

What Did I Say Again? Relating User Needs to Search Outcomes in Conversational Commerce

Kevin Schott

kevin.schott@gesis.org

GESIS – Leibniz Institute for the Social Sciences
Cologne, Germany

Daniel Hienert

daniel.hienert@gesis.org

GESIS – Leibniz Institute for the Social Sciences
Cologne, Germany

Andrea Papenmeier

a.papenmeier@utwente.nl

University of Twente
Enschede, Netherlands

Dagmar Kern

dagmar.kern@gesis.org

GESIS – Leibniz Institute for the Social Sciences
Cologne, Germany

ABSTRACT

Recent advances in natural language processing and deep learning have accelerated the development of digital assistants. In conversational commerce, these assistants help customers find suitable products in online shops through natural language conversations. During the dialogue, the assistant identifies the customer’s needs and preferences and subsequently suggests potentially relevant products. Traditional online shops often allow users to filter search results based on their preferences using facets. Selected facets can also serve as a reminder of how the product base was filtered. In conversational commerce, however, the absence of facets and the use of advanced natural language processing techniques can leave customers uncertain about how their input was processed by the system. This can hinder transparency and trust, which are critical factors influencing customers’ purchase intentions. To address this issue, we propose a novel text-based digital assistant that, in the product assessment step, explains how specific product aspects relate to the user’s previous utterances to enhance transparency and facilitate informed decision-making. We conducted a user study (N=135) and found a significant increase in user-perceived transparency when natural language explanations and highlighted text passages were provided, demonstrating their potential to extend system transparency to the product assessment step in conversational commerce.

CCS CONCEPTS

• **Information systems** → **Users and interactive retrieval**; • **Human-centered computing** → **HCI design and evaluation methods**.

KEYWORDS

Conversational User Interfaces, Conversational Commerce, Product Search, Conversational Search, Chatbot, Explanations

1 INTRODUCTION

In recent years, we have witnessed rapid advances in natural language processing (NLP) and deep learning techniques, resulting in the rise of powerful generative Artificial Intelligence (AI) such as GPT-4 [1]. This has led to the proliferation of conversational user interfaces (CUIs) such as virtual (voice) assistants and chatbots, which allow users to interact with computers in natural human language (NL) [19, 21]. The growing popularity of CUIs can also

be observed in the context of online shopping. *Conversational commerce*, a term which was introduced in 2015 [22, 38] and defined by Balakrishnan et al. as “buying activity by a customer through a digital assistant” [4] has emerged as a new option for companies to sell products to their customers. These digital (shopping) assistants engage customers in a voice or chat dialogue to elicit their needs and preferences and find a suitable product. The conversation mimics a natural dialogue with a human shop assistant intended to gather necessary information and to build trust [37].

The graphical user interfaces of traditional online shops often display facets that allow customers to specify their preferences and filter the search results accordingly. Additionally, selected facets can remind users of the applied filters. In conversational commerce, however, facets are missing, and digital assistants usually employ NLP techniques to interpret the users’ utterances. This can lead to users questioning how the system has processed their input to come up with a recommended product [11, 30].

To address this issue, we investigate how explanations that indicate how user utterances were mapped to product attributes can support users during product assessment in conversational commerce. We understand explanations to mean “[a]ny feature or aspect that enhances the interpretability and transparency of the system, making it more understandable to [users]” [16]. We introduce a text-based digital assistant for laptop search that maps vague user responses regarding different laptop attributes to specific values. Inspired by faceted search, the assistant explains this mapping in the result presentation. We conducted a user study and analyzed participants’ perceptions of two types of result explanations (namely text snippet highlighting and NL explanations) compared to a baseline condition without any explanations. The study’s findings indicate that providing NL explanations together with text snippet highlighting for how user responses were mapped to different attributes of a recommended product can enhance users’ perceived transparency and are deemed helpful by users.

2 RELATED WORK

As the interaction with digital assistants involves aspects of both search and recommendation [43], we discuss related literature from the research fields of information retrieval (IR) and recommender systems (RSs) in addition to the limited research on explanations specific to conversational commerce. Thus, users who search for a suitable product by conversing with a digital assistant receive either

a personalized product suggestion or a list of search results. In the following, we have a closer look at the goals of explanations, their effects on perceived transparency and decision-making, as well as text highlighting and NL explanations as explanatory components.

2.1 Goals of Explanations

Regarding the goals of explanations in RSs, Tintarev [34] describes seven possible aspects: Transparency (explaining how the system works), scrutability (allowing users to tell the system that it is wrong), trustworthiness (increasing users' confidence in the system), effectiveness (helping users make decisions), persuasiveness (convincing users to try or buy), efficiency (helping users to make decisions faster), and satisfaction (improving the ease of use or enjoyment). Focusing on the first aspect, Vorm and Miller [41] propose the five-factor model of transparency in RSs, which includes the factors Qualities of Data (e.g., "What are the sources of data?"), Options (e.g., "Choices known to the system are made available to the user"), User Representation (e.g., "Does the system know and consider the user in its model?"), Social Influence (e.g., "How is the user grouped with others?"), and System Parameters and Logic (e.g., "Information about system logic, reasoning, policies, limitations, etc."). Tintarev and Masthoff [35] note that the terms "explanation" and "justification" are often used interchangeably. However, Vig et al. [40] take a more detailed approach. They distinguish between (result) justification, which explains why a particular item is recommended (e.g., "This movie is a good choice because you liked other movies by this director"), and (system) transparency, which clarifies how the recommendation mechanism operates (e.g., "My algorithm analyzes your viewing history and identifies patterns, such as your preference for this director's movies"). In this work, we focus on system transparency.

2.2 Effects of Explanations on Perceived Transparency and Decision-making

Tsagkias et al. [37] point out that explanations in conversational commerce can increase transparency by enabling users to understand what data from their input is being processed and how the search or recommendation mechanism works. Jin et al. [15] developed the digital shopping assistant "PhoneBot" for mobile phones that provides users with NL explanations that justify product recommendations within the chat. Based on the users' expressed preferences, the assistant explains how well a suggested mobile phone ranks within the online shop's product library. The study shows that these explanations can enhance users' perceived transparency by improving explainability. In a user study with a digital assistant for hotel search, Hernandez-Bocanegra and Ziegler [12] found that enabling users to interact with explanations by requesting individual customer comments can increase their perception of transparency through enhanced perception of explanation quality compared to the sole provision of aggregated customer opinions.

