

ForeSeer: Product Aspect Forecasting Using Temporal Graph Embedding

Zixuan Liu*
zucksliu@cs.fwashington.edu
University of Washington
Seattle, WA, USA

Edward W. Huang
ewhuang@amazon.com
Amazon
Palo Alto, CA, USA

Karthik Subbian
ksubbian@amazon.com
Amazon
Palo Alto, CA, USA

Gaurush Hiranadani
hgaurush@amazon.com
Amazon
Palo Alto, CA, USA

Yi Xu
yxaamzn@amazon.com
Amazon
Seattle, WA, USA

Sheng Wang
swang@cs.washington.edu
University of Washington
Seattle, WA, USA

Kun Qian
qianku@amazon.com
Amazon
Palo Alto, CA, USA

Belinda Zeng
zengb@amazon.com
Amazon
Seattle, WA, USA

ABSTRACT

Developing text mining approaches to mine aspects from customer reviews has been well-studied due to its importance in understanding customer needs and product attributes. In contrast, it remains unclear how to predict the future emerging aspects of a new product that currently has little review information. This task, which we named product aspect forecasting, is critical for recommending new products, but also challenging because of the missing reviews. Here, we propose ForeSeer, a novel textual mining and product embedding approach progressively trained on temporal product graphs for this novel product aspect forecasting task. ForeSeer transfers reviews from similar products on a large product graph and exploits these reviews to predict aspects that might emerge in future reviews. A key novelty of our method is to jointly provide review, product, and aspect embeddings that are both time-sensitive and less affected by extremely imbalanced aspect frequencies. We evaluated ForeSeer on a real-world product review system containing 11,536,382 reviews and 11,000 products over 3 years. We observe that ForeSeer substantially outperformed existing approaches with at least 49.1% AUPRC improvement under the real setting where aspect associations are not given. ForeSeer further improves future link prediction on the product graph and the review aspect association prediction. Collectively, ForeSeer offers a novel framework for review forecasting by effectively integrating review text, product network, and temporal information, opening up new avenues for online shopping recommendation and e-commerce applications.

*This research was conducted during the author's internship at Amazon.



This work is licensed under a Creative Commons Attribution International 4.0 License.

CIKM '23, October 21–25, 2023, Birmingham, United Kingdom.
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0124-5/23/10.
<https://doi.org/10.1145/3583780.3614887>

CCS CONCEPTS

• **Information systems** → **Data mining; Information retrieval; Language models**; • **Computing methodologies** → **Machine learning; Information extraction; Semantic networks; Temporal reasoning; Transfer learning**.

KEYWORDS

Aspect forecasting, textual mining, temporal graph embedding, multi-time forecasting, information extraction, contrastive learning

ACM Reference Format:

Zixuan Liu, Gaurush Hiranadani, Kun Qian, Edward W. Huang, Yi Xu, Belinda Zeng, Karthik Subbian, and Sheng Wang. 2023. ForeSeer: Product Aspect Forecasting Using Temporal Graph Embedding. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*, October 21–25, 2023, Birmingham, United Kingdom. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3583780.3614887>

1 INTRODUCTION

Customer reviews reflect the properties of the product and the needs of the customer on online shopping systems [24, 88]. As a result, one widely studied task is to extract descriptive keywords from customer reviews, such as "not greasy" and "soft inside". These descriptive keywords, which are referred to as aspects in the literature, are later used as key features for various applications that are beneficial to buyers making decisions and sellers improving the market, such as customer behavior prediction [48], sentiment analysis [42], opinion mining [26, 71], product rating [34], and marketing refinement [62]. To tackle this problem, many text mining methods have been proposed to mine such aspects from customer reviews [4, 7, 18, 25, 53, 73, 79, 83, 84].

Despite its usefulness, however, customer reviews are also subjective and sometimes biased, therefore mining effective aspects from reviews for a product requires adequate numbers of reviews. For a new product with limited customer reviews, aspect mining methods will suffer from starting cold and mining inaccurate and noisy aspects. In this paper, we aim at a novel *product aspect forecasting* task, by forecasting the top aspects in the future for products with

inadequate reviews (Fig. 1), i.e., *what aspects will the customer use to describe a new product after six months or three years?* To the best of our knowledge, product aspect forecasting has not been studied in the literature and is one step further than traditional aspect extraction tasks. The most related task is zero-shot aspect-based sentiment analysis [9, 17, 54], but their frameworks are restricted to sentiment analysis or still require multilingual supervision. The recent progress in large language models such as ChatGPT [44] or GPT-4 [2] may also open a door for improved zero-shot aspect extraction performance, but besides the heavy computational costs using them, the limited availability of reviews of new products still restricts the model from forecasting their top aspects in the future.

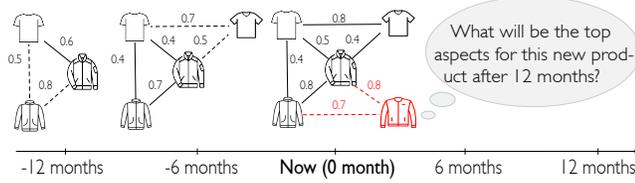


Figure 1: Problem setting of aspect forecasting. The product graph is evolving because new product nodes and edges are added to the graph. The edge weights also change over time.

In this paper, we formulate the product aspect forecasting task as a multi-future top-K aspect retrieving problem. An intuitive solution to the problem is to find similar products with longer histories and infer based on their received feedback, just like a customer would search for the reviews of similar older products that have the same functionality, brand, or same style before making the decision. Motivated by this, in this paper, we propose a three-stage framework ForeSeer to address it. ForeSeer is a flexible textual mining and product embedding framework progressively trained on temporal product graphs. The key idea is to find similar older products from the temporal product graph and propagate potential aspects mined from their reviews to the new item. This enables forecasting currently under-estimated aspects which might be frequently mentioned in the future (Fig. 2). ForeSeer has three core steps, contrastive review-aspect association learning, progressive temporal graph-based product embedding, and aspect-guided product embedding refinement with temporal information. The first step efficiently captures the static semantic relationship between review texts and aspects, and the second step together captures the evolving graphical product similarity information and how products aggregate reviews over time. In the third step, we design a novel product-aspect temporal attention (PATA) module to help adjust the product embeddings guided by aspect embeddings for multi-length future times forecasting.

We evaluated ForeSeer on a real-world customer review dataset, which contains 11,000 products, 11,536,382 reviews, and a product similarity graph constructed using user queries and clicks. Our dataset contains the timestamp for each review and product over three years. We observed substantial improvement against comparison approaches on product aspect forecasting in both settings with accessible estimated or ground truth association. We found the product embeddings derived from ForeSeer show more visible patterns when the timestamp is increasing, suggesting an accurate incorporation of temporal information. In addition to aspect forecasting, ForeSeer achieved prominent performance on multi-time

future link prediction and review-aspect association prediction, demonstrating its wide applicability in modeling temporal review information. ForeSeer is developed as, to our knowledge, the first approach to product aspect forecasting and can be broadly applied to other temporal graph mining tasks.

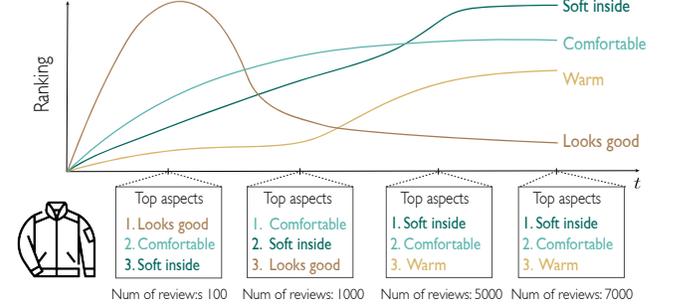


Figure 2: The key idea of ForeSeer is to find and embed similar products dynamically and propagate their top aspects mined from reviews to the new item to forecast future aspects.

2 PRODUCT ASPECT FORECASTING

2.1 Problem Definition: A new task

In principle, products receive and aggregate reviews and aspects are mined from the review text pieces, enabling the products to extract and aggregate aspects associations. We then understand relationships among the three core components, products, aspects, and reviews in detail, and formally define the product aspect mining and forecasting problem as a top- K_I aspect retrieval problem.

2.1.1 Product aspect mining from review texts. A descriptive aspect A_l , is a descriptive textual span extracted from a review text sequence r_j that summarizes a particular attribute or feature of P_i , the product r_j belongs to (examples and illustration in Table 1). Given an aspect list $L = \{A_1, A_2, \dots, A_{|L|}\}$ and r_j , we define the review-aspect association vector $\mathbf{v}_j \in \{0, 1\}^{|L|}$, where $v_j^l = 1$ if r_j and A_l are associated. The aspect list can be obtained by supervised or unsupervised mining from reviews [77, 84]. The longer the list, the more descriptive and diverse aspects can be included in the list.

A product network $\mathcal{G} = \{\mathcal{P}, \mathcal{E}\}$ is given with a set of product nodes \mathcal{P} and an edge set \mathcal{E} representing the similarities between products. \mathbf{A}^w is the weighted adjacency matrix of \mathcal{E} . Each product receives reviews and maintains a review list $R_i = \{r_{i,1}, \dots, r_{i,|R_i|}\}$. Let $\mathbf{v}_{i,k}$ be the review-aspect association vector of $r_{i,k}$ and $\mathbf{V}_i^L = \{\mathbf{v}_{i,k}^L\}_{k=1}^{|R_i|}$ be the collected A_l 's association list of P_i . The product-aspect association vector $\mathbf{u}_i \in \mathbb{R}^{|L|}$ can then be obtained by aggregating all the review-aspect association vectors in the review list through $\text{Agg}(\cdot)$, a permutation-invariant aggregation function:

$$\mathbf{u}_i = [\text{Agg}(V_i^1); \dots; \text{Agg}(V_i^{|L|})]. \quad (1)$$

The binarized top- K_I aspect label vector $\mathbf{y}_i^{K_I} = \mathbf{y}_i$ can then be constructed by ranking \mathbf{u}_i and picking up the top- K_I product-aspect associations of P_i to be positive. Let $\mathbf{x}_i \in \mathbb{R}^{d_p}$ be the node feature of P_i , the product aspect mining task is using \mathbf{x}_i and graphical information in \mathcal{G} such as \mathbf{A}^w to predict \mathbf{y}_i . Note that the number of positive labels can be less than K_I , resulting in the multi-label nature of \mathbf{y}_i .

