

---

# Uncovering Customer Issues through Topological Natural Language Analysis

---

**Shu-Ting Pi**  
Amazon  
Cupertino, CA 95014  
shutingp@amazon.com

**Sidarth Srinivasan \***  
Amazon  
Seattle, WA 98109  
sidarthsri@g.ucla.edu

**Yuying Zhu**  
Amazon  
Seattle, WA 98109  
imyuying@amazon.com

**Michael Yang**  
Amazon  
Seattle, WA 98109  
abyang@amazon.com

**Qun Liu**  
Amazon  
Seattle, WA 98109  
qunliu@amazon.com

## Abstract

E-commerce companies deal with a high volume of customer service requests daily. While a simple annotation system is often used to summarize the topics of customer contacts, thoroughly exploring each specific issue can be challenging. This presents a critical concern, especially during an emerging outbreak where companies must quickly identify and address specific issues. To tackle this challenge, we propose a novel machine learning algorithm that leverages natural language techniques and topological data analysis to monitor emerging and trending customer issues. Our approach involves an end-to-end deep learning framework that simultaneously tags the primary question sentence of each customer’s transcript and generates sentence embedding vectors. We then whiten the embedding vectors and use them to construct an undirected graph. From there, we define trending and emerging issues based on the topological properties of each transcript. We have validated our results through various methods and found that they are highly consistent with news sources.

## 1 Introduction

E-commerce websites handle a vast number of online customer service requests daily. During a typical online customer service interaction, customers first interact with a chatbot which asks them questions to identify their intent. This intent is usually classified based on the product or service that the customer needs assistance with. For instance, an online consumer electronics retailer might use its chatbot to classify requests as relating to cell phones, computers, or home appliances, among others. The chatbot then routes the customer to an agent who specializes in the requested product or service to assist. While the actual business practices among companies may differ, the interaction process between customers and agents is generally similar. Agents usually begin with a greeting and ask for details about the customer’s questions. They then engage in diagnosis and finally conclude with some closing remarks.

Generally, processing customer requests can take several minutes, making it one of the most time-consuming aspects of e-commerce business. Therefore, developing a standardized process to handle specific issues is critical to help customers save considerable time and optimize the available resources of agents. This is especially important during emerging events or sudden surges in cus-

---

\*The work was carried out while SS was an intern at Amazon.

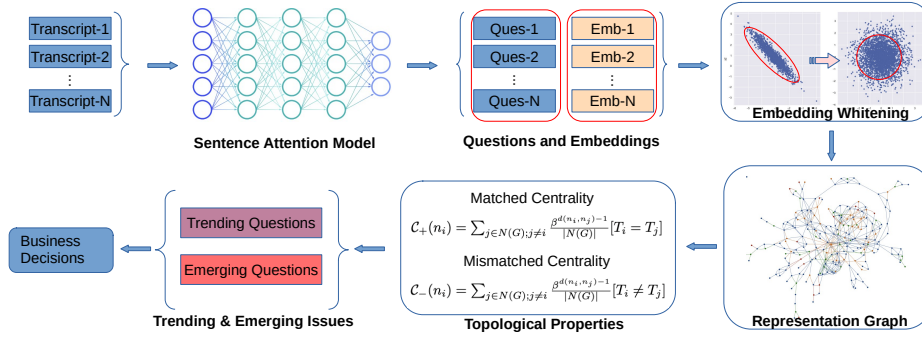


Figure 1: Our proposed workflow involves several steps. Initially, the transcripts are passed to a sentence attention model to extract the primary questions asked by the customers and their corresponding sentence embeddings. The embeddings are then whitened to obtain representations in an isotropic coordinate system. These whitened vectors are then utilized to construct an undirected graph, and their topological properties are calculated to identify both trending and emerging issues.

customer inquiries. By anticipating common issues and developing standardized procedures for agents, businesses can improve their response times, reduce customer frustration, and ultimately enhance customer satisfaction.

We present a novel machine learning framework, as illustrated in Fig. 1, that can detect emerging and trending issues without predefined lists. **The terms "trending" and "emerging" refer to the topics that are most frequently discussed within the current time window and the topics that show the most rapid increase in discussion compared to the previous time window.** Our approach comprises three distinct components. Firstly, a deep learning model is employed, which utilizes an attention mechanism to automatically tag the primary question sentence in each customer’s transcript and generates sentence-level embedding vectors. To improve the performance of cosine similarity, we decouple the covariance matrix to whiten the embedding vectors, bringing the coordinates of the feature space close to an orthonormal basis. We then construct an undirected graph based on cosine similarity. Finally, we analyze the topology of the graph by calculating the centrality of each customer’s question. This enables us to quantify both trending and emerging issues.

