258	0.317	0.272	0.451	0.543	<b>0.266</b>	0.268

**Table 5.** MAE results for cross-domain review rating prediction. BID = best in-domain. Categories are sorted in descending order by number of reviews, with the most popular review categories at the top of the list.

### Cross-Domain Review Rating Prediction

Tables 5 and 6 show results for cross-domain review rating prediction. While experiments are also performed using 10-fold cross-validation, we train the classifier for a particular domain using data from the other 11 datasets, i.e., simulating the scenario where we do not have any labelled data for the target domain. We also include results for the best performance for each review category when we train with data from the same domain, which is represented as BID (best in-domain). This enables us to compare whether and the extent to which the use of out-of-domain data for training can help to improve the performance for each review category.

Results show a remarkable difference between popular and non-popular review categories. For the 6 most popular review categories (books, restaurants, attractions, clothing, homeware, hotels), the best performance is obtained by the best in-domain (BID) classifier, which indicates that for review categories with large amounts of training data, it is better to use in-domain data. However, when we look at the bottom 6 review categories (nightlife, events, casinos, hair salons, resorts, dentists), we observe that the cross-domain classifier leveraging out-of-domain data for training achieves higher performance. While the sole use of word embeddings (w2v) leads to improved performance for the least popular categories, the improvement is even better when we incorporate either phrase embeddings or polarised phrase embeddings. The results are also positive for the w2v+PAPE over the w2v+APE; even if the results are not better in 100% of the cases, w2v+PAPE tends to outperform w2v+APE in most cases, with just a small difference when APE is better, showing that one can rely on PAPE for all cases dealing

RMSE							
	A <sup>3</sup> NCF	BID	w2v	ape	pape	w2v +ape	w2v +pape
books	1.130	<b>1.023</b>	1.354	1.481	1.440	1.307	1.278
restaurants	1.121	<b>0.828</b>	0.891	1.045	1.094	0.870	0.868
attractions	0.902	<b>0.823</b>	1.062	1.155	1.275	1.050	1.034
clothing	1.237	<b>1.001</b>	1.265	1.453	1.418	1.229	1.190
homeware	1.330	<b>1.081</b>	1.443	1.600	1.579	1.417	1.386
hotels	1.073	<b>0.784</b>	0.854	0.980	1.077	0.836	0.844
nightlife	1.359	0.958	0.937	1.133	1.177	0.921	<b>0.912</b>
events	1.402	1.004	0.973	1.148	1.221	0.943	<b>0.936</b>
casinos	1.414	1.118	1.081	1.250	1.322	1.048	<b>1.039</b>
hair salons	1.437	0.993	0.900	1.116	1.246	<b>0.866</b>	0.876
resorts	1.382	1.136	1.064	1.214	1.296	1.024	<b>1.018</b>
dentists	1.513	0.983	0.887	1.181	1.323	<b>0.870</b>	0.873

**Table 6.** RMSE results for cross-domain review rating prediction. BID = best in-domain. Categories are sorted in descending order by number of reviews, with the most popular review categories at the top of the list.

with non-popular domains. Our combining methods (w2v+APE and w2v+PAPE) also outperform the state-of-the-art baseline A<sup>3</sup>NCF for all of the non-popular review domains, with some variation for the popular review domains. This again reinforces the potential of our methods combining phrase embeddings for cross-domain review rating prediction applied to review domains with little training data available. These experiments show a clear shift in the performance of in-domain classifiers when the amount of training data decreases, encouraging the use of out-of-domain data for training in those cases.

Likewise, this motivates the use of the cross-domain classifier for new domains where no labelled data is available for training, for instance because reviews are collected from a website where no ratings are given by the user, such as Facebook, Twitter or comments on blogs.

