# Intent-based Product Collections for E-commerce using Pretrained Language Models

Hiun Kim*, Jisu Jeong*†, Kyung-Min Kim*†, Dongjun Lee‡, Hyun Dong Lee*,
Dongpil Seo*, Jeeseung Han*, Dong Wook Park*, Ji Ae Heo*, Rak Yeong Kim*

*NAVER CLOVA    †NAVER AI LAB    ‡LBox Co., Ltd.

*Abstract*—**Building a shopping product collection has been primarily a human job. With the manual efforts of craftsmanship, experts collect related but diverse products with common shopping intent that are effective when displayed together, e.g., backpacks, laptop bags, and messenger bags for freshman bag gifts. Automatically constructing a collection requires an ML system to learn a complex relationship between the customer's intent and the product's attributes. However, there have been challenging points, such as 1) long and complicated intent sentences, 2) rich and diverse product attributes, and 3) a huge semantic gap between them, making the problem difficult. In this paper, we use a pretrained language model (PLM) that leverages textual attributes of web-scale products to make intent-based product collections. Specifically, we train a BERT with triplet loss by setting an intent sentence to an anchor and corresponding products to positive examples. Also, we improve the performance of the model by search-based negative sampling and category-wise positive pair augmentation. Our model significantly outperforms the search-based baseline model for intent-based product matching in offline evaluations. Furthermore, online experimental results on our e-commerce platform show that the PLM-based method can construct collections of products with increased CTR, CVR, and order-diversity compared to expert-crafted collections.**

*Keywords*-**Product Collections; E-Commerce; Pretrained Language Models;**

## I. INTRODUCTION

Current e-commerce platforms heavily rely on search engines or recommender systems to display their products to customers. However, there have been limitations of such tools. Customers do not precisely know which keywords they have to provide with search engines in advance. Recommender systems only provide multiple products that are just top-k results of computation between user features and product features without explicit consideration of customer's intention [1]–[5]. Thus, unless the customer makes many iterative searches or browsing, their intention may not be satisfied. As a solution, e-commerce platforms regularly display intent-based shopping product collections on their websites to provide convenient shopping experiences to users. The assortment of products with a common theme can reduce the burden of iterative search or browsing by combining coherent but diverse products to fulfill customer's intentions at once. Although there has been rapid progress in

the recommender system domain, making product collection has still been a human task. The quality of outcome is more prone to errors than that of a search engine or recommender system. In other words, it is not yet an alternative solution to overcome the limitations of a current search engine or recommender system because both volumes and diversity of current product collections mostly retained in human scale. Technically, there are three challenging problems to build a product collection automatically. 1) understanding of long and complicated intent sentences. 2) handling rich and diverse attributes of products combined in a collection. 3) closing a vast semantic gap between the customer's intention and the product's attributes. Fig. 1 shows an example of constructing a collection. Typically, a single intent sentence contains multiple intents (e.g., functionality, utility, or style) as well as a product has rich and diverse textual attributes (e.g., 'daily simple', 'H&M', 'basic', or 'Uniqlo') making the matching task complex.

In this paper, we build an intent-based product collection by using a pretrained language model. BERT [6] has powerful language understanding capability by pretrained knowledge and Self-Attention [7] to handle long sentences. Specifically, we adopt Sentence-BERT (SBERT) [8] to fine-tune its parameters for matching textual attributes of products and corresponding intent sentences with triplet loss. We set a concatenation of the title, section name, and date of a product collection as a query and up to eleven textual attributes of products in the collections as positive examples. We use term-based search using the query for gathering hard negative examples. Also, we augment positive pairs by modifying a query to contain category name.

Offline experimental results on five in-house datasets show that our negative sampling and data augmentation techniques improve recall and precision by a large margin. Compare to the baseline search-based model, our model significantly improves the intent-based matching performance. When performing online evaluations in our service, the product collections from the model increase the CTR, CVR, order-diversity by 16%, 29%, 60%, respectively, comparing to the product collections made by human experts. We additionally present an ablation study by different sizes of the SBERT model. Examples of intents and matched products using our

Figure 1: An example of constructing an intent-based product collection for e-commerce.

model and an example of exposed product collection on our service are displayed in Appendix B and C, respectively.

## II. RELATED WORK

### A. Intent-based Product Collections

Knowledge graphs are used to identifying products having similar attributes [9]–[11]. [12] built a huge knowledge graph whose node types are items, e-commerce concepts, primitive concepts, and taxonomy. And for a given specific concept, they collect products that are connected with the concept. Motivated by existing bag-of-words approaches [13], [14], Angelov [15] proposed Top2Vec, a representation learning for topic modeling. However, since their models treat an intent as one object, e.g. a node or a topic, their representation power of intents is limited.

### B. Metric Learning for Product Retrieval

Metric learning learns and measures the similarity of data instances, where a relevant representation of sentences can directly benefit performance [16]. Many representation learning techniques are used for metric learning such as Bilinear model [17], CNNs [18], [19], RNNs [20], and others [21]–[27]. In application to product search, Nigam et al. [28] introduced a shared embedding model for semantic product search. They focus on search queries, such as categories and names, that are typically short thus their approach also has a limitation for understanding and relating long, complicated intents to rich product attributes.

