

Model-agnostic vs. Model-intrinsic Interpretability for Explainable Product Search

Qingyao Ai*
University of Utah
Salt Lake City, UT, USA
aiqy@cs.utah.edu

Lakshmi Narayanan.R*
University of Utah
Salt Lake City, UT, USA
lakshminarayanan.ramasamy@utah.edu

ABSTRACT

Product retrieval systems have served as the main entry for customers to discover and purchase products online. With increasing concerns on the transparency and accountability of AI systems, studies on explainable information retrieval has received more and more attention in the research community. Interestingly, in the domain of e-commerce, despite the extensive studies on explainable product recommendation, the studies of explainable product search is still in an early stage. In this paper, we study how to construct effective explainable product search by comparing model-agnostic explanation paradigms with model-intrinsic paradigms and analyzing the important factors that determine the performance of product search explanations. We propose an explainable product search model with model-intrinsic interpretability and conduct crowdsourcing to compare it with the state-of-the-art explainable product search model with model-agnostic interpretability. We observe that both paradigms have their own advantages and the effectiveness of search explanations on different properties are affected by different factors. For example, explanation fidelity is more important for user's overall satisfaction on the system while explanation novelty may be more useful in attracting user purchases. These findings could have important implications for the future studies and design of explainable product search engines.

CCS CONCEPTS

• **Information systems** → **Retrieval models and ranking**; *Retrieval tasks and goals*; *Evaluation of retrieval results*.

KEYWORDS

Product Search, Attention Mechanism, Search Explanation

ACM Reference Format:

Qingyao Ai and Lakshmi Narayanan.R. 2021. Model-agnostic vs. Model-intrinsic Interpretability for Explainable Product Search. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management (CIKM '21), November 1–5, 2021, Virtual Event, QLD, Australia*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3459637.3482276>

*Equal Contribution

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CIKM '21, November 1–5, 2021, Virtual Event, QLD, Australia

© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-8446-9/21/11...\$15.00
<https://doi.org/10.1145/3459637.3482276>

1 INTRODUCTION

As online marketplaces have gradually dominated the retail market, product retrieval systems such as product search engines and recommendation systems have become the main entry for users to discover products. Meanwhile, with increasing concerns on the transparency and accountability of AI systems, studies on explainable AI have received more attention in both academic communities and industry [18, 24]. Specifically in the domain of e-commerce information retrieval, explainability means the ability of a product retrieval system in providing explanations that allow users to understand, trust, and effectively control the retrieved products. Previous studies have shown that providing recommendation results together with explanations on why the items are retrieved not only increases the conversion rates from clicks to purchases, but also improves user's satisfaction on e-shopping websites [76]. Thus, how to improve the explainability of product retrieval systems has become an important challenge and opportunity for e-commerce.

Interestingly, despite of the extensive studies on explainable product recommendation [11, 29, 43, 58], the effectiveness and the potentials of explainable product search is mostly unexplored. As of today, more than 80% of shoppers find products starting from search online¹, which means that product search is still the most popular method to find products on e-commerce platforms. On the one hand, developing product search engines is similar to developing product recommendation systems from multiple perspectives, including the need of personalization [2], the model of heterogeneous information [72], etc. On the other hand, by explicitly formulating and feeding a query to the systems, user's requirements and expectations for product search engines are significantly different from those for product recommendation. For instance, while it is preferable to recommend PC games to a customer who recently purchased Alienware gaming laptops, it may not be a good idea when the user is searching for "running shoes". Thus, how to retrieve and explain search results based on both the explicit need and implicit preferences of product search users make explainable product search an unique challenge in explainable information retrieval.

Existing studies on explainable AI can be broadly categorized into two directions, namely *model-intrinsic* (or pre-hoc) interpretability and *model-agnostic* (or post-hoc) interpretability [38]. Model-intrinsic interpretability focuses on the construction of transparent AI systems that can explicitly explain its behavior based on its inference process. In contrast, model-agnostic interpretability focuses on explaining model outputs without knowing the internal mechanism of the model. Previous studies on explainable IR have explored both paradigms in document retrieval [21, 55, 56, 64] by creating pre-hoc

¹<https://www.retaildive.com/news/87-of-shoppers-now-begin-product-searches-online/530139/>

or post-hoc explanations with text-matching signals extracted by the retrieval models from query-document pairs. In product search, however, it has been shown that text matching is relatively less important [1, 16] comparing to other information such as knowledge entities and their relationships [26, 39] in determining user’s purchase decisions. Thus, how to create model-intrinsic/agnostic explanations in product search and how those two approaches would benefit or affect the development of explainable product search systems is mostly unknown. To the best of our knowledge, the only study on explainable product search is the Dynamic Relation Embedding Model [3] (DREM) that utilizes product knowledge graph to generate post-hoc result explanations. For evaluation, however, Ai et al. [3] simply use a survey to examine whether users are more likely to purchase after seeing the explanations and conduct no comparison of different explanation methodologies as well as possible factors that affect the effectiveness of search explanations.

To fill in this blank, we propose to construct and train an intrinsic-explainable model for product search with user-interaction data and knowledge graph. Inspired by the Zero Attention mechanism [1], we propose to extend DREM with a Hierarchical Gated Network (HGN) that explicitly construct user representations from items and knowledge entities related to the user’s purchase history. By extracting the attention weights from HGN, our proposed model is capable of generating model-intrinsic explanations for product search results. To understand the advantages and drawbacks of pre-hoc and post-hoc search explanations, we conduct a crowdsourcing study with Amazon Mechanical Turk to evaluate and analyze the performance of model-agnostic and model-intrinsic explanations generated by the original DREM and the proposed DREM with HGN. Experiment results show that model-intrinsic explanations usually could be more informative and reliable while model-agnostic explanations could have better potentials in attracting users to purchase the product. Further, we propose an explanation performance task and build models to explore the possibility of automatically evaluating search explanations without human annotations. Based on feature analysis, we find that the fidelity of search explanations could be more important for user’s overall satisfaction with the search engines while the novelty of the search explanations could be more useful in attracting users to purchase the item.

2 RELATED WORK

There are three lines of studies that are important to our work: Interpretable AI, Explainable IR and Product Search.

Interpretable AI. The research of interpretable and explainable AI is a growing topic as the concerns on transparency and accountability of AI systems have increased dramatically recently [24]. In general, existing studies on interpretable AI can be broadly categorized into two groups, i.e., the studies on explaining machine learning (ML) models based on their internal structures, and the studies on explaining model outputs by treating the ML model as a black box [38]. Examples of the first group including the examination of network neurons and layers [7, 22, 46, 53, 71], the use of attention networks [32, 63, 69], and the design of disentangled model structure and information representations [13, 17, 30, 40]. Examples of the second group including the construction of proxy models with linear classifiers [50], decision trees [52, 77], extracted

rules [4, 23], and salience map [54, 73]. Both paradigms have their own advantages and disadvantages depending on application scenarios, and the field of interpretable AI is still young with numerous new studies and approaches emerging every year [24].

Explainable Recommendation. The studies of explainable retrieval systems have drawn the attention of researchers mainly starting from the last decade. Early IR systems based on term matching are transparent and explainable in nature [49, 51, 60]. However, as more state-of-the-art retrieval systems rely on complex ML and latent representation models [25, 45], interpretability is no longer a minor problem for IR. Most existing studies on explainable IR focus on recommendation tasks [75]. For example, model-based explainable recommendation methods attempt to develop models that generate both recommendations and explanations together [14, 59, 74]. Peake and Wang [48] created post-hoc explanations based on the latent vectors in recommendation models; Zhang et al. [76] explained recommendation results with facets extracted from user reviews. Another line of explainable recommendation research focuses on analyzing the nature of user behaviors to help users better understand recommendations [6, 28, 29]. Bilgic and Mooney [11] used statistical histograms as explanations to help users understand rating distribution; Tintarew and Mashthoff [58] provided user-centered design approaches to analyze the explanation effectiveness.