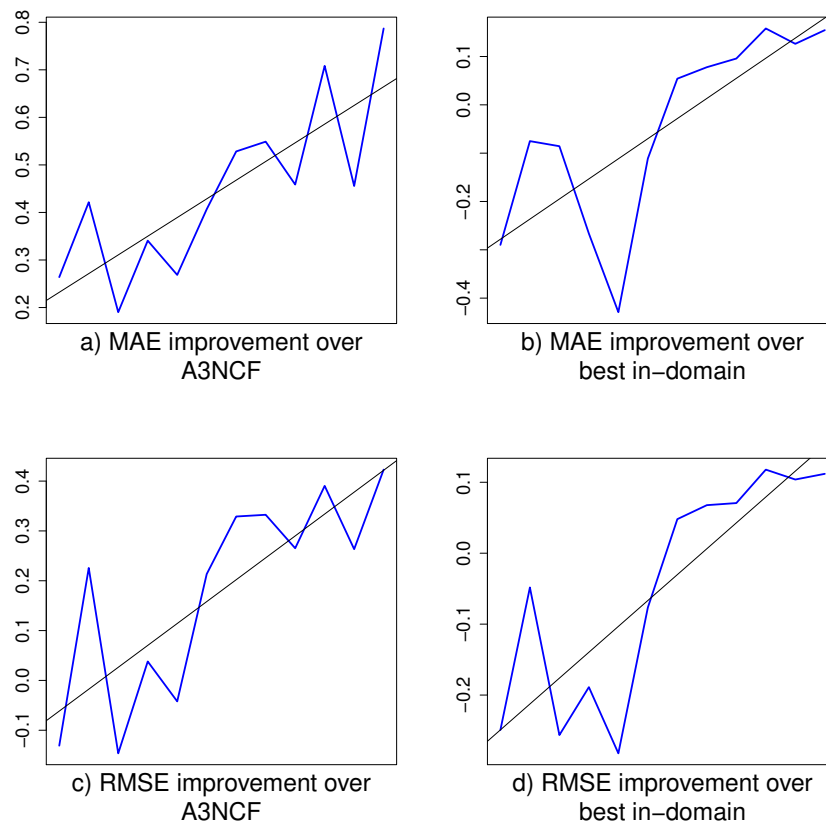
Further to this, Figure 2 shows the relative improvements of our w2v+PAPE cross-domain classifier over the A<sup>3</sup>NCF and the BID classifiers, both for the MAE and RMSE metrics. Review domains are sorted by popularity on the X axis in the sub plots a-d, with the least popular domains on the right. The relative improvement on each review domain is represented by the blue line. The line of best fit (in black) shows the overall tendency of our cross-domain classifier to achieve an increasing relative improvement over other methods as the popularity of the review domain and subsequently the amount of in-domain training data increases. This reinforces our cross-domain classifier as a suitable alternative to achieve the best results on the least popular domains, outperforming any other in-domain alternative.

### Effect of the Size of the Training Data

Since cross-domain experiments benefit from larger training sets, we assess the effect of the size of the training data on the performance of the rating prediction system. We evaluate the performance of cross-domain rating prediction by using different percentages of the training data available, ranging from 1% to 100%, which we then compare with the best in-domain rating prediction system. The subset of the training data is randomly sampled until the percentage is satisfied.

Tables 7 and 8 show the results of using different percentages of the training data for the rating prediction based on MAE and RMSE respectively. Results for each review category is compared with the best in-domain (BID) result, and the best results are highlighted in bold, i.e., either the best in-domain when this is the best, or the specific percentages when the cross-domain prediction system performs best.

As we observed before, in-domain prediction system perform better for the top 6 review categories based on popularity, thanks to the availability of more training data. We are particularly interested in seeing how much data we need with the less popular categories for the cross-domain rating prediction systems to perform better than their in-domain counterparts. Looking at the 6 least popular review categories, we observe that using only 5% of the training data suffices in all cases to outperform the best in-domain system, with the percentage dropping up to 2% for casinos and hair salons and up to 1% for



**Figure 2.** Relative improvement of the w2v+PAPE cross-domain classifier over the A<sup>3</sup>NCF and best in-domain (BID) classifiers displayed by MAE (a/b) and RMSE (c/d). Review domains are sorted by popularity on the X axis in the sub plots a-d, with the least popular domains on the right. The relative improvement on each review domain is represented by the blue line. The line of best fit (in black) shows the overall tendency of our cross-domain classifier to achieve an increasing relative improvement over other methods as the popularity of the review domain and subsequently the amount of in-domain training data increases.

resorts. Note that out of the (up to) 58 million reviews available for cross-domain training, a subset of 5% only represents (up to) 2.9 million reviews, which are easy to obtain through online review platforms.

