

Identifying Shopping Intent in Product QA for Proactive Recommendations

Besnik Fetahu, Nachshon Cohen, Elad Haramaty, Liane Lewin-Eytan, Oleg Rokhlenko, and Shervin Malmasi

Amazon.com, Inc. Seattle, WA, USA

{besnikf,nachshon,eladh,lliane,olegro,malmasi}@amazon.com

Abstract. Voice assistants have become ubiquitous in smart devices allowing users to instantly access information via voice questions. While extensive research has been conducted in question answering for voice search, little attention has been paid on how to enable proactive recommendations from a voice assistant to its users. This is a highly challenging problem that often leads to user friction, mainly due to recommendations provided to the users at the wrong time. We focus on the domain of e-commerce, namely in identifying *Shopping Product Questions* (SPQs), where the user asking a product-related question may have an underlying shopping need. Identifying a user’s shopping need allows voice assistants to enhance shopping experience by determining *when* to provide recommendations, such as *product* or *deal* recommendations, or *proactive shopping actions recommendation*. Identifying SPQs is a challenging problem and cannot be done from question text alone, and thus requires to infer latent user behavior patterns inferred from user’s past shopping history. We propose features that capture the user’s latent shopping behavior from their purchase history, and combine them using a novel Mixture-of-Experts (MoE) model. Our evaluation shows that the proposed approach is able to identify SPQs with a high score of $F1=0.91$. Furthermore, based on an *online* evaluation with real voice assistant users, we identify SPQs in real-time and recommend shopping actions to users to add the queried product into their shopping list. We demonstrate that we are able to accurately identify SPQs, as indicated by the significantly higher rate of added products to users’ shopping lists when being prompted after SPQs vs random PQs.

1 Introduction

Voice assistants, like Alexa or Google Assistant provide ubiquitous services through a variety of devices (e.g. smart speakers, phones, TVs etc.). Users interact with voice-assistants for different purposes [42, 33] such as question answering [11], task completion [9], conversational search [43], entertainment, or control of smart devices. Determining the underlying *user intent* is an active field of research [47, 26, 53], given that new skills are continuously added to such voice assistants.

In this context, making proactive follow-on suggestions is an active area of research [12]. However, key challenges with voice-based conversational recommender systems remain, such as their failure to adapt to evolving user behavior [37]. Additionally, interactions are typically initiated by users, and proactive system recommendations may lead

to increased user friction. Knowing *when* to proactively recommend actions or items to users on the next turn, such as suggesting the right product recommendations [8, 34, 36, 19, 28, 31] or next actions, is tightly dependent on accurately identifying the user’s underlying intent. Correctly identifying this intent can avoid user dissatisfaction by providing recommendations only when necessary. Our work focuses on identifying the right time to make a proactive personalized recommendation in a voice-based conversational system.

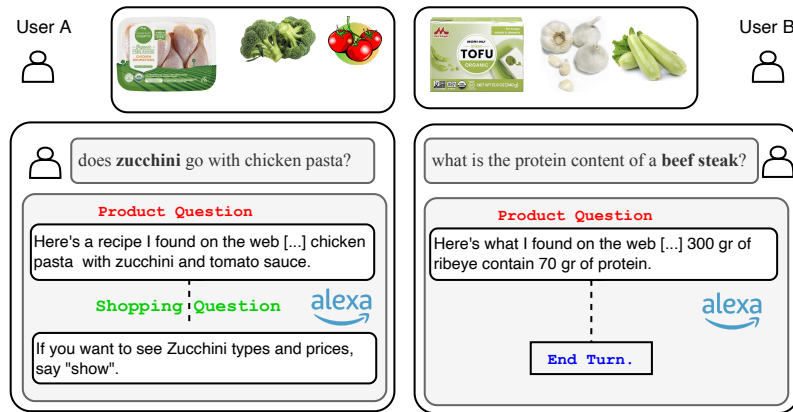


Fig. 1. User A searches for ingredients for a *recipe*. Our approach identifies the user having a shopping need, hence, triggering voice assistant’s shopping recommendations. For User B’s (likely vegetarian) our model predicts no shopping intent, hence, after answering, the turn is over.

E-commerce is an important functionality in voice assistants. In some contexts, the shopping intent can be explicit, such as in product searches (e.g. "Buy me dog food", or "Search for an Apple watch stand"). However, intent cannot always be captured from the question alone, e.g., product-related questions, such as: "What can I use as a substitute for toilet paper?"; "How long do you cook rice?"; "Are dogs allowed to eat tuna?".

The examples show that the user intent is often not clear from the Product Question (PQ) itself. PQs may emerge from an underlying desire to potentially purchase a product (*shopping need*), or only for a *general knowledge need*.

Voice product search [5, 40] and question answering [16, 6, 39, 38] have received attention, while understanding the reasons why users inquire information about certain products, especially in terms of *shopping need*, has seen very little progress. This is the main challenge that we tackle, namely determine whether a user has a shopping need. This allows voice assistants to determine *when* to surface recommendations that improve user’s shopping experience, e.g. after answering a product question, show product recommendations, deals or promo suggestions, or offering additional product details for further examination. In this work, we do not focus on *what* to recommend given the vast literature in product (sequential) recommendation [8, 34, 36, 19, 28, 31].