The article is structured into five sections. Section 2 details the development of a sentence tagging model that automatically identifies the main question in a contact transcript and generates its corresponding sentence embedding vector. In Section 3, we explain how to whiten the embedding vectors to make cosine distance a better metric and how to use them to construct an undirected graph. Section 4 demonstrates how to define trending and emerging issues based on the topological properties of each transcript and validates the results using a variety of methods. Finally, in Section 5, we present our conclusions.

**Related Works** Our work relates to several areas in the literature, including sentence-level attention, sentence tagging, and graph-based clustering. Works related to sentence attention include (Yang et al., 2016), (Nallapati et al., 2016), and (Lin et al., 2016), which apply the idea to document classification, summarization, and text noise reduction, respectively. Regarding sentence tagging, (Collobert and Weston, 2008) and (Santos and Zadrozny, 2014) are works that perform sentence tagging from token-level representation, and have also influenced our approach. Additionally, we reference several works in graph-based clustering, including (Wang et al., 2019) and (Novák et al., 2010), which use this technique to handle repetitive sequences and multiview data. These works have inspired our research in the application of topological data analysis to natural language processing.

## 2 Question Tagging and Sentence Attention Model

We present a deep learning model that can automatically tag the primary question in a contact transcript. This model is a crucial component for detecting both trending and emerging issues, as the identified questions will be utilized in later sections.

## 2.1 The Dataset

There has been a lack of publicly available datasets related to customer service transcripts. To address this gap, we partnered with customer service team to initiate this research, using a dataset from MessageUs, an online chat system that enables customers to communicate with customer service agents. We collected over 500,000 contact transcripts during 2022, recording conversations between customers and agents. Each contact also comes with a unique label of the product or service that the customer and agents discussed. It’s worth noting that the dataset only contains customer text data, with all confidential information, such as names and account details, anonymized to protect privacy before being shared with researchers. Although the dataset is from a specific database, the methodology presented in this article can be applied to other use cases as well.

## 2.2 Sentence Embedding

Our goal is to identify the primary question sentence in a customer-agent contact transcript. We propose two hypotheses: (1) the primary question sentence typically appears in the first few sentences of the customer’s interaction with the agent, and (2) it contains the most relevant information about the product or service being discussed. If these hypotheses hold, we can treat the problem as a machine learning task: **identifying the customer sentences near the agent’s initial response that are most useful in predicting the product or service of the contact for a machine learning classifier**. To achieve this, we need information on attention weights at the sentence level.

We propose a deep learning model, as shown in Fig. 2, to achieve our goal. Unlike traditional text classification models that represent an article as a 2D tensor  $\in R^{\mathbb{N}_{at} \times \mathbb{N}_{we}}$ , where  $\mathbb{N}_{at}$  is the number of tokens in the article and  $\mathbb{N}_{we}$  is the dimension of the word embedding, our approach represents each article as a 3D tensor  $\in R^{\mathbb{N}_{as} \times \mathbb{N}_{st} \times \mathbb{N}_{we}}$ . Here,  $\mathbb{N}_{as}$  is the number of sentences in the article, and  $\mathbb{N}_{st}$  is the number of tokens per sentence. To accommodate varying numbers of sentences and tokens per sentence in each transcript, we use zero-padding to ensure a consistent tensor shape for subsequent processing.

To obtain sentence-level embeddings, we treat each sentence as a temporal slice and apply a time-distributed wrapper to a sequence model  $\Sigma$ , such as BERT or LSTM. This ensures that the model receives only one sentence per time step, allowing us to embed each sentence. The resulting output tensor,  $Q' \in R^{\mathbb{N}_{as} \times \mathbb{N}_{se}}$ , contains  $\mathbb{N}_{as}$  sentence vectors, each with  $\mathbb{N}_{se}$  dimensions.

Our “bag of sentences” model currently does not consider sentence positions, but we have observed that customer questions tend to appear in early sentences during interactions with agents. This suggests that sentence positions can impact attention weights, so we incorporate sentence position information into the model.