### C. Pretrained Language Models for Product Understanding

BERT has been applied for general semantic relevances task [29] and web search [30] for its powerful language understanding capability compared with conventional lexical matching [31]. For e-commerce tasks, BERT has been applied to product specification understanding [32] and user behavior prediction using session log-based product representations [33]. Reimers et al. [8] presented Sentence-BERT (SBERT) for similarity learning using siamese or triplet BERT architecture.

We choose SBERT for our task, because SBERT could understand arbitrary intents and rich product attributes end to end without manual effort of building ontologies. Even if metric learning for BERT has been used for product matching [34], applying SBERT to textual attributes of large, manually crafted collections to facilitate intent-based product retrieval is not intensively studied to the best of our knowledge.

## III. METHOD

### A. Task Description

When shopping products are visible at a glance, it is easy to compare and choose the product for specific themes or purposes. For this reason, many e-commerce platforms, such as Amazon, Shopify, and Zalando, collect the products under certain themes or purposes and expose them to their main page. Because of its importance, this work is usually done by human experts. Human experts figure out various shopping intents of users and make product collections that users can satisfy. However, they spend too much time and have a lot of difficulties choosing the proper products across diverse themes and purposes of shopping. This is the motivation of our work. We have a sentence, which includes users' shopping intent, as an input and then output the collection of products with a common intent. Fig. 1 emphasizes the advantage of intent-based product retrieval. Formally speaking, we have a set of products $P_{full}$ and a shopping intent $I$, then our task is to match a subset of products $P_{sub}$ from $P_{full}$ with $I$.

$$P_{sub} = Matching(P_{full}, I) \qquad (1)$$

### B. Metric Learning with Product Collections Data

We adopted SBERT for implementing $Matching$ function in Eq. (1). We reconstruct existing product collection data created by human operators for metric learning with triplet loss [35].
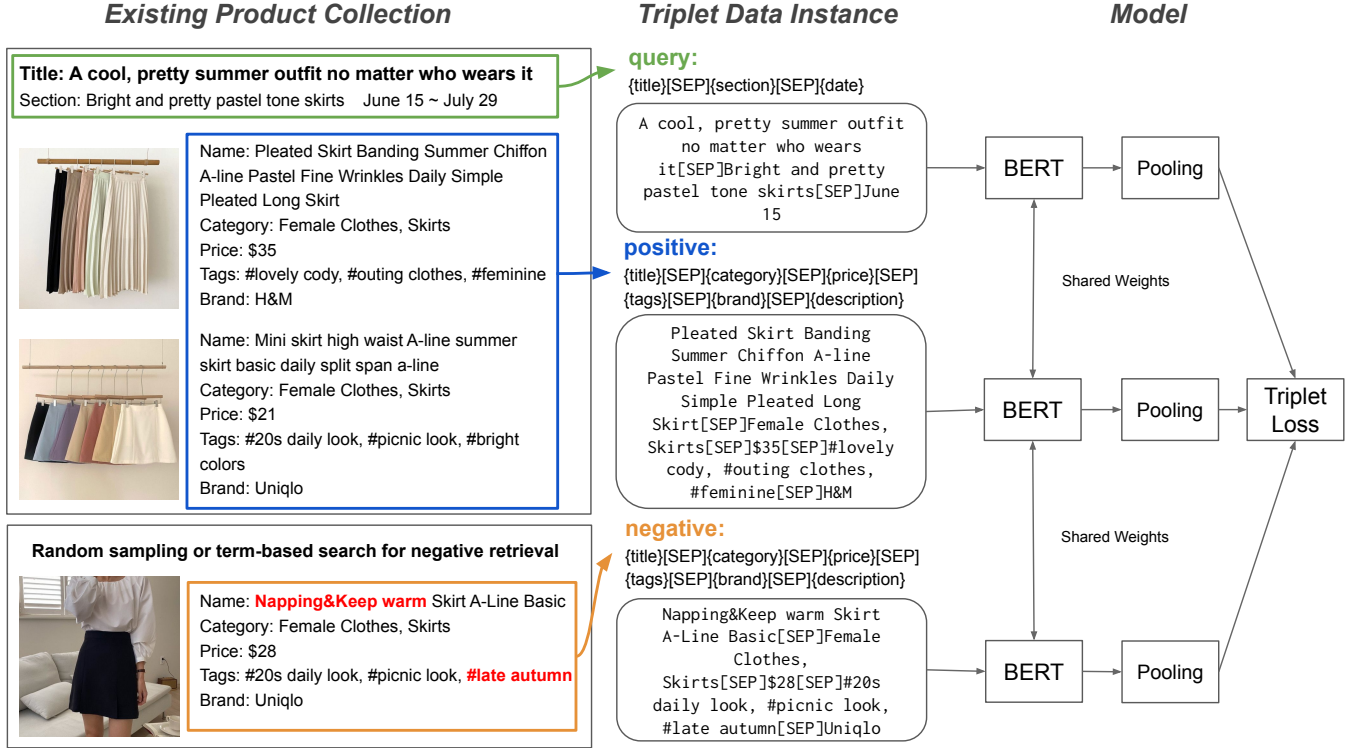
Figure 2: Illustration of data instances and model architecture used for intent-based product retrieval.

*1) Training Objectives and Loss Functions:* The training objective is to create representation for a given intent query and positive products to have high similarity while keeping low similarity between the query and negative product. On the right side of Fig. 2 shows neural network architecture corresponds to anchor, positive, and negative examples.