We also observe that, while performance keeps improving as we increase the size of the training data, it generally shows a tendency to plateau after 25% of the reviews are used for training. This reflects that after about 14.5 million reviews the system has enough training data and can hardly improve its performance.

## CONCLUSIONS

We have proposed a novel method that leverages aspect phrase embeddings for predicting the star rating of online reviews, which is applicable in in-domain and cross-domain settings. We have also shown that our cross-domain approach is effective for making predictions in review domains with a paucity of training data, where training data from other domains can be successfully exploited. Previous work on review rating prediction had been limited to popular reviews domains, such as restaurants or hotels. Our study broadens the findings of previous works by experimenting on 12 different datasets pertaining to 12 different review domains of very different levels of popularity, and collects from different sources including Yelp, Amazon and TripAdvisor. Given that some of these review domains have very limited availability of labelled data for training, our aim has been to propose a cross-domain review rating prediction system that would perform well for those non-popular domains. Likewise, a cross-domain

	MAE							
	1	2	5	10	25	50	75	100
books	0.706	0.682	0.666	0.661	0.663	0.667	0.671	0.672
	BID = <b>0.521</b>							
restaurants	0.571	0.539	0.518	0.505	0.501	0.501	0.501	0.501
	BID = <b>0.466</b>							
attractions	0.615	0.604	0.596	0.594	0.593	0.595	0.596	0.596
	BID = <b>0.549</b>							
clothing	0.718	0.699	0.682	0.671	0.66	0.657	0.656	0.656
	BID = <b>0.518</b>							
homeware	0.869	0.848	0.818	0.800	0.787	0.779	0.777	0.776
	BID = <b>0.543</b>							
hotels	0.531	0.518	0.503	0.496	0.493	0.491	0.491	0.490
	BID = <b>0.441</b>							
nightlife	0.604	0.579	<b>0.553</b>	<b>0.540</b>	<b>0.530</b>	<b>0.526</b>	<b>0.525</b>	<b>0.524</b>
	BID = 0.554							
events	0.601	0.575	<b>0.550</b>	<b>0.537</b>	<b>0.527</b>	<b>0.523</b>	<b>0.522</b>	<b>0.521</b>
	BID = 0.565							
casinos	0.713	<b>0.687</b>	<b>0.657</b>	<b>0.641</b>	<b>0.629</b>	<b>0.625</b>	<b>0.624</b>	<b>0.623</b>
	BID = 0.689							
hair salons	0.413	<b>0.388</b>	<b>0.367</b>	<b>0.358</b>	<b>0.348</b>	<b>0.344</b>	<b>0.342</b>	<b>0.341</b>
	BID = 0.405							
resorts	<b>0.685</b>	<b>0.658</b>	<b>0.632</b>	<b>0.619</b>	<b>0.606</b>	<b>0.604</b>	<b>0.603</b>	<b>0.602</b>
	BID = 0.689							
dentists	0.361	0.332	<b>0.305</b>	<b>0.292</b>	<b>0.278</b>	<b>0.272</b>	<b>0.270</b>	<b>0.268</b>
	BID = 0.317							

**Table 7.** MAE Results for cross-domain review rating prediction with different percentages of the training data. BID = best in-domain. Categories are sorted in descending order by number of reviews, with the most popular review categories at the top of the list.

review rating prediction system can be used to predict ratings of reviews gathered from platforms where users do not assign ratings, such as Facebook or Twitter.

Our review rating prediction system leverages both POS taggers and sentiment lexicons to extract aspect phrases from reviews, i.e., phrases referring to different features of a business. To enable generalisation of aspect phrases to different domains, we make use of universal representations using word embeddings; we propose two different models for feature representations, (1) Aspect Phrase Embeddings (APE), which aggregates all aspect phrases of a review, and (2) Polarised Aspect Phrase Embeddings (PAPE), which considers positive and negative aspect phrases separately to create an embedding representation for each. We compare our results with those of the best-performing classifier by Cheng et al. (2018) and another baseline that uses word embedding representations of the entire review. We developed both in-domain and cross-domain review rating prediction systems following this methodology; this allows us to compare performance on in-domain and cross-domain experiments for different review domains.