In RSs, Pu et al. [29] show that explanations can increase users' perception of transparency, enhancing users' trust and confidence in decision-making. Cramer et al. [7] demonstrate that allowing users to view the criteria based on which an art recommender makes its suggestions can enhance users' perceived and actual understanding of how the system works. In their user study, Gedikli et al. [8]

compared different explanation styles for a movie recommender. They discovered that personalized tag clouds color-coded to represent users' sentiment towards tags such as "politics", "drama", or "classic" can support users in decision-making, ultimately improving recommendation effectiveness and increasing users' perceived level of system transparency. In another study for a movie recommender, Vig et al. [40] provided users with community tags for a recommended movie and predicted their perceived relevance and preference for those tags. They show that this type of explanation can improve users' perceived justification for ratings predicted by the RS and support their decision-making. In the context of music recommendation, Millicamp et al. [23] found that explanations could enhance users' perception of the recommendations' effectiveness and improve users' understanding. Furthermore, a user study by Verbert et al. [39] demonstrates that visually explaining and allowing users to explore publication recommendations in the context of academic conferences can increase their effectiveness. Investigating personalized NL explanations in a news RS, ter Hoeve et al. found that they may increase users' trust [33].

In the research area of IR, Khurana et al. [17] show that providing NL explanations for why a chatbot did not understand a user's query or does not know what to do next in the context of spreadsheet applications can enhance users' perception of usefulness, transparency, and trust. Additionally, Papenmeier and Topp [27] show that providing transparency by backchanneling how the users' input was processed by the system can improve their perception of a chatbot's competence as well as their conversational engagement in conversational commerce. Toader et al. [36] demonstrate that improving users' perceived competence of a chatbot can enhance their trust in the system.

2.3 Text Highlighting as an Explanatory Component in Information Retrieval

The research field of explainable information retrieval (ExIR) aims to make IR systems more transparent and trustworthy [3]. Text highlighting is used as a simple tool to support users in evaluating the relevance of results in web and scholarly search engines. Searchers rely on descriptions like titles, snippets, and URLs to decide whether they should click and read a particular result [42]. Users' query terms are highlighted in titles and snippets to give them feedback on where on a website their search terms were found [14]. The highlighting of text passages as direct answers to user queries is also used in "feature snippets", which are integrated into Google's search results [32]. Additionally, highlighting can be used to identify relevant passages directly within documents [44].

Highlighting is also used for local explanations in machine learning. Thus, feature attribution by color highlighting is a common method for explaining the predictions of classification models [18, 31]. Recent work investigates how feature attribution can be applied in ranking tasks by highlighting snippets of search queries and document texts [3].

2.4 Natural Language Explanations

Natural language (NL) is the most human-like way for systems to provide explanations [2, 5]. Such explanations are interpretable

by users with various backgrounds and levels of technical knowledge [20]. Juneja et al. [16] conducted a study on search engine explanations with users lacking AI and IR expertise. They found that the participants did not require explanations from a search engine for single, unambiguous answers such as specific dates or values. However, for more complex web search tasks such as product search, the participants deemed concise and easy-to-read NL explanations useful. The authors also note that both text highlighting (in “feature snippets”) and NL explanations are used by Google’s search engine to clarify the search results that were retrieved. Users can access the latter by clicking on the three-dots icon next to a search result in Google. Additionally, Gkatzia et al. [9] demonstrate that providing NL explanations for uncertain data such as weather data on a user interface can aid users in decision-making by enhancing their comprehension and informedness, as opposed to solely presenting graphical representations of the data.

3 RESEARCH QUESTION AND HYPOTHESES

While extensive research has been conducted on explaining the output of RSs, research has not yet established a systematic understanding of the effects of explanations within the result presentation in conversational commerce. In particular, the use of advanced NLP techniques to process user input poses new challenges in terms of transparency. In conversational commerce, a one-to-one mapping of the user’s sometimes vague utterances to predefined product attributes is not always obvious compared to what users are used to from traditional faceted search.

Our research is guided by the following research question:

RQ: How can explanations that indicate how user utterances were mapped to product attributes support users during the product assessment step in conversational commerce?

Inspired by common practices in IR, we investigate two different approaches:

- (1) **Highlighting (H):** The gathered user preferences are mapped to product attributes, which are then highlighted in the result presentation (see Fig. 2a).
- (2) **Highlighting + Explanation (H+E):** In addition to the highlighting, the mapping is explained in natural language (see Fig. 2b).

To answer our research question, we conducted an online user study using a between-subjects design with three conditions (H, H+E, B = Baseline without any explanations). To limit complexity and inspired by Google’s use of highlighting in conjunction with NL explanations (see Section 2.4), we decided not to include a condition in which only NL explanations are provided. We assume that explanations positively affect perceived system transparency, usefulness, and users’ confidence in their purchase decisions. Table 1 presents our hypotheses and the underlying assumptions.

4 STUDY DESIGN

4.1 Apparatus - Digital Product Advisor

We developed a rule-based digital assistant for laptop search, as shown in Fig. 1. The underlying dataset was collected from Amazon¹ and contains 3,638 laptop descriptions. The digital assistant

¹<https://www.amazon.com/>

queries users about their preferences for different laptop aspects in a predefined order that aims to mimic the questioning behavior of a human shop assistant: Purpose, price, screen size, hard drive storage, RAM size, and battery life. The user’s response regarding the intended purpose gets mapped to either basic, advanced, or gaming and is used to provide recommendations for the subsequent attributes. After each attribute, the system continuously filters the dataset for suitable laptops by mapping the user’s responses to specific values. Research has shown that the vagueness of user queries in the context of product search varies to a high degree [26]. Therefore, our assistant uses a BERT model and an LSTM model to identify vague user statements (e.g., “rather small” or “not too big”) and map them to specific value ranges (e.g., 12 to 14 inches) for the respective product attributes. Especially when vague user input is provided, we expect explanations to help users understand why a particular product is recommended.

After inquiring about all the described laptop aspects, our system presents a suitable laptop to the user. If multiple laptops in the dataset match the user’s preferences, the one with the highest average user rating is suggested. However, if only one fitting laptop remains earlier in the conversation, the assistant ends the conversation and moves on to the result presentation. To simulate a first suggestion from a human shop assistant, we have chosen to display only one result. Our setup focuses on providing explanations for this initial suggestion. The suitable laptop is presented as a typical online shop result page with a title, user rating, product image, and product description.