Existing product collections consist of triplet anchor (query), and positive products as in the left side of Fig. 2. For negative products, we sampled randomly or retrieved from the term-based search results.

For a given query $Q$, a positive product $P$ and a negative product $N$, let $e_Q$, $e_P$, and $e_N$ be an embedding vector of $Q$, $P$, and $N$, respectively, after applying BERT and pooling. For triplet loss for metric learning, we use Euclidean norm and triplet margin to 1 followed by [8]. The formula is described as follows:

$$Loss(Q, P, N) = ReLU(\|e_Q - e_P\| - \|e_Q - e_N\| + Margin)$$

### C. Constructing BERT Input of Intents and Products

A single product collection is decomposed to 50∼100 triplet positive pairs for training.

*1) Intent Query:* The intent query is composed of a title, section name (name of products group in the product collection), and start date of the product collection. For example, title of product collection in Fig. 2 'A cool, pretty summer outfit no matter who wears it' is an intent that describes the functionality, utility, or style. Products inside

that product collection have high conformance for that intent. The section name 'Bright and pretty pastel tone skirts' of the product collection involves both an intent and product category (i.e. skirts). If category information exists in the section name, product collection needs to be composed of products from that category. The product collection's start date of exposure on the e-commerce platform is also used as part of a query. This information indicates seasonality. For example, 'June 15' is added to the query for the model to possibly learn representations for relevant products of early summer (e.g. skirts with light color). The elements of the query are concatenated with [SEP] token.

*2) Positive and Negative Products:* The product is represented as a sentence containing up to eleven textual product attributes such as title, categories, price, tags, or brand concatenated with [SEP] token. The positive product is derived from the existing product collection with corresponding intent query. The negative product is randomly sampled or retrieved with BM25 term-based search [31] by the query text. The purpose of the BM25 term-based search is to gather hard negatives effectively [36].

### D. Joint Learning with Category-Wise Data Augmentation

*1) Overlapped Intents across Categories:* Products in intent-based collections usually share the same or similar product category. For instance, Fig. 1 shows the intent of 'A cool, pretty summer skirts no matter who wears it'. This

**Original Product Collection**

| Title: summer outfit<br>Section: pastel tone skirts<br><br>P1 skirts (category)<br>P2 skirts<br>P3 pants<br>P4 shirts<br>P5 skirts | **Original Anchor and Positive**<br><br>`summer outfit[SEP]pastel tone skirts - P1 skirts`<br>`summer outfit[SEP]pastel tone skirts - P2 skirts`<br>`summer outfit[SEP]pastel tone skirts - P3 pants`<br>`summer outfit[SEP]pastel tone skirts - P4 shirts`<br>`summer outfit[SEP]pastel tone skirts - P4 skirts` |
|---|---|

**Augmented Product Collections by Product Category**

| Title: summer outfit<br>Section: skirts<br><br>P1 skirts<br>P2 skirts<br>P5 skirts | **Category-aware Anchor and Positive - skirts**<br><br>`summer outfit[SEP]skirts    -    P1 skirts`<br>`summer outfit[SEP]skirts    -    P2 skirts`<br>`summer outfit[SEP]skirts    -    P5 skirts` |
|---|---|
| Title: summer outfit<br>Section: pants<br><br>P3 pants | **Category-aware Anchor and Positive - pants**<br><br>`summer outfit[SEP]pants    -    P3 pants` |
| Title: summer outfit<br>Section: shirts<br><br>P4 shirts | **Category-aware Anchor and Positive - shirts**<br><br>`summer outfit[SEP]shirts    -    P4 shirts` |

Figure 3: Example of category-wise data augmentation.

intent contains category information 'skirts'. Since we are training with many existing product collections datasets, there will be product collections entitled 'A cool, pretty summer pants no matter who wears it', or 'A cool, pretty summer shirts no matter who wears it'. Therefore, the model has a chance to match the wrong products (pants and shirts) along with skirts since their pure intent sentences are similar.

In addition, experts also include different category products at the mid or later part of product collections for complementary products displays (e.g. displaying skirts and pants together). However, our training data do not consider the significance of the relative order of products in a collection. Thus, queries may have wrong category information to the portion of products in a collection (e.g. skirts query mapped to pants products). A model trained with this data can potentially result in categorically wrong product retrieval.

*2) Category-Wise Data Augmentation:* To alleviate the categorical correctness issue, we perform simple data augmentation as visualized in Fig. 3. 1) Randomly select a specific ratio of product collections from all product collections in the training data. 2) Transform a single product collection that has multiple product categories into multiple product collections with a single product category. 3) Replace a section name in the query of transformed product collections by common category name of products. This simple augmentation forces model to more aware category information in the query, and it significantly improves categorical precision, compared to the tradeoff made in the recall. The detailed experimental result will be discussed in Section IV.

Table I: Statistics of product collections for training data.

| Data | Counts |
|---|---|
| Collections | 88,371 |
| Sections in Collections | 139,794 |
| Products in Collections | 8,790,251 |
| Avg. Product per Collection | 99.46 |
| Avg. Product per Section | 62.88 |
| Total Product Categories | 1,154 |

### E. Model and Learning

*1) Models:* Pretrained language model BERT is used for making product representations for various e-commerce tasks [32]–[34], and it plays a baseline model for the modern intent understanding of natural languages in the dialog systems domain [37], [38]. Furthermore, BERT is plausible for joint learning [37], which is helpful for the training model with augmented data aiming for both intent-based retrieval as well as categorically correct retrieval. For these reasons, we adopt BERT as a base module for understanding products and intents.