Explainable Search. Search is fundamentally different from recommendation as user intents are explicitly expressed with queries. Different from explainable recommendation, the studies on explainable search mostly focuses on the domain of ad-hoc retrieval, i.e., retrieving text documents such as news articles or web pages based on user’s query. For example, Zeon Trevor et al. [21] proposes to use DeepSHAP [41] to explain the outputs of neural retrieval models; Verma and Ganguly [65] explore different sampling methods to build explanation models for a given retrieval model and proposes a couple of metrics to evaluate the explanations based on the terms in queries and documents. Unfortunately, those methods are not applicable to product search as they are purely designed for text retrieval and text matching signals are relatively unimportant [1, 16] compared to other information such as entity relationships and user purchase history in determining user’s purchases. As for how to create result explanations with heterogeneous entity and information in product search, to the best of our knowledge, the only study on this topic is proposed by Ai et al. [3] that construct a dynamic relation embedding model to incorporate product knowledge graph and use it to explain product search results. However, they only conducted a laboratory study to examine the effectiveness of the explanations generated by their model and did no comparison and study on different explanation methodologies as well as what factors are important for product search explanations.

Product Search. Early studies on product search focus on retrieving products based on structured product facets such as brands and categories [19, 20, 37]. However, as there exists a significant vocabulary gap between user queries and product descriptions [47, 61], state-of-the-art approaches usually conduct product search in latent space with deep learning techniques [10, 27, 68]. For example, Bi et al. [8, 9] extract fine-grained review information with embedding networks; Guo et al. [26] model long/short term user preferences with attention networks over user query history. There are also considerable studies on extracting ranking features and applying

learning-to-rank methods for product search [5, 15, 31, 34, 70]. In this paper, our main focus is not to build the state-of-the-art product search models but to explore how to build effective search explanations to better improve user experience.

3 METHODOLOGY

In this section, we describe our proposed method for explainable product search. We start from introducing the framework of latent product retrieval models, the structure of the state-of-the-art explainable product search model (i.e., DREM) for model-agnostic explanation, and then propose a hierarchical gated network (HGN) to extend DREM for model-intrinsic search explanations.

3.1 Latent Product Retrieval Framework

As discussed previously, the goal of product search is to retrieve products according to user’s needs so that we can maximize the average transaction rate (i.e., user purchases) in search sessions. Usually, this means ranking and showing products to users according to their probabilities to be purchased [2, 62]. Different from traditional IR tasks such as ad-hoc retrieval, information in product search is often stored in heterogeneous forms and classic retrieval models based on text matching often performs suboptimal in practice [3, 27, 47]. Therefore, the state-of-the-art methods in product search often build retrieval models in latent spaces by representing and matching queries, users, and items with latent vectors.

In general, user’s purchase decisions are affected by two factors [2, 3, 27]: (1) the explicit purchase intents in the current session, which are usually expressed by user’s queries, and (2) the implicit preferences over product properties (e.g., colors and brands), which are usually inferred from user’s historical behaviors (e.g., previous purchases). Formally, let $\mathbf{q}, \mathbf{u}, \mathbf{i} \in \mathbb{R}^\alpha$ be the α dimensional latent representations of the search query, the user’s personal preferences, and the item, respectively. Following previous studies [2, 26], we model the probability of an item i being purchased by a user u after submitting a query q with a latent generative model as

$$P(i|u, q) = \frac{\exp(\mathbf{i} \cdot \mathbf{S}_{uq})}{\sum_{i' \in I} \exp(\mathbf{i}' \cdot \mathbf{S}_{uq})}, \quad \mathbf{S}_{uq} = \mathbf{q} + \mathbf{u} \quad (1)$$

where I is the universal set of candidate items, and \mathbf{S}_{uq} is the latent representation of the user’s purchase intent in search, which could be modeled as the linear combination of \mathbf{q} and \mathbf{u} .²

Under this formulation, the representations of queries, users, and items can be directly optimized for product search by maximizing the log likelihood of observed user purchases in search defined as

$$\mathcal{L} = \log \prod_{u, q, i} P(i|u, q) = \sum_{u, q, i} (\mathbf{i} \cdot (\mathbf{q} + \mathbf{u}) - \log \sum_{i' \in I} \exp(\mathbf{i}' \cdot (\mathbf{q} + \mathbf{u}))) \quad (2)$$

While directly computing \mathcal{L} is prohibitive due to the softmax function and the large number of items in I , there are many effective and mature solutions built with approximation algorithms such as hierarchical softmax and negative sampling [44]. In this paper, we adopt the negative sampling strategy that approximate the denominator of softmax function by randomly sampling negative samples from I . Therefore, the key problem of product search in latent space is how to construct the representations of queries, users, and items.

²For simplicity, we ignore the discussions of more complicated models for \mathbf{S}_{uq} as it is not the focus of this paper.

3.2 Dynamic Relation Embedding Model and Model-agnostic Explanations

To the best of our knowledge, the first model proposed for explainable product search is the Dynamic Relation Embedding Model (DREM) [3]. In order to utilize heterogeneous data and knowledge for product search and explanations, Ai et al. [3] proposed to build a latent dynamic knowledge graph that jointly encodes the relationships between queries, users, items, as well as product-related knowledge entities. Specifically, the construction of DREM and search explanations include two parts: the modeling of entity relationships, and the extraction of explainable knowledge path between users and retrieved items.

3.2.1 Product Knowledge Graph and Query Modeling. Product search is different from product recommendation as the relevance and relationships between users and items could vary based on user’s information need expressed in the search query. To model both the static relationships between knowledge entities and the dynamic relationships between users, queries, and items, Ai et al. [3] propose to adopt the TransE models [12] for product search and treat *Search&Purchase* as a special relationship that translates users to items. Formally, let $(h, r, t) \in \mathcal{G}$ be a relation triple with head entity h , relation r , and tail entity t (e.g., *iPhone is Produced_by Apple*) in observed data \mathcal{G} . Then DREM defines a linear translation function and a latent generative model to model (h, r, t) as

$$P(t|h, r) = \frac{\exp(\mathbf{t} \cdot (\mathbf{h} + \mathbf{r}))}{\sum_{t' \in T} \exp(\mathbf{t}' \cdot \mathbf{h} + \mathbf{r})} \quad (3)$$

where T is the universal set of possible tail entity t , and $\mathbf{h}, \mathbf{r}, \mathbf{t} \in \mathbb{R}^\alpha$ are the embedding representations of the head entity, the relation, and the tail entity, respectively. In other words, *the entity h can be translated to entity t through relation r with probability $P(t|h, r)$.*

As the relationship between users and items (i.e., *Search&Purchase*) varies according to different search queries, Ai et al. [3] propose to create a dynamic relation embedding by encoding query string with a non-linear project function ϕ as

$$\mathbf{q} = \phi(\{w_q | w_q \in q\}) = \tanh(W \cdot \frac{\sum_{w_q \in q} w_q}{|q|} + b) \quad (4)$$

where w_q and $\mathbf{w}_q \in \mathbb{R}^\alpha$ are query words and their corresponding embedding representations, $W \in \mathbb{R}^{\alpha \times \alpha}$ and $b \in \mathbb{R}^\alpha$ are model parameters, and $\mathbf{q} \in \mathbb{R}^\alpha$ is the query embedding as well as the relation embedding of *Search&Purchase*. The probability of a user u searched and purchased an item i is then computed in the same way with other relation triples as shown in Eq. (3).

To optimize the embedding representations of all entities and relations for product search, DREM directly maximizes the log likelihood of all observed relation triples as

$$\begin{aligned} \mathcal{L} &= \sum_{(u, q, i)} \log P(i|u, q) + \sum_{(h, r, t) \in \mathcal{G}} \log P(t|h, r) \\ &\approx \sum_{(u, q, i)} \log \sigma((\mathbf{u} + \mathbf{q}) \cdot \mathbf{i}) + k \cdot \mathbb{E}_{i' \sim P_i} [\log \sigma(-(\mathbf{u} + \mathbf{q}) \cdot \mathbf{i}')] \\ &\quad + \sum_{(h, r, t) \in \mathcal{G}} \log \sigma((\mathbf{h} + \mathbf{r}) \cdot \mathbf{t}) + k \cdot \mathbb{E}_{t' \sim P_t} [\log \sigma(-(\mathbf{h} + \mathbf{r}) \cdot \mathbf{t}')] \end{aligned} \quad (5)$$

where $\sigma(x)$ is the sigmoid function (i.e., $\sigma(x) = \frac{1}{1+e^{-x}}$) and we apply a negative sampling strategy with sample size k . P_i is defined as a uniform item noisy distribution and P_l is defined as a frequency-based entity noisy distribution [3, 62].