Figure 1 shows a hypothetical user interaction with a voice assistant. Depending on the user’s interests, shopping recommendations may differ. For instance, for *User A*, based on their purchase history we can infer that they may be interested in purchasing a queried product, recommending shopping related actions, and thus ease the shopping process for the user. For *User B* on the other hand, who has vegetarian preferences, the intent is general knowledge need.

Identifying Shopping Product Questions (SPQ) in voice search poses several challenges. Contrary to shopping intent prediction in e-commerce sites [10, 30, 15], signals such as click-rate, browse time, hardware related gestures are not present. Furthermore, shopping need detection in voice search is inherently harder given the multi-purpose use of voice assistants, compared to a restricted use of e-commerce sites.

We propose an approach based on Graph Attention Networks (GAT) [49] to identify shopping need. Questions are considered nodes and are connected if they share the same product. As input features, we consider features that are geared at capturing diverse and latent aspects of shopping need, such as *product information*, *user past purchasing behavior*, and the *question* itself. Finally, based on the mixture-of-experts approach (MoE) [44], we propose a mechanism to compute a joint node representation from the diverse features.

Experimental evaluation on more than 370k voice assistant user questions shows that our proposed approach allows us to achieve highly accurate results in distinguishing SPQs. Furthermore, we carry out an online experiment with real voice assistant users, where we identify SPQs in real-time, and recommend users to add the queried product into their *shopping lists* following a SPQ, identified with high accuracy.¹ The approach presented in this paper through detection of SPQs allows voice assistants to recommend personalized responses according to user’s interests and needs, such as “*check price*”, “*show deals*”, “*add to cart*”, “*compare products*” etc. Our main contributions in this work are as follows:

- A new problem definition for identifying SPQs that enable voice assistants to determine when to recommend to users.
- An approach to identify SPQs in voice search, using a novel way to combine diverse features through Mixture of Experts stemming from various user signals and containing different feature types.
- As part of these experiments, we perform detailed ablations showing the impact of the different features in identifying queries with shopping need, highlighting the latent nature of users shopping behavior, and thus opening directions for improving users shopping experience with voice assistants.
- Detailed offline and online experimental evaluation on voice assistant users and user queries.

2 Related Work

Users Needs in E-commerce. Informational need in e-commerce and voice shopping is typically expressed by asking questions about products. This domain is also referred

¹ As future work we plan to experiment with different types of recommended shopping actions and assess the impact on the user experience.

to as Product Question Answering (PQA), and has attracted much interest in the past years [6, 7, 23, 55, 14, 41]. Some research in the PQA domain has also been conducted with respect to the voice medium [16, 6, 39, 38], where queries have different characteristics, and where cross-lingual and speech-to-text components add another layer of complexity. Most of the works in the PQA domain focus on different answering approaches and on different types of questions [41, 52, 35, 54, 24, 25]. Here, we do not focus on the answering side at all, but rather focus on the informational need behind product questions.

Transactional need in e-commerce providers has been investigated, mainly through purchase intent prediction [10, 30, 15, 2, 45], where behavior features correspond to browse or hardware related (touch interaction on smartphones) [15]. Esmeli et al. [10] extract features such as product click rate, view time, number of visits, to predict purchase intent. Our key differences with these works is that we tackle the problem of voice modality, where voice assistants are used for highly diverse purposes [33, 42], which is not the case of dedicated e-commerce sites, where the primary purpose is for shopping. Secondly, features that can be extracted from Web sites are not available for voice assistants, making the problem of predicting shopping needs of voice users more challenging. The work of [18] explores the problem of predicting the purchase rate of a ranking model in a voice setting, where there is an extreme position bias towards the first offer. This work lies in the domain of voice product search, where the user’s intent is clearly transactional – unlike our work, which aims to identify an underlying shopping need in the voice PQA domain.

Product Search and Recommendations. There has been a large body of research in product search and product recommendations [3, 27, 29, 57, 50, 56], which are crucial in the shopping domain, for providing relevant products, and recommendations reflecting users’ preferences. The voice medium has also attracted research in this context [5, 40], as shopping queries and behavioral shopping patterns in voice are different than in web [20]. Learning product embeddings is another challenge that has received much attention [17, 46, 48], as having relevant product embeddings allows for more accurate recommendations and better user experience. Our goal is different in that, the product of interest is explicit in the user query. Our goal is to *understand* the underlying user’s intent, rather than to provide the most relevant recommendation. Nevertheless, advanced product embeddings used for product recommendation might help in capturing the customer affinity to a specific product and improve further the shopping need classification.

Behavioral Features in Search and Recommendation. The works in [1, 21] propose the integration of implicit user feedback for ranking Web search results. Although these works tackle a different problem and are not comparable to our work, we share a similar goal in that we obtain user behavior features (e.g. purchase history) that can serve to learn a latent user shopping intent behavioral model.

Works on session based recommendation [36, 19, 28, 31] propose approaches for incorporating diverse features to facilitate recommendation within a given user session (e.g. music search). At its core it predicts the next item with which the user will interact. These works are complementary to ours as they address the question on *what* to recom-

mend, whereas we focus on determining *when* to trigger recommendations by the voice assistants recommender systems.

