

DAGFiNN: A Conversational Conference Assistant

IVICA KOSTRIC, KRISZTIAN BALOG, TØLLØV ALEXANDER ARESVIK, NOLWENN BERNARD, EYVINN THU DØRHEIM, PHOLIT HANTULA, SANDER HAVN-SØRENSEN, RUNE HENRIKSEN, HENGAMEH HOSSEINI, EKATERINA KHLYBOVA, WERONIKA LAJEWSKA, SINDRE EKRHEIM MOSAND, and NARMIN ORUJOVA, University of Stavanger, Norway

DAGFiNN is a conversational conference assistant that can be made available for a given conference both as a chatbot on the website and as a Furhat robot physically exhibited at the conference venue. Conference participants can interact with the assistant to get advice on various questions, ranging from where to eat in the city or how to get to the airport to which sessions we recommend them to attend based on the information we have about them. The overall objective is to provide a personalized and engaging experience and allow users to ask a broad range of questions that naturally arise before and during the conference.

CCS Concepts: • **Information systems** → **Recommender systems**; *Users and interactive retrieval*.

Additional Key Words and Phrases: Conversational AI; Conversational Recommender System; Digital Assistant;

ACM Reference Format:

Ivica Kostric, Krisztian Balog, Tølløv Alexander Aresvik, Nolwenn Bernard, Eyvinn Thu Dørheim, Pholit Hantula, Sander Havn-Sørensen, Rune Henriksen, Hengameh Hosseini, Ekaterina Khlybova, Weronika Lajewska, Sindre Ekrheim Mosand, and Narmin Orujova. 2022. DAGFiNN: A Conversational Conference Assistant. In *Sixteenth ACM Conference on Recommender Systems (RecSys '22)*, September 18–23, 2022, Seattle, WA, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3523227.3551467>

1 INTRODUCTION

Conversational systems for information access have received much attention in recent years [1, 6], with conversational recommendation being one of the main tasks [5, 9]. A conversational recommender system (CRS) is a multi-turn, interactive system that can elicit user preferences in real-time using natural language [10]. Common research threads related to CRS include question-based user preference elicitation [11], multi-turn conversational recommendation strategies [12], natural language understanding and generation [14], and evaluation and user simulation [20]. With some exceptions [2, 17], the vast majority of work uses text as the sole modality and targets a single domain. Arguably, in many real-world contexts, a rich user experience would encompass multiple modalities and domains.

Specifically, in this paper we study the problem of conversational recommendations in the context of an academic conference. To mimic the capabilities of a person sitting behind the information desk, a CRS needs to have basic conversational skills as well as knowledge about things that conference participants would typically ask about, such as recommendations for places to eat and to visit, and conference schedule. A person behind the counter would also personalize responses, remember past interactions, and use additional digital/printed material to help convey information (e.g., show directions to a restaurant on a map). The system we present here, DAGFiNN, represents a step towards the digital version of such conference assistant.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

Manuscript submitted to ACM

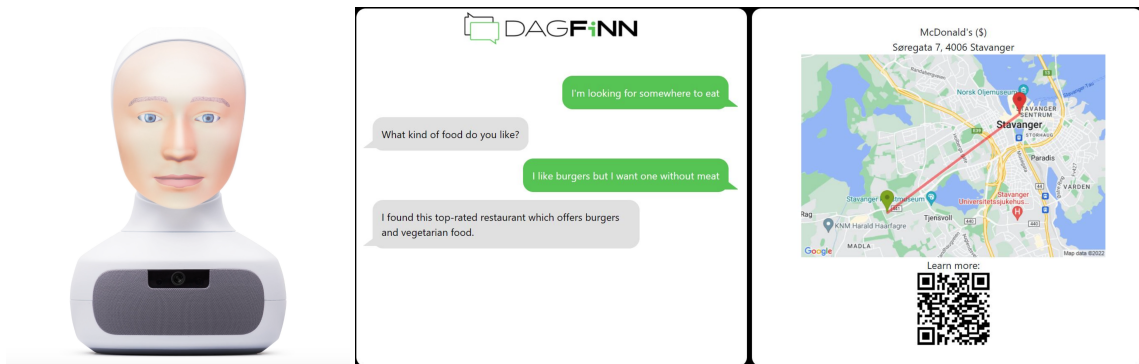


Fig. 1. Furhat (Left) with the external screen (Right), showing an excerpt from a POI recommendation dialogue.