3.2.2 Post-hoc Search Explanations. With the latent knowledge graph learned from observed product purchases and meta data, Ai et al. [3] argues that DREM is capable of creating post-hoc search explanations for each item retrieved for a user-query pair. Specifically, as all relations and entities are encoded in the latent space with the TransE models defined in Eq. (3), one can infer an arbitrary item from a user-query pair by finding a set of relations and intermediate entities that translate the joint representation of user and query (i.e., $S_u q$) to the item representation (i.e., i). Then, the path from the user to the item can be used to create an explanation of why the item is relevant to the user’s search intent.

Let $\{r_u^j\}$ (where r_u^0 is *Search&Purchase*) and $\{r_i^m\}$ be two sequences of relations that finally translate a user u and an item i to an entity in entity space Ω_e . DREM defines a soft matching path between u and i through $e \in \Omega_e$ with score:

$$\begin{aligned} M(e|u, i) &= \log(P(e|u, \{r_u^j\})P(e|i, \{r_i^m\})) \\ &= \log P(e|u) + \log P(e|i) \\ &= \log \frac{\exp(\mathbf{e}_u \cdot \mathbf{e} - \gamma j)}{\sum_{e' \in \Omega_e} \exp(\mathbf{e}_u \cdot \mathbf{e}')} + \log \frac{\exp(\mathbf{e}_i \cdot \mathbf{e} - \gamma m)}{\sum_{e' \in \Omega_e} \exp(\mathbf{e}_i \cdot \mathbf{e}')} \end{aligned} \quad (6)$$

where $\mathbf{e}_u = \mathbf{u} + \sum_j r_u^j$, $\mathbf{e}_i = \mathbf{i} + \sum_m r_i^m$, and γ is a hyper-parameter³. While the soft matching score of a path does not have any meanings to users, Ai et al. [3] argues that it indicates the model’s confidence on the path. Thus, they sort all potential inference path from user-query pair to a target item with the soft matching score and directly create a search explanation using simple templates and the relations/entities on the path to explain why i is retrieved for u by q . For example, suppose that there is a path from user u to Apple Pencil with query “*tablet*” and relation *Brought_Together*, then we can explain why we retrieved Apple Pencil with a post-hoc explanation as “Apple Pencil is retrieved because it is frequently *Brought_Together* with products retrieved by query “*tablet*””.

3.3 Hierarchical Gated Network and Model-intrinsic Explanations

While DREM can provide post-hoc explanations to search results with inference paths on knowledge graph, the retrieval process of the model is simply ranking items according to the dot product between user-query pair and item representation in the latent space, which are not necessarily correlated to the generated explanations. In practice, we may prefer a transparent retrieval model that could provide direct explanations to its inference process for many reasons such as model reliability and result accountability [38]. Inspired by the Zero Attention Mechanism (ZAM) [1], in this paper, we propose an extension to DREM to enhance its interpretability and enable it to provide model-intrinsic search explanations.

3.3.1 Attention Network with Gates. ZAM is first proposed to conduct selective personalization in product search [1]. The idea of ZAM is to relax the assumption of traditional attention mechanism

by allowing the network to attend none input data when the query is not relevant to any input vectors. Let \mathbf{q} be the query vector and X be the input vectors of an attention network, then ZAM computes the output \mathbf{y} by attending \mathbf{q} to both X and a zero vector $\mathbf{0}$ as

$$\mathbf{y} = \sum_{x \in X} \frac{\exp(f(\mathbf{q}, \mathbf{x}))}{\exp(f(\mathbf{q}, \mathbf{0})) + \sum_{x' \in X} \exp(f(\mathbf{q}, \mathbf{x}'))} \mathbf{x} \quad (7)$$

where $\mathbf{0}$ is a vector with all elements equal to 0, and $f(\mathbf{q}, \mathbf{x})$ is the attention function that computes the attention score of x with q .

By adding $\mathbf{0}$ to the attention network, ZAM naturally creates a gate that controls whether the output vector of the attention network would be fed into downstream applications or not. Let \mathbf{a}_X be the vector of $\{f(\mathbf{q}|\mathbf{x})|x \in X\}$, then ZAM can be reformulated as

$$\mathbf{y} = \frac{\exp(\mathbf{a}_X)}{\exp(f(\mathbf{q}, \mathbf{0})) + \exp^+(\mathbf{a}_X)} \cdot X \quad (8)$$

where $\exp^+(\mathbf{a}_X)$ is the element-wise sum of $\exp(\mathbf{a}_X)$. Thus, the output \mathbf{y} would be influenced by the input X only when the aggregated attention of X is significantly larger than a threshold $f(\mathbf{q}, \mathbf{0})$.

3.3.2 User Modeling in Hierarchy. We now describe how we extend DREM to a transparent product search model with the idea of ZAM. The construction of a interpretable retrieval model with heterogeneous product knowledge involves two questions: (1) how to model user preferences in a specific knowledge domain according to the current search query, and (2) how to jointly combine user preferences in each knowledge domain to retrieve items for the current search query. To solve these questions, we proposed to build a Hierarchical Gated Network (HGN) to model user preferences in search with knowledge entities associated to user’s purchase history. A illustration of the DREM with HGN is shown in Figure 1.

Formally, let Ω_e^u be the set of entities with type e associated to a user u . For example, Ω_e^u could be the items or brands purchased by u . For each knowledge domain e , we compute a latent embedding \mathbf{u}_e for u in e by attending each entity with the current query q as

$$\mathbf{u}_e = \sum_{e \in \Omega_e^u} \frac{\exp(f_e(\mathbf{q}, \mathbf{e}))}{\exp(f_e(\mathbf{q}, \mathbf{0})) + \sum_{e' \in \Omega_e^u} \exp(f_e(\mathbf{q}, \mathbf{e}'))} \mathbf{e} \quad (9)$$

where $f_e(\mathbf{q}, \mathbf{e})$ is a simple attention function defined as

$$f(\mathbf{q}, \mathbf{e}) = (\mathbf{e} \cdot \tanh(\mathbf{W}_e^f \cdot \mathbf{q} + \mathbf{b}_e)) \cdot \mathbf{W}_e^h \quad (10)$$

where $\mathbf{W}_e^h \in \mathbb{R}^\beta$, $\mathbf{W}_e^f \in \mathbb{R}^{\alpha \times \beta \times \alpha}$, $\mathbf{b}_e \in \mathbb{R}^{\alpha \times \beta}$, and β is a hyper-parameter that controls the number of the attention heads.

To further aggregate \mathbf{u}_e from each domain to create the final user embedding \mathbf{u} , we apply another layer of zero attention above all knowledge domains. Let $\Omega^u = \{\Omega_e^u\}$, then

$$\mathbf{u} = \sum_{\Omega_e^u \in \Omega^u} \frac{\exp(f_u(\mathbf{q}, \mathbf{u}_e))}{\exp(f_u(\mathbf{q}, \mathbf{0})) + \sum_{\Omega_e^u \in \Omega^u} \exp(f_u(\mathbf{q}, \mathbf{u}_e))} \mathbf{u}_e \quad (11)$$

where $f_u(\mathbf{q}, \mathbf{u}_e)$ is another attention function with similar form of Eq. (10) but different set of parameters.

Intuitively, the idea of HGN is to construct a hierarchical zero attention network that aggregates fine-grained user preferences from each knowledge domain to a unified user vector based on the search query. For parameter optimization, we simply follow the methodology of DREM introduced in Section 3.2.1 and replace the original user vector \mathbf{u} with the new user vector constructed by

³Please refer to the original paper [3] for more details.

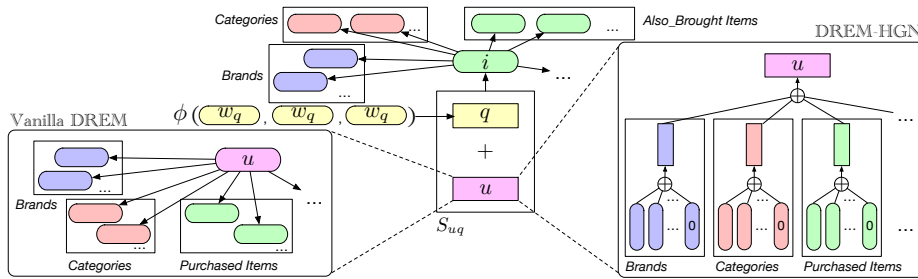


Figure 1: An illustration of the vanilla DREM and DREM with HGN. Different types of entities are colored differently. Squashed rectangles are vectors randomly initialized and learned in training, and rectangles are vectors encoded from other vectors.