For full model, we adopt SBERT [8], a siamese and triplet network version of the BERT, to project intent and product representation in the same vector space. We used custom pretrained BERT weight (110M, 12L, 768 output dim.) for SBERT. We used the Wordpiece tokenizer [39] for both pretraining and fine-tuning the BERT model to understand intent and product attributes with semantically relevant subwords.

*2) Architectural Analysis for Intent-based Product Retrieval:* BERT recognizes `[SEP]` token to separate different attribute types of a product. After the model ingests the triplet dataset, Self-Attention [7] captures a latent representation of query ($e_Q$) and text information of product ($e_P$). Self-Attention helps to make relevant representations for long and complicated intent queries. It also helps to make relevant representations for arbitrary length product attributes by handling long-range dependencies. Layer normalization [40] in Self-Attention alleviates potential length difference between query and product where the concatenated product attributes are usually a few times longer than an intent query. We used the default `MEAN`-strategy pooling in SBERT to create output representation. For production, we perform a cosine similarity search on top of $e_Q$ and precomputed $e_P$.

## IV. EXPERIMENTS

### A. Training and Evaluation

*1) Training Data:* We extracted product collections for the fashion category published from in 2.5 years period. The statistics of extracted product collections is stated in Table I. Due to the operation policy of our platforms, information of some old products are not available. Thus, only currently

Table II: Dataset configuration and statistics of models.

| Model | Negative Sampling | | Cat.-Wise Data Aug. | Products for Positive Pairs | Products for Negative Pairs | Total Triplet Data Instances |
|---|---|---|---|---|---|---|
| | Product Cat. | BM25 | | | | |
| $PR_{easy0}$ | 25 | 0 | 0% | | 968,202 | 73,752,535 |
| $PR_{hard0}$ | 10 | 15 | 0% | | 6,477,564 | 62,307,766 |
| $PR_{hard15}$ | 10 | 15 | 15% | 968,202 | 6,363,214 | 67,266,919 |
| $PR_{hard40}$ | 10 | 15 | 40% | | 6,060,951 | 78,870,990 |
| $PR_{hard55}$ | 10 | 15 | 55% | | 5,976,943 | 87,502,182 |

Table III: The number of products belonging to each category in the evaluation data.

| Category | Product Counts |
|---|---|
| Underw. | 148,118 |
| Bags | 212,531 |
| Accs. | 476,229 |
| Goods | 225,799 |
| Shoes | 326,109 |



Figure 4: The numbers of positive products belong to the top 100 most frequent categories.

accessible products ('Products for Positive Pairs' in Table II) are used for training. The category distribution of products is depicted in Fig. 4.

With the same query and positive products from product collection data, we apply different negative sampling settings and category-wise data augmentation ratios to make five models. Table II depicts models with dataset configuration and statistics. The model name represents the negative sampling strategy and the ratio of category-wise data augmentation. For a single positive product, we use two types of negative samples. First, random sampling from the same categories, and second, sampling by BM25 Search using the query from product collections. We denote $_{hard}$ in the model name when involves BM25-based negative samples and $_{easy}$ for random sampling. We augment collections category-wise at a specific ratio from all product collections. We denote the augmentation percentage to the model name (e.g. $_{40}$). The model names will connect evaluation results to the corresponding dataset configuration.

### B. Metric Analysis

*1) Evaluation Data and Protocol:* The evaluation procedure for product retrieval systems is usually divided into two-phase: 1) indexing embeddings from evaluation products (ideally, same as online products) 2) making query embedding and retrieve products from evaluation indices. For queries, we used the same product collection data used for training to make queries for evaluation. For products, we used extract products with specific customer review counts

(in this case same or larger than 1) from online products. In our e-commerce platform, the number of online products in the major categories reaches tens of millions, it is not feasible to perform the evaluation on top of entire online products. Hence we used this filtering condition to make a small product volume for evaluation while maintaining extensive coverage of products across product collections. We only used product collections contains five or more evaluation products for recall evaluation. The statistics of evaluation data are depicted in Table III.

*2) Evaluation Metrics:* In a single product collection dataset, we evaluate how well trained SBERT model retrieves and restores original products by a given query (i.e. recall to measure intent conformance of retrieval result). If data is category-wise augmented, we perform an additional evaluation for how well trained SBERT model retrieves products with the same category mentioned in a given query (i.e. precision to measure categorical correctness of retrieved result). Formally speaking, For a set $P_{gt}$ of evaluation products in the specific product collections, a recall is defined as follows:

$$recall = \frac{\text{the number of retrieved products} \in P_{gt}}{\text{the number of } P_{gt}}$$

For a set $C_{gt}$ of products with ground truth category in query, a precision is defined as follows:

$$precision = \frac{\text{the number of retrieved products} \in C_{gt}}{\text{the total number of retrieved products}}$$

In measuring recall and precision, the number of retrievals performed for evaluations is different by the ratio of category-wise augmentation of each model training (0∼55% in this paper). However, we assume the relative difficulties of retrieval tasks for all model training are mostly the same because only queries used for training will be used for evaluation.

We train the model up to 400,000 steps and the evaluation interval is 20,000 steps. The training batch size is 55. We used P40 GPU for most model training. We used Faiss [41] for similarity search for evaluation.