3 Data Analysis

3.1 Dataset and Labels

Since we cannot observe the true user intent, we rely on their purchase behavior: if the user asked a question and then purchased the referred product, we consider it as having a shopping need.

A shopping intent does not often result in an instantaneous purchase. Selecting the time window for the purchase information, from the time the user issued a PQ up to the actual purchase, represents an interesting tradeoff. Longer periods tend to better capture the hidden purchase intent, however at the same time pose the risk of labeling unrelated utterances with shopping intent. We analyzed different time windows, including one hour/day/week/two weeks, and 28 days, and found that a period of 28 days increased coverage by 187% while increasing noise by 8% only² compared to a period of one week. Given a PQ, we checked the future user purchases. If the referred product – or a similar product, sharing the same category – was purchased within 28 days, we label the PQ as SPQ, otherwise NSPQ.

We collected a dataset of 374k PQs from a leading voice assistant, evenly distributed among SPQs and NSPQs³, split into 306k/34k/37k subsets for train/validation/test set, respectively. For each question, we collected its text, voice assistant’s response, the product and its category, and user past purchases.

3.2 Data Insights

To understand correlated signals with SPQs, we perform a correlation analysis (Pearson’s correlation coefficient), between the question type, their timing, and the searched product.

Questions. The relation between the user’s shopping need and the question text is involved. We found that questions asked from users about vitamins or pet-food have a high shopping intent (88%, 82%), while questions about televisions or cellphones have low shopping intent (only 16%, 18%). Figure 2 shows the histogram of the SPQ rate for the 50 most popular product categories. We see that different product types have different SPQ rates compared to the average 50%. There are product types where most of the queries are general knowledge need (4 categories with SPQ rate below 0.25%), while there are product types where most questions have a shopping need (5 categories with SPQ rate above 0.75%). For the rest, more than half of the popular categories

² We estimated noise by scrambling the utterances and purchases and measure how often an utterance is coincidentally followed by a purchase.

³ The sample was designed to have an even distribution of NSPQs and SPQs, as such it is not representative of the overall traffic and none of the results reported here should be extrapolated to the entire traffic

have SPQ rate above 60%. This demonstrates that the product category, which can be inferred from the text of the question, is indicative of SPQs.

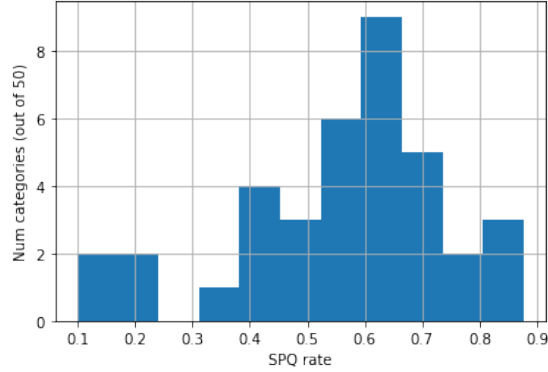


Fig. 2. SPQ rate per product type (out of 50 types with at least 1000 PQs, which represent 29.7% of the entire dataset).

However, while a shopping need is correlated with some textual features of the questions, this correlation is not natural nor intuitive. To demonstrate this, we sampled 10K questions and presented them to annotators. The annotator was presented with the question, the voice assistant response, and was required to estimate whether the user will likely purchase the referred product. To demonstrate this, we labelled 10K utterances, asking annotators to predict whether the user will purchase the product. We saw a very low correlation of 0.05 between human evaluation and future purchases. This indicates that from text alone, humans cannot predict whether a question has a shopping intent. An example of this phenomena, e.g., “*what do fish eat?*”, which to annotators seems like a general knowledge need question, while later purchases of fishing materials demonstrate the latent shopping need behind this question.

Users. The user properties also have an involved relation with the shopping intent. On the one hand, we found that shopping need is not a user property. We took two random PQs from each user, and measured the correlation between the first one being an SPQ and the second being a NSPQ. To avoid two questions referring to the same product category, which could lead to a false correlation, we de-duplicated PQs from the same user about the same product category. The correlation was only $r = 0.14$, showing that the same user may ask questions that have or do not have a shopping need.

On the other hand, we found a relation between the user’s shopping history and the shopping need. For each PQ, we checked if the referred product was previously purchased by the user (28 days prior). This property has a strong correlation of $r = 0.67$ with the SPQ property, which is based on a future purchase. This shows that when the user asks about a product category that was previously purchased, they are more likely to purchase again.

While there is strong correlation between SPQs and preceding purchases, this might simply reflect user tendency to repurchase products, and be unrelated to the question. To refute this possibility, we measured the correlation of purchasing the referred product during two consecutive 28-days periods, before the actual user’s question (PQ). We found that the correlation is significantly weaker when compared to consecutive periods coming before and after the user’s PQ, showing that the PQ plays an important role. More specifically, given a PQ referring to product P from the user’s shopping history, we consider four 28-days time periods T_0, T_{-1}, T_{-2} and T_{-3} , where T_0 is the time period coming after asking the question, T_{-1} is the time period coming before the question, and consecutively $T_{-3} < T_{-2} < T_{-1}$. The correlation between purchasing product P during T_0 and purchasing P during T_{-1} is measured to be $r = 0.67$. If this correlation was due to the periodic nature of purchasing P , we would expect a similar correlation of purchasing P also during T_{-2} and T_{-3} . However, the correlation was $r = 0.18$, showing that the high correlation of purchasing P between times T_0 and T_{-1} is strongly connected to the PQ, which indicates its shopping intent.