We tested two different explanation variants to present the results showing how a user’s utterances relate to different aspects of the recommended laptop alongside a baseline condition (B) without any explanation. In the first variant (H), the aspects the user was asked about are highlighted in color (see Fig. 2a). Given the common usage of highlighting (see Section 2.3), we adopted it as a straightforward tool in our user study to investigate its effects on user perception in the context of conversational commerce. In the second variant (H+E), our assistant offers NL explanations in the form of preference repetition to the user in speech bubbles (see Fig. 2b). The explanation consists of a citation of the user’s response regarding a laptop aspect they were asked about (“You said: [...]”) and a description of how it impacted the assistant’s search (“So I searched for [...]”). An example would be “You said: ‘The laptop should be rather small and portable as I move around a lot at work’. So I searched for: 15 inches or less.” Inspired by faceted search and search engine user interfaces, we deliberately integrate several aspect-specific explanations into the result presentation instead of providing a single, summative explanation within the chat (like in Jin et al.’s [15] “PhoneBot”, see Section 2.2). This way, users can find the explanations exactly next to the product attributes they are referring to.

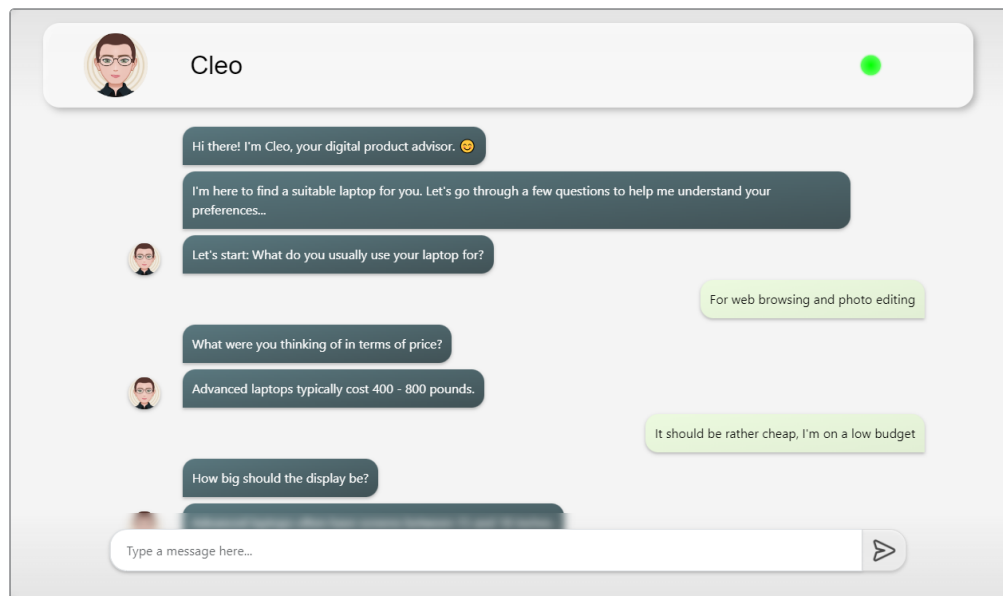
4.2 Procedure and Scenario

We set up an online study on SoSci Survey². Participants first provide consent and answer closed questions about their demographics. They then rate their domain knowledge for laptops on a scale from 1 (no knowledge/non-expert) to 7 (high knowledge/expert) before

²<https://www.sosicurvey.de/en/index>

Table 1: Research hypotheses and underlying assumptions. Dependent variables are formatted in bold.


Assumption	Hypothesis
Jin et al. [15] show that explanations provided within the chat positively influence users' perceived transparency in conversational commerce. Similarly, in traditional RSs, literature shows that explanations also enhance perceived transparency [7, 8, 29, 40] (see Section 2.2). We assume that those findings likewise hold true for the result presentation in conversational commerce.	H 1. <i>Users who are provided with H+E have a higher perception of transparency than those with H only. Users assigned to B have the lowest perception of transparency.</i>
Vig et al. [40] and Verbert et al. [39] show that explanations can increase decision effectiveness in RSs by allowing users to make more informed decisions. We assume that enabling users to make more informed decisions also increases their decision-making confidence after assessing a recommended item in conversational commerce. As highlighting can be categorized as a type of explanation [16], albeit more rudimentary than NL explanations, we assume that it leads to a higher confidence than the baseline condition.	H 2. <i>Users who are provided with H+E have the highest confidence in their decision. The lowest confidence is reported by users in B.</i>
Pu et al. [29] show that by improving perceived transparency, confidence, and trust, explanations can increase users' purchase intentions in RSs. Furthermore, Zhou [45] found that by increasing customers' trust, perceived transparency can improve their purchase intentions in online stores. We assume that these findings can be transferred to or are applicable to the result presentation in conversational commerce.	H 3. <i>Users who are provided with H+E have a higher purchase intention compared to both H and B.</i>
Explanations were shown to increase decision-making effectiveness in RSs [39, 40]. We assume that these findings can be transferred to the result presentation in conversational commerce. However, while effectiveness is specifically related to decision-making (as also described by Tintarev [34]), see Section 2.1), we use the related variable "perceived usefulness" to examine the users' perception of the assistant as a whole. We posit that enhanced decision-making effectiveness also leads to an increase in the perceived usefulness of the assistant. This aligns with research by Ing and Ming [13] who use shopping effectiveness as one of several items to measure the perceived usefulness of blogger recommendations.	H 4. <i>Users who are provided with H+E perceive the digital advisor as most useful. B is perceived the least useful.</i>
Khurana et al. [17] show that NL explanations of errors provided by a chatbot in a spreadsheet tool enhance users' perception of usefulness. We assume that this finding can be applied to the result presentation in conversational commerce.	H 5. <i>H+E is perceived as more useful than H.</i>

**Figure 1: Interface for our digital assistant**

HP laptop, 14.0" display Windows 10 8GB RAM 128GB storage, 3.5GHz 2 cores AMD Ryzen 3 processor, 10.0 hours h battery life

From HP
£ 324

★★★★★ 5 (based on 2 reviews)



Brand	HP
Operating System	Windows 10
Screen Size	14.0 inches
Hard Drive Type	SSD
Hard Drive Size	128 gigabytes
RAM Size	8 gigabytes
Processor Manufacturer	AMD
Processor Cores	2
Processor Speed	3.5 gigahertz
Graphics Card	Integrated
Casing Color	silver
Weight	3.24 pounds


(a) Highlighting (H)

HP laptop, 14.0" display Windows 10 8GB RAM 128GB storage, 3.5GHz 2 cores AMD Ryzen 3 processor, 10.0 hours h battery life

From HP
£ 324

You said: "The laptop should be rather cheap". So I searched for: 450 pounds or less.