To do this, we adapt the idea of position embedding used in many language models for tokens, but apply it to sentences. We assign each sentence an index, ranging from  $-\mathbb{N}_{as}$  to  $+\mathbb{N}_{as}$ , representing the number of sentences between the current sentence and the agent’s first response. For instance, an index of -5 indicates that the sentence is five steps before the agent’s first sentence, while +5 indicates five steps after. We shift the indices by  $\mathbb{N}_{as}$ , resulting in an allowed index range of 0 to  $+2\mathbb{N}_{as}$ , with the sentence having an index of  $\mathbb{N}_{as}$  being the agent’s first sentence to avoid negative indices. For a sentence with index  $i$ , the  $p$ -th component of the position embedding vector  $E_i$  is given by  $E_i(2p)$  and  $E_i(2p + 1)$ :

$$E_i(2p) = \sin(i/10000^{2p/d_{pos}}) \tag{1}$$

$$E_i(2p + 1) = \cos(i/10000^{2p/d_{pos}}) \tag{2}$$

where  $d_{pos}$  is the dimension of embedding vector. By adding  $Q'_i$  and  $E_i$ , we get the final sentence embedding vector  $Q_i = Q'_i + E_i$ .

## 2.3 Sentention Attention

We define sentence level attention(Vaswani et al., 2017):

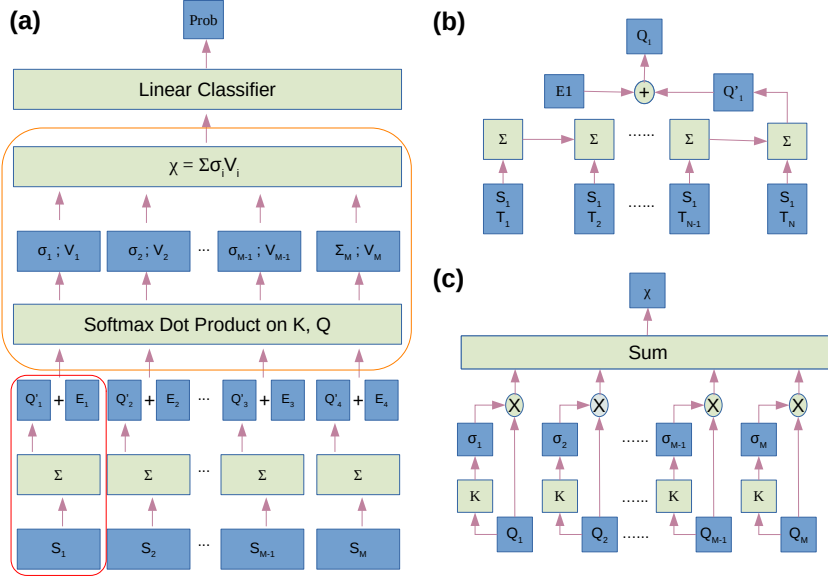


Figure 2: The Sentence Attention Model. The model consists of blue blocks, representing tensors, and green blocks representing operators. (a) The neural network is comprised of sentence tensors,  $S_i$ , and a sequence model,  $\Sigma$ , which outputs  $Q'_i$ . The position embedding vector  $E_i$  is combined with  $Q'_i$  to create the sentence embeddings  $Q_i$ . Finally, a linear classifier predicts the product/service. (b) The red block in (a) is described in detail. Tensor notation  $(S_1; T_n)$  refers to the  $n$ -th token in sentence  $S_1$ . The sequence model  $\Sigma$  processes each word in a sentence using a time-distributed wrapper to handle multiple sentences. (c) The orange block in (a) is explained. A dense layer,  $K$ , with a softmax activation function is applied to all sentence embedding vectors  $Q_i$  (via a time-distributed wrapper) to calculate attention scores  $\sigma_i$ . Note that  $Q_i$  is equivalent to  $V_i$ .

$$Attention(\vec{K}, \hat{Q}, \hat{V}) = \vec{\chi} = softmax\left(\frac{\hat{Q} \cdot \vec{K}}{\sqrt{d_k}}\right)^T \cdot \hat{V} = \vec{\sigma} \cdot \hat{V} \quad (3)$$

, where  $\hat{Q} \in R^{N_{as} \times N_{se}}$  is a query tensor (i.e. the sentence embedding),  $\vec{K} \in R^{N_{se}}$  is a key vector and  $\hat{V} \in R^{N_{as} \times N_{se}}$  is a value tensor. We can generally define  $\hat{Q} = \hat{V}$  without any impact on model performance. As a result, the attention value  $\vec{\chi} \in R^{N_{se}}$  is a vector and the sum of its elements must equal 1, due to the application of the softmax function.