Table 1: Three review-aspect association examples are presented to show the different difficulty levels mining these aspects. The mining difficulty for an aspect depends on its relationship to the related corpus (More challenging from top to bottom).

Relation with corpus	Product type	Aspect	Related corpus	Raw review texts
Exact match	Table	Very sturdy	very sturdy	This table is very sturdy .
Synonym	Candle	Looks great	looks perfect	I bought this product for my wife and it looks perfect for me.
Summary	Carpet cleaning	Great carpet machine	did a great job cleaner cleaning my carpet	This Compact Carpet Cleaner did a great job cleaning my carpet!

Table 2: Head and tail aspect examples with frequency information over a three-year period.

Head aspect	3-year frequency	Tail aspect	3-year frequency
Looks great	1,955,762	Buttons are good quality	145
Works great	1,796,908	Great fit and finish	145
Very sturdy	1,199,984	Simple plug and play	145

2.1.2 Predicting future aspect trends for products. In practice, products are released at different times and start receiving reviews over time. Consider in a discrete time period $\mathcal{T} = \{0, \dots, T\}$, the product network is evolving and $\mathcal{G}_t = \{\mathcal{P}_t, \mathcal{E}_t\}$ is the product network snapshot at time $t \in \mathcal{T}$. $P_i \in \mathcal{P}_{t > t_i^{start}}$ only occurs after its proposed time t_i^{start} in the monotonically increasing product node list \mathcal{P}_t over time. The product similarities edge set \mathcal{E}_t and its weighted adjacency matrix A_t^w are also changing as t changes. Products begin receiving reviews once become available to customers, therefore the corresponding review lists $R_i(t; t_i^{start}) = R_{i,t}$ are also expanding over time. Any review $r_{i,k}$ in $R_{i,t} = \{r_{i,1}, \dots, r_{i,|R_{i,t}|}\}$ also has a proposed time $t_{i,k}^{prop}, t_i^{start} \leq t_{i,k}^{prop} \leq t$. For simplicity, assume $R_{i,t}$ is ordered by ascending review proposed time. Now V_i^l will be $V_{i,t}^l$ with reviews up to $R_{i,t}$, and the product-aspect association vector becomes $\mathbf{u}_{i,t}$ by replacing the V_i^l to $V_{i,t}^l$ in Eqn (1).

Let the above t be the present time, and $\mathbf{x}_{i,t}$ be the node feature of P_i in \mathcal{G}_t . Suppose at a future time $t' = t + \Delta t$, the product aspect vector and the top- K_l aspect label vector is $\mathbf{u}_{i,t'}$ and $\mathbf{y}_{i,t'}$. The Δt -future product aspect forecasting problem at time t is to predict $\mathbf{y}_{i,t'}$ using $\mathbf{x}_{i,t}$ and \mathcal{G}_t . Moreover, we are interested in forecasting top product aspects $\mathbf{y}_{i,t+\Delta t_n}$ at multiple future with $N_{\Delta T}$ different time lengths $\Delta T \triangleq \{\Delta t_n\}_{n=1}^{N_{\Delta T}}$ simultaneously using information till t . In practice, the accessible information is till the present t (referred to as 'now'). Yet setting 't at now' to an earlier stage will therefore meet with the old product that is earlier 'new', together with a smaller product network that has fewer products with shorter review lists.

2.2 Motivations and challenges of the problem

2.2.1 Motivation. We illustrate the motivation of the product aspect forecasting problem. In a large e-commerce system, existing products keep receiving reviews from buyers, and sellers continue to release more and more new products (Fig. 1). The high-level target of this problem is predicting the customer feedback for products in the long future: e.g., *what will be the top aspects of a new product six months or three years later?* This problem might be easier to answer for products with long histories - they have many customer reviews, thus simply extracting all aspects till now and counting the most frequent ones can already obtain statistically significantly satisfying results. However, for the larger and more important group of newly released products with limited reviews, such a strategy will give a much less reliable aspect forecasting, as aspects that currently have low frequency might be frequently mentioned in future reviews. On one hand, customers cannot trust products with less accurate aspects prediction; on the other hand, sellers worry about long-history products monopolizing the market. Therefore,

a robust and reliable solution clearly relieves the dilemma that new products have limited reviews to help forecast on both the seller and customer side despite the non-trivial problem nature.

2.2.2 Challenges. Forecasting aspects for 'new' products can not be effectively achieved by trivially aggregating mined results from aspect mining methods, as the number of reviews is limited besides their subjectiveness. Also, the scale of reviews can be very large and aspect list length L can also be large in practice, leading to high sparsity of positive associations and extremely imbalanced aspect frequencies (Table 2). Therefore the top-K aspect retrieving formulation also introduces association sparsity and extremely imbalanced aspect frequencies, making it more challenging.

Another challenge is that the 'new' products are distributed at different time points, resulting in unaligned review list series evolving processes. It is because reviews will keep occurring for products once they are released, and all the products will receive different reviews at different times. Since forecasting aspect trends for the 'new' products with limited reviews is more important and challenging, trivially training a graph neural network from a single product graph is also not effective. The evolving product-review associations and product similarities need to be efficiently extracted, aggregated, and embedded across the whole time range to avoid over-fitting effects.

The third challenge is that the aspect evolving trends will be quite different in shorter and longer time periods, as $\mathbf{y}_{i,t'}$ will reflect the genuine converged customer feedback when $\Delta t \rightarrow \infty$ and the recent popular thoughts from customers when Δt is small. Therefore predicting top aspect trends for products at different Δt future indicates capturing different stages of product evolving from the entire 'new' (with limited reviews) to 'old' (with adequate reviews) evolving process from time to time. Trivially using one single product embedding at current time t without further future-length aware adjustment will result in unsatisfactory performance.

3 METHODS

We propose ForeSeer, a textual mining and product embedding framework progressively trained on temporal product graph to predict the dynamic product embeddings at a future time point $t + \Delta t$ using all information till the current time point t . ForeSeer then exploits the product embeddings to predict the top aspects of that future time. ForeSeer is a three-stage framework (Fig. 3). In the first stage, we exploit contrastive learning to co-embed reviews and aspects, leveraging the aspect semantic information to acquire better review embeddings. We resample review-aspect associations to increase the presence of rare aspects so that the review embeddings can be less biased to frequently-appeared aspects. In the second stage, we progressively train the product graph embedding network on the temporal product graph. This helps ForeSeer see more new products and the new reviews they received, as well as evolving product similarity information, resulting in an efficient product embedding alignment. In the third stage, we develop a

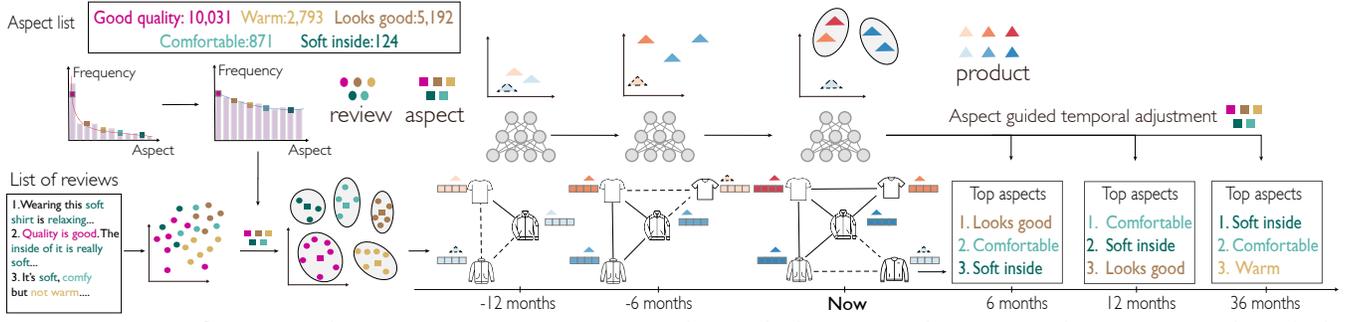


Figure 3: ForeSeer first resamples review-aspect associations and co-embed aspects and reviews. It then progressively embeds products leveraging temporal product graphs. Finally, it exploits temporal information to adjust product and aspect embeddings to provide multi-time future forecasting.

novel Product-Aspect Temporal Attention (PATA) module to further adjust product embeddings with the help of learned aspect embeddings. PATA module helps to forecast future product-aspect associations with explicit consideration of temporal information, thus providing time-sensitive forecasting.

