




# ClayBot: Increasing Human-Centricity in Conversational Recommender Systems

Fouad Zablith<sup>(✉)</sup> 

Olayan School of Business, American University of Beirut,  
P.O. Box 11-0236, Riad El Solh, 1107 2020 Beirut, Lebanon  
fouad.zablith@aub.edu.lb

**Abstract.** Conversational recommender systems are increasingly studied to provide more fine-tuned recommendations based on user preferences. However, most existing product recommendation approaches in online stores are designed to interact with people through questions that mainly focus on products or their attributes, and less on buyers' core purchase needs. This work proposes ClayBot, a novel conversational recommendation agent, which aims to capture people's intents and recommend products based on the jobs or actions that their buyers aim to do. Interactions with ClayBot are guided by an openly accessible knowledge graph, which connects a sample of computing products to the actions annotated in product reviews. A demonstration of ClayBot is presented as an Amazon Alexa Skill to showcase the feasibility of handling more human-centered interactions in the product recommendation and explanation process.

**Keywords:** Conversational Recommender Systems · Knowledge Graphs · Human-centered · Artificial Intelligence · Product reviews

## 1 Introduction

Conversational recommender systems (CRS) are increasingly gaining research visibility in artificial intelligence (AI). This rising attention can be attributed to several factors, including the challenges in eliciting users' preferences and needs, as well as justifying why does a person prefer an item versus another [4].


In real life, human interaction and questions usually drive the recommendation process and help with discovering people's preferences to reach a recommended product [4]. Asking questions is one of the key components in CRS. In this domain, researchers have been studying "when to ask" and "what to ask" to people, and their approaches mainly focus on formulating questions either around products, or product attributes [3].

While such questions can be effective for experienced buyers, they may be hard to answer by people who are less informed about how products or their attributes can contribute to fulfilling their needs. This work investigates how we can make conversational product recommendation systems more human-centered and aligned with users' needs.

## 2 Human-Centered Product Recommendations in CRS

Clayton Christensen famously revealed that “when we buy a product, we essentially ‘hire’ it to get a job done” [2]. This core notion of the Jobs Theory reflects that in most situations, people don’t perceive products as a set of attributes, but as a means to perform a certain job. Jobs to be done are often articulated in the form of actions. Designers are usually trained to be sensitive to such potential actions when developing human-centered products. Potential actions that a person is aiming to do may be used as a proxy for further understanding the person’s needs and uncovering persona-centered problems in the design thinking process [6]. For example, people may decide to buy a particular laptop to support them while studying, or a certain tablet to be able to stream their favorite TV series in bed. A substantial number of buyers generously share such experiences in product reviews, which often include a description of the actions and jobs to be done [1].

This work aims to shift the focus of questions asked in the majority of existing CRS approaches from a product and product’s attributes angle [3], to a human-centered perspective by formulating questions and understanding intents around the jobs or actions that buyers are aiming to fulfill. Drawing on previous efforts to construct a knowledge graph that connects actions detected in product reviews to their products and related entities [8], we propose a framework that leverages the knowledge graph data connections to guide the interaction during the product recommendation process. Figure 1 provides an overview of the CRS framework.



**Fig. 1.** CRS framework overview for recommending products based on buyers’ intents.

The framework includes three layers namely the interaction, processing, and data layers. Following buyers’ intents to the CRS, the *Intent-SPARQL Translator* component is responsible for converting the intents into SPARQL queries. The translator is designed to include semantic and pattern-based placeholders to identify the actions and other elements mentioned in the buyer’s intents. For example, if a buyer articulates that “she would like a device for drawing,” then the translator automatically detects *drawing* as a potential action and integrates it in the SPARQL queries. The queries are then passed to the *SPARQL endpoint* to extract from the knowledge graph the relevant triples needed for the recommendation process. The *Recommender Engine* traverses the triples extracted

from the SPARQL endpoint and uses them to rank products and to aggregate relevant action and product-related entities—reviews, images, and others. The *Voice and Visual Presentation* component presents the recommended results to the user through voice and visual features.