Our model's attention vector  $\vec{\chi}$  has a simple interpretation. As shown in Fig. 2, each sentence in a transcript is represented as a vector  $\vec{V}_i$  (the  $i$ -th row of the 2D tensor  $\hat{V}$ ). The attention vector  $\vec{\chi}$  is a linear combination of all sentences, computed as  $\vec{\chi} = \sum_i \sigma_i \vec{V}_i$ , subject to the constraint that  $\sum_i \sigma_i = 1$ . We then pass  $\vec{\chi}$  through a fully connected layer with a softmax activation function to predict the product or service associated with the transcript. The attention weights  $\sigma_i$  reflect the importance of each sentence in determining the product or service, with higher weights assigned to sentences containing more critical information.

## 2.4 Experiments

Our model was trained on 500,000 transcripts with 152 classes to predict the product or service discussed in each contact. We used DistilBERT (Sanh et al., 2020) as the embedding model  $\Sigma$ , with an output dimension of 768. To prepare the transcripts for training, we padded each one with zeros to create 64 sentences, each with 128 words.

After training, we calculated attention weights  $\sigma_i$  for each customer sentence in a transcript. To identify the primary question, we assumed that it is the sentence with the highest attention weight that is

$N$  steps before or after the agent’s first sentence. In testing on 4,000 human-annotated transcripts, we found that  $N = 2$  (i.e.,  $\pm 2$  steps) yielded the best results, correctly identifying the primary question sentence in 84.3% of the total transcripts. Fig. 1 in the Appendix illustrates an example of how our sentence attention model tags the primary question sentence.

To evaluate the impact of sentence position embedding on performance, we compared the model with and without this feature. Although the difference in accuracy was minimal (83.4%), we observed that sentence position embedding resulted in higher attention weights on the primary question sentence. Further investigation is needed to explore the effects of position embedding.

### 3 Topological Natural Language Analysis

Our objective is to identify both emerging and trending topics among the questions gathered by the model presented in Sec. 2. However, conventional clustering methods face several challenges in achieving this goal, including sensitivity to hyperparameters, scalability issues with large numbers of classes and samples, and lack of flexibility in defining distances.

Moreover, identifying emerging topics involves changes in the volume of a topic between time windows, and conventional clustering methods cannot determine whether two clusters in different datasets are related to the same topics. It is even possible for an emerging topic to be present in the current time window without appearing in the previous time window. Additionally, conventional clustering methods may treat emerging topics as noise due to their much smaller volume than trending topics.

To overcome these challenges, we propose a topological-analysis-based method robust to hyperparameter selection and can quantitatively detect both emerging and trending topics between different time windows.

#### 3.1 Sentence Embedding Whitening

Assuming we have gathered a substantial number of customer’s primary questions and their associated embedding vectors  $Q'_t{}^i$ , where  $Q'$  represents the output embedding vectors of  $\Sigma$  without the addition of position embedding,  $i = 1 \sim N$  corresponds to the  $i$ -th transcript, and  $t$  signifies the tagged sentence. Provided that we can define a metric or distance to express the semantic similarity between two questions, it becomes possible to construct a graph in which an edge connects two questions if their cosine distance falls below a specified threshold.

However, cosine distance may not be an appropriate metric for measuring text similarity. In fact, several experiments have shown that cosine similarity is not suitable for use with BERT-based representations and its performance in many similarity tasks is inferior to that of traditional embedding methods (Reimers and Gurevych, 2019) such as GloVe (Pennington et al., 2014) or Word2Vec (Mikolov et al., 2013). This is mainly due to the fact that cosine distance assumes an orthonormal coordinate system, which is not the case for most pre-trained sequence models. Various approaches have been proposed to address this issue, such as BERT-flow (Li et al., 2020) or representation whitening (Jianlin Su, 2021), to ensure that data distributions tend to be isotropic, a property that an independent basis set should have. Representation whitening, a post-processing method that does not require model training, is particularly effective in calibrating sentence representations. Hence, we adopt this strategy to improve the performance of our sentence representation.