3.1 A base model

We first introduce a base model that extends aspect mining methods for aspect forecasting and later contrast this base model with ForeSeer to clarify the key technical ideas of ForSeer. Let $\mathbf{R} = \{r_j\}_{j=1}^{|\mathbf{R}|}$ be the list of reviews collected from all products. The base model builds a BERT-based textual encoder that is pretrained on all reviews with the mask language modeling (MLM) objective [10]. The model can then be used to embed r_j as $\mathbf{h}_j = \text{BERT}(r_j)$. At present time t ('now'), the learned review embedding \mathbf{h}_j can then be aggregated to the product P_i they belong to as \mathbf{f}_i via an aggregator function. Let R_i be the review list of P_i at time t and $r_{i,k} \in R_i$ be one of the review of P_i , the aggregation process is:

$$\mathbf{f}_i = \text{Agg}(\{\mathbf{h}_{i,k} | r_{i,k} \in R_i\}). \quad (2)$$

The base model then leverages $\mathbf{x}_i = [\mathbf{f}_i; \mathbf{SE}_i]$ as a node feature of P_i and trains a graph neural network on the last observed product graph \mathcal{G}_{last} to further get a product embedding Z_{last} . Here, \mathbf{SE}_i is the static product-specific embedding that encodes other product-related information, such as seller descriptions, which will not change over time. Let A_{last}^w be the weighted adjacency matrix of \mathcal{G}_{last} , and X be the product node feature matrix of all products that make up \mathbf{x}_i , where $P_i \in \mathcal{G}_{last}$. Let σ be an activation function and $H^{(0)} = X$, for $H^{(l_p)}$, the output of the l_p -th GNN layer we have:

$$H^{(l_p)} = \sigma([A_{last}^w H^{(l_p-1)} W^{l_p}]). \quad (3)$$

Let $Z_{last} = H^{(l_p)}$ be the resulting embedding matrix of all products after the final layer, its i -th row \mathbf{z}_i is the embedding of P_i and is used to get the prediction of top aspects in the future. Suppose we are interested in multiple futures $\Delta T \triangleq \{\Delta t_n\}_{n=1}^{N_{\Delta T}}$, we use $N_{\Delta T}$ classifier heads $\{\text{CLS}_{P \rightarrow A, n}\}_{n=1}^{N_{\Delta T}}$ to map Z_{last} and different Δt_n to the prediction $\hat{\mathbf{y}}_{i, t+\Delta t_n}$ for different future $t + \Delta t_n$. They are finally optimized with a multi-time multi-label cross-entropy objective:

$$\hat{\mathbf{y}}_{i, t+\Delta t_n} = \text{CLS}_{P \rightarrow A, n}(\mathbf{z}_i, \Delta t_n),$$

$$\mathcal{L}_{P \rightarrow A}^{(i)} = \frac{1}{N_{\Delta T} \cdot |\mathbf{L}|} \sum_{n=1}^{N_{\Delta T}} \sum_{l=1}^{|\mathbf{L}|} -y_{i, t+\Delta t_n}^l \log \hat{y}_{i, t+\Delta t_n}^l. \quad (4)$$

3.2 Mining aspect to enhance review embedding

The major limitation of the base model is it only pre-trains the language model using review texts instead of clearly capturing the review-aspect association. Because reviews are usually long and subjective (Table 1), it is very inefficient to capture signals of potential aspects from the embedding of the entire review text. Also, to have more descriptive and diverse aspects, the aspect list can be very long. Under this situation, the positive aspect association in review texts will be very sparse, and the aspect distribution can be really imbalanced and long-tailed. We thus propose to refine the review embedding by fine-tuning the review encoder to mine aspect associations. The key intuition is to embed the review focusing on representing sparse aspect associations.

Specifically, given a review text r_j , we fine-tune its embedding \mathbf{h}_j to capture the review-aspect association vector $\hat{\mathbf{v}}_j$ with a classifier head $\text{CLS}_{r \rightarrow a}(\cdot)$ and a multi-label cross-entropy loss objective:

$$\hat{\mathbf{v}}_j = \text{CLS}_{r \rightarrow a}(\mathbf{h}_j); \quad \mathcal{L}_{cls, r \rightarrow a}^{(j)} = \frac{1}{|\mathbf{L}|} \sum_{l=1}^{|\mathbf{L}|} -v_j^l \log \hat{v}_j^l. \quad (5)$$

Note that here the association between the review texts and the resulted aspects will not change with time; it is therefore compatible to combine this objective with the MLM objectives [10] in the BERT pre-training stage. The resulting prediction of the association of the review aspect $\hat{\mathbf{v}}$ can be viewed as part of the refined review embeddings. For instance, the updated representation $\mathbf{h}'_j = [\mathbf{h}_j; \hat{\mathbf{v}}_j]$ can be the concatenation of BERT embedding and predicted association vectors. Having it can thus better capture sparse aspect association information during review aggregation using Eqn. (2).

3.3 Aspect-guided review embedding matching

However, the limitation of the above refinement is that it has no idea that some aspects might describe similar meanings. It thus lacks the ability to efficiently capture the tail aspects that might have a similar meaning to the top frequent aspects but are rare and more diverse. Incorporating class names is useful when classes have their own description [40]. Therefore, our next key improvement is to jointly embed reviews and aspects through cross-domain contrastive learning. The key intuition is that by also embedding aspects, tail aspects can be closer to head popular aspects if they have similar semantic meanings. The resulting aspect embeddings can therefore guide two review embeddings to be closer if they point to aspects with similar meanings. Moreover, if we treat the base

model as a review feature extractor, training a cascaded contrastive learner will enable an effective way to train multiple review feature extractors and incorporate resampling and ensemble techniques.

Specifically, we train multiple feature extractors by instantiating M base aspect extraction replicas with different reweighting factors $\alpha_1, \dots, \alpha_M$ in addition to the standard aspect extraction instance ($\alpha_0 = 0$). We use $\mathbf{h}_j^{(\alpha_m)}$ and $\hat{\mathbf{v}}_j^{(\alpha_m)}$ to denote the trained features output by the m -th replica follows Eqn (5). We then downsample frequent aspects to encourage the model to focus on review-aspect associations from infrequent aspects. For replica m , we first construct a bin for each aspect by assigning reviews that have positive associations with that aspect. Here, a review might be allocated to multiple aspect bins. We thus apply the resampling strategy hierarchically, which chooses the l -th aspect bin with probability p_l by reweighting and then normalizing the importance of the aspect based on the reciprocal of the α_m -reweighted aspect frequency:

$$p_l = \frac{w_l^{(m)}}{\sum_l w_l^{(m)}}, w_l^{(m)} = \left(\frac{1}{\text{freq}(A_l)}\right)^{\alpha_m}, \alpha_m \geq 0. \quad (6)$$

The reweight factor α_m leads the model to focus more on the aspects of tail (large α_m) or head (small α_m). $\alpha_m = 0$ corresponds to the standard aspect extractor that conduct uniform sampling. Then, a review r_j is sampled from the chosen aspect bin with an assigned importance weight based on the number of review-aspect associations and the reweighted aspect importance.

This resampling strategy enables the possibility to ensemble the features from all M feature extractors and perform mixture-of-experts (MoE) learning in the contrastive learner. We can now exploit contrastive learning to co-embed reviews and aspects so that the embedding of reviews can pay more attention to rare aspects. Specifically, we first build review and aspect embedding networks E_R and E_A with the same output dimension d . E_R encodes the embeddings of the review r_j and the predicted associations from all replicas of the extraction of aspects simultaneously and matches them into a calibrated embedding of the review \mathbf{z}_j^{re} :

$$\mathbf{z}_j^{re} = E_R(\text{CONCAT}([\mathbf{h}_j; \hat{\mathbf{v}}_j; \mathbf{h}_j^{(\alpha_1)}; \hat{\mathbf{v}}_j^{(\alpha_1)}; \dots; \mathbf{h}_j^{(\alpha_M)}; \hat{\mathbf{v}}_j^{(\alpha_M)}])). \quad (7)$$

For aspect A_l , we use SpanBERT [22] to get the pre-trained features and feed it into E_A to get the aspect embedding \mathbf{z}_l^{asp} :

$$\mathbf{z}_l^{asp} = E_A(\text{SpanBERT}(A_l)). \quad (8)$$

We then utilize the aspect embeddings \mathbf{z}_l^{asp} of A_l to guide the optimization of the calibrated review embedding \mathbf{z}_j^{re} of review r_j using the inner product of \mathbf{z}_j^{re} and \mathbf{z}_l^{asp} :

$$\begin{aligned} s_j^l &= (\mathbf{z}_j^{re})^T \mathbf{z}_l^{asp}; Z^{asp} = [\mathbf{z}_1^{asp}; \dots; \mathbf{z}_{|L|}^{asp}]; \\ s_j &= (\mathbf{z}_j^{re})^T Z^{asp} = [s_j^1; \dots; s_j^{|L|}] \in \mathbb{R}^{|L|}. \end{aligned} \quad (9)$$

We finally employ a multi-label cross-domain contrastive loss similar to NT-Xent loss [56] to maximize positive association and minimize negative association denoted by \mathbf{v}_j :

$$\mathcal{L}_{CL}(s_j, \mathbf{v}_j; \tau) = -\log \sum_l \frac{v_j^l \exp(s_j^l / \tau)}{\sum_k \exp(s_j^k / \tau)}. \quad (10)$$

By jointly learning an informative, low-dimensional embedding space for reviews and aspects, the resulting \mathbf{s}_j has two advantages.

First, it will efficiently ensemble learnings from multiple aspect extraction replicas, thus more effectively capturing tail aspects. Second, by minimizing Eqn (10), \mathbf{s}_j will encode the ranking information of the aspects for r_j . This indicates a more flexible way to retrieve K_a positive associations. By simply tuning K_a , more (recall-inclined) or less (precision-inclined) positive associations can be given, resulting in a more flexible aspect extraction module.

3.4 Progressive temporal graph-based product embedding

The major limitation of the graph-based product embedding module of the base model is that it is only trained on the last observed product graph snapshot with the product features at that time. In fact, only a small ratio of products are ‘new’ in the last observed graph snapshot. This leads the model to have severe overfitting effects. We thus propose to progressively train the model on the temporal product graph snapshot series rather than only on the last observed graph. The key intuition is that the model can not only see more ‘new’ products during training but also be aware of the other temporal evolving information, including product similarities and how a product becomes from ‘new’ to ‘old’ by receiving more and more reviews.