To the best of our knowledge, DAGFiNN is a first of its kind, being an open-source, multi-domain, multi-modal conversational recommender system. The specific research challenges we address are as follows:

- Provide a holistic user experience by providing recommendations across multiple domains, basic conversational capabilities, and personalization integrated into a single system. Currently, recommenders are built for point-of-interest (POI) and conference content, but the set of domains can be extended in a modular fashion.
- Support multiple user interfaces (UI), specifically, a Web chat and a Furhat robot with an optional external screen, and use rich responses customized to each user interface to effectively convey information.
- Utilize multiple input modalities for the Furhat robot. Furhat is a social robot with a microphone, speaker, and camera, which makes it possible to recognize users that interacted with the robot previously.

The DAGFiNN project is the result of the combined effort of multiple student teams, each responsible for a specific component. DAGFiNN has been field-tested at an international conference, which we briefly report on in Section 5. The source code, more detailed descriptions of the skills, and a demonstration video are available at <https://dagfinn.ai>.

2 RELATED WORK

Recommender systems have been extensively studied for the academic [7] and point-of-interest (POI) [13, 16] domains. There is a rich body of recommendation methods to build on. However, our focus in this work is on the conversational user experience; more advanced recommendation methods may be incorporated in the future. Notably, the type of items that can be recommended is slightly different in our case. For example, the majority of the research on academic recommendation focuses on scientific literature recommendation [3, 7, 15], while in our work, similarly to [18], the focus is on recommending conference sessions. We focus not only on recommending POIs, but also on how to get there.

There are currently only a few open-source CRS prototypes available. CRSLab [21] and Macaw [19] are general frameworks, while VoteGoat [4] and IAI MovieBot [8] are specific to the movies domain. None of these systems has support for rich multi-modal interactions across multiple user interfaces.

3 CONVERSATIONAL CONFERENCE ASSISTANT

DAGFiNN is a conversational conference assistant that can be made available both as a chatbot on the conference website and as a Furhat robot physically exhibited at the conference venue. Conference participants can interact with DAGFiNN to get advice on a variety of questions, from point-of-interest (POI) recommendations to suggestions




Fig. 2. DAGFiNN architecture. Left: Clients with different interfaces; Web chat and Furhat robot with external screen. Right: Server.

on which sessions to attend at the conference. DAGFiNN has several skills, modular in design, making it easy to add new ones or remove the ones that are no longer needed. As its **Core** skill, it is equipped with basic “chit-chat” conversational capabilities, to handle questions like *Who are you?* or *What’s the weather like in Seattle?* It uses simple rules to determine the appropriate response for some anticipated questions. These feature relies mostly on templated responses to frequently asked queries and handles out-of-scope utterances.

In order to provide a personalized experience, there is an additional personalization layer to allow for tailoring responses to individuals. In its current form, DAGFiNN can recognize users interacting with the Furhat robot by scanning a QR code from their badge or, if permitted, by taking pictures of users’ faces during their interaction.

In terms of recommendation skills, we target two domains that represent some of the most common information needs that arise in a conference context: POI and conference content. DAGFiNN can elicit user preferences via multi-turn interaction to narrow the search space and can handle follow-up questions about items. The **POI recommender** answers queries about restaurants, museums, and other activities in the vicinity of the conference venue, like *Where can I go running?* or *Do you know of any good Indian restaurants?* Once the user has made a selection, they can ask follow-up questions, e.g., *How do I get there?* The collection of items is extracted from TripAdvisor and Google into a manually curated database, with details on name, category (e.g., restaurant, bar, museum, park), price range, address, rating, etc. The skill is triggered by asking for a recommendation on one of the categories. After a preference elicitation step, where preferences are represented internally as lists of liked and disliked keywords, a recommendation is made on each turn from the set of top-rated matching items. When the user is satisfied with the recommendation, they can further ask questions about the recommended item or inquire about transport options.

The **conference content recommender** skill can provide details on the conference programme (keynotes, tutorials, workshops, sessions) and authors, and can give recommendations on which session to attend based on elicited interests. Example queries include *Who are the conference’s keynote speakers?*, *Can you recommend me a session to attend?* and *What is the next session about?* The items (sessions, speaker, schedule, and room information) are obtained from the conference website.