★★★★★



Brand	HP
Operating System	Windows 10
Screen Size	14.0 inches
Hard Drive Type	SSD
Hard Drive Size	128 gigabytes
RAM Size	8 gigabytes
Processor Manufacturer	AMD
Processor Cores	2
Processor Speed	3.5 gigahertz
Graphics Card	Integrated
Casing Color	silver

You said: "I'm looking for a smaller laptop". So I searched for: 15 inches or less.

You said: "Not too much". So I searched for: 1000 GB or less.

You said: "I guess 8 is fine then". So I searched for: 8 GB.

(b) Highlighting plus natural language explanations (H+E)

Figure 2: Explanation variants

being presented with the scenario and task description (adapted from [25]):

“Imagine that your laptop stopped working, and you are now searching for a new one. The advisor will inquire about your preferences and requirements to suggest a suitable laptop for you. Your task will be to interact with the advisor, examine the recommended laptop and decide whether you would buy it.”

Participants are then randomly assigned and directed to one of the three interface variants (B, H, H+E), receive a reminder of the scenario, and start interacting with the digital assistant (see Fig. 1) until they receive their laptop recommendation (as shown in Fig. 2 according to their condition). Upon scrolling down to view the complete product overview that includes all product aspects that the participants were asked about, a button labeled “I have made a decision regarding the purchase of the laptop. Take me to the next step!” appears on the interface. After clicking this button,

participants are asked about the likeliness of buying the recommended laptop and subsequently about their confidence in their previous answer. In case of low likeliness to buy, they are asked for additional feedback with an open question. Thereafter, participants are redirected to our questionnaire and answer open and closed questions about their experience with the digital assistant. Our institute's ethics committee has confirmed the study's ethical clearance.

4.3 Measurements

This section presents the dependent variables that were tested and the types of data that were gathered. Tables provide insight into how we measured the dependent variables.

4.3.1 Measures for Perceived Transparency, Confidence, and Purchase Intention. We took measures for the dependent variables confidence and perceived transparency. Additionally, we measured purchase intention, i.e., the users' likeliness to buy the recommended laptop. All mentioned variables were measured with 5-point Likert scale questions or statements (cp. Table 2). The statements for PT1, and PT2 are inspired by previous work from Jin et al. [15], while the one for PT3 is inspired by Hellman et al. [10]. Peters [28] inspired the statement for confidence.

4.3.2 Measures for Usefulness. We also measured users' perceived usefulness of the digital assistant for all conditions (statement inspired by Pu et al. [29]) as well as the usefulness for H and H+E in particular. For the interface variants H and H+E, we asked additional specific questions that were not applicable for the baseline condition. We took measures for the perceived usefulness of the highlighting and the explanations, respectively. For H+E, we also measured the explanations' understandability (cp. Table 3). We again used a five-point Likert scale with statement and questions.

4.3.3 Qualitative User Feedback. As the quantitative variables do not give us detailed information about why people like some things and dislike others, we included some open-ended questions to gather more individual feedback. This also helps us to identify areas for improvement for future work. All open questions are listed in Table 4. Q5 was asked after the perceived usefulness rating for H and H+E (cp. Table 3).

4.4 Participants

We recruited 135 participants (68 male, 66 female, 1 inter/non-binary) through the online crowdsourcing platform Prolific³. During the course of the study, 14 participants either chose to withdraw or had their participation automatically canceled by Prolific after 56 minutes. In these cases, Prolific automatically recruited new participants until a total of 135 study completions was reached. They had a mean domain knowledge of 4.4, with a standard deviation of 1.5. A Kruskal-Wallis test showed that there was no significant difference ($p = 0.3584$) in domain knowledge between conditions. The participants (ages 18-91, $M = 40.9$, $STD = 13.5$) were randomly assigned to one of the three conditions ($B = 47$, $H = 43$, $H+E = 45$). To be eligible, they had to be located in the UK, speak English as their

primary language, and not have any language difficulties. Participation in the study took approximately 11 minutes per participant ($M = 10.53$ min, $STD = 5.22$ min). Each participant who completed the study received a compensation of 2.25 GBP, corresponding to an hourly rate of 12.27 GBP.

5 RESULTS

This section presents our findings and the data collected during our user study. First, we describe the results of our statistical analyses and evaluate our hypotheses. We then present the qualitative insights we gathered from the participants' answers to the open-ended questions.

5.1 Statistical Analysis

To test for significant differences between conditions, we used the Kruskal-Wallis test with a Dunn's test as post-hoc analysis, applying the Bonferroni correction to account for the multiple testing bias. To compare only two conditions, we used the Mann-Whitney U test. We used non-parametric tests because the survey data was not normally distributed. For analyzing correlations, we used the measure of Spearman's rho. Means and standard deviations for our dependent variables are presented in Table 5.

5.1.1 Effects of Explanations on Perceived Transparency, Confidence, and Purchase Intention. The Kruskal-Wallis tests showed no significant differences ($p > 0.05$) between conditions for PT2 (**H1** rejected for PT2), confidence (**H2** rejected), purchase intention (**H3** rejected), and perceived usefulness (**H4** rejected) (cp. Table 6). Significant differences between groups were found for PT1, PT3, and PT4. Post-hoc Dunn's tests were conducted to determine which specific conditions differed significantly. For PT1 (see Fig. 3a), H+E was rated significantly higher than both B ($p < 0.001$) and H ($p < 0.01$). For PT3 (see Fig. 3b), H+E was also rated significantly higher than B ($p < 0.001$) and H ($p < 0.01$). Finally, for PT4 (see Fig. 3c), H+E received a significantly higher score than B ($p < 0.01$). However, no significant difference was found between H+E and H ($p = 0.0527$). Thus, **H1** can be partially accepted for PT1, PT2, and PT4.

5.1.2 Perceived Usefulness and Understandability. A Kruskal-Wallis test revealed no significant differences ($p > 0.05$) between conditions for perceived usefulness (cp. Table 6). A Mann-Whitney U test further showed no significant difference in perceived usefulness (H, H+E) between H and H+E ($W = 1033.5$, $p = 0.5621$). Thus, **H5** has to be rejected. For explanation understandability, the mean was 4.62 and the standard deviation 0.53 (cp. Table 5).

5.1.3 Effects of Domain Knowledge on Confidence. As reported above, we did not find a significant main effect of decision-making confidence between conditions. However, when categorizing all participants in two groups of below-average (1-3) and above-average (5-7) domain knowledge, a Mann-Whitney U test showed that the confidence of participants with above-average domain knowledge was significantly higher ($p < 0.01$) than of those with below-average domain knowledge. We also found a weak positive correlation ($r = 0.3024$) between domain knowledge and confidence.