Without loss of generality, we now let $t \in \mathcal{T}$ be a time point we want to train our model at and let the graph be the snapshot \mathcal{G}_t at that time. To enable progressive training, multiple fixations are needed. First, the feature vectors of product nodes need to be aware of the evolving temporal information, such as the related information of the progressive product-aspect associations until time t . Given a review $r_{i,k} \in R_{i,t}$, the review list of P_i at time t , we aggregate the learnt embedding $\mathbf{z}_{i,k}$ and the predicted review-aspect association $\mathbf{s}_{i,k}$ for all reviews in $R_{i,t}$ to acquire the product feature $\mathbf{f}_{i,t}$ of P_i at time t :

$$\mathbf{f}_{i,t} = \text{Agg}(\{\text{CONCAT}([\mathbf{z}_{i,k}; \mathbf{s}_{i,k}]) | r_{i,k} \in R_{i,t}\}). \quad (11)$$

We then incorporate the related temporal information of P_i at time t . Similar to the positional embedding [61], the 2d-dimensional temporal embedding of a discrete-time index t_0 is:

$$\text{Emb}_{t_0} = [\sin(\omega_1 t_0), \cos(\omega_1 t_0), \dots, \sin(\omega_d t_0), \cos(\omega_d t_0)]. \quad (12)$$

For P_i at time t , there are three pieces of temporal information that need to be incorporated: its proposed time t_i^{start} , the training time t and the gap between these two times $t_i^{gap} = t - t_i^{start}$. The resulting temporal embedding therefore has three sub-parts:

$$T_{i,t}^{emb} = \text{CONCAT}([\text{Emb}_{t_i^{start}}, \text{Emb}_t, \text{Emb}_{t_i^{gap}}]). \quad (13)$$

We then compose the overall node feature vector $\mathbf{x}_{i,t}$ for P_i at training time t with three parts of features mentioned before:

$$\mathbf{x}_{i,t} = \text{CONCAT}([\mathbf{f}_{i,t}; T_{i,t}^{emb}; \mathbf{SE}_i]). \quad (14)$$

Let \mathbf{A}_t^w be the normalized weighted adjacency matrix of \mathcal{G}_t , to enable progressive training, we sequentially train the graph neural network at a subset of graph snapshots $\{\mathcal{G}_{t_1}, \dots, \mathcal{G}_{t_Q} | t_q \in \mathcal{T}\}$. We further refine the GNN layer that contains an extra multi-layer perceptron (MLP) module that focuses on learning product-specific information. Specifically, let $W_{t_{q-1}}^l$ and $W_{MLP, t_{q-1}}^l$ be the weight matrix of l -th GNN layer and the extra MLP module after training

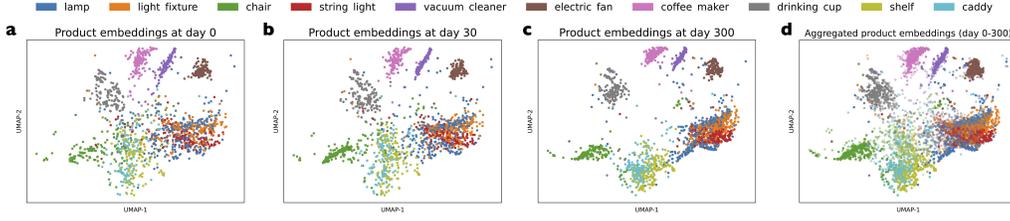


Figure 4: UMAP plots visualizing product embeddings at day 0 (a), day 30 (b), day 300 (c), and aggregation of product embeddings at all of the three time points (d).

the network on $\mathcal{G}_{t_{q-1}}$ and let $H_{t_q}^{(0)} = \mathbf{x}_{i,t_q}$, the product embedding process in Eqn (3) will be updated at t_q as:

$$H_{t_q}^{(l_p)} = \sigma(\text{CONCAT}([A_{t_q}^w H_{t_q}^{(l_{p-1})} W_{t_{q-1}}^{l_p}; W_{MLP,t_{q-1}}^{l_p} H_{t_q}^{(l_{p-1})}])). \quad (15)$$

At present time t , we take $Z_t^P = H_t^{(L_p)}$ got on \mathcal{G}_t as the resulting product embedding matrix. By progressively updating the network, ForeSeer can see more new products and how they ‘grow’ from ‘new’ to ‘old’, which is helpful in efficiently aligning the product embedding.

3.5 PATA: Aspect-guided product embedding adjustment for multi-future forecasting

One limitation of the base model is it uses only one product embedding for all Δt_n future predictions without any refinement when capturing future product-aspect associations. We thus develop a novel multi-head product-aspect temporal attention (PATA) module to adjust product embedding based on learned aspect embedding Z^{asp} . Our intuition is to adjust product embedding with a designated attention module guided by aspect embedding. This novel design empowers ForeSeer with the ability to predict both short and long-future time ranges with lower computational costs.

Specifically, given $N_{\Delta T}$ target futures, for each Δt_n , we set up a query network $EP_{\Delta t_n}^Q$ and a key network $EA_{\Delta t_n}^K$ to get the product query matrix from Z_t^P and aspect key matrix from Z^{asp} :

$$\begin{aligned} Q_{t+\Delta t_n}^P &= EP_{\Delta t_n}^Q(Z_t^P). \\ K_{t+\Delta t_n}^{asp} &= EA_{\Delta t_n}^K(Z^{asp}) \end{aligned} \quad (16)$$

We then calculate cross-attention score $Attn_{t+\Delta t_n}$ between product query and aspect key for Δt_n future:

$$Attn_{t+\Delta t_n} = \text{Softmax}\left(\frac{Q_{t+\Delta t_n}^P (K_{t+\Delta t_n}^{asp})^T}{\sqrt{d}}\right) \in \mathbb{R}^{|\mathcal{P}_t| \times |L|}. \quad (17)$$

We then set up a shared value network EV^{val} that takes Z_t^P as input and outputs value matrix Val_t^P . Let $val_{i,t}$ and $attn_{i,t+\Delta t_n}$, the i -th row of Val_t^P and $Attn_{t+\Delta t_n}$, be the value and attention score vector for P_i , we obtain the final product embedding for Δt_n future as:

$$z_{i,t+\Delta t_n}^{final} = attn_{i,t+\Delta t_n} \circ val_{i,t}. \quad (18)$$

We next exploit the final product embedding to get the final prediction using Eqn (4). With the above design, we can acquire different product embedding for different Δt_n future prediction tasks. Using a shared value network is especially efficient when $N_{\Delta T}$ is large and $d \ll |L|$. This also regularizes the model by sharing weights for learning multiple imbalanced top- K_l aspect retrieving tasks.

4 EXPERIMENT

In this section, we evaluated ForeSeer on a large scale e-commerce dataset with a variety of tasks, aiming to answer the following research questions (RQs):

- **RQ1:** How does ForeSeer perform on product aspect forecasting without or with genuine annotated aspects for reviews?
- **RQ2:** Is ForeSeer sensitive to aspects with long trends?
- **RQ3:** How does different components contributes to the success of ForeSeer on capturing review-aspect associations?
- **RQ4:** Is the PATA module in ForeSeer able to capture temporal information for different lengths of future?
- **RQ5:** How does the product embedding captured by ForeSeer evolve over time?

4.1 Experimental setup

4.1.1 Dataset. We evaluated our ForeSeer framework on a large-scale e-commerce dataset that contains a dynamic product network with 11,536,382 time-stamped review events in a three-year period. The product network is a series of homogeneous dynamic product similarity graphs with 1,096 daily time-stamped snapshots. The final graph has 11,000 product nodes from 418 different product types based on their name, high-level property, and usage, such as "Lamp", "TABLE", and "Caddy". The unpublished product nodes will not be added to the graph until their proposed times are reached. The average start date of products is 77. The first snapshot has 999,614 edges and the final snapshot has 2,400,404 edges. Each edge represents a similarity ranging from 0 and 1 between the two products. We calculate the similarities based on multiple factors such as user click information. We obtained an aspect list with 10,000 popular aspects extracted via sequential extraction tools [76, 77] and 30,217,638 review-aspect associations from these reviews. We constructed the top- K_l ($K_l = 10$) aspects label for every product node at all timesteps by aggregating all the review-aspect associations it has and picking up the top-ranked aspects after normalization. Example head and tail aspects with their frequency over three years are shown in Table 2. While the top 3 aspects have more than 1,000,000 counts, the tail aspects only have 150 counts from more than 11 million review texts, demonstrating extreme aspect imbalances.

4.1.2 Tasks. We studied ForeSeer with three challenging tasks: single review-aspect association prediction, multi-time aspect forecasting, and multi-time link prediction.

Single review-aspect association prediction: We model the single review-aspect association prediction as a multi-label classification problem. We split 80% of the reviews as the training set and leave the remaining 20% as the test set.

Multi-time aspect forecasting: To further assess the effectiveness of the learned review associations, we tested the multi-time aspect

Table 3: Performance of ForeSeer and comparison approaches on aspect forecasting at three different time gaps (3 months, 6 months, 3 years) under settings of with annotated aspects and without annotated aspects.