Our experiments show that both of our methods leveraging phrase embeddings lead to improvements over the rest of the baseline methods, both in the in-domain and the cross-domain settings. Interestingly, a comparison of results for these two experiment settings shows that performance scores are higher for the in-domain classifier when we make predictions on the most popular domains, however the cross-domain classifier leads to substantial improvements for the least popular domains. Our results indicate that a classifier trained from in-domain data is more suitable for popular review domains, whereas unpopular review domains can be improved when out-of-domain data is used for training along with our aspect phrase embedding representation. Further looking into the amount of data needed for training a cross-domain rating prediction system, we observe that a small sample of about 5% of the data available (i.e., fewer than 3 million reviews) for the other domains is enough to outperform baseline methods.

Our work defines the state-of-the-art and the first approach to cross-domain review rating prediction,

	RMSE							
	1	2	5	10	25	50	75	100
books	1.311	1.281	1.262	1.258	1.264	1.270	1.276	1.278
	BID = <b>1.023</b>							
restaurants	0.976	0.925	0.895	0.876	0.869	0.869	0.869	0.868
	BID = <b>0.828</b>							
attractions	1.056	1.040	1.029	1.027	1.028	1.032	1.033	1.034
	BID = <b>0.823</b>							
clothing	1.259	1.236	1.217	1.206	1.194	1.191	1.190	1.190
	BID = <b>1.001</b>							
homeware	1.484	1.467	1.435	1.415	1.400	1.389	1.387	1.386
	BID = <b>1.081</b>							
hotels	0.907	0.889	0.865	0.854	0.848	0.845	0.844	0.844
	BID = <b>0.784</b>							
nightlife	1.022	0.987	<b>0.951</b>	<b>0.933</b>	<b>0.919</b>	<b>0.914</b>	<b>0.913</b>	<b>0.912</b>
	BID = 0.958							
events	1.048	1.013	<b>0.977</b>	<b>0.959</b>	<b>0.945</b>	<b>0.939</b>	<b>0.937</b>	<b>0.936</b>
	BID = 1.004							
casinos	1.146	<b>1.115</b>	<b>1.078</b>	<b>1.058</b>	<b>1.045</b>	<b>1.040</b>	<b>1.039</b>	<b>1.039</b>
	BID = 1.118							
hair salons	1.015	<b>0.969</b>	<b>0.929</b>	<b>0.910</b>	<b>0.890</b>	<b>0.882</b>	<b>0.878</b>	<b>0.876</b>
	BID = 0.993							
resorts	<b>1.119</b>	<b>1.085</b>	<b>1.054</b>	<b>1.038</b>	<b>1.024</b>	<b>1.019</b>	<b>1.018</b>	<b>1.018</b>
	BID = 1.136							
dentists	1.058	1.004	<b>0.953</b>	<b>0.925</b>	<b>0.896</b>	<b>0.884</b>	<b>0.878</b>	<b>0.873</b>
	BID = 0.983							

**Table 8.** RMSE Results for cross-domain review rating prediction with different percentages of the training data. BID = best in-domain. Categories are sorted in descending order by number of reviews, with the most popular review categories at the top of the list.

expanding the hitherto limited research focusing on in-domain settings for popular domains. Research in cross-domain review rating prediction is still in its infancy, and we hope that our work will encourage further research in the task. Future work on cross-domain review rating prediction could further focus on the detection of implicit aspect phrases for improving performance, as well as in capturing cultural differences in writing reviews, given that users from different countries may express themselves differently when writing positive or negative reviews. To further test additional methods to cross-domain rating prediction, our plans for future work include testing methods for transfer learning and domain adaptation.

## REFERENCES

- Bird, S. and Loper, E. (2004). NLTK: the natural language toolkit. In *Proceedings of the ACL 2004 Interactive poster and demonstration sessions*, page 31. Association for Computational Linguistics.
- Catherine, R. and Cohen, W. (2017). Transnets: Learning to transform for recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*, pages 288–296. ACM.
- Cheng, Z., Ding, Y., He, X., Zhu, L., Song, X., and Kankanhalli, M. S. (2018). A3NCF: An adaptive aspect attention model for rating prediction. In *IJCAI*, pages 3748–3754.
- Fan, M. and Khademi, M. (2014). Predicting a business star in Yelp from its reviews text alone. *arXiv preprint arXiv:1401.0864*.
- Ganu, G., Elhadad, N., and Marian, A. (2009). Beyond the stars: improving rating predictions using review text content. In *WebDB*, volume 9, pages 1–6. Citeseer.
- Goss, J. (2017). *A Melting Pot of Cuisines: Examining the Relationship Between Restaurant Ethnicities and Food Safety Inspection Scores*. PhD thesis, Georgetown University.
- He, R. and McAuley, J. (2016). Ups and downs: Modeling the visual evolution of fashion trends with