HGN. Through this way, we can easily track down the usage of each knowledge entity in product search and create a higher-level of transparency and interpretability to DREM.

3.3.3 Pre-hoc Search Explanations. The advantage of HGN-based DREM is its ability to create model-intrinsic explanations. Attention network is explainable by nature as the importance of input data is directly reflected by their attention weights in model outputs. With the help of HGN, we can not only infer the importance of each user-associated entity in building the final retrieval model (i.e., S_{uq} in Eq. (1)), but also distinguish how much utility is obtained from understanding user’s preferences over retrieved items or the general relevance/popularity between items and search queries.

For pre-hoc search explanations, the attention weights in HGN can be split into two parts. The first part is the attention score of each user-associated knowledge domain and entity. Let $A_{\Omega_e^u}^e$ be the attention weight that entity e received within domain Ω_e^u , and $A_{\Omega_e^u}^u$ be the attention weight that domain Ω_e^u received in search, then

$$\begin{aligned} A_{\Omega_e^u}^e &= \frac{\exp(f_e(q, e))}{\exp(f_e(q, 0)) + \sum_{e' \in \Omega_e^u} \exp(f_e(q, e'))} \\ A_{\Omega_e^u}^u &= \frac{\exp(f_u(q, \Omega_e^u))}{\exp(f_u(q, 0)) + \sum_{\Omega_e^u \in \Omega^u} \exp(f_u(q, \Omega_e^u))} \end{aligned} \quad (12)$$

Intuitively, $A_{\Omega_e^u}^u$ is the importance of domain Ω_e^u in building the final user model u , and $A_{\Omega_e^u}^e$ is the importance of entity e in the domain. To explain the behavior of the product search model with these information, we can adopt simple templates to generate user readable search explanations with the attention weights. For example, for a specific user-query pair, if the attention weight of the domain *Brand* is 0.5, and *Apple* is the entity that received the highest attention within *Brand*, we can generate a pre-hoc search explanations as “this product are retrieved 50% because of the *Brand* of products previously purchased by the user, such as *Apple*”.

The second part of the attention weights in HGN is the attention on the zero vector. As depicted in Figure 1, HGN allows the model to attend to a zero vector when aggregating the information extracted from each knowledge domain with weight A_0^u as

$$A_0^u = \frac{\exp(f_u(q, 0))}{\exp(f_u(q, 0)) + \sum_{\Omega_e^u \in \Omega^u} \exp(f_u(q, \Omega_e^u))} \quad (13)$$

Particularly, we apply a negative sampling strategy to maximize the probability of observed purchase $P(i|u, q)$, which has been proven to be equivalent to factorizing the pointwise mutual information between S_{uq} and i [36]. According to Eq. (1) and (13), when A_0^u is close to 1, S_{uq} would downgrade to q and the final retrieval

model is essentially retrieving items according to their mutual information with the query. From this perspective, the weight of the zero vector in HGN can be seen an indicator of the importance of item popularity under the query in the generation of the final ranked list. Therefore, given a particular A_0^u , we could explain the results retrieved by HGN as “this product is retrieved $A_0^u\%$ because of its popularity under the query”.

4 RETRIEVAL EXPERIMENTS

In general, the evaluation of an explainable product search model involves two parts: (1) the evaluation of retrieval performance in terms of retrieving items that are most likely to be purchased by users, and (2) the evaluation of explanation effectiveness in terms of illustrating the connections between users, queries, and retrieved items as well as increasing the conversion rates from search to purchase. In this section, we focus on the first part and introduce our settings and results in retrieval experiments.

4.1 Experimental Setup

The goal of retrieval experiments is to evaluate the effectiveness of product search models in retrieving relevant items for user-query pairs. To this end, we conduct experiments on a well-established product search dataset and implement a couple of state-of-the-art baselines to analyze the performance of DREM with HGN.

4.1.1 Dataset. Our testbed is a well-established Amazon product search benchmark datasets [2, 26, 62]. The dataset contains user’s purchases, reviews, and queries in a variety of categories as well as detailed descriptions and meta data of a large number of items on Amazon⁴. Specifically, we conduct experiments on three categories, i.e., *Electronics*, *Health&PersonalCare*, and *Sports&Outdoors*, and use the 5-core data where each user/item has at least 5 reviews [2, 26].

Other than review text, to incorporate rich product meta data for product search, we also consider five types of entity relationships in our experiments. They are *Also_bought*: users who purchased item i_1 has also purchased item i_2 ($i_1 \rightarrow i_2$); *Also_viewed*: users who viewed item i_1 also viewed item i_2 ($i_1 \rightarrow i_2$); *Bought_together*: item i_1 was purchased together with item i_2 in a single transaction ($i_1 \rightarrow i_2$); *Brand*: item i has brand b ($i \rightarrow b$); and *Category*: item i has category c ($i \rightarrow c$). More data statistics can be found in Table 1.

4.1.2 Baselines. Other than the naive DREM proposed by Ai et al. [3], in our experiments, we include six state-of-the-art product search baselines including classic retrieval models such as

⁴Please refer to [2, 26, 62] for the details of Amazon search datasets.

Table 1: Statistics for the 5-core data.

	<i>Electronics</i>	<i>Health&PersonalCare</i>	<i>Sports&Outdoors</i>
Vocabulary size	142,922	38,772	32,386
Number of reviews	1,689,188	346,355	296,337
Number of users	192,403	38,609	35,598
Number of items	63,001	18,534	18,357
Number of brands	3,525	3,855	2,412
Number of categories	983	861	1,443
<i>Also_bought</i> per item	36.70 \pm 38.56	63.03 \pm 35.36	75.18 \pm 31.98
<i>Also_viewed</i> per item	4.36 \pm 9.44	15.43 \pm 9.35	14.46 \pm 12.24
<i>Bought_together</i> per item	0.59 \pm 0.72	0.86 \pm 0.77	0.83 \pm 0.76
<i>Brand</i> per item	0.47 \pm 0.50	0.76 \pm 0.43	0.67 \pm 0.47
<i>Category</i> per item	4.39 \pm 0.95	4.20 \pm 0.93	4.82 \pm 1.33
Number of reviews	1,275,432/413,756	261,281/85,074	224,807/71,530
Number of user-query pairs	1,204,928/5,505	232,187/207	214,919/1,739
Relevant items per pair	1.12 \pm 0.48/1.01 \pm 0.09	1.13 \pm 0.47/1.00 \pm 0.00	1.12 \pm 0.45/1.01 \pm 0.13

- **QL**: the query-likelihood model [49] that ranks items according to the log likelihood of queries in the unigram language model built with item descriptions and reviews.
- **BM25**: the classic probabilistic model proposed by Robertson and Walker [51] built on the item’s descriptions and reviews.
- **LTR**⁵: a learning-to-rank model built with LambdaMART.

and latent product search baselines such as

- **LSE**: the Latent Semantic Entity model [62] that ranks items based on the similarity of queries and items in latent spaces.
- **HEM**: the Hierarchical Embedding Model [2] that personalizes product search results with a latent retrieval framework.
- **ZAM**: the original Zero Attention Model [1] that conducts selective personalization in product search.

4.1.3 Implementation and Evaluation Details. Following previous studies [3, 26], we partition the data in each product category by randomly hiding 30% user purchases from the training process and use them as the testing data. We randomly select 30% queries as the test queries and match users with queries extracted from their purchase history to form training and testing user-query pairs. An item is considered relevant to a user-query pair when it is relevant to the query and has been purchased by the user. More information about our data partition can be found in Table 1.

For implementation details, we follow the settings proposed by Ai et al. [3] by building QL and BM25 with galago⁶, building LTR with ranklib⁷, and tuning the Dirichlet smoothing parameter μ in QL from 1000 to 3000, the scoring parameter k and b in BM25 from 0.5 to 4 and 0.25 to 1, respectively. The number of trees and leaf in LambdaMART model used in LTR are set as 1000 and 10, and we tune the learning rate from 0.01 to 0.1. For latent product retrieval models such as LSE, HEM, ZAM, the vanilla DREM [3] and our extended DREM with HGN (DREM-HGN), we use Adagrad [42] with batch size 64 to optimize the latent vectors and set the sample size of negative sampling as 5. We clipped batch gradients with norm 5 to avoid unstable updates and train each model for 20 epochs by gradually decrease the learning rate from 0.5 to 0 (note that most models converge after 10 epochs). For fair comparison, we fixed the personalization weight η in HEM, ZAM, DREM, and DREM-HGN as 0.5 (which results in Eq. (1)) and the size of all latent vectors as 100 (i.e., $\alpha = 100$). We acknowledge that having larger vector size could boost the performance of some latent product search models,

⁵We extract ranking features for LTR following the same method used by Ai et al. [3], which is ignored in this paper due to page limit.