The core idea of representation whitening is to find a linear transformation that sets the mean of the set  $\{Q'_t{}^i\}_{i=1 \sim N}$  to zero and the covariance matrix to an identity matrix. To achieve this, we first calculate the mean  $\mu$  and covariance matrix  $K$ . As we aim for a diagonalized covariance matrix, we can compute the unitary matrix  $U$  and singular values matrix  $A$  using singular value decomposition (SVD) (Halko et al., 2009):  $U, A, U^\dagger = SVD(K)$ . In this case, the whitening matrix  $W$  that diagonalizes the covariance matrix  $K$  is:

$$W^\dagger K W = I \quad ; \quad W = U \sqrt{A}^{-1}$$

Once we obtain the whitening matrix  $W$ , the whitened sentence embedding vectors become:

$$z_i = (Q'_t{}^i - \mu)W$$

Under this transformation, the mean of  $\{z_i, i = 1 \sim N\}$  equals zero, and the covariance matrix being equal to an identity matrix  $I$  is guaranteed. Although we lack labels to evaluate the performance of whitened sentence embedding vectors in capturing similarity, we manually inspected a small portion of the data with cosine similarity. Whitened vectors  $z_i$  indeed appear more reasonable than the original ones. From this point, we will use  $z_i$  as the embedding vector of the tagged question sentence in the  $i$ -th transcript unless specified otherwise.

### 3.2 Undirected Representation Graph and Centrality

Let us assume that we have gathered primary customer questions and their corresponding sentence embeddings,  $z_i$ , over two time periods,  $T_0$  and  $T_1$ . Each sentence can be represented as a node in an undirected graph, with edges connecting nodes whose cosine similarity surpasses a specified threshold,  $\alpha$ . It is important to note that this graph is constructed using data from both time periods,  $T_0$  and  $T_1$ . In this graph, a node with a higher number of neighbors indicates a greater number of similar questions.

Centrality is a significant topological property of a node, which describes its importance within a graph. Various types of centrality exist, each applicable to different scenarios. In this case, we define two new types of centrality that modify the decay centrality. Assuming we have an undirected graph represented by an adjacency matrix  $A$ , where  $A(i, j) = 1$  if nodes  $\cos(z_i, z_j) \geq \alpha$ , the decay centrality of a node  $n_i$  in a graph  $G$  is defined as (van Steen, 2010):

$$\mathcal{C}(n_i) = \sum_{j \in N(G); j \neq i} \frac{\beta^{d(n_i, n_j) - 1}}{|N(G)|}$$

Here,  $N(G)$  refers to the set of nodes in the graph, and  $|N(G)|$  is the total number of nodes in the graph. This normalization factor ensures that the centrality  $\mathcal{C}(n_i)$  is independent of the graph's size. The attenuation factor  $\beta$  is typically selected such that  $0 < \beta < 1$ . The topological distance  $d$  between nodes  $i$  and  $j$  is the graph distance, not the Euclidean distance  $|z_i - z_j|$ . The numerator is given by an exponent  $d(n_i, n_j) - 1$ , which results in  $\beta^{d(n_i, n_j) - 1} = 1$  when nodes  $i$  and  $j$  are directly connected. A question with higher decay centrality means the graph has more similar questions.

Given that time is an additional property of each node, we modify the definition above and propose two new types of centrality:

- matched decay centrality:

$$\mathcal{C}_+(n_i) = \sum_{j \in N(G); j \neq i} \frac{\beta^{d(n_i, n_j) - 1}}{|N(G)|} [W_i = W_j]$$

- mismatched decay centrality:

$$\mathcal{C}_-(n_i) = \sum_{j \in N(G); j \neq i} \frac{\beta^{d(n_i, n_j) - 1}}{|N(G)|} [W_i \neq W_j]$$

In these equations, the brackets  $[\cdot \cdot \cdot]$  represent Iverson brackets, yielding 1 if the statement inside is true and 0 if false.  $W_i$  refers to the time window that node  $i$  belongs to.  $\mathcal{C}_+$  considers contributions only when two nodes belong to the same time window, while  $\mathcal{C}_-$  accounts for contributions from different windows, decaying exponentially as the distance increases. It is important to note that node  $j$  can reach node  $i$  through intermediate nodes in either the same or different time windows. Furthermore, we have  $\mathcal{C}(n_i) = \mathcal{C}_+(n_i) + \mathcal{C}_-(n_i)$ . This definition splits the decay centrality into two terms based on the time window. We can effectively identify trending and emerging issues by utilizing matched and mismatched decay centrality.