- one-class collaborative filtering. In *proceedings of the International Conference on World Wide Web*, pages 507–517. International World Wide Web Conferences Steering Committee.
- Hu, M. and Liu, B. (2004). Mining and summarizing customer reviews. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 168–177. ACM.
- Huang, J., Rogers, S., and Joo, E. (2014). Improving restaurants by extracting subtopics from yelp reviews. *iConference 2014 (Social Media Expo)*.
- Jin, Z., Li, Q., Zeng, D. D., Zhan, Y., Liu, R., Wang, L., and Ma, H. (2016). Jointly modeling review content and aspect ratings for review rating prediction. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 893–896. ACM.
- Kaplan, A. M. and Haenlein, M. (2010). Users of the world, unite! the challenges and opportunities of social media. *Business horizons*, 53(1):59–68.
- Laddha, A. and Mukherjee, A. (2018). Aspect opinion expression and rating prediction via lda-crf hybrid. *Natural Language Engineering*, 24(4):611–639.
- Lei, X., Qian, X., and Zhao, G. (2016). Rating prediction based on social sentiment from textual reviews. *IEEE Transactions on Multimedia*, 18(9):1910–1921.
- Leung, C. W., Chan, S. C., and Chung, F.-I. (2006). Integrating collaborative filtering and sentiment analysis: A rating inference approach. In *Proceedings of the ECAI 2006 Workshop on Recommender Systems*, pages 62–66.
- Li, F., Liu, N., Jin, H., Zhao, K., Yang, Q., and Zhu, X. (2011). Incorporating reviewer and product information for review rating prediction. In *Proceedings of the International Joint Conference on Artificial Intelligence*, volume 11, pages 1820–1825.
- Liu, B. and Zhang, L. (2012). A survey of opinion mining and sentiment analysis. In *Mining text data*, pages 415–463. Springer.
- McAuley, J., Targett, C., Shi, Q., and Van Den Hengel, A. (2015). Image-based recommendations on styles and substitutes. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 43–52. ACM.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, pages 3111–3119.
- Mousavizadeh, M., Koohikamali, M., and Salehan, M. (2015). The effect of central and peripheral cues on online review helpfulness: A comparison between functional and expressive products. In *Proceedings of the International Conference on Information Systems*.
- Ning, X., Yac, L., Wang, X., Benatallah, B., Dong, M., and Zhang, S. (2018). Rating prediction via generative convolutional neural networks based regression. *Pattern Recognition Letters*.
- Pang, B. and Lee, L. (2005). Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*, pages 115–124. Association for Computational Linguistics.
- Pang, B., Lee, L., and Vaithyanathan, S. (2002). Thumbs up?: sentiment classification using machine learning techniques. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 79–86. Association for Computational Linguistics.
- Perrin, A. (2017). Social media usage: 2005–2015. Technical report, Pew Research Center.
- Poria, S., Cambria, E., Hazarika, D., Majumder, N., Zadeh, A., and Morency, L.-P. (2017). Context-dependent sentiment analysis in user-generated videos. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 873–883.
- Qu, L., Ifrim, G., and Weikum, G. (2010). The bag-of-opinions method for review rating prediction from sparse text patterns. In *Proceedings of the 23rd International Conference on Computational Linguistics*, pages 913–921. Association for Computational Linguistics.
- Quan-Haase, A. and Young, A. L. (2010). Uses and gratifications of social media: A comparison of facebook and instant messaging. *Bulletin of Science, Technology & Society*, 30(5):350–361.
- Salehan, M. and Kim, D. J. (2016). Predicting the performance of online consumer reviews: A sentiment mining approach to big data analytics. *Decision Support Systems*, 81:30–40.
- Salehan, M., Zhang, S., and Aghakhani, N. (2017). A recommender system for restaurant reviews based on consumer segment. In *Proceedings of the Americas Conference on Information Systems*.
- Seo, S., Huang, J., Yang, H., and Liu, Y. (2017a). Interpretable convolutional neural networks with dual local and global attention for review rating prediction. In *Proceedings of the Eleventh ACM Conference*