4 Identifying SPQs

To identify SPQs we present a model $SPQI$ with three components: (a) features of different types, (b) feature aggregation through mixture of experts [44], and (c) a Graph Attention Network.

Textual Features. The PQ is an important intent indicator, containing key aspects that are asked about a product (e.g., *price*, *delivery time*, etc.). Additionally, voice assistant’s response to the question contains information if the question was understood and answered. Furthermore, when a customer asks a question the voice assistant responds with either an *answer* or indicates that it *does not understand* the utterance. For cases where the voice assistant provides an answer, the user can continue their conversation by either asking additional questions regarding the product or make a shopping decision. We encode textual features extracted with a RoBERTa model [32], and use the [CLS] token for representation.

Product Features. An important aspect of a PQ is the queried product itself. Products have a rich structure, and users may have shopping preferences towards seemingly unrelated products or specific categories, e.g., because they share a similar price range or a similar theme, e.g., “*zucchini*” and “*broccoli*”, or “*healthy food*”. We use the product and the product category embeddings⁴ that are trained from scratch to capture shopping need related features.

We compute the product p , product category c , and parent category \hat{c} of c embeddings. The features are considered independently from each other⁵ and allow the model to learn how to leverage them during training.

Behavioral Features. To identify shopping intent, the user’s shopping behavioral patterns, extracted from the purchase history, are key. The purchase history is a sparse feature, consisting of the tuples $\langle \text{user}, \text{product} \rangle$. We pre-train a skip-gram based embeddings [4] for better generalization. We use the tuple $\langle u, p \rangle$ to train a model that predicts

⁴ Products are organized in a taxonomic product category graph

⁵ As future work we foresee integrating graph embeddings for categories and products.

if the product p will be purchased by user u , where $u, p \in \mathbb{R}^{50}$ represent the product and user vectors. The resulting embedding have two highly desirable properties: (1) create user vectors that are similar according to purchasing patterns, and (2) product vectors that are similar according to the users that co-purchased them.

Using the pre-trained model, for an input question, from the list of purchases made by the user in previous 28 days \mathbf{H} , we compute: (i) average dot/cosine similarity, maximum dot/cosine similarity, and the sum of dot/cosine similarity.

4.1 Learning Joint Question Representations via MoE

An important question is how to combine the various feature types (textual, numerical, categorical). Typically, such features are combined using either max or average pooling mechanisms [13, 51]. One drawback of this is that they do not allow the models to learn weights for each feature type according to their impact on the classification task.

We propose the use of mixture-of-experts [44], which allows the dynamic mixing of the different question feature types. Let $\mathbf{F} \in \mathbb{R}^{f \times n}$ be the concatenated set of features. For each feature dimension in \mathbf{F} , MoE weighs the importance of the f individual feature types as follows.

$$\text{MoE} = \mathbf{W}'_e (\mathbf{F}\mathbf{W}_e + c_1)^T + c_2; \quad \lambda_{f_j, z} = \frac{\text{MoE}_{f_j, z}}{\sum_{f' \in |f|} \text{MoE}_{f', z}} \quad (1)$$

where $\mathbf{W}_e \in \mathbb{R}^{n \times i}$, $\mathbf{W}'_e \in \mathbb{R}^{n \times i}$, represent trainable parameters used to compute feature type weights. Finally, $\lambda_{f_j, z} \in \mathbf{\Lambda}$ represents the feature weight for dimension z and feature type f_j , while $\mathbf{\Lambda} \in \mathbb{R}^{f \times n}$ represents all the feature weights.

The final joint question representation that is passed onto the GAT framework, is computed by performing the Hadamard product between each feature vector and the computed MoE weights: $\mathbf{h} = \sum_{f' \in |f|} \mathcal{F}_{f'} \circ \mathbf{\Lambda}_{f'}$.

4.2 Question Graph Construction & Classification

Considering questions as independent instances has the disadvantage of not being able to leverage similar *questions* or *users* that may have similar purchasing behavior. Hence, we formulate the SPQ identification problem as a node classification task using GAT [49]. GAT allow to create soft dependencies between questions (if they share the same product), and as such, influences the prediction of shopping need for a question, by taking into account questions about the same product from other users. An important aspect of utilizing GATs, is the shape of the graph. As nodes are the user questions, represented by the joint representation \mathbf{h} , and nodes are connected if they share the same queried *product*.

We train a GAT model to compute the question representations: $\hat{\mathbf{h}}_q = \text{GAT}(\mathbf{h}_q)$, which is used to obtain the SPQ probability as $s_q = \sigma(\hat{\mathbf{h}}_q \cdot \theta)$ to classify into either an SPQ or a NSPQ.

5 Experiments

We discuss the offline and online evaluation results, and further introduce baselines for the offline experimentation phase.