## 4 Trending and emerging Issues

### 4.1 Topological Perspective

We are only interested in the nodes in the current time window, so we define trending and emerging issues from a topological perspective as follows:




Figure 3: The centrality distribution of Tablet in Feb 2023 is depicted in panels (a)-(f), where we compare the cosine similarity thresholds of  $\alpha = 0.8$  and  $\alpha = 0.6$ . As for panel (g), the graph built using a few selected samples clearly demonstrates that similar sentences tend to cluster together, resulting in high centrality for nodes around the cluster centers.

- Trending issues are identified as the customer’s primary questions with a large trending score:  $S_t = C_+$ .
- Emerging issues are identified as the customer’s primary questions with a large emerging score:  $S_e = \tanh^2(C_+/\gamma)[C_+ - C_-]/C$ .

The definitions for the trending and emerging scores are simple to understand. The trending score, denoted by  $S_t$ , counts the number of similar questions in the graph that exist in the current time window. On the other hand, the emerging score, given by  $[C_+ - C_-]/C$ , represents the difference in centrality contributed from the previous time window and the current time window. It’s important to note that  $C_+ + C_- = C$ . As a result, the emerging score lies between -1 and +1, where -1 and +1 correspond to the centrality entirely coming from the previous and current time window respectively, while 0 indicates equal contributions from both time windows. To avoid identifying small-size clusters that may not be significant for a business that receives a large volume of customer contacts, we introduce an additional filtering factor,  $\tanh^2(C_+/\gamma)$ , where  $\gamma$  is a parameter that controls the strength of the filter. This factor applies a weight to the emerging score, such that  $\tanh^2(C_+/\gamma)$  drops to 0 when  $C_+ \ll \gamma$ , and saturates to 1 when  $C_+ \gg \gamma$ .

## 4.2 Experiment

We constructed a graph using the MessageUs transcript data by applying a cosine similarity threshold of  $\alpha = 0.7$ , an attenuation factor of  $\beta = 0.5$ , and a filter factor of  $\gamma = 0.1 \times \text{Max}\{\mathcal{C}_+^i; i = 1 \sim N\}$ , where  $N$  is the total number of nodes. This filter factor was set to 10% of the maximum number of  $C_+$  values, and it helped to exclude small clusters that were unlikely to be significant. We used topological data analysis to identify the most prominent clusters, and we found that the topics of these clusters were highly consistent, even though the exact values of  $C_+$  and  $C_-$  were sensitive to the chosen hyperparameters. Varying the hyperparameters affected only the sizes of the clusters and did not significantly alter the topics we discovered, which demonstrates the robustness of the topology analysis approach.

Fig. 3 shows the distribution of  $C_+$ ,  $C_-$ , and  $S_e$  for Tablet-related transcripts, using Jan 2023 and Feb 2023 as the previous and current time windows, respectively (around 100K data points). We compare the distributions for a cosine similarity threshold of  $\alpha = 0.8$  (panels (a)-(c)) and  $\alpha = 0.6$  (panels (d)-(f)). Decreasing  $\alpha$  results in a more compact graph and a more Gaussian-like distribution of centrality. We observe that the emerging score distribution can be separated into two parts: a Gaussian-like distribution around zero due to the filter factor applied to nodes with  $C_+ \ll \gamma$ , and a long-tail region due to nodes with  $C_+ \gg \gamma$ . The filter factor helps to focus on the outliers, and the nodes in the long-tail region are generally similar, with high-score nodes unlikely to become low-score nodes due to changes in hyperparameters.

It is important to note that similar sentences tend to form clusters. As a result, the neighbors of a high centrality node also tend to have high centrality, as shown in Fig. 3(g). To avoid locating the

same cluster multiple times, it is necessary to ensure that two centers are sufficiently far apart. To achieve this, we first designate the node with the largest  $S_t$  (for trending) or  $S_e$  (for emerging) as the cluster center and its surrounding neighbors with graph distance  $\leq 3$  as the cluster members. Next, we search for the node with the next-largest score at least a graph distance of 4 away from any known clusters as the next cluster center. We repeat this procedure until we have obtained the desired number of clusters.

Table.1 of the appendix provides examples of sentences from the top-1 trending and emerging clusters of Tablet in February 2023. As one can see, all the customer’s questions within each cluster are very similar, demonstrating the effectiveness of our approach in cosine similarity and topology-based topic detection.