- 484 on Recommender Systems, pages 297–305. ACM.
- 485 Seo, S., Huang, J., Yang, H., and Liu, Y. (2017b). Representation learning of users and items for review
- 486 rating prediction using attention-based convolutional neural network. In *3rd International Workshop on*
- 487 *Machine Learning Methods for Recommender Systems (MLRec)(SDM'17)*.
- 488 Sparks, B. A. and Browning, V. (2011). The impact of online reviews on hotel booking intentions and
- 489 perception of trust. *Tourism management*, 32(6):1310–1323.
- 490 Taboada, M., Brooke, J., Tofiloski, M., Voll, K., and Stede, M. (2011). Lexicon-based methods for
- 491 sentiment analysis. *Computational linguistics*, 37(2):267–307.
- 492 Tang, D., Qin, B., Liu, T., and Yang, Y. (2015). User modeling with neural network for review rating
- 493 prediction. In *Proceedings of IJCAI*, pages 1340–1346.
- 494 Titov, I. and McDonald, R. (2008). Modeling online reviews with multi-grain topic models. In *Proceedings*
- 495 *of the 17th international conference on World Wide Web*, pages 111–120. ACM.
- 496 Tripathy, A., Agrawal, A., and Rath, S. K. (2016). Classification of sentiment reviews using n-gram
- 497 machine learning approach. *Expert Systems with Applications*, 57:117–126.
- 498 Vargas, F. A. and Pardo, T. A. S. (2018). Aspect clustering methods for sentiment analysis. In *International*
- 499 *Conference on Computational Processing of the Portuguese Language*, pages 365–374. Springer.
- 500 Wang, B., Chen, B., Ma, L., and Zhou, G. (2019). User-personalized review rating prediction method
- 501 based on review text content and user-item rating matrix. *Information*, 10(1):1.
- 502 Wang, B., Xiong, S., Huang, Y., and Li, X. (2018). Review rating prediction based on user context and
- 503 product context. *Applied Sciences*, 8(10):1849.
- 504 Wang, H., Lu, Y., and Zhai, C. (2010). Latent aspect rating analysis on review text data: a rating regression
- 505 approach. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery*
- 506 *and Data Mining*, pages 783–792. ACM.
- 507 Wang, Y., Huang, M., Zhao, L., and Zhu, X. (2016). Attention-based LSTM for aspect-level sentiment
- 508 classification. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*,
- 509 pages 606–615.
- 510 Xing, S., Liu, F., Wang, Q., Zhao, X., and Li, T. (2019). A hierarchical attention model for rating
- 511 prediction by leveraging user and product reviews. *Neurocomputing*, 332:417–427.
- 512 Ye, Q., Zhang, Z., and Law, R. (2009). Sentiment classification of online reviews to travel destinations by
- 513 supervised machine learning approaches. *Expert systems with applications*, 36(3):6527–6535.
- 514 Zhang, L., Wang, S., and Liu, B. (2018a). Deep learning for sentiment analysis: A survey. *Wiley*
- 515 *Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(4):e1253.
- 516 Zhang, Q. and Qu, Y. (2013). Topic-opposite sentiment mining model for online review analysis. *Journal*
- 517 *of Frontiers of Computer Science and Technology*, 7(7):620–629.
- 518 Zhang, S., Salehan, M., Leung, A., Cabral, I., and Aghakhani, N. (2018b). A recommender system for
- 519 cultural restaurants based on review factors and review sentiment. In *Proceedings of the Americas*
- 520 *Conference on Information Systems*.
- 521 Zhao, G., Qian, X., and Xie, X. (2016). User-service rating prediction by exploring social users' rating
- 522 behaviors. *IEEE transactions on multimedia*, 18(3):496–506.
- 523 Zhou, X., Wan, X., and Xiao, J. (2015). Representation learning for aspect category detection in online
- 524 reviews. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 417–424.