Δt	Method	Without annotated aspects				With annotated aspects			
		AUPRC	AUROC	Kappa	Max F1	AUPRC	AUROC	Kappa	Max F1
3 months	Frequency baseline	21.0 ± 5.7	84.8 ± 4.1	28.6 ± 6.2	37.4 ± 5.6	91.2 ± 4.2	96.9 ± 2.0	86 ± 5.4	90.0 ± 4.2
	MLP	13.1 ± 7.2	89.8 ± 3.2	4.5 ± 4.3	25.1 ± 6.5	14.6 ± 7.7	91.3 ± 2.9	4.9 ± 4.1	28.5 ± 7.1
	LSTM	17.0 ± 6.0	93.7 ± 2.5	6.5 ± 5.1	29.6 ± 6.1	17.4 ± 6.1	93.8 ± 2.6	7.3 ± 5.5	30.1 ± 6.2
	GRU	17.3 ± 6.2	93.6 ± 2.6	7.0 ± 5.5	30.1 ± 5.5	18.2 ± 6.2	94.0 ± 2.6	5.9 ± 5.3	31.0 ± 6.3
	GNN	19.2 ± 6.6	95.1 ± 3.8	25.1 ± 8.1	32.3 ± 6.8	20.9 ± 6.5	97.2 ± 3.9	24.9 ± 7.0	34.2 ± 6.6
	ForeSeer (Ours)	84.1 ± 7.2	99.9 ± 0.2	75.8 ± 7.9	83.9 ± 5.9	80.2 ± 8.4	99.9 ± 0.1	72.3 ± 8.2	80.0 ± 7.2
6 months	Frequency baseline	20.2 ± 5.6	84.3 ± 4.2	27.5 ± 6.2	36.1 ± 5.6	87.4 ± 5.1	96.4 ± 2.2	81.8 ± 6.0	86.3 ± 4.9
	MLP	11.1 ± 7.1	88.4 ± 3.5	4.3 ± 4.1	23.5 ± 6.7	13.9 ± 7.5	90.9 ± 2.7	5.1 ± 4.0	27.6 ± 6.8
	LSTM	14.9 ± 5.7	93 ± 2.8	6.6 ± 5.2	27.5 ± 6.0	15.3 ± 5.8	93.0 ± 2.9	7.2 ± 5.5	27.8 ± 6.1
	GRU	15.4 ± 5.9	92.9 ± 2.8	6.7 ± 5.3	28.1 ± 6.2	16.2 ± 6.0	93.3 ± 5.6	28.9 ± 6.3	26.3 ± 5.3
	GNN	18.1 ± 6.5	93.8 ± 3.5	23.0 ± 7.7	29.4 ± 6.7	19.3 ± 6.4	95.9 ± 3.3	22.9 ± 6.9	31.9 ± 6.6
	ForeSeer (Ours)	81.8 ± 7.7	99.8 ± 0.2	74.2 ± 8.1	82.2 ± 6.3	82.3 ± 8.6	99.7 ± 0.3	73.6 ± 8.5	82.6 ± 6.9
3 years	Frequency baseline	18.6 ± 5.4	83.8 ± 4.2	25.4 ± 6.1	34.0 ± 5.6	81.0 ± 6.1	96.2 ± 2.3	75.8 ± 6.5	81.2 ± 5.3
	MLP	10.2 ± 4.9	86.3 ± 3.0	3.1 ± 3.0	21.2 ± 5.3	11.7 ± 5.3	89.7 ± 4.2	5.5 ± 5.2	23.3 ± 6.1
	LSTM	11.4 ± 5.0	90.3 ± 3.3	5.8 ± 4.9	23.4 ± 5.7	12.0 ± 5.2	90.3 ± 3.5	6.6 ± 5.3	24.1 ± 6.0
	GRU	12.1 ± 5.2	90.3 ± 3.3	6.0 ± 5.1	24.4 ± 5.9	13.1 ± 5.4	90.7 ± 6.7	5.9 ± 5.1	25.1 ± 6.0
	OA-Mine	5.9 ± 3.9	78.6 ± 5.9	3.4 ± 2.7	6.1 ± 2.6	5.9 ± 3.9	78.6 ± 5.9	3.4 ± 2.7	6.1 ± 2.6
	GNN	16.1 ± 6.3	92.1 ± 3.1	20.9 ± 7.1	28.6 ± 6.6	17.3 ± 6.4	93.0 ± 2.7	20.6 ± 6.9	29.7 ± 6.5
ForeSeer (Ours)	68.5 ± 8.9	98.9 ± 1.1	65.6 ± 8.5	73.2 ± 6.9	83.6 ± 7.9	99.9 ± 0.1	74.7 ± 7.7	82.8 ± 6.4	

forecasting task with two settings: (i) ground truth association accessible, (ii) only learned approximate association accessible. For this task, we chose to predict $\Delta t = 3, 6$ months and 3 years (always predicting the last observed label) at every test timestep t . We use the label at the final time if the target future time is out of range. We aimed to test the inductive performance by randomly splitting 10% nodes out as the test set.

Multi-time link prediction: For multi-time link prediction, we chose to predict the link at the current time, 180 days later, and at the final time at every test time t . The positive edge ratio is 50%.

4.1.3 Comparison approaches. There is no existing framework that directly handles the multi-time aspect forecasting problem. We therefore designed and compared our method with four types of baselines. **Frequency baseline** directly aggregates the association learnt in section 3.2 at time t and uses it as the prediction. It can not incorporate other information and the performance will highly depend on the quality of learned aspect associations from reviews. We aim to assess the effectiveness of the multi-time forecasting PATA module. **MLP baseline** directly predicts the multi-time future with product features without graphical information. **Recurrent-based baselines (GRU, LSTM)** predict multi-time future step-by-step by using an autoregressive network structure and also ignore graphical similarity information. We thus aim to assess the importance of using graphical information. **Graph-based baseline (GNN)** leverages only the final graphical information \mathcal{G}_{last} . We aim to assess the effectiveness of our progressive training process in multiple Δt -future. **Weakly supervised baseline (OA-Mine)** directly mines candidate aspects from reviews and aggregates the candidates as the final predictions. We aim to assess the effectiveness of learned review-aspect associations and the importance of using learning-based models as this baseline does not leverage review-aspect association information.

For multi-time link prediction tasks for different Δt -future, we only compare our method with graph-based baselines, since other baselines can not handle this task. We built an extra lightweight GNN for multi-task link prediction for our methods and a GNN

baseline with the same network structure. For our method, we used features learned from our multi-head product aspect temporal attention module for different futures as the input for different future time predictions. For the graph-based baseline, we used the same features for different future time predictions.

4.1.4 Implementation details. For the aggregation function $Agg(\cdot)$ used to construct future aspect labels, we used $sum(\cdot)$ with per-aspect z-score normalization operation. Specifically, for each aspect, we got the sum of its occurred association for every product at the end of the time period and calculated the mean and standard deviation of them. We then performed z-score normalization on the sum of the associations of this aspect for every product at time t . We used $mean(\cdot)$ as the aggregator to get the aggregated feature $f_{i,t}$ of product P_i at time t . We maintained both the predicted logits and the binarized association prediction from $s_{i,k}$. We incorporated a pre-trained product embedding obtained from a multi-modality neural model that encodes information from the product title and descriptions to construct the auxiliary product embedding SE_i .

For simplicity, we only trained one resampling replica with $\alpha = 0.5$ and match its embedding with embedding learned by standard aspect extraction replica ($\alpha = 0$). We used AdamW as the optimizer and 10^{-4} as the learning rate to train the aspect-guided embedding match model for 10 epochs on four 16GB Nvidia-V100 GPUs and set the number of retrieved positive associations $K_a = 5$ for better precision performance. We used a two-layer MLP with ReLU for all the embedding networks and a linear projection layer for all query and key blocks. All the embedding, query, and key vector output dimensions are 100. To avoid the over-smoothing problem, we used a GNN with one designed graph neural layer and two cascaded MLP layers. We used the same network structure for the GNN baseline. For the MLP baseline, we used a two-layer MLP with a hidden size of 128, for GRU and LSTM baseline we set the number of the layer as 2 with a hidden size of 128. For the lightweight GNN for multi-task link prediction task, we used a standard GraphSAGE layer and one cascaded MLP layer. We used Adam as the optimizer and 10^{-3} as the learning rate and trained all aspect forecasting and link prediction tasks for 200 epochs on a 16GB Nvidia-V100 GPU.

Table 4: Case study showing the aspect we predicted for skin moisturizer. Our method identified time-sensitive phrases (red) while the frequency baseline failed to identify them.

Product type	Skin moisturizer
Count baseline prediction with one month reviews	Smooth, Feels great, Product is great, Clean, Light weight, Soothing, Love this product, Dries quickly, Works great, Great price
ForeSeer prediction with one month reviews	Smooth, Clean, Great price, Product is great, Soothing, A little goes a long way , Sensitive skin , Light weight, Not greasy , Good quality
Final Label (Ground truth)	Smooth, Love the texture, Light weight, Product is great, Soothing, Great price, A little goes a long way , Sensitive skin , Fragrance free, Not greasy

4.2 Experimental results

4.2.1 Improvement on aspect forecasting without annotated aspects (RQ1). We first investigated the performance of ForeSeer on aspect forecasting when aspects are not annotated in the review (Table 3). We found that our method achieved the best performance on all three pieces of time gaps (3 months, 6 months, and 3 years), indicating that our temporal graph embedding framework can effectively model the dynamics of aspects and products. For instance, ForeSeer obtains 63.1% and 49.1% improvements on 90-day gap forecasting and 3-year gap forecasting, respectively. We found that the frequency baseline didn't perform well on this task, especially when the time gap is larger, further confirming the importance of aligning products temporally. We found that the GNN baseline has undesirable performance, demonstrating the effectiveness of progressive training on avoiding over-fitting. We found that the MLP, GRU, and LSTM baselines performed badly on all future forecasting, necessitating the importance of graphical information. We also found that two temporal baselines outperform the MLP baseline, suggesting the importance to adjust product embedding for multi-time forecasting. We found that OA-Mine, one of the state-of-the-art approaches in aspect mining, obtained a less prominent performance, indicating that aspect mining methods cannot be applied to this novel aspect forecasting task. Finally, we noticed that graph-based approaches GNN, in general, performed better than methods that do not consider graphs, necessitating the consideration of graph dynamics in this task.

4.2.2 Improvement on aspect forecasting with annotated aspects (RQ1, RQ2). We next evaluated an easier setting where aspects are annotated in each review. In this setting, the frequency baseline achieved very good performances since it simply counts the aspects in current reviews and uses it to predict future aspects. For aspects that are not time-sensitive, the frequency baseline can be regarded as an upper bound. Nevertheless, we found that ForeSeer still achieved good performance in this setting, where ForeSeer showed comparable performance with the frequency baseline on a 6-month gap and even outperformed the frequency baseline on a 3-year gap. All other comparison approaches obtained a much less promising performance. This indicates that ForeSeer can obtain a comparable performance with the frequency baseline on aspects that are not time-sensitive and substantially better performance on aspects that are time-sensitive. To further examine this (RQ2), we illustrated one case study showing how ForeSeer can effectively recognize long-term aspects (Table 4). ForeSeer corrected forecasted time-sensitive aspects such as "a little goes a long way" and "sensitive skin", while the frequency baseline failed to identify, reassuring ForeSeer's ability to forecast long-term future aspects.