Since no quantitative metrics are available to directly validate our model’s performance in detecting trending and emerging issues, we rely on our intuition to assess its effectiveness. With regards to trending issues, we are confident that our model is accurate, as we have checked the trending issues from November 2022 to February 2023 across more than 10 products such as Kindle, Echo, Music, eBook, and Prime Video. We found that all the top trending topics are very stable and consistent with our business intuition. The most common issues were related to returns, refunds, or canceling subscriptions, which accounted for the majority of the volume.

To validate the emerging issues, we collected the top three emerging topics for over 10 product lines between November 2022 and February 2023, excluding some clusters with a relatively small  $S_e$  or unclear topics, resulting in 62 emerging issues. We validated whether a topic is emerging by manually selecting a few keywords that best represented the topic and comparing their term frequency changes between the two months to check for a significant change or by checking if we could find the issue from online news sources or related discussion forums. Based on our survey, we found that about 70% of the identified topics could easily be classified as emerging. During our validation, we found that many emerging topics were well-covered in online news sources. In Table 2 of the appendix, we provide a few examples of emerging topics found in Feb 2023 and our validation results, and most of the topics align well with the news sources.

## 5 Conclusion

In summary, we have presented a unique machine learning framework for extracting customers’ trending and emerging issues. Our work starts with an attention-based deep learning model that tags customers’ primary questions and generates corresponding sentence embeddings simultaneously. We then transform the sentence embeddings into an isotropic coordinate system using whitening techniques to improve the cosine similarity performance. Finally, we apply topological natural language analysis methods to analyze the centrality of each question, enabling us to identify trending and emerging issues.

Our work makes a significant contribution by demonstrating the application of a sentence-level attention mechanism in conversational transcripts, an area that has been understudied. We combine this mechanism with topological data analysis to extract useful information for a real-world problem. In the future, we aim to further explore topological properties of natural language understanding and their applications in e-commerce.

## References

- Collobert, R. and Weston, J. (2008). A unified architecture for natural language processing: deep neural networks with multitask learning. *ICML '08: Proceedings of the 25th international conference on Machine learning*, 81:160–167.
- Halko, N., Martinsson, P.-G., and Tropp, J. A. (2009). Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *arXiv*, page 0909.4061.
- Jianlin Su, Jiarun Cao, W. L. Y. O. (2021). Whitening sentence representations for better semantics and faster retrieval. *arXiv*, page 2103.15316.
- Li, B., Zhou, H., Junxian He, M. W., Yang, Y., and Li, L. (2020). On the sentence embeddings from pre-trained language models. *arXiv*, page 2011.05864.

- Lin, Y., Shen, S., Liu, Z., Luan, H., and Sun, M. (2016). Neural relation extraction with selective attention over instances. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 1:2124–2133.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv*, page 1301.3781.
- Nallapati, R., Zhou, B., dos santos, C. N., and Caglar Gulcehre, B. X. (2016). Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv*, page 1602.06023.
- Novák, P., Neumann, P., and Macas, J. (2010). Graph-based clustering and characterization of repetitive sequences in next-generation sequencing data. *BMC Bioinformatics*, 11:378.
- Pennington, J., Socher, R., and Manning, C. (2014). Glove: Global vectors for word representation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, page 1532–1543.
- Reimers, N. and Gurevych, I. (2019). Sentencebert: Sentence embeddings using siamese bert-networks. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, page 3982–3992.
- Sanh, V., Debut, L., and Julien Chaumond, T. W. (2020). Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv*, (1910):01108.
- Santos, C. D. and Zadrozny, B. (2014). Learning character-level representations for part-of-speech tagging. *Proceedings of the 31st International Conference on Machine Learning, PMLR*, 32:1818–1826.
- van Steen, M. (2010). An introduction to graph theory and complex networks.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, H., Yang, Y., and Liu, B. (2019). Gmc: Graph-based multi-view clustering. *IEEE Transactions on Knowledge and Data Engineering*, 32:1116–1129.
- Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., and Hovy, E. (2016). Hierarchical attention networks for document classification. *Proceedings of NAACL-HLT 2016*, pages 1480–1489.