⁶<https://sourceforge.net/p/lemur/wiki/Galago/>

⁷<https://sourceforge.net/p/lemur/wiki/RankLib/>

especially those using rich product knowledge and information (e.g., DREM and DREM-HGN) [3]. However, this is not the focus of this paper and we ignore the tuning of α so that the embeddings learned by different models have comparable dimensionalities.

We adopt mean average precision (MAP), mean reciprocal rank (MRR) and normalized discounted cumulative gain (NDCG) to evaluate the performance of product search models. For each user-query pair, we retrieve 100 items among all candidate items in each dataset to generate the rank list and compute MAP and MRR accordingly. We also report NDCG with cutoff 10 and 50. Significant tests are measured by the Fisher randomization test [57] with $p < 0.05$.

4.2 Retrieval Results

The results of our retrieval experiments are shown in Table 2. Similar to previous studies [2, 62], we observed that latent product search models usually perform better than classic retrieval models constructed based on text matching signals. For example, our best latent product search baseline (i.e., the vanilla DREM) has significantly outperformed QL and BM25 on all the datasets. This is reasonable as previous studies have observed significant vocabulary gap between queries and item descriptions [47, 62], and users often purchase items that “seems” irrelevant to their submitted query in text [16]. After incorporating more complex behavior features such as item popularity, the LTR baseline has managed to outperform DREM on *Sports&Outdoors*, but still performs worse than DREM on *Electronics* and *Health&PersonalCare*.

Among all latent product search models, the non-personalized baseline (i.e., LSE) performs the worst, which demonstrates the importance of personalization in product search. In the results of personalized product search models, we observed that DREM has significantly outperformed other baselines with large improvements from 25% to 50%. This indicates that incorporating rich information from product knowledge graph and meta data is indeed helpful in improving the effectiveness of product search. Further, our proposed model (i.e., DREM-HGN) has achieved the best performance in our experiments. It has outperformed all the baselines significantly and achieved 5.6%, 53.6%, and 28.2% MRR improvements over the vanilla DREM on *Electronics*, *Health&PersonalCare*, and *Sports&Outdoors*, respectively. This indicates that user representations encoded from product knowledge with the hierarchical gated network is more useful than the user embedding learned by DREM from randomly initialized vectors.

While the results of DREM-HGN outperforming the vanilla DREM is not surprising as the former has incorporated a complicated attention network to model query-specific user preferences, the main advantage of HGN is its transparency and model-intrinsic interpretability that allows the generation of pre-hoc explanation for product search. To compare the post-hoc and pre-hoc search explanations generated by DREM and DREM-HGN in detail, we further conduct a series of explanation evaluation and analysis.

5 EXPLANATION EVALUATION

In practice, both pre-hoc and post-hoc explanations have their unique advantages for IR. For example, pre-hoc search explanations are considered more reliable as they are directly inferred from the

Table 2: The retrieval performance of product search models. Best performance in each model group are highlighted in bold. * and † denote significant improvements over the best classic retrieval baselines and latent search baselines, respectively.

Model	Electronics				Health&PersonalCare				Sports&Outdoors			
	MAP	MRR	NDCG@10	NDCG@50	MAP	MRR	NDCG@10	NDCG@50	MAP	MRR	NDCG@10	NDCG@50
QL	0.166	0.164	0.187	0.210	0.063	0.063	0.059	0.097	0.068	0.067	0.080	0.117
BM25	0.216	0.213	0.227	0.270	0.076	0.076	0.088	0.131	0.092	0.091	0.097	0.142
LTR	0.216	0.216	0.230	0.303	0.060	0.060	0.055	0.113	0.109 [†]	0.109 [†]	0.120 [†]	0.166
LSE	0.108	0.108	0.137	0.183	0.016	0.016	0.000	0.113	0.015	0.015	0.019	0.040
HEM	0.156	0.156	0.182	0.197	0.157 [*]	0.157 [*]	0.146 [*]	0.200 [*]	0.075	0.075	0.086	0.128
ZAM	0.115	0.115	0.130	0.162	0.208 [*]	0.208 [*]	0.244 [*]	0.263 [*]	0.074	0.075	0.087	0.153
Vanilla DREM	0.231 [*]	0.232 [*]	0.268 [*]	0.314 [*]	0.349 [*]	0.349 [*]	0.378 [*]	0.429 [*]	0.099	0.099	0.113	0.180 [*]
DREM-HGN	0.244 ^{*†}	0.245 ^{*†}	0.275 ^{*†}	0.339 ^{*†}	0.536 ^{*†}	0.536 ^{*†}	0.556 ^{*†}	0.588 ^{*†}	0.126 ^{*†}	0.127 ^{*†}	0.141 ^{*†}	0.215 ^{*†}

structure of the retrieval model. In contrast, post-hoc search explanations are more flexible as it neither enforces the retrieval model to have intrinsic interpretability nor requires access to the internal structure and data flow of the model. To the best of our knowledge, DREM is the only explainable models for product search in the literature. Ai et al. [3] has conducted a laboratory user study and show that the post-hoc search explanations extracted by DREM are useful in attracting users to purchase the items. However, how this is achieved or what factors are important for the quality of search explanations are mostly unexplored. In this section, we conduct experiments to evaluate and compare the pre-hoc explanations created by DREM-HGN with the post-hoc explanations created by the vanilla DREM for product search. Specifically, we want to study and shed some lights on the following research questions:

RQ1: Which types of search explanations do users prefer? Model-intrinsic ones or model-agnostic ones?

RQ2: What factors are important for product search explanations?

5.1 Experimental Setup

We design and conduct a crowdsourcing experiment on Amazon Mechanical Turk⁸ (AMT) to evaluate the pre-hoc and post-hoc search explanations generated by DREM-HGN and DREM.

5.1.1 Explanation Generation. As described in Section 3, both DREM and DREM-HGN create explanations with templates using knowledge entities and relations extracted by the models. For fair comparison, we adopt a single set of explanation templates for both DREM and DREM-HGN. For example, given a specific relation triple such as (item, Brand, Apple) extracted by the vanilla DREM, we would create an explanation as “This product was retrieved because the user often buys products with **brands** such as **Apple**”. Also, as DREM-HGN relies on the attention weights extracted from HGN to explain its behavior, we add the corresponding information in the template and create explanations such as “This product was retrieved **50%** because the user often buys products with **brands** such as **Apple**”.

To avoid the randomness in search explanations and to improve the robustness of crowdsourcing, we increase the redundancy of our experiment by allowing each model to provide a group of explanations instead of a single one. Specifically, we extracted and grouped the top-3 search explanations extracted by DREM and DREM-HGN and allow each explanation to include at most 3 relevant knowledge entities. Instead of requiring each crowdsourcing worker to annotate each search explanation, we let workers to annotate the search explanations in groups so that the final results would be

influenced less by the quality variance of explanations provided by each model. We refer to the group of explanations generated by DREM and DREM-HGN as Model-Agnostic Explanation (MAE) and Model-Intrinsic Explanation (MIE), respectively.

5.1.2 Annotation Strategy. Previous studies [3, 67, 74] evaluated product recommendation and search explanations mainly from three perspectives: (1) whether the explanation has provided more relevant information about the item and the query, or *Informativeness*; (2) whether the explanation is useful in attracting the user to purchase the item, or *Usefulness*; and (3) whether providing the explanation would increase user’s satisfaction for the service provided by the product search engine, or *Satisfaction*.

In this paper, we adopt the same strategy to evaluate the performance of search explanations. However, instead of requiring crowdsourcing workers to directly annotate each explanation with a 5-level score [67], we propose to conduct pairwise comparisons for the explanations generated by DREM and DREM-HGN for each user-query-item triple and let the workers to annotate their pairwise preferences only. Pairwise preferences have been proven to be much more robust and reliable comparing to pointwise relevance judgements in IR [33]. Through this way, we hope to improve the quality of our crowdsourcing experiments as well as exploring the possibility of building automatic search explanation evaluation models for product search, which is further discussed in Section 5.3.