*3) BM25 Search Baseline:* Table IV represents the recall and precision performance of the baseline BM25 Search. For query in BM25 Search model evaluation, we additionally

Table IV: Recall and precision evaluation results using BM25 Search.

| | Recall@100 - BM25 Search | | | | | | Precision@100 - BM25 Search | | | | | |
| Underw. | Bags | Accs. | Goods | Shoes | Avg. | | Underw. | Bags | Accs. | Goods | Shoes | Avg. |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 0.0378 | 0.0341 | 0.0211 | 0.0403 | 0.0273 | 0.03212 | | 0.5363 | 0.6172 | 0.7797 | 0.1662 | 0.7261 | 0.5651 |

Table V: Recall evaluation results of models at 400,000 training steps. The best score is **bold** and second best score is underlined among models with hard negatives and category-wise data augmentation ($PR_{hard15}$, $PR_{hard40}$, $PR_{hard55}$).

| Model | Recall@100 - SBERT 12L | | | | | | Recall@100 - SBERT 6L | | | | | |
| | Underw. | Bags | Accs. | Goods | Shoes | Avg. | Underw. | Bags | Accs. | Goods | Shoes | Avg. |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| $PR_{easy0}$ | 0.5141 | 0.4840 | 0.5122 | 0.5032 | 0.4156 | 0.4858 | 0.4733 | 0.4163 | 0.4920 | 0.4983 | 0.4044 | 0.4569 |
| $PR_{hard0}$ | 0.5830 | 0.4493 | 0.5776 | 0.5983 | 0.4433 | 0.5303 | 0.5505 | 0.4106 | 0.5440 | 0.5698 | 0.3786 | 0.4907 |
| $PR_{hard15}$ | **0.5847** | **0.4374** | **0.5724** | **0.5874** | **0.4331** | **0.5230** | <u>0.5859</u> | **0.4026** | **0.5535** | **0.5781** | <u>0.4088</u> | **0.5058** |
| $PR_{hard40}$ | <u>0.5816</u> | <u>0.3601</u> | <u>0.5564</u> | <u>0.5583</u> | <u>0.4192</u> | <u>0.4951</u> | <u>0.5851</u> | <u>0.3773</u> | <u>0.5340</u> | <u>0.5359</u> | **0.4151** | <u>0.4895</u> |
| $PR_{hard55}$ | 0.5659 | 0.3459 | 0.5215 | 0.5286 | 0.3930 | 0.4710 | **0.6169** | 0.3154 | 0.5227 | 0.5262 | 0.3914 | 0.4745 |

Table VI: Precision evaluation results of models at 400,000 training steps. The best score is **bold** and second best score is underlined.

| Model | Precision@100 - SBERT 12L | | | | | | Precision@100 - SBERT 6L | | | | | |
| | Underw. | Bags | Accs. | Goods | Shoes | Avg. | Underw. | Bags | Accs. | Goods | Shoes | Avg. |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| $PR_{hard15}$ | 0.2877 | 0.6101 | 0.5334 | **0.1810** | 0.4928 | 0.4210 | 0.2861 | **0.6589** | 0.5382 | **0.1821** | 0.5506 | 0.4432 |
| $PR_{hard40}$ | <u>0.4711</u> | **0.6717** | <u>0.6288</u> | 0.1547 | <u>0.6832</u> | <u>0.5219</u> | <u>0.4529</u> | <u>0.6579</u> | <u>0.6695</u> | 0.1536 | <u>0.6698</u> | <u>0.5207</u> |
| $PR_{hard55}$ | **0.5305** | <u>0.6447</u> | **0.6929** | <u>0.1715</u> | **0.7625** | **0.5604** | **0.5073** | 0.6572 | **0.7060** | <u>0.1715</u> | **0.6831** | **0.5450** |

filter uninformative words in the title (e.g. 'collection') since some of our product collections are entitled with the word 'collection'. We also exclude the date information for the query, since many collections share the same date and potentially lead to noise in retrieval results. We tokenize query and product attributes by space. All other details on evaluation are the same as SBERT models.

*4) Metric Overview:* Table V and Table VI show recall and precision performance of SBERT models in various product categories. The metrics of models at each training interval are depicted in Appendix A. Since models require optimal recall and precision performance for our task, we marked the best score and second best score only among $PR_{hard15}$, $PR_{hard40}$, $PR_{hard55}$ models (i.e. models with hard negative samples and category-wise data augmentation). We conduct experiments in SBERT 12L and 6L model, here we focused on the performance of the 12L model, a comparison between the 12L and 6L model will be discussed in Section V.

*5) Recall:* Table V depicts recall performances. $PR_{hard0}$ is the best by adopting hard negative sampling without category-wise data augmentation, nearly 9.2% recall improvements for between $PR_{easy0}$ and $PR_{hard0}$. This means hard negative sampling using the query of product collection is effective for triplet metric learning. We expect these hard negative examples not only considering product category but also considering similar but undesirable product attributes such as specific material or color differences. This helps the model to learn rich and unique product representations of the given intent query.

$PR_{hard15}$ and $PR_{hard40}$ are the second and third best performances, respectively. This means category-wise data augmentation degrades recall performance. Since category information is more common compared to the intent sentence in the dataset, adopting category information directly to query (intent) embedding forces model to create more general query representations. Recall depends on unique representations of queries, thus general query representations degrade recall performances.