5.1 Baselines and Ablations

RoBERTa-Baseline. This baseline highlights the difficulty of identifying shopping intent from text alone. We consider two variations: (1) `RoBERTa-Question`, where only the question text is used for training and inference, and (2) `RoBERTa-Text`, where additionally voice assistant’s answer is used for training and inference.

MLP-Baseline. All the computed features are pushed through an MLP layer for classification. We consider several ablations (`MLP-Question` and `MLP-Text` are identical to the `RoBERTa` baselines) of this baseline.

SPQI Setup and Ablations. We consider different model ablations using different feature subsets, and we distinguish between `SPQI-MoE` and `SPQI-CONCAT`, with the difference being on how the diverse features are aggregated with `MoE` or simply concatenated. `SPQI` has four layers of graph convolutions [22], and each node has 1024 dimensions.

6 Evaluation Results

6.1 Offline Evaluation

Model	Behavior			Product			Text+Product			Text+Behavior			Product+Behavior			Full		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1
MLP-Concat	0.827	0.910	0.866	0.758	0.755	0.757	0.758	0.766	0.762	0.856	0.916	0.885*	0.876	0.881	0.878	0.875	0.881	0.878
MLP-MoE	0.907	0.886	0.896	0.756	0.758	0.757	0.761	0.730	0.744	0.895	0.921	0.907*	0.903	0.908	0.906	0.903	0.914	0.908
SPQI-Concat	0.903	0.884	0.893*	0.760	0.757	0.759	0.760	0.772	0.766	0.898	0.893	0.895	0.890	0.904	0.895*	0.895	0.909	0.902*
SPQI-MoE	0.909	0.890	0.900*	0.762	0.764	0.763*	0.760	0.793	0.776*	0.914	0.885	0.900	0.903	0.910	0.907	0.903	0.917	0.910*

Table 1. SPQ detection performance of competing approaches. Significant results (as per paired t-test) are marked with * (p -value < 0.05). `SPQI-MoE` ablations are compared to `MLP-MoE`, and `SPQI-Concat` vs. `MLP-Concat`. For the sake of clarity, we omit from the table the results for `RoBERTa-Question` and `RoBERTa-Text`, given that they are the lowest performing with $F1=0.694$ and $F1=0.699$, respectively. On the same feature set, `SPQI-MoE-Query` and `SPQI-MoE-Text` obtain significantly higher results with $F1=0.777$ and $F1=0.774$, respectively.

Model Performance. Table 1 shows the results of identifying SPQs. Our approach, `SPQI-MoE` achieves the best results with $F1 = 0.91$. `SPQI-Concat-Full` has a 1% drop in F1. This shows that the flexibility of `MoE` to dynamically decide per question, which features are important in predicting user’s shopping need.

Comparing `SPQI-Full-MoE` against `MLP-MoE-Full`, the use of GATs provides an additional advantage in terms of F1. This shows that even for the same feature set and way to combine them, GATs ability to consider neighbouring nodes for determining shopping need is helpful. This in itself is intuitive and has been widely studied in recommender systems (e.g. collaborative filtering), where tying users with similar search patterns can be helpful in recommendations. In our case, according to how sub-graphs are constructed for classification, our models exploit information coming from users who share similar purchasing patterns and ask similar questions. Finally, `RoBERTa`

baselines obtain the lowest performance, showing that identifying SPQs cannot be done from the question text alone.

In nearly all the cases, combining the question features using the MoE approach, allows the models to feature concatenation. Furthermore, computing the final question representation using GATs, where neighbouring nodes can have an impact on its representation, allows the model to obtain more accurate intent classification results.

Although the SPQI models are larger (in parameters) than their MLP counterparts, the performance decrease is not due to model size, as the comparison between SPQI-MoE and SPQI-Concat shows that SPQI-MoE achieves better performance.

Ablations. The user behavioral features achieve the biggest performance improvement. SPQI-Text+Behavior is the most accurate model with $P = 0.914$, which is 1% higher than SPQI-Full-MoE. However, in terms of recall we get a drop of 3%. The addition of product embedding features, allows the model to improve its recall by covering cases of users that may have a sparse purchase history or that few purchased products within the same category as the queried product. Interestingly, none of the feature combinations manage to achieve the full performance of SPQI-MoE, where we draw conclusions that while the precision may be higher on certain feature subsets (e.g. SPQI-MoE-Text+Behavior), the combination of all the features allows to obtain maximum coverage.

Similarly, here too, the SPQ-MoE and SPQ-Concat ablations perform better than most of their counterparts. In only one case the MLP obtain significantly better performance.

In terms of architecture, SPQI-MoE-Text obtains an increase of $F1 = +7\%$ over RoBERTa. This improvement can be attributed to the use of GATs. Additionally, we do not see any significant difference between the models when trained only on the users questions, as opposed to both the question and the voice assistant answers. One reason may be that answers are usually handled by Q&A systems that do not provide personalized answers to the users.

Overall, only MLP-MoE-Text+Behavior achieves results that are significantly better than those of SPQI-MoEText+Behavior. The performance decrease comes from lower recall.