5.1.3 Data Sampling. Our crowdsourcing dataset is sampled from the retrieval experiment dataset of *Electronics*. *Electronics* is one of the most popular product categories on Amazon. Products in *Electronics* usually have less complicated knowledge structures (e.g., less entity relations per item as shown in Table 1) and are more familiar to workers on AMT. Specifically, we randomly sampled 101 user-query pairs from the test data of *Electronics* where both DREM and DREM-HGN achieved MRR scores greater or equal to 0.1. For fairness, we extracted the user-query-item triples to explain by pairing user-query pairs with the item purchased by the user in the corresponding session. Thus, all sampled items are indeed purchased by the user and AMT workers only need to judge which explanations can better explain the user’s purchase in the search session. Specifically, we recruited three workers per case, and applied a voting process to assign the final labels.

5.1.4 UI Design. Figure 2 provides an illustration of the UI we used for the crowdsourcing experiments. On the top of the UI, we provided a variety of information related to the current item, including product links, images, titles, descriptions, the search query, and the recent purchases and reviews of the current user. In the

⁸<https://requester.mturk.com/>



Figure 2: An illustration of the crowdsourcing UI.

center of the UI, we implemented a tab-based frame that allows workers to navigate and annotate the informativeness, usefulness, and satisfaction of search explanations. In each tab, we provided and asked workers to read a detailed instruction on the annotation process on the left, and show the groups of explanations created by DREM (i.e., MAE) and DREM-HGN (i.e., MIE) in the middle. We also provided the original links on Amazon to all items and entities shown in the product descriptions, user reviews, or search explanations. To avoid unnecessary biases in the annotation process, we anonymized MAE and MIE by randomly assigning them as “Group A” and “Group B”. Workers only need to click the buttons on the right to indicate which explanation group provide better search explanations: “Group A”, “Group B”, both, or none. We eyeballed the collected data and manually filtered out workers with unreasonable behaviors. The source code of our models, experiment platforms, and all crowdsourcing data can be found in links below^{9,10}.

5.2 Crowdsourcing Results

To answer RQ1, we show the results of our crowdsourcing experiment in Table 3. As shown in table, most workers found that the model-intrinsic explanations provided by DREM-HGN (i.e., MIE) are preferable over the model-agnostic explanations provided by DREM (i.e., MAE) from the perspectives of *Informativeness* and *Satisfaction*. This is not surprising as MIE provides more information about the actual inference process of the retrieval model (e.g., the attention weights), which makes it more reliable and trustworthy to users. However, in terms of *Usefulness*, we do not observe any significant differences between MIE and MAE. The overall scores of MAE on *Usefulness* is slightly higher than those for MIE. In fact, the Fleiss Kappa κ [66] of binary classification (MIE is better or not) on *Usefulness* is -0.03 , which is much lower than those for *Informativeness* (0.1) and *Satisfaction* (0.11). One possible reason is that *Usefulness* – whether the explanations can attract the user to purchase the item – is a subjective question which varies significantly based on user’s preferences and worker’s opinions. In contrast, the questions of whether the explanations provide more information (i.e., *Informativeness*) or increase user’s satisfaction on the product search service (i.e., *Satisfaction*) are objective in spite of whether the user purchases the item or not.

To analyze the relation between *Informativeness*, *Usefulness* and *Satisfaction*, we compute the Pearson Coefficient for each pair of

Table 3: Crowdsourcing results for explanation evaluation.

	<i>Informativeness</i> ($\kappa = 0.10$)	<i>Usefulness</i> ($\kappa = -0.03$)	<i>Satisfaction</i> ($\kappa = 0.11$)
MIE wins	55%	43%	51%
MAE wins	37%	47%	35%
Equal	8%	11%	14%

Table 4: Descriptions of search explanation features.

Performance

MRR: the MRR of the retrieval model.

log_purchase_prob: $P(i|u, q)$ in DREM or DREM-HGN.

Fidelity

exist_confidence: model confidence on the existence of the relations/entities used in explanations ($M(e|u, i)$ for MAE or 1 for MIE).
existence_rate: percentage of relations/entities used in explanations that are actually observed in the dataset.

Novelty

entity_iuf: inverse user frequency of entities used in explanations.
entity_iif: inverse item frequency of entities used in explanations.
user_entity_mutual_info: mutual information between users and entities used in explanations in observed data.
item_entity_mutual_info: mutual information between items and entities used in explanations in observed data.
relation_info_entropy: the entropy of the distribution of users who have the relations used in explanations.

labels. The coefficients are 0.483 for (*Informativeness*, *Usefulness*), 0.457 for (*Usefulness*, *Satisfaction*), and 0.494 for (*Informativeness*, *Satisfaction*). Interestingly, we observe that the coefficient between *Usefulness* and *Satisfaction* is the lowest among all pairs. This may indicate that user’s satisfaction on search engines and search explanations is not directly related to whether the explanations encourage the purchases of the items. Even when a user decide not to purchase an item after seeing the explanations, they may still feel satisfied if the explanations have helped them make more informed decisions. The high coefficient between *Informativeness* and *Satisfaction* can also serve as a side evidence for this phenomenon.

5.3 Performance Prediction and Analysis

To answer RQ2, we propose a performance prediction task for explainable product search by training a classification model to infer user’s preferences over search explanations. By creating such performance prediction models, we want to explore the possibility of evaluating product search explanations without involving human in the loop and analyze the importance of different explanation properties with respect to the effectiveness of search explanations.

5.3.1 Feature Design. The goal of the classification model is to predict user’s preferences over an arbitrary pair of result explanations in product search. To this end, we extract three groups of features to represent each search explanation in the feature space. They are (1) *Performance* features, which indicate the retrieval effectiveness of the explainable search model that creates the explanation; (2) *Fidelity* features, which indicate whether the information provided by the explanation is correct or trustworthy; and (3) *Novelty* features, which indicate whether the information shown in the explanation is novel or surprising to the user, query, or item. For each feature, we compute the maximum, minimum, and mean scores of the entities in each explanation and concatenate the features of all explanations in each group to form the feature vector of the group. In total, we have 59 features for each explanation group. Detailed information about each feature can be found in Table 4.

⁹<https://github.com/utahIRlab/AMTurk-Product-Search-Explanation-Evaluation>

¹⁰<https://github.com/utahIRlab/drem-attention>




Figure 3: Feature Importance in the GBDT model for explanation performance prediction.

Table 5: Explanation performance prediction results.

	Total	Correct	Type-1-error	Type-2-error
<i>Informativeness</i>	202	125	16	61
<i>Usefulness</i>	202	127	24	51
<i>Satisfaction</i>	202	127	28	47

To form the feature vector of a pair of explanation groups and avoid introducing biases to the experiment, we concatenated the features of MIE and MAE in both forward and backward orders and create two data points in the classification task for each user-query-item triple. Therefore, we have 202 pair of input data and corresponding labels in the performance prediction task.

5.3.2 Experiment Setup. We build the performance prediction model with GBDT [35] in LightGBM¹¹. We conducted a 5-fold cross validation to predict the pairwise preference of *Informativeness*, *Usefulness* and *Satisfaction*. For each GBDT, we tuned the maximum tree depth from 5 to 20, leaves number from 10 to 30, minimum leaf data from 10 to 50, and learning rate from 0.1 to 0.5. The final results are aggregated over all the test folds in cross validation.

5.3.3 Prediction Results. Table 5 depicts the results of our explanation performance prediction experiment. *Correct* represents the pairs of explanations where the model has correctly predicted their pairwise preferences; *Type-1-error* refers to the pairs where MIE and MAE are equally good while the model predicted that one is better than the other; and *Type-2-error* are cases where MIE or MAE is better than the other while the model predicted otherwise.

Can we predict the performance of product search explanations without human annotations? The overall preference prediction accuracy is 61.9% for *Informativeness* and 62.9% for *Usefulness* and *Satisfaction*. If we allow Type-1-errors, the accuracy would further increase to 69.8%, 74.8%, and 76.7%, respectively. While such results are far from perfect, they are much better than those produced by a random model and show that it is possible infer user’s preferences on product search explanations from their feature representations. This could serve as an evidence to the potential of automatic product search explanation evaluation in future studies.