## A Appendix: Example of Transcript

Speaker	Attention Weight	Text
Bot	0.000	Hi thanks for contacting customer service. I'm service chat helper. Type a little about what you'd like help with so I can get you to the right place.
Customer	0.179	I would like to know when was I billed last month for my monthly subscription and what was the total price?
Bot	0.000	Sure I will get someone who can help you with that.
Customer	0.029	Ok thanks !
Bot	0.000	Ok, I'll get someone to help you here through chat.
Customer	0.027	Ok thanks
Bot	0.000	If you have details you think would help the associate. Type them here. An associate will join the chat.
Agent	0.000	Hello my name is David. I'll be glad to help you today.
Customer	0.253	Hi I would like to know how much was my monthly renewal plan?
Agent	0.000	I see your membership is active and it's due on 10 <sup>th</sup> next month.
Customer	0.200	I saw that I will be billed 5 days ago for my renewal. The charge was \$99 and I was not sure what it is.
Agent	0.200	Please allow me a minute to check that.
Customer	0.020	Ok thanks

Figure 1: An illustration of the functioning of the question tagging model. The model calculates the attention score,  $\sigma_i$ , for each sentence in a transcript. The sentence preceding and following the agent's first sentence, with the highest score is predicted as the customer's primary question, i.e. the highlighted sentence. Please note that the sentences provided are not derived from real data but were formulated by our researchers.

## B Appendix: Trending & Emerging Issues on Tablet

Trending Issues	Emerging Issues
I need to utilize the warranty for my children's tablet due to a cracked screen.	Whenever I access my kids' profile on the tablet, it consistently displays an error message. However, switching to the parent profile resolves the issue.
I want to check the warranty for my son's cracked screen on his tablet.	The tablet used by my children encounters an error message, exclusively on the kids' profile, while the adult profiles function properly.
I own a tablet designed for kids, and my son has cracked its screen. I require a replacement through the attached warranty.	While accessing my child's profile, an error message appears, but I can switch to the adult profile without any issues.
I need to ascertain the warranty status of my tablet.	The tablet fails to load the kids' profile, even after attempting a reset and factory reset. The "oops something went wrong" message persists, while the adult profile functions normally.
My daughter has damaged her kids' tablet screen. Can we file a claim to have it repaired?	When I attempt to access my child's page on the tablet, my information loads, but Kid's profile continually displays the "oops something went wrong" message.

Table 1: Our examination of the Tablet's top-1 trending and emerging clusters within the chosen timeframe revealed that customers frequently reported cracked screens as a prevalent issue. As a result, they sought warranty replacements or repairs. Additionally, an emerging issue was identified concerning problems with the kids' profile. These findings indicate the efficacy of our methodology, as the sentences within the trending and emerging topics extracted from the same cluster exhibit notable similarities. Please note that the sentences provided are not derived from real data but were formulated by our researchers.

## C Appendix: Selected Emerging Issues on Various Products

Product	Cluster Center Sentence	Topic Summarization	Validation
Tablet	an error message, stating "oops something went wrong," is encountered when attempting to access my child's profile.	Customers have reported encountering an error message when attempting to enter child's profile.	We found this issue was actively discussed in a popular electronic device forum during the specific timeframe.
Music	I experienced an error code that prevented me from playing any music from any playlist.	Customers have reported encountering an exception error with a specific code while playing music.	We compared the probability of finding the keywords "exception/play/error" together in a transcript and observed a 30% increase from the previous month.
Game & Software	I am extremely disappointed with the significant delay in receiving my preordered Hogwarts Legacy. Despite the expected arrival date being today, I have been informed that it will not be shipped until the 23rd. This level of delay is highly unsatisfactory and unacceptable.	Customers have reported that their preordered game, Hogwarts Legacy, has been delayed.	We confirmed that the official release date of Hogwarts Legacy was about the same time period, which explains why the related issue emerged.
Live Games	I'm curious if subscribing to MLB TV can access to both MLB Network shows and games, or only the games?	Customers have raised questions about watching MLB games.	We confirmed that the 2023 MLB season commenced in late February, which explains why the related issues emerged during that month.
TV Stick	The screen displays a message stating "system recovery: your device will restart in a few minutes and should resume normal operation. If it fails to restart..."	Customers have reported that their devices received updates but failed to restart properly.	We checked a popular device forum and confirmed that this issue has been widely discussed.

Table 2: The table presents a selection of emerging issues identified in selected timeframe, categorized by product line. The first column displays the corresponding question sentences that represent the cluster center node. In the second column, we summarize the topic of each cluster. The third column shows the evidence that confirms these issues as emerging. Please note that the sentences provided are not derived from real data but were formulated by our researchers.