In summary, the ablations show that our contributions are twofold. First, the proposed features that capture user’s shopping need are highly suitable for voice search. Second, the proposed MoE for integrating the different features into a joint question representation, and the use of GATs for the intent classification task, with consistent improvement over the competing approaches (e.g., MLP-Concat/MoE). Note that the MLP baseline represents a strong baseline, since it is trained on the proposed features and uses the MoE.

6.2 Online Recommendation Evaluation

The SPQ definition in Section 3 is only a proxy of shopping intent. PQs can theoretically be unrelated to the associated purchase. Our online evaluation, with real voice assistant users, verifies that the identified SPQs by our SPQI-MoE have a shopping intent. For this experiment, we split the users into two groups: 1) $T1$ where the users, according to our model, are recommended a shopping action to add a product into their shopping

list, and 2) C where all the users issuing a PQ are prompted to add the product into their shopping list.

$T1$ obtains 81.5% higher F1 score than C in identifying SPQs, and additionally users on $T1$ were 79.5% more likely to add the product into their shopping list when compared to users in C . This experiment shows that our approach is able to identify user’s shopping need, validated by users adding the queried product into their shopping list. Note that, not adding a product to the shopping list does not indicate the contrary, namely, the lack of shopping intent.

7 Conclusion

We presented an approach to make proactive shopping recommendations in a leading industry voice assistant. By identifying shopping need, we allow the voice assistant to accurately recommend to its users relevant products or proactively suggest shopping actions that enhance their shopping experience.

We proposed a set of features that capture shopping intent from various perspectives, and relied on graph attention networks and on a novel mechanism to combine features for classifying a PQ as an SPQ. Experimental evaluations show that our proposed features and the way we encode user question through GATs, yield significant improvements over text classification approaches, achieving an increase in $F1$ of +21%. The experiments confirm the highly latent nature of shopping intent. This work, unlike others, explores users’ shopping intent in voice, not only in the context of product search, but also in the context of product question answering. Furthermore, focusing on informational intent and exploring whether the underlying need behind is related to general knowledge or to shopping.

References

1. Agichtein, E., Brill, E., Dumais, S.T.: Improving web search ranking by incorporating user behavior information. In: SIGIR (2006). <https://doi.org/10.1145/1148170.1148177>, <https://doi.org/10.1145/1148170.1148177>
2. Ahmadvand, A., Kallumadi, S., Javed, F., Agichtein, E.: Jointmap: Joint query intent understanding for modeling intent hierarchies in e-commerce search. In: SIGIR (2020). <https://doi.org/10.1145/3397271.3401184>, <https://doi.org/10.1145/3397271.3401184>
3. Bi, K., Ai, Q., Croft, W.B.: Learning a fine-grained review-based transformer model for personalized product search. In: SIGIR (2021). <https://doi.org/10.1145/3404835.3462911>, <https://doi.org/10.1145/3404835.3462911>
4. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. *Trans. Assoc. Comput. Linguistics* **5**, 135–146 (2017), <https://transacl.org/ojs/index.php/tacl/article/view/999>
5. Carmel, D., Haramaty, E., Lazerson, A., Lewin-Eytan, L., Maarek, Y.: Why do people buy seemingly irrelevant items in voice product search?: On the relation between product relevance and customer satisfaction in ecommerce. In: WSDM (2020). <https://doi.org/10.1145/3336191.3371780>, <https://doi.org/10.1145/3336191.3371780>