What properties are important for effective product search explanations? Intuitively, MIE tends to have higher fidelity because it is directly inferred from the model’s internal structure, and MAE may have higher novelty as DREM could extract unobserved entity relations based on soft matching [3]. As MIE is better than MAE on *Informativeness* and *Satisfaction* while MAE is slightly better than MIE on *Usefulness*, one may expect that fidelity would be more important for the *Informativeness* and *Satisfaction* of product search explanation while novelty may be more important for *Usefulness*.

To examine this hypothesis, we plot the aggregated feature importance in the GBDTs from cross validation based on the total gains of splits which use the feature. As shown in Figure 3, we observe that fidelity features such as *exist_confidence* are the most important features for *Informativeness* and *Satisfaction*, which indicates that the reliability of result explanations are important for user’s overall satisfaction with the explainable product search engine. In contrast, novelty features such as the inverse user/item frequency and the mutual information between user/item and knowledge entities used in explanations are more important for the prediction of *Usefulness*. From this perspective, it seems that users are more likely to purchase an item if the search engine could provide some interesting and novel explanations to why it retrieves the corresponding item.

Also, in our experiments, we observe that performance features (i.e., MRR and *log_purchase_prob*) have shown no or minor effect on user’s preference over MIE or MAE. This may because we have filtered out test cases where DREM and DREM-HGN have significant retrieval performance differences (e.g., we only sampled cases where both DREM and DREM-HGN have $MRR \geq 0.1$ for crowd-sourcing), but it may also indicate that users are not sensitive to the retrieval performance of a product search engine when judging the quality of search explanations. In other words, building explainable product search models that provide effective explanations requires us to rethink of the model design from different perspectives but not simply focusing on the optimization of retrieval performance.

6 CONCLUSION AND FUTURE WORK

In this paper, we present the first study that compares model-intrinsic and model-agnostic explanations for explainable product search. Specifically, we propose a hierarchical gated network as an extension to the state-of-the-art explainable product search model (i.e., DREM), and then conduct a series of experiments to compare and analyze the effectiveness of the post-hoc and pre-hoc search explanations generated by the vanilla DREM and DREM with HGN.

We acknowledge that there are still many limitations of this study such as the template-based explanation generation, the systematic bias introduced by UI and the differences between AMT workers and real product search users. In future, we will seek for opportunities of online experiments with real product search engines to further analyze the effectiveness of product search explanations and validate the observations in this paper.

7 ACKNOWLEDGMENTS

This work was supported in part by the School of Computing, University of Utah and in part by NSF IIS-2007398. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

¹¹<https://github.com/microsoft/LightGBM>

REFERENCES

- [1] Qingyao Ai, Daniel Hill, Vishy Vishwanathan, and W Bruce Croft. 2019. A Zero Attention Model for Personalized Product Search. In *Proceedings of the 28th ACM international conference on information and knowledge management*. ACM.
- [2] Qingyao Ai, Yongfeng Zhang, Keping Bi, Xu Chen, and W Bruce Croft. 2017. Learning a hierarchical embedding model for personalized product search. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 645–654.
- [3] Qingyao Ai, Yongfeng Zhang, Keping Bi, and W Bruce Croft. 2019. Explainable Product Search with a Dynamic Relation Embedding Model. *ACM Transactions on Information Systems (TOIS)* (2019).
- [4] Robert Andrews, Joachim Diederich, and Alan B Tickle. 1995. Survey and critique of techniques for extracting rules from trained artificial neural networks. *Knowledge-based systems* 8, 6 (1995), 373–389.
- [5] Kamelia Aryafar, Devin Guillory, and Liangjie Hong. 2017. An ensemble-based approach to click-through rate prediction for promoted listings at Etsy. In *Proceedings of the ADKDD'17*. 1–6.
- [6] Krisztian Balog and Filip Radlinski. 2020. Measuring recommendation explanation quality: The conflicting goals of explanations. In *Proceedings of the 43rd International ACM Sigir Conference on Research and Development in Information Retrieval*. 329–338.
- [7] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. 2017. Network dissection: Quantifying interpretability of deep visual representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6541–6549.
- [8] Keping Bi, Qingyao Ai, and W. Bruce Croft. 2021. Learning a Fine-Grained Review-Based Transformer Model for Personalized Product Search. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (Virtual Event, Canada) (SIGIR '21)*. Association for Computing Machinery, New York, NY, USA, 123–132. <https://doi.org/10.1145/3404835.3462911>
- [9] Keping Bi, Qingyao Ai, Yongfeng Zhang, and W Bruce Croft. 2019. Conversational Product Search Based on Negative Feedback. In *Proceedings of the 28th ACM international conference on information and knowledge management*. ACM.
- [10] Keping Bi, Choon Hui Teo, Yesh Dattatreya, Vijai Mohan, and W Bruce Croft. 2019. A Study of Context Dependencies in Multi-page Product Search. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 2333–2336.
- [11] Mustafa Bilgic and Raymond J Mooney. 2005. Explaining recommendations: Satisfaction vs. promotion. In *Beyond Personalization Workshop, IUI*, Vol. 5. 153.
- [12] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems* 26 (2013), 2787–2795.
- [13] Christopher P Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. 2017. Understanding disentangling in *beta*-VAE. *Proceedings of the 2017 NIPS Workshop on Learning Disentangled Representations* (2017).
- [14] Robin Burke. 2002. Hybrid recommender systems: Survey and experiments. *User modeling and user-adapted interaction* 12, 4 (2002), 331–370.
- [15] David Carmel, Elad Haramaty, Arnon Lazerson, and Liane Lewin-Eytan. 2020. Multi-Objective Ranking Optimization for Product Search Using Stochastic Label Aggregation. In *Proceedings of The Web Conference 2020*. 373–383.
- [16] David Carmel, Elad Haramaty, Arnon Lazerson, Liane Lewin-Eytan, and Yoelle Maarek. 2020. Why Do People Buy Seemingly Irrelevant Items in Voice Product Search? On the Relation between Product Relevance and Customer Satisfaction in ECommerce. In *Proceedings of the 13th International Conference on Web Search and Data Mining (Houston, TX, USA) (WSDM '20)*. Association for Computing Machinery, New York, NY, USA, 79–87. <https://doi.org/10.1145/3336191.3371780>
- [17] Henriette Cramer, Vanessa Evers, Satyan Ramlal, Maarten Van Someren, Lloyd Rutledge, Natalia Stash, Lora Aroyo, and Bob Wielinga. 2008. The effects of transparency on trust in and acceptance of a content-based art recommender. *User Modeling and User-Adapted Interaction* 18, 5 (2008), 455.
- [18] Mengnan Du, Ninghao Liu, and Xia Hu. 2018. Techniques for interpretable machine learning. *Commun. ACM* (2018).
- [19] Huizhong Duan, ChengXiang Zhai, Jinxing Cheng, and Abhishek Gattani. 2013. A probabilistic mixture model for mining and analyzing product search log. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. 2179–2188.
- [20] Huizhong Duan, ChengXiang Zhai, Jinxing Cheng, and Abhishek Gattani. 2013. Supporting keyword search in product database: a probabilistic approach. *Proceedings of the VLDB Endowment* 6, 14 (2013), 1786–1797.
- [21] Zeon Trevor Fernando, Jaspreet Singh, and Avishek Anand. 2019. A study on the Interpretability of Neural Retrieval Models using DeepSHAP. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1005–1008.
- [22] Jonathan Frankle and Michael Carbin. 2019. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *ICLR* (2019).
- [23] LiMin Fu. 1994. Rule generation from neural networks. *IEEE Transactions on Systems, Man, and Cybernetics* 24, 8 (1994), 1114–1124.
- [24] Leilani H Gilpin, David Bau, Ben Z Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. 2018. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*. IEEE, 80–89.
- [25] Jiafeng Guo, Yixing Fan, Liang Pang, Liu Yang, Qingyao Ai, Hamed Zamani, Chen Wu, W Bruce Croft, and Xueqi Cheng. 2019. A deep look into neural ranking models for information retrieval. *arXiv preprint arXiv:1903.06902* (2019).
- [26] Yangyang Guo, Zhiyong Cheng, Liqiang Nie, Yinglong Wang, Jun Ma, and Mohan Kankanhalli. 2019. Attentive long short-term preference modeling for personalized product search. *ACM Transactions on Information Systems (TOIS)* 37, 2 (2019), 1–27.
- [27] Yangyang Guo, Zhiyong Cheng, Liqiang Nie, Xin-Shun Xu, and Mohan Kankanhalli. 2018. Multi-modal preference modeling for product search. In *Proceedings of the 26th ACM international conference on Multimedia*. 1865–1873.
- [28] Jonathan Lee Herlocker. 2000. *Understanding and improving automated collaborative filtering systems*. Citeseer.
- [29] Jonathan L Herlocker, Joseph A Konstan, and John Riedl. 2000. Explaining collaborative filtering recommendations. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work*. ACM, 241–250.
- [30] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. 2017. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. *ICLR* 2, 5 (2017), 6.
- [31] Yujing Hu, Qing Da, Anxiang Zeng, Yang Yu, and Yinghui Xu. 2018. Reinforcement learning to rank in e-commerce search engine: Formalization, analysis, and application. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 368–377.
- [32] Sarthak Jain and Byron C Wallace. 2019. Attention is not explanation. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics* (2019).
- [33] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2017. Accurately interpreting clickthrough data as implicit feedback. In *ACM SIGIR Forum*, Vol. 51. Acm New York, NY, USA, 4–11.
- [34] Shubhra Kanti Karmaker Santu, Parikshit Sondhi, and ChengXiang Zhai. 2017. On application of learning to rank for e-commerce search. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 475–484.
- [35] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. Lightgbm: A highly efficient gradient boosting decision tree. In *Advances in neural information processing systems*. 3146–3154.
- [36] Omer Levy and Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. *Advances in neural information processing systems* 27 (2014), 2177–2185.
- [37] Lipyew Lim, Haixun Wang, and Min Wang. 2013. Semantic queries by example. In *Proceedings of the 16th International Conference on Extending Database Technology*. ACM, 347–358.
- [38] Zachary C Lipton. 2018. The mythos of model interpretability. *Commun. ACM* 61, 10 (2018), 36–43.
- [39] Shang Liu, Wanli Gu, Gao Cong, and Fuzheng Zhang. 2020. Structural Relationship Representation Learning with Graph Embedding for Personalized Product Search. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 915–924.
- [40] Francesco Locatello, Stefan Bauer, Mario Lucic, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. 2018. Challenging common assumptions in the unsupervised learning of disentangled representations. *Proceedings of the 36th International Conference on Machine Learning* (2018).
- [41] Scott M Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. *Advances in Neural Information Processing Systems* 30 (2017), 4765–4774.
- [42] Liangchen Luo, Yuanhao Xiong, Yan Liu, and Xu Sun. 2019. Adaptive gradient methods with dynamic bound of learning rate. *arXiv preprint arXiv:1902.09843* (2019).
- [43] Julian McAuley and Jure Leskovec. 2013. Hidden factors and hidden topics: understanding rating dimensions with review text. In *Proceedings of the 7th ACM conference on Recommender systems*. ACM, 165–172.
- [44] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [45] Bhaskar Mitra, Nick Craswell, et al. 2018. An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval* 13, 1 (2018), 1–126.
- [46] Anh Nguyen, Alexey Dosovitskiy, Jason Yosinski, Thomas Brox, and Jeff Clune. 2016. Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. In *Advances in Neural Information Processing Systems*. 3387–3395.
- [47] Petteri Nurmi, Emil Lagerspetz, Wray Buntine, Patrik Floréen, and Joonas Kukkonen. 2008. Product retrieval for grocery stores. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information*