4.2.3 Improvement on predicting review-aspect association with embedding matching (RQ3). Next, we examined the importance of the

aspect-guided embedding matching strategy by performing an ablation study on the review-aspect association prediction (Fig. 5). We observed the limited precision performance of using either with or without resampling multi-label classification instances, leading to poor ability or focusing too much on capturing tail aspect associations. Instead, we observed substantial improvements from our aspect-guided embedding framework by successfully exploiting the advantage of two multi-label classification instances. For example, the precision of our precision-inclined embedding matching model ($K_a = 5$) is 22% and 34% higher than the with or without resampling multi-label classification instances. The superior precision performance suggests the learned review embeddings guided by aspects concentrate less-biased features for both head and tail aspects, which is critical to their success.

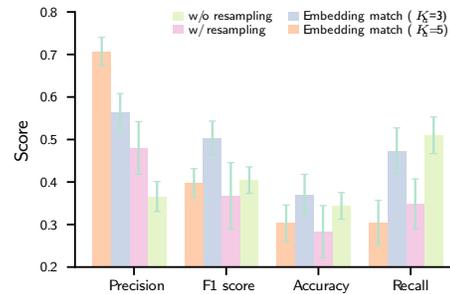


Figure 5: Bar plot showing the performance of our embedding match approach and comparison approaches on review-aspect association prediction evaluated using precision, F1 score, accuracy, and recall.

4.2.4 Improvement on predicting future links (RQ4). After confirming the superior performance of ForeSeer on aspect forecasting, we next investigated whether ForeSeer's PATA module is able to capture temporal information for multiple futures by letting it predict product-product edges that might emerge in the multiple future times (Table 5). We observed a consistent improvement of our methods over different time gaps compared to GNN. We adopted a multi-task prediction setting for efficient inference and training where predictions for different time gaps are treated as different tasks. Specifically, we found that the GNN baseline cannot generalize well on the intermediate-long future range (6 months) as it didn't consider the dynamics when simultaneously predicting multiple time gaps. In contrast, our model addressed this by explicitly modeling the temporal dynamics of products, confirming that ForeSeer can effectively capture both product and aspect dynamics.

4.2.5 Contrasting static and temporal embeddings (RQ5). We further visualized the dynamic product embedding space colored by their product types at different time steps to assess how ForeSeer captures the temporal pattern (Fig. 4). As the time step becomes

Table 5: Comparison of ForeSeer and GNN on predicting future links at three different time points (current, 6 months, 3 years).

Δt -future	Approaches	Accuracy	Precision	Recall	F1 score	AUPRC	AUROC
Current (0 day)	Baseline GNN	88.32 ± 0.54	85.29 ± 0.37	91.28 ± 0.72	88.13 ± 0.61	94.64 ± 0.88	94.73 ± 0.76
	ForeSeer (Ours)	92.64 ± 0.53	92.68 ± 0.35	92.66 ± 0.69	92.65 ± 0.58	97.54 ± 0.83	97.66 ± 0.75
6 months	Baseline GNN	18.36 ± 0.32	10.93 ± 0.24	89.02 ± 0.72	19.49 ± 0.23	89.13 ± 0.84	44.79 ± 0.24
	ForeSeer (Ours)	85.12 ± 0.51	92.96 ± 0.38	91.87 ± 0.78	92.37 ± 0.55	92.21 ± 0.87	74.91 ± 0.49
3 years	Baseline GNN	80.28 ± 0.48	77.95 ± 0.35	87 ± 0.63	82.03 ± 0.77	92.24 ± 0.88	89.34 ± 0.81
	ForeSeer (Ours)	84.25 ± 0.51	92.95 ± 0.44	87.31 ± 0.88	86.39 ± 0.64	94.08 ± 0.96	91.75 ± 1.02

larger, we observed more visible clustering patterns among products. This suggests that the quality of the product embeddings becomes better as the data is accumulated by modeling the temporal dynamics. We also noticed the product embedding similarity reflects edges in the product-product graphs ("shelf" and "caddy", "lamp", "light fixture" and "string light"), indicating these product embeddings successfully encode graphical information.

5 RELATED WORKS

Product aspect mining. Few work that tries to forecast aspects for products and most previous works only aim at product aspect mining tasks. Mining product aspects from large-scale commercial data has been a well-studied problem [7, 15, 25, 53, 77]. Previous works target the problem using rule-based [7, 25, 53], supervised learning based extractor [4, 18] or propagation-based methods [35, 46]. These works are less generalizable to product aspect forecasting, as they require either domain-specific features or supervision signals from downstream tasks such as sentiment or opinion labels. Recently, other works similarly extract descriptive product attributes from seller-provided product profiles by using sequential supervised textual span labeling information [29, 64, 70, 76, 93] or distant supervision [65, 83]. OA-Mine [84] leverage weakly-supervised seed set to mine aspects from product titles. These methods need either more costly massive sequential labeled datasets or predesigned hierarchical attribute taxonomy, which is hard to obtain for customer-centered aspects. Zero-shot aspect extraction is another recently emerged related topic that aims to extract aspects in new domains without annotated data by leveraging transfer learning [17], natural language inference [54], document sentiment classification [9] and the recent emerged super-scale ChatGPT [44] and GPT-4 [2]. However, the zero-shot indicates that no prior knowledge of the new domain is needed instead of the number of reviews being limited for products. They are thus not directly applicable to forecasting tasks. In summary, all the methods above focus on text sequence level extraction and do not efficiently aggregate learned attributes at the product level. They are thus useful to help construct the aspect list for product aspect forecasting tasks, but are not able to forecast future aspect trends. Unlike these methods, our methods are the first framework that can forecast aspects that might be mentioned in future reviews for a new product with a limited number of reviews.

Label-guided classification. Label-guided text classification has been studied in fields such as social recommendation and document classification [6, 32]. It has been proven to be beneficial for classification performance [37, 68]. While these methods focus on classification, our method provides an efficient cross-domain contrastive learning framework that can easily ensemble these models as expert multi-label classification instances. Compared to [63], which also proposes to mix multiple encoder instances, our representation learning scheme allows the model to also produce informative reviews and aspect representations to help downstream tasks.

Temporal graph learning. Temporal graph-based representation learning has been intensively studied [14, 49, 74]. Previous works either incline on graph structure [28, 43, 52, 87, 91] or temporal dynamics [21, 38, 41, 59, 60, 72]. The dynamic graph structure is extracted by graph adjacency dynamics [12, 30], skip-gram-based modeling [11, 39, 95], leveraging clique information [16, 81, 90], step-by-step embedding updating [3, 13, 31, 33, 47, 55, 92]. These methods only focus on capturing implicit or short graphical dynamics changing and can not efficiently model large-scale review-aspect associations and evolving product-review associations. Some methods are designed specifically for heterogeneous networks [19, 66, 82, 85, 86]. Temporal-dynamic focused methods leverage recurrent-based models [1, 8, 23, 67, 94], continuous point process modeling [5, 50] and self-supervised graph representation learning [57, 58]. These methods focus on modeling interactive events between nodes, but are not suitable for modeling dynamic product similarities and massive product-review associations. Some methods leverage self-attention, hierarchical, or temporal attention [20, 36, 51, 69, 75, 78, 89], but only for temporal information aggregation, while our PATA module leverages attention to predict multiple time range future simultaneously. Other recent works such as EvolveGCN [45] and Roland [80] incorporate recurrent modules and efficient graph learning. These frameworks can be easily adapted to our progressive training framework, while we aim to clearly show that progressive training is the key success to in aspect forecasting with a simple yet clear formulation. JODIE [27] is the closest work that can also predict multiple time-range futures with a time projection module. However, it is based on interaction events and is designed for heterogeneous user-item networks, making it impractical for our use case. In summary, most of the existing works can not be directly applied to the product aspect forecasting task because of the training efficiency and inability to handle large-scale reviews.

6 CONCLUSION

In this paper, we have studied a novel task of product aspect forecasting, which aims to predict aspects that users might mention in future reviews. We have proposed a novel framework ForeSeer to solve this task by dynamically embedding products, reviews, and aspects. We have evaluated our method on a large-scale real-world dataset and observed ForeSeer's superior performance on aspect forecasting and link prediction. In the future, we are interested in boosting ForeSeer with more advanced progressive training strategies. We are interested in applying ForeSeer to other temporal graph embedding frameworks, such as modeling biological signaling pathways. We are further interested in exploring how ForeSeer can assist the classic task of aspect mining, allowing us to apply ForeSeer to advance a larger number of e-commerce applications, such as integrating it with recommendation systems to enhance the product suggestions based on anticipated aspect preferences.