### 3 Demo: Alexa, Ask ClayBot to Recommend a Product

To showcase the feasibility of the proposed approach, this part presents ClayBot, a conversational product recommendation agent. We explore an initial implementation of ClayBot as an Amazon Alexa Skill<sup>1</sup>. The Amazon Alexa Skills platform provides an environment with features that ease the development and deployment of conversational agents. Such features include for example the creation of user intent templates, handling voice and visual interfaces, and customizing slot types that we used in the Intent-SPARQL Translator, among others. Figure 2 shows a sample voice interaction between a person and ClayBot, coupled with visual feedback through the Alexa app<sup>2</sup>.

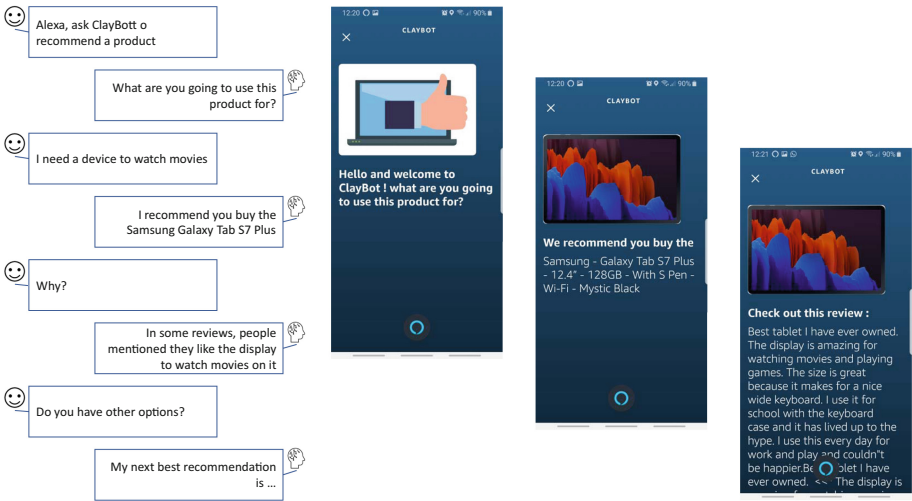


Fig. 2. ClayBot prototype running as an Amazon Alexa Skill.

ClayBot currently relies on a knowledge graph that covers annotated reviews of computing products. Around 3,000 product reviews were annotated as part of another project effort, and the data is accessed through an openly accessible SPARQL endpoint [8]. The knowledge graph includes information pertaining to

<sup>1</sup> ClayBot Alexa Skill page: <https://www.amazon.com/dp/B0BX6LQQT7>.

<sup>2</sup> A video recording of the demo featuring the discussed example in Fig. 2 is available at: [https://youtu.be/ZiILD\\_f51MQ](https://youtu.be/ZiILD_f51MQ).

the actions found in the reviews extracted from a retailer’s website that follows the [schema.org](https://schema.org) vocabulary, the valence (i.e., whether the reviewers were positive or negative with the action they were trying to do while using the device), and additional contextual information such as the role of agents and the environment in which the product was used.

A person can invoke the skill by telling Alexa to ask ClayBot to recommend a product. Following this invocation, ClayBot responds with what the person is going to use this product for. This response aims to shift the focus of the interaction on the job or action intended by the buyer, rather than on the traditional products and attributes focus. This design also limits the possibility of the buyer having generic intents, which may drift the conversation to a more generic intent classification task [5, 7]. In this example, the person may reply that they would need a device to watch movies. At this level, ClayBot will initiate the Intent-SPARQL Translator component to try to match the action of *watching* with the knowledge graph data. In the current version of ClayBot, the matching is performed by defining in the Alexa app slots of type *Action*. The slot types are then detected through text patterns in the intents. For example, a pattern of *text* + *<to>* or *<for>* followed by an *<action>*. All products that support this action will be fetched from the knowledge graph using SPARQL, along with other data needed by the recommender engine to rank the available products. The ranking of products relies on the relative ratio of positive versus negative review annotations related to the action in focus, as a comparative base among the products.