The left side of Table IV depicts recall performance of baseline BM25 Search. Since BM25 Search only considers word-level matching, it fails to map long and complicated queries to rich product attributes, and introduces very low recall performance.

*6) Precision:* Table VI shows precision performance. We observe the percentage of category-wise augmented data correlated to precision performance $PR_{hard55} > PR_{hard40} > PR_{hard15}$. The best model matches product from the correct category by 56% of product collections, nearly 33.2% precision improvements between $PR_{hard15}$ and $PR_{hard55}$, which means joint learning with category-wise positive pair augmentations helps the model to consider category-based precision along with intent matching.

The right side of Table IV depicts precision performance of baseline BM25 Search. In many cases, category information is explicitly stated in product attributes, thus word-level matching can be enough to capture. The performance of BM25 Search is similar to $PR_{hard55}$, in other words,

pretrained language models successfully satisfy categorical conditions while maintaining significantly improved recall performance.

*7) Optimal Fine-Tuned Model for Production:* We choose $PR_{hard40}$ for our services because the precision score and the recall score are balanced. Even if the recall of $PR_{hard40}$ is 5.3% less than the one of $PR_{hard15}$, the precision of $PR_{hard40}$ is 24% higher than the one of $PR_{hard15}$. Compare to BM25 Search, the precision of $PR_{hard40}$ is slightly lower. However, the recall of it is outperformed significantly. Note that, in a service environment, we apply post-processing for the more improved precision performance for SBERT models.

## V. ABLATION STUDY

Along with 12L (110M) model, we conduct an ablation study with the 6L (68M) model to analyze how the number of transformer layers and model capacity affects recall and precision performance.

We choose the 12L model for our services by the reliable recall and precision performance in both intent-based and category-based product retrieval. For example, the 6L model is better on the precision performance of $PR_{hard15}$, but the 12L model is better on recall performance on a similar margin. This is meaningful since the increase in the recall is generally harder than the increase in the precision in our task in terms of retrieval search space. In addition, the 6L model shows slightly better recall performance to the 12L model on $PR_{hard55}$, but precision performance under the 12L model, which implies that adding more transformer layers is helpful for precision performance. Furthermore, for $PR_{hard40}$, both recall and precision of the 12L model is higher than those of the 6L model.

## VI. ONLINE PERFORMANCE

In our service, we collect products under certain themes and show them to customers. Currently, making the collections of products is done by dozens of professional human operators. They first think of the theme from their own knowledge, and then repeatedly use lexical search to gather related products. We apply our model to automatically retrieve relevant products from a given theme query.

In this section, we compare the online performances of manually created product collections and product collections crafted using our model. For each product collection, we compute CTR (Click Through Rate), CVR (Conversion Rate), and the order-diversity. And then we compute the relative score as follows:

$$relative\ score = \frac{\text{avg. score of collections using our model}}{\text{avg. score of collections made by experts}}$$

Here, CTR is the number of product clicks over the number of product collection views and CVR is the number of product purchases over the number of product collection views. The order-diversity, which can measure the diversity



Figure 5: Daily relative performance of our model. The x-axis means the date in June 2021. The orange dash line means the average relative CTR, CVR, and order-diversity. Overall, comparing to the product collections made by human operators, the product collections made by our model increase the CTR, CVR, order-diversity by 16%, 29%, 60%, respectively.

of purchased items, is the number of purchased products over the number of products in the product collection. Thus, we can say that if the order-diversity of a product collection is large, then it contains a lot of quality products that are good enough to be purchased among web-scale products.

Comparing to the product collections made by humans, the product collections made by our model increase the CTR, CVR, order-diversity by 16%, 29%, 60%, respectively, in June 2021. The daily relative performance is illustrated in Fig. 5. Using pretrained language models, it is possible to learn the product features and collect related products with a common intent well. For these reasons, our model consistently outperforms human operators, who are prone to make the product collection with their individual preferences

and interests.

Note that when we make a product collection, we first retrieve a bunch of products and reorder them. Human experts usually focus on popularity or the number of reviews. In our model, we run a simple linear regression model with similar features for reordering.

## VII. CONCLUSION AND FUTURE WORK

In this paper, we create an intent-based product collection using SBERT that well performs in offline and online evaluation results to enable a better shopping experience for e-commerce platform users. We adopt SBERT for triplet metric learning for creating intent-based product collections by the deep understanding of intent and related product attributes. We enhance the base model by hard negative sampling for improved intent-based recall performance and category-wise positive pair augmentation for improved category-based precision performance. We analyze effective training data construction options for fine-tuning, and ablation study to identify optimal pretrained model choices for production. Furthermore, we measure online performance on CTR, CVR, order-diversity, and conclude that our model can create better intent-based product collections compare to human operators. Currently, we are training and inferencing our model for most e-commerce categories including digital, living, or food. In the future, we will conduct more detailed experiments on other categories. We will also explore approaches for more effective model training and more personalized collections.

## ACKNOWLEDGMENT

## REFERENCES

[1] J. B. Schafer, J. A. Konstan, and J. Riedl, "E-commerce recommendation applications," *Data mining and knowledge discovery*, vol. 5, no. 1, pp. 115–153, 2001.