REFERENCES

- [1] Alex Beutel, Paul Covington, Sagar Jain, Can Xu, Jia Li, Vince Gatto, and Ed H Chi. 2018. Latent cross: Making use of context in recurrent recommender systems. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. 46–54.
- [2] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712* (2023).
- [3] Lei Cai, Zhengzhang Chen, Chen Luo, Jiaping Gui, Jingchao Ni, Ding Li, and Haifeng Chen. 2021. Structural temporal graph neural networks for anomaly detection in dynamic graphs. In *Proceedings of the 30th ACM international conference on Information & Knowledge Management*. 3747–3756.
- [4] Abir Chakraborty. 2022. Aspect Based Sentiment Analysis Using Spectral Temporal Graph Neural Network. *arXiv preprint arXiv:2202.06776* (2022).
- [5] Xiaofu Chang, Xuqin Liu, Jianfeng Wen, Shuang Li, Yanming Fang, Le Song, and Yuan Qi. 2020. Continuous-time dynamic graph learning via neural interaction processes. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 145–154.
- [6] Eli Chien, Wei-Cheng Chang, Cho-Jui Hsieh, Hsiang-Fu Yu, Jiong Zhang, Olga Milenkovic, and Inderjit S Dhillon. 2021. Node Feature Extraction by Self-Supervised Multi-scale Neighborhood Prediction. *arXiv preprint arXiv:2111.00064* (2021).
- [7] Anjali Dadhich, Blessy Thankachan, and Blessy Thankachan. 2021. Social & Juristic challenges of AI for Opinion Mining Approaches on Amazon & Flipkart Product Reviews Using Machine Learning Algorithms. *SN Computer Science* 2, 3 (2021), 1–21.
- [8] Hanjun Dai, Yichen Wang, Rakshit Trivedi, and Le Song. 2016. Deep coevolutionary network: Embedding user and item features for recommendation. *arXiv preprint arXiv:1609.03675* (2016).
- [9] Pengfei Deng, Jianhua Yuan, Yanyan Zhao, and Bing Qin. 2022. Zero-shot Aspect-level Sentiment Classification via Explicit Utilization of Aspect-to-Document Sentiment Composition. *arXiv preprint arXiv:2209.02276* (2022).
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [11] Lun Du, Yun Wang, Guojie Song, Zhicong Lu, and Junshan Wang. 2018. Dynamic Network Embedding: An Extended Approach for Skip-gram based Network Embedding. In *IJCAI*, Vol. 2018. 2086–2092.
- [12] Palash Goyal, Sujit Rokka Chhetri, and Arquimedes Canedo. 2020. dyngraph2vec: Capturing network dynamics using dynamic graph representation learning. *Knowledge-Based Systems* 187 (2020), 104816.
- [13] Palash Goyal, Nitin Kamra, Xinran He, and Yan Liu. 2018. Dyngem: Deep embedding method for dynamic graphs. *arXiv preprint arXiv:1805.11273* (2018).
- [14] Jin Guo, Zhen Han, Zhou Su, Jiliang Li, Volker Tresp, and Yuyi Wang. 2022. Continuous Temporal Graph Networks for Event-Based Graph Data. *arXiv preprint arXiv:2205.15924* (2022).
- [15] Mingqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*. 168–177.
- [16] Hong Huang, Zixuan Fang, Xiao Wang, Youshan Miao, and Hai Jin. 2020. Motif-Preserving Temporal Network Embedding. In *IJCAI*. 1237–1243.
- [17] Soufian Jebbara and Philipp Cimiano. 2019. Zero-shot cross-lingual opinion target extraction. *arXiv preprint arXiv:1904.09122* (2019).
- [18] A Jeyapriya and CS Kanimozhi Selvi. 2015. Extracting aspects and mining opinions in product reviews using supervised learning algorithm. In *2015 2nd international conference on electronics and communication systems (ICECS)*. IEEE, 548–552.
- [19] Yugang Ji, MingYang Yin, Yuan Fang, Hongxia Yang, Xiangwei Wang, Tianrui Jia, and Chuan Shi. 2021. Temporal heterogeneous interaction graph embedding for next-item recommendation. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 314–329.
- [20] Pengfei Jiao, Xuan Guo, Xin Jing, Dongxiao He, Huaming Wu, Shirui Pan, Maoguo Gong, and Wenjun Wang. 2021. Temporal network embedding for link prediction via vae joint attention mechanism. *IEEE Transactions on Neural Networks and Learning Systems* (2021).
- [21] Guangyin Jin, Min Wang, Jinlei Zhang, Hengyu Sha, and Jincui Huang. 2022. STGNN-TTE: Travel time estimation via spatial-temporal graph neural network. *Future Generation Computer Systems* 126 (2022), 70–81.
- [22] Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. 2020. Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics* 8 (2020), 64–77.
- [23] Seyed Mehran Kazemi, Rishab Goel, Sepehr Eghbali, Janahan Ramanan, Jaspreet Sahota, Sanjay Thakur, Stella Wu, Cathal Smyth, Pascal Poupard, and Marcus Brubaker. 2019. Time2vec: Learning a vector representation of time. *arXiv preprint arXiv:1907.05321* (2019).
- [24] Sung Guen Kim and Juyoung Kang. 2018. Analyzing the discriminative attributes of products using text mining focused on cosmetic reviews. *Information Processing & Management* 54, 6 (2018), 938–957.
- [25] Michael Y Kpibbaareh, Wei-Ping Wu, Brighter Agyemang, Charles R Haruna, and Tandoh Lawrence. 2022. A Generic Graph-Based Method for Flexible Aspect-Opinion Analysis of Complex Product Customer Feedback. *Information* 13, 3 (2022), 118.
- [26] KL Santhosh Kumar, Jayanti Desai, and Jharna Majumdar. 2016. Opinion mining and sentiment analysis on online customer review. In *2016 IEEE International Conference on Computational Intelligence and Computing Research (ICCIC)*. IEEE, 1–4.
- [27] Srijan Kumar, Xikun Zhang, and Jure Leskovec. 2019. Predicting dynamic embedding trajectory in temporal interaction networks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 1269–1278.
- [28] Jundong Li, Harsh Dani, Xia Hu, Jiliang Tang, Yi Chang, and Huan Liu. 2017. Attributed network embedding for learning in a dynamic environment. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 387–396.
- [29] Ning Li, Chi-Yin Chow, and Jia-Dong Zhang. 2020. EMOVA: A semi-supervised end-to-end moving-window attentive framework for aspect mining. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 811–823.
- [30] Taisong Li, Jiawei Zhang, S Yu Philip, Yan Zhang, and Yonghong Yan. 2018. Deep dynamic network embedding for link prediction. *IEEE Access* 6 (2018), 29219–29230.
- [31] Hao Liu, Shuwang Zhou, Changfang Chen, Tianlei Gao, Jiyong Xu, and Minglei Shu. 2022. Dynamic knowledge graph reasoning based on deep reinforcement learning. *Knowledge-Based Systems* 241 (2022), 108235.
- [32] Jingzhou Liu, Wei-Cheng Chang, Yuexin Wu, and Yiming Yang. 2017. Deep learning for extreme multi-label text classification. In *Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval*. 115–124.
- [33] Xi Liu, Ping-Chun Hsieh, Nick Duffield, Rui Chen, Muhe Xie, and Xidao Wen. 2019. Real-time streaming graph embedding through local actions. In *Companion Proceedings of The 2019 World Wide Web Conference*. 285–293.
- [34] Yao Liu, Cuiqing Jiang, and Huimin Zhao. 2019. Assessing product competitive advantages from the perspective of customers by mining user-generated content on social media. *Decision Support Systems* 123 (2019), 113079.
- [35] Guangnan Lu, Jiecheng Li, and Jian Wei. 2022. Aspect sentiment analysis with heterogeneous graph neural networks. *Information Processing & Management* 59, 4 (2022), 102953.
- [36] Yuanfu Lu, Xiao Wang, Chuan Shi, Philip S Yu, and Yanfang Ye. 2019. Temporal network embedding with micro-and macro-dynamics. In *Proceedings of the 28th ACM international conference on information and knowledge management*. 469–478.
- [37] Qianwen Ma, Chunyuan Yuan, Wei Zhou, and Songlin Hu. 2021. Label-specific dual graph neural network for multi-label text classification. In *Proceedings of the 59th Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 3855–3864.
- [38] Yao Ma, Ziyi Guo, Zhaocun Ren, Jiliang Tang, and Dawei Yin. 2020. Streaming graph neural networks. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 719–728.
- [39] Sedigheh Mahdavi, Shima Khoshraftar, and Aijun An. 2018. dynode2vec: Scalable dynamic network embedding. In *2018 IEEE international conference on big data (Big Data)*. IEEE, 3762–3765.
- [40] Yu Meng, Yunyi Zhang, Jiabin Huang, Chenyan Xiong, Heng Ji, Chao Zhang, and Jiawei Han. 2020. Text classification using label names only: A language model self-training approach. *arXiv preprint arXiv:2010.07245* (2020).
- [41] Shengjie Min, Zhan Gao, Jing Peng, Liang Wang, Ke Qin, and Bo Fang. 2021. STGSN—A Spatial-Temporal Graph Neural Network framework for time-evolving social networks. *Knowledge-Based Systems* 214 (2021), 106746.
- [42] Divya Mittal and Shiv Ratan Agrawal. 2022. Determining banking service attributes from online reviews: text mining and sentiment analysis. *International Journal of Bank Marketing* (2022).
- [43] Giang Hoang Nguyen, John Boaz Lee, Ryan A Rossi, Nesreen K Ahmed, Eunye Koh, and Sungchul Kim. 2018. Continuous-time dynamic network embeddings. In *Companion proceedings of the web conference 2018*. 969–976.
- [44] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35 (2022), 27730–27744.
- [45] Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao Scharidl, and Charles Leiserson. 2020. Evolvegcn: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 5363–5370.
- [46] Guang Qiu, Bing Liu, Jiajun Bu, and Chun Chen. 2011. Opinion word expansion and target extraction through double propagation. *Computational linguistics* 37,