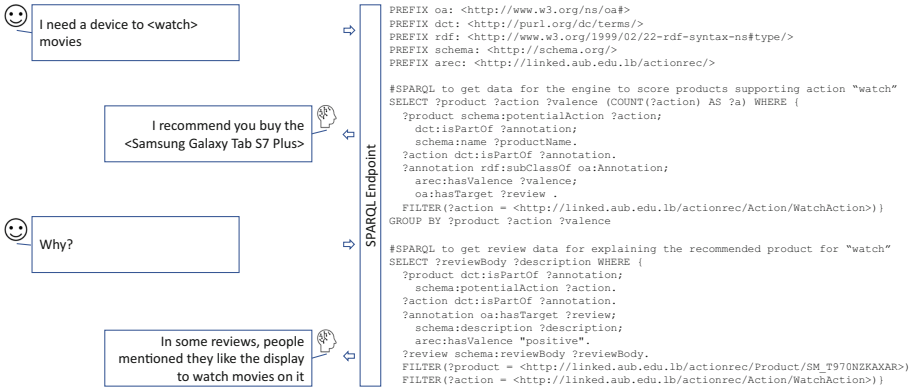
ClayBot then communicates the recommended product to the user and waits for the person’s next instruction. At this stage, the person may ask for explanation why this product was recommended, ask for another recommendation, restart the interaction with another objective, or simply end the conversation. When asked for other options, ClayBot recommends the next best product ranked by the recommender engine. As explanation, ClayBot is currently able to read and visualize a sample of related reviews from the knowledge graph. Figure 3 shows two examples of SPARQL queries. The first query extracts the knowledge graph data needed by the engine for the computation of the products’ scores relative to the action detected in the buyer intent. The second query extracts the review data related to the recommended product and action in focus<sup>3</sup>.

## 4 Conclusion

This paper demonstrated a novel attempt to offer more human-centered conversational product recommendation systems. The approach aims to increase human centricity by focusing the agent-people interactions on the jobs and actions that buyers are aiming to fulfill through the products they are seeking.

This work may benefit from several future research paths. First, the recommendation and explanation process can be extended to cover more elements

<sup>3</sup> The queries can be tested on the following SPARQL endpoint:  
<https://linked.aub.edu.lb/actionrec/sparql>.



**Fig. 3.** SPARQL examples to get knowledge graph data relevant to a buyer's intents.

captured in the knowledge graph (e.g., environment of use, product features, and others). Second, it is beneficial to investigate a wider range of questions and their impact on the conversation between a person and the recommender agent. Third, it is valuable to evaluate users' perception of their interaction with the proposed approach compared to existing product recommendation efforts.

This research provides the following contributions: first, it helps CRS with better uncovering buyers' needs that often revolve around performing certain jobs and recommending products accordingly; second, it demonstrates the potential of leveraging knowledge graphs to provide guided and explainable interactions between people and voice-enabled AI recommendation agents.

**Acknowledgments.** This work was partially supported by the University Research Board of the American University of Beirut. Special thanks to Rayan Al Arab for his support in developing the tools.

## References

- Christensen, C., Hall, T., Dillon, K., Duncan, D.S.: *Competing Against Luck: The Story of Innovation and Customer Choice*. HarperBusiness, New York (2016)
- Christensen, C., Hall, T., Dillon, K., Duncan, D.S.: Know your customers' jobs to be done. *Harv. Bus. Rev.* **94**(9), 54–62 (2016)
- Gao, C., Lei, W., He, X., de Rijke, M., Chua, T.S.: Advances and challenges in conversational recommender systems: a survey. *AI Open* **2**, 100–126 (2021)
- Jannach, D., Pu, P., Ricci, F., Zanker, M.: Recommender systems: trends and frontiers. *AI Mag.* **43**(2), 145–150 (2022)
- Moradizyveh, S.: Intent Recognition in Conversational Recommender Systems (2022). [arXiv:2212.03721](https://arxiv.org/abs/2212.03721) [cs]
- Norman, D.: *The Design of Everyday Things*. Basic Books, New York (2013)
- Schuurmans, J., Frasincar, F.: Intent classification for dialogue utterances. *IEEE Intell. Syst.* **35**(1), 82–88 (2020)
- Zablith, F.: ActionRec: toward action-aware recommender systems on the web. In: *International Semantic Web Conference Demos - CEUR Proceedings*, vol. 2980 (2021)