³<https://www.prolific.com/>

Table 2: Dependent variables measured with 5-point Likert scale questions or statements

Dependent variable	Question/Statement	Likert Scale Options
Perceived transparency	The product advisor explained why the product was recommended to me. (PT1)	Strongly disagree (1) – Strongly agree (5)
	I understood why the product was recommended to me. (PT2)	Strongly disagree (1) – Strongly agree (5)
	The product advisor provided information about how my preferences were considered. (PT3)	Strongly disagree (1) – Strongly agree (5)
	How easy was it for you to understand why the product was recommended to you? (PT4)	Very difficult (1) – Very easy (5)
Confidence	How confident are you in your decision?	Not at all confident (1) – Extremely confident (5)
Purchase intention	How likely is it that you would buy this laptop?	Extremely unlikely (1) – Extremely likely (5)

Table 3: 5-point Likert scale statements or questions for the interface variants H and H+E

Dependent variable	Question/Statement	Likert Scale Options
Perceived usefulness	The advisor helped me find a suitable product.	Strongly disagree (1) – Strongly agree (5)
Perceived usefulness (H, H+E)	How useful did you find the [highlighting/explanations] while forming an opinion on the recommended product?	Not at all useful (1) – Extremely useful (5)
Understandability (only for H+E)	How easy was it for you to understand the language of the explanations?	Very difficult (1) – Very easy (5)




Figure 3: Significant Kruskal-Wallis test results for PT1, PT3, and PT4

5.2 Qualitative User Feedback

We analyzed participants’ responses to the open-ended questions (cp. Table 4) through thematic analysis [6], using codes to inductively identify patterns and common themes. Regarding participants’ responses to Q1, it seems that the laptop recommendation is perceived as slightly less confusing in H and H+E. 79% (H) and 71% (H+E) of the participants answered this question with ‘nothing,’ in contrast to 53% in the baseline condition. The other statements in all conditions, among the single statements, are about complaining

about incorrect processing of user input (6%) and the lack of alternative laptop choices (7%). In the baseline group, 4% of the participants explicitly asked for an explanation of why the recommended item was suggested.

Feedback on what stood out to participants (Q2) can be divided into feedback on the conversation and feedback on the result or the presentation of the result. The speed of the conversation was praised by 12% of all participants; 10% liked the suggestions and explanations given by the digital advisor during the conversation. 4% liked the conversation in general, while 2% complained about

Table 4: Open-ended questions for eliciting qualitative data

Open-ended questions
What specifically confused you about the recommendation? (Q1)
Which specific aspect(s) from your interaction with the product advisor stood out to you? (Q2)
What (other) explanations would you like to be provided by the advisor? (Q3)
What specifically helped you understand the recommendation? (Q4)
Please justify your rating [for highlighting/explanation usefulness] above. (Q5)

not being able to ask questions, and 4% would have liked to be asked more questions. Regarding the results and presentation of the results, again, 3% complained about the lack of choice. In both the H and the H+E conditions, 13% of participants explicitly stated that they liked the recommended item (e.g., participant P103 wrote “that [the advisor was] able to find a product that fit all my needs”). In the baseline condition, zero participants left an explicit comment about this when asked what stood out to them. 51% of the participants in the baseline condition stated that the layout and the presentation itself helped them to understand the recommended item Q4. In the H and H+E conditions, this was mentioned less (H: 30%, H+E: 11%). They were more likely to find the highlighting (H: 28%, H+E: 9%) and the explanations helpful (H+E: 64%). P22, for example, explained that they liked “The pop up boxes explaining why the product had been chosen for me”.

When asked for (further) explanations (Q3), B: 25%, H: 42% and H+E: 27% answered “none”. In the B and H conditions, 23% and 26%, respectively, explicitly asked for an explanation as to why the recommended item was chosen (e.g., P23: “Would like to understand why this specific product was chosen [...]”). 11% of the participants in the H+E condition asked for more choices and explanations of their differences. Justifying why a result was chosen contributed to perceived usefulness (Q5) for 22% of participants in the H+E condition (e.g., P140: “It justifies why the product was selected for me. It allows me to see how and why the product may or may not be suitable for my needs”). On the other hand, 15% of the participants stated that the NL explanations described the obvious. Highlighting was mentioned as useful by 40% in the H condition. 21% liked it because it drew attention (e.g., P125: “Helps to see if it met my requirements”). Highlighting was not explicitly mentioned at all in the H+E condition.

Across conditions, 30% of participants responded that it is either unlikely or extremely unlikely that they would buy the laptop recommended to them. 37% of them noted that they would have preferred a different operating system and/or laptop brand, while 27% mentioned other product specifications would prevent them from buying (e.g., P141: “Because of the graphics card”). Additionally, 24% stated that the digital assistant did not understand their preferences correctly (e.g., P23: “Because I asked for a 17 inch laptop, and the

screen size is 16 inches”). Two participants (5%) mentioned that the laptop’s aesthetics did not appeal to them, while P10 stated that they “feel as if the interaction was not real enough and [doesn’t] give [them] the most confidence in choosing a laptop to buy”. B: 21%, H: 13%, and H+E: 8% mentioned a lack of product information displayed on the interface.

6 DISCUSSION

The results indicate that the NL explanations provided by our digital assistant can support users during their assessment of a recommended product by significantly improving their perception of transparency. Our findings are in line with those from the RS domain [7, 8, 23, 29, 40], which show that explanations can enhance transparency as perceived by users (see Section 2.2).

According to qualitative feedback, users found the NL explanations to be particularly helpful for understanding why an item was recommended (see Section 5.2). This suggests that the explanations improved the effectiveness of the recommendation, which Tintarev and Masthoff [34] defined as one goal of explanations in RSs (see Section 2.1). Additionally, participants’ feedback indicates that reducing their confusion about the recommendation was one factor that enhanced their perception of transparency. We assume that the personalized explanations improved the transparency factors “User Representation” and “System Parameters and Logic” defined by Vorm and Miller [41] (see Section 2.1). Meanwhile, highlighting did not significantly affect any of the dependent variables. Qualitative feedback indicates that it did not assist participants in deducing why a recommended item was selected.