6. Carmel, D., Lewin-Eytan, L., Maarek, Y.: Product question answering using customer generated content-research challenges. In: SIGIR (2018)
7. Chen, S., Li, C., Ji, F., Zhou, W., Chen, H.: Review-driven answer generation for product-related questions in e-commerce. In: WSDM (2019)
8. Chen, W., He, M., Ni, Y., Pan, W., Chen, L., Ming, Z.: Global and personalized graphs for heterogeneous sequential recommendation by learning behavior transitions and user intentions. In: RecSys (2022). <https://doi.org/10.1145/3523227.3546761>, <https://doi.org/10.1145/3523227.3546761>
9. Choi, J.I., Kuzi, S., Vedula, N., Zhao, J., Castellucci, G., Collins, M., Malmasi, S., Rokhlenko, O., Agichtein, E.: Wizard of tasks: A novel conversational dataset for solving real-world tasks in conversational settings. In: Proceedings of the 29th International Conference on Computational Linguistics. pp. 3514–3529. International Committee on Computational Linguistics, Gyeongju, Republic of Korea (Oct 2022), <https://aclanthology.org/2022.coling-1.310>
10. Esmeli, R., Bader-El-Den, M., Abdullahi, H.: Towards early purchase intention prediction in online session based retailing systems. *Electronic Markets* pp. 1–19 (2020)
11. Faustini, P., Chen, Z., Fetahu, B., Rokhlenko, O., Malmasi, S.: Answering unanswered questions through semantic reformulations in spoken QA. In: Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 5: Industry Track). pp. 729–743. Association for Computational Linguistics, Toronto, Canada (Jul 2023). <https://doi.org/10.18653/v1/2023.acl-industry.70>, <https://aclanthology.org/2023.acl-industry.70>
12. Fetahu, B., Faustini, P., Castellucci, G., Fang, A., Rokhlenko, O., Malmasi, S.: Follow-on question suggestion via voice hints for voice assistants. *EMNLP 2023* (2023)
13. Gao, J., Lyu, T., Xiong, F., Wang, J., Ke, W., Li, Z.: MGNN: A multimodal graph neural network for predicting the survival of cancer patients. In: SIGIR (2020). <https://doi.org/10.1145/3397271.3401214>, <https://doi.org/10.1145/3397271.3401214>
14. Gao, S., Ren, Z., Zhao, Y., Zhao, D., Yin, D., Yan, R.: Product-aware answer generation in e-commerce question-answering. In: Proc. of WSDM. pp. 429–437 (2019)
15. Guo, L., Hua, L., Jia, R., Zhao, B., Wang, X., Cui, B.: Buying or browsing?: Predicting real-time purchasing intent using attention-based deep network with multiple behavior. In: SIGKDD (2019). <https://doi.org/10.1145/3292500.3330670>, <https://doi.org/10.1145/3292500.3330670>
16. Guy, I.: Searching by talking: Analysis of voice queries on mobile web search. In: Perego, R., Sebastiani, F., Aslam, J.A., Ruthven, I., Zobel, J. (eds.) SIGIR (2016). <https://doi.org/10.1145/2911451.2911525>, <https://doi.org/10.1145/2911451.2911525>
17. Hao, J., Zhao, T., Li, J., Dong, X.L., Faloutsos, C., Sun, Y., Wang, W.: P-companion: A principled framework for diversified complementary product recommendation. In: CIKM (2020)
18. Haramati, E., Fairstein, Y., Lazerson, A., Lewin-Eytan, L.: External evaluation of ranking models under extreme position-bias. In: WSDM (2022)
19. Hidasi, B., Karatzoglou, A.: Recurrent neural networks with top-k gains for session-based recommendations. In: CIKM (2018). <https://doi.org/10.1145/3269206.3271761>, <https://doi.org/10.1145/3269206.3271761>
20. Ingber, A., Lazerson, A., Lewin-Eytan, L., Libov, A., Osherovich, E.: The challenges of moving from web to voice in product search. In: Proceedings of the 1st International Workshop on Generalization in Information Retrieval (GLARE 2018) (2018), <http://glare2018.dei.unipd.it/paper/glare2018-paper5.pdf>

21. Joachims, T.: Optimizing search engines using clickthrough data. In: SIGKDD
22. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: ICLR (2017), <https://openreview.net/forum?id=SJU4ayYgl>
23. Kulkarni, A., Mehta, K., Garg, S., Bansal, V., Rasiwasia, N., Sengamedu, S.: Productqna: Answering user questions on e-commerce product pages. In: Companion Proceedings of WWW (2019)
24. Lai, T., Bui, T., Li, S., Lipka, N.: A simple end-to-end question answering model for product information. In: Proc. of ECONLP. pp. 38–43 (2018)
25. Lai, T.M., Bui, T., Lipka, N., Li, S.: Supervised transfer learning for product information question answering. In: Proc. of ICMLA. pp. 1109–1114. IEEE (2018)
26. Le, D.T., Weber, V., Bradford, M.: Combining semantic search and twin product classification for recognition of purchasable items in voice shopping. In: 4th Workshop on e-Commerce and NLP (Aug 2021)
27. Le, T.H., Lauw, H.W.: Explainable recommendation with comparative constraints on product aspects. In: WSDM (2021)
28. Li, J., Ren, P., Chen, Z., Ren, Z., Lian, T., Ma, J.: Neural attentive session-based recommendation. In: CIKM (2017). <https://doi.org/10.1145/3132847.3132926>, <https://doi.org/10.1145/3132847.3132926>
29. Li, Y., Chen, W., Yan, H.: Learning graph-based embedding for time-aware product recommendation. In: CIKM (2017)
30. Ling, C., Zhang, T., Chen, Y.: Customer purchase intent prediction under online multi-channel promotion: A feature-combined deep learning framework. IEEE Access **7**, 112963–112976 (2019)
31. Liu, Q., Zeng, Y., Mokhosi, R., Zhang, H.: STAMP: short-term attention/memory priority model for session-based recommendation. In: SIGKDD (2018). <https://doi.org/10.1145/3219819.3219950>, <https://doi.org/10.1145/3219819.3219950>
32. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized BERT pretraining approach. CoRR **abs/1907.11692** (2019), <http://arxiv.org/abs/1907.11692>
33. Lopatovska, I., Rink, K., Knight, I., Raines, K., Cosenza, K., Williams, H., Sorsche, P., Hirsch, D., Li, Q., Martinez, A.: Talk to me: Exploring user interactions with the amazon alexa. Journal of Librarianship and Information Science **51**(4), 984–997 (2019)
34. Ma, R., Liu, N., Yuan, J., Yang, H., Zhang, J.: CAEN: A hierarchically attentive evolution network for item-attribute-change-aware recommendation in the growing e-commerce environment. In: RecSys (2022). <https://doi.org/10.1145/3523227.3546773>, <https://doi.org/10.1145/3523227.3546773>
35. McAuley, J., Yang, A.: Addressing complex and subjective product-related queries with customer reviews. In: Proc. of WWW. pp. 625–635 (2016)
36. Meng, W., Yang, D., Xiao, Y.: Incorporating user micro-behaviors and item knowledge into multi-task learning for session-based recommendation. In: SIGIR (2020). <https://doi.org/10.1145/3397271.3401098>, <https://doi.org/10.1145/3397271.3401098>
37. Rafailidis, D., Manolopoulos, Y.: The technological gap between virtual assistants and recommendation systems. CoRR **abs/1901.00431** (2019), <http://arxiv.org/abs/1901.00431>
38. Ramnath, K., Sari, L., Hasegawa-Johnson, M., Yoo, C.: Worldly wise (wow) - cross-lingual knowledge fusion for fact-based visual spoken-question answering. In: NAACL-HLT (2021). <https://doi.org/10.18653/v1/2021.naacl-main.153>, <https://doi.org/10.18653/v1/2021.naacl-main.153>