- 1 (2011), 9–27.
- [47] Mahmudur Rahman, Tanay Kumar Saha, Mohammad Al Hasan, Kevin S Xu, and Chandan K Reddy. 2018. Dylink2vec: Effective feature representation for link prediction in dynamic networks. *arXiv preprint arXiv:1804.05755* (2018).
- [48] Jyoti Rana, Loveleen Gaur, and KC Santosh. 2022. Classifying Customers' Journey from Online Reviews of Amazon Fresh via Sentiment Analysis and Topic Modelling. In *2022 3rd International Conference on Computation, Automation and Knowledge Management (ICCAKM)*. IEEE, 1–6.
- [49] Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and Michael Bronstein. 2020. Temporal graph networks for deep learning on dynamic graphs. *arXiv preprint arXiv:2006.10637* (2020).
- [50] Sina Sajadmanesh, Sogol Bazargani, Jiawei Zhang, and Hamid R Rabiee. 2019. Continuous-time relationship prediction in dynamic heterogeneous information networks. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 13, 4 (2019), 1–31.
- [51] Aravind Sankar, Yanhong Wu, Liang Gou, Wei Zhang, and Hao Yang. 2018. Dynamic graph representation learning via self-attention networks. *arXiv preprint arXiv:1812.09430* (2018).
- [52] Simon Schaefer, Daniel Gehrig, and Davide Scaramuzza. 2022. AEGNN: Asynchronous Event-based Graph Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12371–12381.
- [53] Richa Sharma, Shweta Nigam, and Rekha Jain. 2014. Mining of product reviews at aspect level. *arXiv preprint arXiv:1406.3714* (2014).
- [54] Lei Shu, Hu Xu, Bing Liu, and Jiahua Chen. 2022. Zero-Shot Aspect-Based Sentiment Analysis. *arXiv preprint arXiv:2202.01924* (2022).
- [55] Uriel Singer, Ido Guy, and Kira Radinsky. 2019. Node embedding over temporal graphs. *arXiv preprint arXiv:1903.08889* (2019).
- [56] Kihyuk Sohn. 2016. Improved deep metric learning with multi-class n-pair loss objective. *Advances in neural information processing systems* 29 (2016).
- [57] Li Sun, Junda Ye, Hao Peng, and Philip S Yu. 2022. A Self-supervised Riemannian GNN with Time Varying Curvature for Temporal Graph Learning. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 1827–1836.
- [58] Sheng Tian, Ruofan Wu, Leilei Shi, Liang Zhu, and Tao Xiong. 2021. Self-supervised representation learning on dynamic graphs. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 1814–1823.
- [59] Rakshit Trivedi, Hanjun Dai, Yichen Wang, and Le Song. 2017. Know-evolve: Deep temporal reasoning for dynamic knowledge graphs. In *international conference on machine learning*. PMLR, 3462–3471.
- [60] Rakshit Trivedi, Mehrdad Farajtabar, Prasenjeet Biswal, and Hongyuan Zha. 2019. Dyrep: Learning representations over dynamic graphs. In *International conference on learning representations*.
- [61] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [62] Sanjeev Verma, Rohit Sharma, Subhamay Deb, and Debojit Maitra. 2021. Artificial intelligence in marketing: Systematic review and future research direction. *International Journal of Information Management Data Insights* 1, 1 (2021), 100002.
- [63] Jiyao Wang, Zijie Chen, Yang Qin, Dengbo He, and Fangzhen Lin. 2023. Multi-Aspect co-Attentional Collaborative Filtering for extreme multi-label text classification. *Knowledge-Based Systems* 260 (2023), 110110.
- [64] Qifan Wang, Li Yang, Bhargav Kanagal, Sumit Sanghai, D Sivakumar, Bin Shu, Zac Yu, and Jon Elsas. 2020. Learning to extract attribute value from product via question answering: A multi-task approach. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 47–55.
- [65] Xin Wang, Jiaxiu He, David J Curry, and Jun Hyun Ryoo. 2022. Attribute embedding: Learning hierarchical representations of product attributes from consumer reviews. *Journal of Marketing* 86, 6 (2022), 155–175.
- [66] Xiao Wang, Yuanfu Lu, Chuan Shi, Ruijia Wang, Peng Cui, and Shuai Mou. 2020. Dynamic heterogeneous information network embedding with meta-path based proximity. *IEEE Transactions on Knowledge and Data Engineering* (2020).
- [67] Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alexander J Smola, and How Jing. 2017. Recurrent recommender networks. In *Proceedings of the tenth ACM international conference on web search and data mining*. 495–503.
- [68] Lin Xiao, Xin Huang, Boli Chen, and Liping Jing. 2019. Label-specific document representation for multi-label text classification. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*. 466–475.
- [69] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. 2020. Inductive representation learning on temporal graphs. *arXiv preprint arXiv:2002.07962* (2020).
- [70] Huimin Xu, Wenting Wang, Xinnian Mao, Xinyu Jiang, and Man Lan. 2019. Scaling up open tagging from tens to thousands: Comprehension empowered attribute value extraction from product title. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 5214–5223.
- [71] Kaiquan Xu, Stephen Shaoyi Liao, Jiexun Li, and Yuxia Song. 2011. Mining comparative opinions from customer reviews for competitive intelligence. *Decision support systems* 50, 4 (2011), 743–754.
- [72] Mengjia Xu, Apoorva Vikram Singh, and George Em Karniadakis. 2022. DynG2G: An Efficient Stochastic Graph Embedding Method for Temporal Graphs. *IEEE Transactions on Neural Networks and Learning Systems* (2022).
- [73] Xueke Xu, Xueqi Cheng, Songbo Tan, Yue Liu, and Huawei Shen. 2013. Aspect-level opinion mining of online customer reviews. *China Communications* 10, 3 (2013), 25–41.
- [74] Guotong Xue, Ming Zhong, Jianxin Li, Jia Chen, Chengshuai Zhai, and Ruochen Kong. 2022. Dynamic network embedding survey. *Neurocomputing* 472 (2022), 212–223.
- [75] Hansheng Xue, Luwei Yang, Wen Jiang, Yi Wei, Yi Hu, and Yu Lin. 2021. Modeling dynamic heterogeneous network for link prediction using hierarchical attention with temporal rnn. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 282–298.
- [76] Jun Yan, Nasser Zalmout, Yan Liang, Christian Grant, Xiang Ren, and Xin Luna Dong. 2021. Adatag: Multi-attribute value extraction from product profiles with adaptive decoding. *arXiv preprint arXiv:2106.02318* (2021).
- [77] Li Yang, Qifan Wang, Zac Yu, Anand Kulkarni, Sumit Sanghai, Bin Shu, Jon Elsas, and Bhargav Kanagal. 2022. Mave: A product dataset for multi-source attribute value extraction. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*. 1256–1265.
- [78] Luwei Yang, Zhibo Xiao, Wen Jiang, Yi Wei, Yi Hu, and Hao Wang. 2020. Dynamic heterogeneous graph embedding using hierarchical attentions. In *European Conference on Information Retrieval*. Springer, 425–432.
- [79] Jisu Yi and Yun Kyung Oh. 2022. The informational value of multi-attribute online consumer reviews: a text mining approach. *Journal of Retailing and Consumer Services* 65 (2022), 102519.
- [80] Jiakuan You, Tianyu Du, and Jure Leskovec. 2022. ROLAND: graph learning framework for dynamic graphs. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2358–2366.
- [81] Wencao Yu, Wei Cheng, Charu C Aggarwal, Kai Zhang, Haifeng Chen, and Wei Wang. 2018. Netwalk: A flexible deep embedding approach for anomaly detection in dynamic networks. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 2672–2681.
- [82] Ling Zhan and Tao Jia. 2022. Coarsas2hvc: Heterogeneous information network embedding with balanced network sampling. *Entropy* 24, 2 (2022), 276.
- [83] Hanlei Zhang, Hua Xu, Ting-En Lin, and Rui Lyu. 2021. Discovering new intents with deep aligned clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 14365–14373.
- [84] Xinyang Zhang, Chenwei Zhang, Xian Li, Xin Luna Dong, Jingbo Shang, Christos Faloutsos, and Jiawei Han. 2022. OA-Mine: Open-World Attribute Mining for E-Commerce Products with Weak Supervision. In *Proceedings of the ACM Web Conference 2022*. 3153–3161.
- [85] Yao Zhang, Yun Xiong, Xiangnan Kong, and Yangyong Zhu. 2017. Learning node embeddings in interaction graphs. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 397–406.
- [86] Zhen Zhang, Jiajun Bu, Zhao Li, Chengwei Yao, Can Wang, and Jia Wu. 2021. TigeCMN: On exploration of temporal interaction graph embedding via Coupled Memory Neural Networks. *Neural Networks* 140 (2021), 13–26.
- [87] Yifeng Zhao, Xiangwei Wang, Hongxia Yang, Le Song, and Jie Tang. 2019. Large Scale Evolving Graphs with Burst Detection. In *IJCAI*. 4412–4418.
- [88] Yabing Zhao, Xun Xu, and Mingshu Wang. 2019. Predicting overall customer satisfaction: Big data evidence from hotel online textual reviews. *International Journal of Hospitality Management* 76 (2019), 111–121.
- [89] Li Zheng, Zhenpeng Li, Jian Li, Zhao Li, and Jun Gao. 2019. AddGraph: Anomaly Detection in Dynamic Graph Using Attention-based Temporal GCN. In *IJCAI*. 4419–4425.
- [90] Lekui Zhou, Yang Yang, Xiang Ren, Fei Wu, and Yueting Zhuang. 2018. Dynamic network embedding by modeling triadic closure process. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [91] Dingyuan Zhu, Peng Cui, Ziwei Zhang, Jian Pei, and Wenwu Zhu. 2018. High-order proximity preserved embedding for dynamic networks. *IEEE Transactions on Knowledge and Data Engineering* 30, 11 (2018), 2134–2144.
- [92] Linhong Zhu, Dong Guo, Junming Yin, Greg Ver Steeg, and Aram Galstyan. 2016. Scalable temporal latent space inference for link prediction in dynamic social networks. *IEEE Transactions on Knowledge and Data Engineering* 28, 10 (2016), 2765–2777.
- [93] Tiangang Zhu, Yue Wang, Haoran Li, Youzheng Wu, Xiaodong He, and Bowen Zhou. 2020. Multimodal joint attribute prediction and value extraction for e-commerce product. *arXiv preprint arXiv:2009.07162* (2020).
- [94] Yu Zhu, Hao Li, Yikang Liao, Beidou Wang, Ziyu Guan, Haifeng Liu, and Deng Cai. 2017. What to Do Next: Modeling User Behaviors by Time-LSTM. In *IJCAI*, Vol. 17. 3602–3608.
- [95] Yuan Zuo, Guannan Liu, Hao Lin, Jia Guo, Xiaoqian Hu, and Junjie Wu. 2018. Embedding temporal network via neighborhood formation. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 2857–2866.