Although participants from H+E reported finding the NL explanations helpful in the open-ended questions, the quantitative findings show that they do not improve users’ confidence in their decision-making. Thus, our study on conversational commerce could not transfer Pu et al.’s [29] finding of a significantly positive effect of perceived transparency on users’ confidence in RSs to the result presentation in conversational commerce. Qualitative user feedback suggests that users’ confidence was influenced by several confounding factors, including reliance on or preference for a certain operating system or brand, the inability to compare results due to the suggestion of only one laptop, unsatisfactory or missing product information, and qualitative aspects such as aesthetic preferences. The desire of participants to compare multiple items can be attributed to the “Options” factor in Vorm and Miller’s [41] five-factor model of transparency in RSs (see Section 2.1). Analysis of user input showed that participants who mentioned that the NL explanations stated the obvious mostly responded with concrete values rather than vague statements. This suggests that our product advisor was particularly helpful to those participants who did not have a precise idea of the laptop they wanted. Conversely, this indicates that it may be sufficient to show explanations for user utterances in which vagueness was detected.

While the presence of NL explanations did not significantly affect users’ perceived understanding of why a product was recommended (PT2), it did significantly improve the perceived ease of understanding (PT4). This suggests that while NL explanations may not enhance users’ confidence in their understanding of the recommendation, they do make the process of understanding why a

Table 5: Mean and standard deviation for each dependent variable. Variables for which significant differences were found are formatted in bold.

	Baseline		Highlighting		H+E	
	Mean	SD	Mean	SD	Mean	SD
PT1	3.04	1.00	3.26	1.20	3.96	0.77
PT2	4.00	0.81	4.09	1.09	4.29	0.63
PT3	3.30	1.08	3.42	1.22	4.13	0.76
PT4	3.96	0.69	4.00	0.90	4.38	0.72
Confidence	3.79	0.86	4.02	0.80	4.00	0.90
Purchase intention	3.19	1.10	3.23	1.41	3.36	1.23
Perceived usefulness	3.70	0.98	3.60	1.28	3.89	0.91
Perceived usefulness (H, H+E)	-	-	3.70	1.24	3.67	1.02
Understandability (H+E)	-	-	-	-	4.62	0.53

Table 6: Results of the Kruskal-Wallis tests comparing B, H, and H+E. Significant results are formatted in bold.

Dependent variable	p-value (adjusted)	Test Statistic (H)	Effect Size (η^2)
PT1	0.00009	18.6911	0.1300
PT2	0.1518	3.7708	0.0130
PT3	0.0004	15.7999	0.1000
PT4	0.0091	9.3828	0.0560
Confidence	0.3474	2.1148	0.0009
Purchase intention	0.7729	0.5152	-0.0110
Perceived usefulness	0.6294	0.926	-0.0081

product was recommended easier. We assume that the explanations may serve to reduce the cognitive effort required to understand recommendations.

While some participants in H and H+E explicitly mentioned in the open questions that they liked the recommended item, none mentioned this in the baseline condition. This would be consistent with the persuasive effects of explanations, as described by Masthoff and Tintarev [34] (see Section 2.1). However, we could not underscore this finding through a statistically significant difference in participants' purchase intention. We assume this to be a consequence of the explanations' focus on system transparency rather than result justification (see Section 2). Furthermore, we assume that the usability limitations of our digital assistant, such as the lack of choices and users' inability to ask questions or adjust their preferences, had a halo effect on the participants' emotional responses [24], which may have confounded their assessment of both decision-making confidence and purchase intention. We will address this issue in future work by improving the digital assistant based on our participants' feedback.

7 CONCLUSION AND FUTURE WORK

With this study, we explored how explanations can support customers during product assessment in the context of conversational commerce. Focusing on the laptop domain, we investigated users' perceptions when provided with text snippet highlighting and NL in addition to this highlighting as explanation variants for recommended products. The explanations aimed at compensating for the absence of facets and account for the use of advanced NLP

techniques in conversational commerce. Our study suggests that providing NL explanations and text snippet highlighting for a recommended product supports users during item assessment by enhancing their perception of transparency. Furthermore, NL explanations were deemed helpful by users. On the other hand, highlighting without NL explanations could not support users. These findings can assist online retailers in promoting transparency during the product evaluation step when developing user interfaces for this growing field of e-commerce.

In future work, we plan to investigate how users can be further supported in their result assessment within conversational commerce. As many participants of our study wished to be presented with several results to compare rather than just one, we will expand our user interface by adding more results and exploring how to display result lists and explain differences between items to further assist users in evaluating products. Additionally, we will consider our findings that users with limited domain knowledge have less decision-making confidence than those with above-average domain knowledge. We will investigate how to personalize conversational interactions to provide appropriate support to users based on their domain knowledge. As our study had limitations in this regard, users' decision-making confidence could be retested with a more advanced digital assistant that enables them to state their preferences for additional product aspects. Furthermore, as some participants complained about incorrect input processing, a logical next step is to leverage the scrutiny goal of explanations by allowing users to continue the conversation beyond the initial result presentation.

This would enable them to clarify any misunderstandings and adjust their preferences. Lastly, we did not examine the isolated effects of NL explanations on user perception, as this was not the focus of our study. Future research could investigate the independent effects of explanation formats (highlighting vs. standalone NL) to gain a more comprehensive understanding.

ACKNOWLEDGMENTS

This work was supported by the German Research Foundation (DFG) as part of the "VACOS 2" project (no. 388815326). We thank our study participants and our anonymous reviewers for their helpful feedback. We would also like to thank Niklas Kerkfeld for his assistance with the technical implementation.