39. Ravichander, A., Dalmia, S., Ryskina, M., Metze, F., Hovy, E.H., Black, A.W.: Noiseqa: Challenge set evaluation for user-centric question answering. In: EACL (2021), <https://aclanthology.org/2021.eacl-main.259/>
40. Rhee, C.E., Choi, J.: Effects of personalization and social role in voice shopping: An experimental study on product recommendation by a conversational voice agent. *Comput. Hum. Behav.* **109**, 106359 (2020). <https://doi.org/10.1016/j.chb.2020.106359>, <https://doi.org/10.1016/j.chb.2020.106359>
41. Rozen, O., Carmel, D., Mejer, A., Mirkis, V., Ziser, Y.: Answering product-questions by utilizing questions from other contextually similar products. In: Proc. of NAACL. pp. 242–253 (2021)
42. Rzepka, C.: Examining the use of voice assistants: A value-focused thinking approach (2019)
43. Salle, A., Malmasi, S., Rokhlenko, O., Agichtein, E.: Studying the Effectiveness of Conversational Search Refinement Through User Simulation. In: Proceedings of 2021 European Conference on Information Retrieval (ECIR 2021). pp. 587–602. Springer (2021). https://doi.org/10.1007/978-3-030-72113-8_39
44. Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q.V., Hinton, G.E., Dean, J.: Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In: ICLR (2017), <https://openreview.net/forum?id=BlcKMDqIlg>
45. Shen, Y., Yan, J., Yan, S., Ji, L., Liu, N., Chen, Z.: Sparse hidden-dynamics conditional random fields for user intent understanding. In: WWW (2011). <https://doi.org/10.1145/1963405.1963411>, <https://doi.org/10.1145/1963405.1963411>
46. Tagliabue, J., Yu, B., Bianchi, F.: The embeddings that came in from the cold: Improving vectors for new and rare products with content-based inference. In: RecSys (2020). <https://doi.org/10.1145/3383313.3411477>, <https://doi.org/10.1145/3383313.3411477>
47. Tyagi, A., Sharma, V., Gupta, R., Samson, L., Zhuang, N., Wang, Z., Campbell, B.: Fast intent classification for spoken language understanding systems. In: IEEE ICASSP (2020). <https://doi.org/10.1109/ICASSP40776.2020.9054240>, <https://doi.org/10.1109/ICASSP40776.2020.9054240>
48. Vasile, F., Smirnova, E., Conneau, A.: Meta-prod2vec: Product embeddings using side-information for recommendation. In: RecSys (2016). <https://doi.org/10.1145/2959100.2959160>, <https://doi.org/10.1145/2959100.2959160>
49. Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. In: ICLR (2018), <https://openreview.net/forum?id=rJXMpikCZ>
50. Wang, J., Sarwar, B., Sundaresan, N.: Utilizing related products for post-purchase recommendation in e-commerce. In: RecSys (2011). <https://doi.org/10.1145/2043932.2043995>, <https://doi.org/10.1145/2043932.2043995>
51. Wei, Y., Wang, X., Nie, L., He, X., Hong, R., Chua, T.: MMGCN: multi-modal graph convolution network for personalized recommendation of micro-video. In: ACM MM. ACM (2019)
52. Xu, H., Xie, S., Shu, L., Philip, S.Y.: Dual attention network for product compatibility and function satisfiability analysis. In: AAAI. vol. 32 (2018)
53. Yang, J., Drake, T., Damianou, A.C., Maarek, Y.: Leveraging crowdsourcing data for deep active learning an application: Learning intents in alexa. In: WWW (2018). <https://doi.org/10.1145/3178876.3186033>, <https://doi.org/10.1145/3178876.3186033>
54. Yu, J., Zha, Z.J., Chua, T.S.: Answering opinion questions on products by exploiting hierarchical organization of consumer reviews. In: Proc. of EMNLP. pp. 391–401 (2012)

55. Zhang, S., Lau, J.H., Zhang, X., Chan, J., Paris, C.: Discovering relevant reviews for answering product-related queries. In: ICDM'19 (2019)
56. Zhao, J., Chen, H., Yin, D.: A dynamic product-aware learning model for e-commerce query intent understanding. In: CIKM (2019)
57. Zheng, L., Noroozi, V., Yu, P.S.: Joint deep modeling of users and items using reviews for recommendation. In: WSDM (2017). <https://doi.org/10.1145/3018661.3018665>, <https://doi.org/10.1145/3018661.3018665>