[2] K. Wei, J. Huang, and S. Fu, "A survey of e-commerce recommender systems," in *2007 international conference on service systems and service management*. IEEE, 2007, pp. 1–5.

[3] S. Sivapalan, A. Sadeghian, H. Rahnama, and A. M. Madni, "Recommender systems in e-commerce," in *2014 World Automation Congress (WAC)*. IEEE, 2014, pp. 179–184.

[4] K. Shin, H. Kwak, K.-M. Kim, M. Kim, Y.-J. Park, J. Jeong, and S. Jung, "One4all user representation for recommender systems in e-commerce," *arXiv preprint arXiv:2106.00573*, 2021.

[5] S. Jung, Y.-J. Park, J. Jeong, K.-M. Kim, H. Kim, M. Kim, and H. Kwak, "Global-local item embedding for temporal set prediction," in *Fifteenth ACM Conference on Recommender Systems*, ser. RecSys '21, 2021, p. 674–679.

[6] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jun. 2019, pp. 4171–4186.

[7] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30, 2017.

[8] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," 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)*, Nov. 2019, pp. 3982–3992.

[9] Q. Ai, V. Azizi, X. Chen, and Y. Zhang, "Learning heterogeneous knowledge base embeddings for explainable recommendation," *Algorithms*, vol. 11, no. 9, 2018.

[10] R. Catherine, K. Mazaitis, M. Eskenazi, and W. Cohen, "Explainable entity-based recommendations with knowledge graphs," *arXiv preprint arXiv:1707.05254*, 2017.

[11] X. L. Dong, X. He, A. Kan, X. Li, Y. Liang, J. Ma, Y. E. Xu, C. Zhang, T. Zhao, G. Blanco Saldana *et al.*, "Autoknow: Self-driving knowledge collection for products of thousands of types," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ser. KDD '20, 2020, p. 2724–2734.

[12] X. Luo, L. Liu, Y. Yang, L. Bo, Y. Cao, J. Wu, Q. Li, K. Yang, and K. Q. Zhu, "Alicoco: Alibaba e-commerce cognitive concept net," in *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, ser. SIGMOD '20, 2020, p. 313–327.

[13] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman, "Indexing by latent semantic analysis," *Journal of the American society for information science*, vol. 41, no. 6, pp. 391–407, 1990.

[14] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *the Journal of machine Learning research*, vol. 3, pp. 993–1022, 2003.

[15] D. Angelov, "Top2vec: Distributed representations of topics," *arXiv preprint arXiv:2008.09470*, 2020.

[16] M. KAYA and H. BİLGE, "Deep metric learning: A survey," *Symmetry*, vol. 11, no. 9, 2019.

[17] W. Wu, Z. Lu, and H. Li, "Learning bilinear model for matching queries and documents," *J. Mach. Learn. Res.*, vol. 14, no. 1, p. 2519–2548, Jan. 2013.

[18] B. Hu, Z. Lu, H. Li, and Q. Chen, "Convolutional neural network architectures for matching natural language sentences," in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, Eds., vol. 27, 2014.

[19] Y. Shen, X. He, J. Gao, L. Deng, and G. Mesnil, "A latent semantic model with convolutional-pooling structure for information retrieval," in *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, ser. CIKM '14, 2014, p. 101–110.

[20] H. Palangi, L. Deng, Y. Shen, J. Gao, X. He, J. Chen, X. Song, and R. Ward, "Deep sentence embedding using long short-term memory networks: Analysis and application to information retrieval," *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 24, no. 4, p. 694–707, Apr. 2016.

[21] P.-S. Huang, X. He, J. Gao, L. Deng, A. Acero, and L. Heck, "Learning deep structured semantic models for web search using clickthrough data," in *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management*, ser. CIKM '13, 2013, p. 2333–2338.

[22] J. Guo, Y. Fan, Q. Ai, and W. B. Croft, "A deep relevance matching model for ad-hoc retrieval," in *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, ser. CIKM '16, 2016, p. 55–64.

[23] K. Hui, A. Yates, K. Berberich, and G. de Melo, "PACRR: A position-aware neural IR model for relevance matching," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Sep. 2017, pp. 1049–1058.

[24] ——, "Co-pacrr: A context-aware neural ir model for ad-hoc retrieval," in *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, ser. WSDM '18, 2018, p. 279–287.

[25] ——, "Re-pacrr: A context and density-aware neural information retrieval model," *arXiv preprint arXiv:1706.10192*, 2017.

[26] B. Mitra, E. Nalisnick, N. Craswell, and R. Caruana, "A dual embedding space model for document ranking," *arXiv preprint arXiv:1602.01137*, 2016.

[27] B. Mitra, F. Diaz, and N. Craswell, "Learning to match using local and distributed representations of text for web search," in *Proceedings of the 26th International Conference on World Wide Web*, ser. WWW '17, 2017, p. 1291–1299.

[28] P. Nigam, Y. Song, V. Mohan, V. Lakshman, W. A. Ding, A. Shingavi, C. H. Teo, H. Gu, and B. Yin, "Semantic product search," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ser. KDD '19, 2019, p. 2876–2885.

[29] Y. Liu, W. Lu, S. Cheng, D. Shi, S. Wang, Z. Cheng, and D. Yin, "Pre-trained language model for web-scale retrieval in baidu search," *arXiv preprint arXiv:2106.03373*, 2021.