REFERENCES

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023). <https://doi.org/10.48550/arXiv.2303.08774>
- [2] Jose M Alonso, Senén Barro, Alberto Bugarin, Kees van Deemter, Claire Gardent, Albert Gatt, Ehud Reiter, Carles Sierra, Mariët Theune, Nava Tintarev, et al. 2020. Interactive natural language technology for explainable artificial intelligence. In *International Workshop on the Foundations of Trustworthy AI Integrating Learning, Optimization and Reasoning*. Springer, 63–70. https://doi.org/10.1007/978-3-030-73959-1_5
- [3] Avishek Anand, Lijun Lyu, Maximilian Idahl, Yumeng Wang, Jonas Wallat, and Zijian Zhang. 2022. Explainable Information Retrieval: A Survey. *arXiv preprint arXiv:2211.02405* (2022). <https://doi.org/10.48550/arXiv.2211.02405>
- [4] Janarthanan Balakrishnan and Yogesh K. Dwivedi. 2021. Conversational commerce: entering the next stage of AI-powered digital assistants. *Annals of Operations Research* (12 Apr 2021). <https://doi.org/10.1007/s10479-021-04049-5>
- [5] Erik Cambria, Amir Hussain, Catherine Havasi, and Chris Eckl. 2009. Common sense computing: From the society of mind to digital intuition and beyond. In *Biometric ID Management and Multimodal Communication: Joint COST 2101 and 2102 International Conference, BioID_MultiComm 2009, Madrid, Spain, September 16-18, 2009. Proceedings 2*. Springer, 252–259. https://doi.org/10.1007/978-3-642-04391-8_33
- [6] Victoria Clarke and Virginia Braun. 2017. Thematic analysis. *The Journal of Positive Psychology* 12, 3 (2017), 297–298. <https://doi.org/10.1080/17439760.2016.1262613>
- [7] 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 (01 Nov 2008), 455–496. <https://doi.org/10.1007/s11257-008-9051-3>
- [8] Fatih Gedikli, Dietmar Jannach, and Mouzhi Ge. 2014. How should I explain? A comparison of different explanation types for recommender systems. *International Journal of Human-Computer Studies* 72, 4 (2014), 367–382. <https://doi.org/10.1016/j.ijhcs.2013.12.007>
- [9] Dimitra Gkatzia, Oliver Lemon, and Verena Rieser. 2016. Natural Language Generation enhances human decision-making with uncertain information. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Katrin Erk and Noah A. Smith (Eds.). Association for Computational Linguistics, Berlin, Germany, 264–268. <https://doi.org/10.18653/v1/P16-2043>
- [10] Marco Hellmann, Diana Hernandez Bocanegra, and Jürgen Ziegler. 2022. Development of an Instrument for Measuring Users' Perception of Transparency in Recommender Systems. In *Workshops at the International Conference on Intelligent User Interfaces (IUI) 2022: Proceedings of the IUI 2022 Workshops: APEX-UI, HAI-GEN, HEALTHI, HUMANIZE, TeXSS, SOCIALIZE (CEUR Workshop Proceedings, Vol. 3124)*, Alison Smith-Renner and Ofra Amir (Eds.). RWTH Aachen, Aachen, 156–165. <https://doi.org/10.17185/dupublico/75905>
- [11] 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 (Philadelphia, Pennsylvania, USA) (CSCW '00)*. Association for Computing Machinery, New York, NY, USA, 241–250. <https://doi.org/10.1145/358916.358995>
- [12] Diana C. Hernandez-Bocanegra and Jürgen Ziegler. 2023. Explaining Recommendations through Conversations: Dialog Model and the Effects of Interface Type and Degree of Interactivity. *ACM Trans. Interact. Intell. Syst.* 13, 2, Article 6 (apr 2023), 47 pages. <https://doi.org/10.1145/3579541>
- [13] Grace Phang Ing and Ting Ming. 2018. Antecedents of consumer attitude towards blogger recommendations and its impact on purchase intention. *Asian Journal of Business and Accounting* 11, 1 (2018), 293–323. <https://doi.org/10.22452/ajba.vol11no1.10>
- [14] Tereza Iofciu, Nick Craswell, and Milad Shokouhi. 2009. Evaluating the impact of snippet highlighting in search. *Understanding the User-Logging and Interpreting User Interactions in Information Search and Retrieval (UIIR-2009)* 44 (2009).
- [15] Yucheng Jin, Li Chen, Wanling Cai, and Xianglin Zhao. 2024. CRS-Que: A User-centric Evaluation Framework for Conversational Recommender Systems. *ACM Trans. Recomm. Syst.* 2, 1, Article 2 (mar 2024), 34 pages. <https://doi.org/10.1145/3631534>
- [16] Prerna Juneja, Wenjuan Zhang, Alison Marie Smith-Renner, Hemank Lamba, Joel Tetreault, and Alex Jaimes. 2024. Dissecting users' needs for search result explanations. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA.) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 841, 17 pages. <https://doi.org/10.1145/3613904.3642059>
- [17] Anjali Khurana, Parsa Alamzadeh, and Parmit K. Chilana. 2021. ChatrEx: Designing Explainable Chatbot Interfaces for Enhancing Usefulness, Transparency, and Trust. In *2021 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 1–11. <https://doi.org/10.1109/VL/HCC51201.2021.9576440>
- [18] Sunnie S. Y. Kim, Nicole Meister, Vikram V. Ramaswamy, Ruth Fong, and Olga Russakovsky. 2022. HIVE: Evaluating the Human Interpretability of Visual Explanations. In *Computer Vision – ECCV 2022*, Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (Eds.). Springer Nature Switzerland, Cham, 280–298. https://doi.org/10.1007/978-3-031-19775-8_17
- [19] Kate Lister, Tim Coughlan, Francisco Iniesto, Nick Freear, and Peter Devine. 2020. Accessible conversational user interfaces: considerations for design. In *Proceedings of the 17th International Web for All Conference (Taipei, Taiwan) (W4A '20)*. Association for Computing Machinery, New York, NY, USA, Article 5, 11 pages. <https://doi.org/10.1145/3371300.3383343>
- [20] Ettore Mariotti, Jose M Alonso, and Albert Gatt. 2020. Towards harnessing natural language generation to explain black-box models. In *2nd Workshop on Interactive Natural Language Technology for Explainable Artificial Intelligence*. 22–27.
- [21] Michael F. McTear. 2017. The Rise of the Conversational Interface: A New Kid on the Block?. In *Future and Emerging Trends in Language Technology. Machine Learning and Big Data*, José F Quesada, Francisco-Jesús Martín Mateos, and Teresa López Soto (Eds.). Springer International Publishing, Cham, 38–49. https://doi.org/10.1007/978-3-319-69365-1_3
- [22] Chris Messina. 2024. Conversational commerce. When the point of sale comes to the... | by Chris Messina | Chris Messina | Medium. <https://medium.com/chris-messina/conversational-commerce-92e0bcfc3ff>. Accessed: 2024-02-15.
- [23] Martijn Millecamp, Nyi Nyi Htun, Cristina Conati, and Katrien Verbert. 2019. To explain or not to explain: the effects of personal characteristics when explaining music recommendations. In *Proceedings of the 24th International Conference on Intelligent User Interfaces (Marina del Ray, California) (IUI '19)*. Association for Computing Machinery, New York, NY, USA, 397–407. <https://doi.org/10.1145/3301275.3302313>
- [24] Michael Minge and Manfred Thüring. 2018. Hedonic and pragmatic halo effects at early stages of User Experience. *International Journal of Human-Computer Studies* 109 (2018), 13–25. <https://doi.org/10.1016/j.ijhcs.2017.07.007>
- [25] Andrea Papenmeier, Dagmar Kern, Daniel Hienert, Alfred Sliwa, Ahmet Aker, and Norbert Fuhr. 2021. Starting Conversations with Search Engines - Interfaces that Elicit Natural Language Queries. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval (Canberra ACT, Australia) (CHIIR '21)*. Association for Computing Machinery, New York, NY, USA, 261–265. <https://doi.org/10.1145/3406522.3446035>
- [26] Andrea Papenmeier, Alfred Sliwa, Dagmar Kern, Daniel Hienert, Ahmet Aker, and Norbert Fuhr. 2020. 'A Modern Up-To-Date Laptop'-Vagueness in Natural Language Queries for Product Search. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. 2077–2089. <https://doi.org/10.1145/3357236.3395489>
- [27] Andrea Papenmeier and Elin Anna Topp. 2023. Ah, Alright, Okay! Communicating Understanding in Conversational Product Search. In *Proceedings of the 5th International Conference on Conversational User Interfaces* (Eindhoven, Netherlands) (CUI '23). Association for Computing Machinery, New York, NY, USA, Article 49, 5 pages. <https://doi.org/10.1145/3571884.3604318>
- [28] Megan AK Peters. 2022. Confidence in Decision-Making. In *Oxford Research Encyclopedia of Neuroscience*. <https://doi.org/10.1093/acrefore/9780190264086.013.371>
- [29] Pearl Pu, Li Chen, and Rong Hu. 2011. A user-centric evaluation framework for recommender systems. In *Proceedings of the Fifth ACM Conference on Recommender Systems* (Chicago, Illinois, USA) (RecSys '11). Association for Computing Machinery, New York, NY, USA, 157–164. <https://doi.org/10.1145/2043932.2043962>