- retrieval. 781–782.
- [48] Georgina Peake and Jun Wang. 2018. Explanation mining: Post hoc interpretability of latent factor models for recommendation systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 2060–2069.
- [49] Jay Michael Ponte and W Bruce Croft. 1998. *A language modeling approach to information retrieval*. Ph.D. Dissertation. University of Massachusetts at Amherst.
- [50] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. Why should I trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. ACM, 1135–1144.
- [51] Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends® in Information Retrieval* 3, 4 (2009), 333–389.
- [52] Gregor PJ Schmitz, Chris Aldrich, and Francois S Gouws. 1999. ANN-DT: an algorithm for extraction of decision trees from artificial neural networks. *IEEE Transactions on Neural Networks* 10, 6 (1999), 1392–1401.
- [53] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. 2014. CNN features off-the-shelf: an astounding baseline for recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. 806–813.
- [54] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034* (2013).
- [55] Jaspreet Singh and Avishek Anand. 2019. Exs: Explainable search using local model agnostic interpretability. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. 770–773.
- [56] Jaspreet Singh and Avishek Anand. 2020. Model agnostic interpretability of rankers via intent modelling. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 618–628.
- [57] Mark D Smucker, James Allan, and Ben Carterette. 2007. A comparison of statistical significance tests for information retrieval evaluation. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*. 623–632.
- [58] Nava Tintarev and Judith Masthoff. 2007. Effective explanations of recommendations: user-centered design. In *Proceedings of the 2007 ACM conference on Recommender systems*. ACM, 153–156.
- [59] Nava Tintarev and Judith Masthoff. 2007. A survey of explanations in recommender systems. In *2007 IEEE 23rd international conference on data engineering workshop*. IEEE, 801–810.
- [60] Nava Tintarev and Judith Masthoff. 2015. Explaining recommendations: Design and evaluation. In *Recommender systems handbook*. Springer, 353–382.
- [61] Aaron Van den Oord, Sander Dieleman, and Benjamin Schrauwen. 2013. Deep content-based music recommendation. In *Advances in neural information processing systems*. 2643–2651.
- [62] Christophe Van Gysel, Maarten de Rijke, and Evangelos Kanoulas. 2016. Learning latent vector spaces for product search. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*. 165–174.
- [63] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [64] Manisha Verma and Debasis Ganguly. 2019. LIRME: Locally Interpretable Ranking Model Explanation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris, France) (SIGIR'19). Association for Computing Machinery, New York, NY, USA, 1281–1284. <https://doi.org/10.1145/3331184.3331377>
- [65] Manisha Verma and Debasis Ganguly. 2019. LIRME: Locally Interpretable Ranking Model Explanation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 1281–1284.
- [66] Anthony J Viera, Joanne M Garrett, et al. 2005. Understanding interobserver agreement: the kappa statistic. *Fam med* 37, 5 (2005), 360–363.
- [67] Nan Wang, Hongning Wang, Yiling Jia, and Yue Yin. 2018. Explainable Recommendation via Multi-Task Learning in Opinionated Text Data. *SIGIR* (2018).
- [68] Qi Wang, Xinchun Liu, Wu Liu, An-An Liu, Wenyin Liu, and Tao Mei. 2020. MetaSearch: Incremental Product Search via Deep Meta-Learning. *IEEE Transactions on Image Processing* 29 (2020), 7549–7564.
- [69] Sarah Wiegrefe and Yuval Pinter. 2019. Attention is not not Explanation. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing* (2019).
- [70] Chen Wu, Ming Yan, and Luo Si. 2017. Ensemble methods for personalized e-commerce search challenge at CIKM Cup 2016. *arXiv preprint arXiv:1708.04479* (2017).
- [71] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. 2014. How transferable are features in deep neural networks?. In *Advances in neural information processing systems*. 3320–3328.
- [72] Hamed Zamani and W Bruce Croft. 2020. Learning a Joint Search and Recommendation Model from User-Item Interactions. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 717–725.
- [73] Matthew D Zeiler and Rob Fergus. 2014. Visualizing and understanding convolutional networks. In *European conference on computer vision*. Springer, 818–833.
- [74] Yongfeng Zhang. 2016. Explainable Recommendation: Theory and Applications. *PhD thesis* (2016).
- [75] Yongfeng Zhang and Xu Chen. 2020. Explainable recommendation: A survey and new perspectives. *Foundations and Trends in Information Retrieval* 14, 1 (2020), 1–101.
- [76] Yongfeng Zhang, Guokun Lai, Min Zhang, Yi Zhang, Yiqun Liu, and Shaoping Ma. 2014. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. ACM, 83–92.
- [77] Jan Ruben Zilke, Eneldo Loza Mencia, and Frederik Janssen. 2016. DeepRED–Rule extraction from deep neural networks. In *International Conference on Discovery Science*. Springer, 457–473.