[30] Y. Sun, S. Wang, Y. Li, S. Feng, H. Tian, H. Wu, and H. Wang, "Ernie 2.0: A continual pre-training framework for language understanding," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, 2020, pp. 8968–8975.

[31] S. E. Robertson and S. Walker, "Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval," in *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR '94, 1994, p. 232–241.

[32] K. Roy, S. Shah, N. Pai, J. Ramtej, P. Nadkarni, J. Banerjee, P. Goyal, and S. Kumar, "Using large pretrained language models for answering user queries from product specifications," in *Proceedings of The 3rd Workshop on e-Commerce and NLP*, Jul. 2020, pp. 35–39.

[33] J. Tagliabue, F. Bianchi, and B. Yu, "BERT goes shopping: Comparing distributional models for product representations," in *Proceedings of The 4th Workshop on e-Commerce and NLP*, Aug. 2021, pp. 1–12.

[34] J. Tracz, P. I. Wójcik, K. Jasinska-Kobus, R. Belluzzo, R. Mroczkowski, and I. Gawlik, "BERT-based similarity learning for product matching," in *Proceedings of Workshop on Natural Language Processing in E-Commerce*, Dec. 2020, pp. 66–75.

[35] M. Schultz and T. Joachims, "Learning a distance metric from relative comparisons," in *Advances in Neural Information Processing Systems*, S. Thrun, L. Saul, and B. Schölkopf, Eds., vol. 16, 2004.

[36] L. Xiong, C. Xiong, Y. Li, K.-F. Tang, J. Liu, P. Bennett, J. Ahmed, and A. Overwijk, "Approximate nearest neighbor negative contrastive learning for dense text retrieval," *arXiv preprint arXiv:2007.00808*, 2020.

[37] Q. Chen, Z. Zhuo, and W. Wang, "Bert for joint intent classification and slot filling," *arXiv preprint arXiv:1902.10909*, 2019.

[38] C.-S. Wu, S. C. Hoi, R. Socher, and C. Xiong, "TOD-BERT: Pre-trained natural language understanding for task-oriented dialogue," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Nov. 2020, pp. 917–929.

[39] M. Schuster and K. Nakajima, "Japanese and korean voice search," in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2012, pp. 5149–5152.

[40] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.

[41] J. Johnson, M. Douze, and H. Jégou, "Billion-scale similarity search with gpus," *IEEE Transactions on Big Data*, pp. 1–1, 2019.

[42] B. Kim, H. Kim, S.-W. Lee, G. Lee, D. Kwak, D. H. Jeon, S. Park, S. Kim, S. Kim, D. Seo, H. Lee, M. Jeong, S. Lee, M. Kim, S. H. Ko, S. Kim, T. Park, J. Kim, S. Kang, N.-H. Ryu, K. M. Yoo, M. Chang, S. Suh, S. In, J. Park, K. Kim, H. Kim, J. Jeong, Y. G. Yeo, D. Ham, D. Park, M. Y. Lee, J. Kang, I. Kang, J.-W. Ha, W. Park, and N. Sung, "What changes can large-scale language models bring? intensive study on HyperCLOVA: Billions-scale korean generative pre-trained transformers," *arXiv preprint arXiv:2109.04650*, 2021.

Average Recall@100 (SBERT 12L)



Average Precision@100 (SBERT 12L)

Intent

Title: *Light running shoes that are comfortable to wear in summer*
Section name: *{same as title}*
Date: *Aug. 08*

| Key Product Summary for Intent | Key Product Summary for Intent |
|---|---|
| Light and mesh shoes | Comfortable to wear without laces |
| Mash and lightweight materials | Air cushion & lightweight fitness shoes |
| Socks sneakers | Dial sneakers lightweight and cool mesh |

Intent

Title: *Cool and light summer vacation dress*
Section name: *dress*
Date: *Aug. 08*

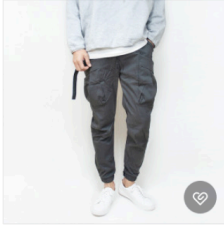| Key Product Summary for Intent | Key Product Summary for Intent |
|---|---|
| Open shoulder dress | dot pattern & loose fit |
| Wrinkles & loose fit | Cool and light cotton, big size, loose fit |
| Cotton dress big size Bright and neon colors | Big size loose fit dating look |

꾸안꾸 스타일의
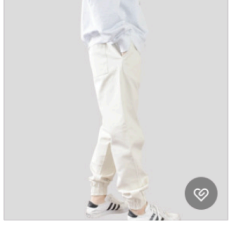남자 조거팬츠
21.09.08. ~ 21.09.21.

조거팬츠,조거바지,남자조거바지추천 19건

50% **29,900원**
빅사이즈 남자 조커팬츠 면 스판 밴딩 건빵 바지 카고 조거 팬츠...
♡ 194 | 💬 282

62% **33,380원**
빅사이즈 조거 벨크로 테크웨어 바지 냉장고 남자 카고 조커 팬...
♡ 26 | 💬 193

52% **18,000원**
(재고정리)남자 봄 가을 스판 빅사이즈 통큰 와이드 편한 조거 ...

50% **43,900원**
롱 발목 밴드 스판 카고 검정 조거 팬츠

The title is "kkuankku style men's jogger pants" in English. Here, 'kkuankku' is a Korean abbreviation that means 'effortlessly chic'. Note that the title is generated by Korean GPT-3 HyperCLOVA [42].