- [30] Pearl Pu, Li Chen, and Rong Hu. 2012. Evaluating recommender systems from the user's perspective: survey of the state of the art. *User Modeling and User-Adapted Interaction* 22, 4–5 (March 2012), 317–355. <https://doi.org/10.1007/s11257-011-9115-7>
- [31] Tim Schrills and Thomas Franke. 2020. Color for Characters - Effects of Visual Explanations of AI on Trust and Observability. In *Artificial Intelligence in HCI*, Helmut Degen and Lauren Reinerman-Jones (Eds.). Springer International Publishing, Cham, 121–135. https://doi.org/10.1007/978-3-030-50334-5_8
- [32] Artur Strzelecki and Paulina Rutecka. 2020. Direct answers in Google search results. *IEEE Access* 8 (2020), 103642–103654. <https://doi.org/10.1109/ACCESS.2020.2999160>
- [33] Maartje Ter Hoeve, Mathieu Heruer, Daan Odijk, Anne Schuth, and Maarten de Rijke. 2017. Do news consumers want explanations for personalized news rankings. In *FATREC Workshop on Responsible Recommendation Proceedings*. <https://doi.org/10.18122/B24D7N>
- [34] Nava Tintarev. 2007. Explanations of recommendations. In *Proceedings of the 2007 ACM Conference on Recommender Systems* (Minneapolis, MN, USA) (*RecSys '07*). Association for Computing Machinery, New York, NY, USA, 203–206. <https://doi.org/10.1145/1297231.1297275>
- [35] Nava Tintarev and Judith Masthoff. 2015. *Explaining Recommendations: Design and Evaluation*. Springer US, Boston, MA, 353–382. https://doi.org/10.1007/978-1-4899-7637-6_10
- [36] Diana Cezara Toader, Grația Boca, Rita Toader, Mara Măcelaru, Cezar Toader, Diana Ighian, and Adrian T Rădulescu. 2019. The effect of social presence and chatbot errors on trust. *Sustainability* 12, 1 (2019), 256. <https://doi.org/10.3390/su12010256>
- [37] Manos Tsagkias, Tracy Holloway King, Surya Kallumadi, Vanessa Murdock, and Maarten de Rijke. 2021. Challenges and research opportunities in eCommerce search and recommendations. *SIGIR Forum* 54, 1, Article 2 (feb 2021), 23 pages. <https://doi.org/10.1145/3451964.3451966>
- [38] Sven Tuzovic and Stefanie Paluch. 2018. *Conversational Commerce – A New Era for Service Business Development?* Springer Fachmedien Wiesbaden, Wiesbaden, 81–100. https://doi.org/10.1007/978-3-658-22426-4_4
- [39] Katrien Verbert, Denis Parra, Peter Brusilovsky, and Erik Duval. 2013. Visualizing recommendations to support exploration, transparency and controllability. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces* (Santa Monica, California, USA) (*IUI '13*). Association for Computing Machinery, New York, NY, USA, 351–362. <https://doi.org/10.1145/2449396.2449442>
- [40] Jesse Vig, Shilad Sen, and John Riedl. 2009. Tagsplanations: explaining recommendations using tags. In *Proceedings of the 14th International Conference on Intelligent User Interfaces* (Sanibel Island, Florida, USA) (*IUI '09*). Association for Computing Machinery, New York, NY, USA, 47–56. <https://doi.org/10.1145/1502650.1502661>
- [41] E. S. Vorm and A. D. Miller. 2018. Assessing the value of transparency in recommender systems: An end-user perspective. In *Proceedings of the 5th Joint Workshop on Interfaces and Human Decision Making for Recommender Systems co-located with ACM Conference on Recommender Systems* (Vancouver, Canada), Vol. 2225. 61–68. <http://ceur-ws.org/Vol-2225/> Retrieved from <https://hdl.handle.net/1805/19751>.
- [42] Hui Zhang et al. 2018. Beyond query-oriented highlighting: Investigating the effect of snippet text highlighting in search user behavior. *Computational Intelligence and Neuroscience* 2018 (2018). <https://doi.org/10.1155/2018/7836969>
- [43] Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W. Bruce Croft. 2018. Towards Conversational Search and Recommendation: System Ask, User Respond. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (Torino, Italy) (*CIKM '18*). Association for Computing Machinery, New York, NY, USA, 177–186. <https://doi.org/10.1145/3269206.3271776>
- [44] Xi Zheng, Akanksha Bansal, and Matthew Lease. 2017. Bullseye: structured passage retrieval and document highlighting for scholarly search. In *Proceedings of the Australasian Computer Science Week Multiconference* (Geelong, Australia) (*ACSW '17*). Association for Computing Machinery, New York, NY, USA, Article 32, 4 pages. <https://doi.org/10.1145/3014812.3014846>
- [45] Feng Zhou. 2015. The Effect Of Perceived Transparency, Perceived Reputation And Perceived Security On Consumer Trust And Purchase Intention. In *ICEB 2015 Proceedings*. Hong Kong, SAR China, 416–425. <https://aisel.aisnet.org/iceb2015/80>