4 ARCHITECTURE

DAGFiNN follows a client-server architecture; the main components are shown in Figure 2. The server is built on top of the open-source Rasa conversational framework.¹ Rasa handles dialogue management by classifying intents and then invoking the respective skills to generate a response. Skills are implemented using a combination of rules, stories, and forms, which represent the various (increasingly complex) ways to provide training data in Rasa to train the dialogue management model. We implemented a logging system, which stores all conversations in a database for later analysis.

¹<https://rasa.com/open-source/>




Fig. 3. Left: Distributions of conversation lengths. Right: Distributions of turns per skill.

Rasa also features connector classes to many different messaging platforms (e.g., Slack, Telegram). The notion of a *channel* in Rasa refers to a communication path between a server and a client (e.g., Rest API, socket.io), on top of which we can develop custom features for specific platforms.

DAGFiNN supports two types of user interfaces at the moment: a Web chat and Furhat robot with an optional external screen. A JavaScript Web chat widget is provided out of the box by Rasa that can be added to any webpage. The Furhat robot talks to the server using a custom Furhat ‘skill’ written in Kotlin and using the socket.io library. Similarly, the external screen uses the same socket.io channel to connect to the same custom connector in Rasa server, thereby the conversation is synchronized between the two devices.

Responses may be customized for different channels, to fit the modalities available on that channel. Finally, the personalization layer handles the identification of users based on facial recognition or QR code, displayed on the participant’s badge. Before taking pictures, for compliance with GDPR, DAGFiNN always asks permission to try to identify users. Knowing the user’s identity allows for the development of additional features in the future.

5 FIELD STUDY

DAGFiNN has been featured as a digital receptionist at the ECIR’22 international conference, where participants could interact with it to get recommendations related to places to visit (POIs) and to the content of the conference. It has offered an engaging experience, especially for POI recommendations, and has received overwhelmingly positive feedback. Figure 3 (Left) shows the distribution of conversation lengths; we observe that most conversations last 3-6 turns. On the right, we show the distribution of intents across skills. Unsurprisingly, the most commonly triggered skill was Core, as most conversations include intents like greetings and good-byes, which are a part of the Core skill.

While users in general liked the recommendations they received, some were probing the abilities with more perplexing requests like *Where to go for a swim?* or *I am interested in seeing the fjords*. Also, given the Furhat robot interface, some people seemed to be more interested in a social chitchat-like experience than inquiries about actual information needs.

6 CONCLUSION AND FUTURE WORK

We have presented DAGFiNN, an open-source conversational conference assistant. It has been developed using the task-oriented open-source framework Rasa, and run on a social Furhat robot along with an external display, allowing for both spoken conversations and visual information. DAGFiNN has been field-tested at an international conference and has received positive feedback. Nevertheless, it is still in a research prototype phase and we plan to expand its functionality in the future. Some ideas include adding new skills, getting implicit feedback through facial expressions, supporting multi-party dialogues, and extending the personalization aspect. Given that DAGFiNN has been developed by a group of students as part of their bachelor/master theses, we are also interested in utilizing it as an educational tool and as a platform for other student projects in the future.

REFERENCES

- [1] Avishek Anand, Lawrence Cavedon, Hideo Joho, Mark Sanderson, and Benno Stein. 2020. Conversational Search (Dagstuhl Seminar 19461). *Dagstuhl Reports* 9, 11 (2020), 34–83.
- [2] David Arteaga, Juan Arenas, Freddy Paz, Manuel Tupia, and Mariuxi Bruzza. 2019. Design of Information System Architecture for the Recommendation of Tourist Sites in the City of Manta, Ecuador through a Chatbot. In *14th Iberian Conference on Information Systems and Technologies (CISTI '19)*. 1–6.
- [3] Xiaomei Bai, Mengyang Wang, Ivan Lee, Zhuo Yang, Xiangjie Kong, and Feng Xia. 2019. Scientific Paper Recommendation: A Survey. *IEEE Access* 7 (2019), 9324–9339.
- [4] Jeffrey Dalton, Victor Ajayi, and Richard Main. 2018. Vote Goat: Conversational Movie Recommendation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '18)*. 1285–1288.
- [5] Chongming Gao, Wenqiang Lei, Xiangnan He, Maarten de Rijke, and Tat-Seng Chua. 2021. Advances and Challenges in Conversational Recommender Systems: A Survey. *AI Open* 2 (2021), 100–126.
- [6] Jianfeng Gao, Michel Galley, Lihong Li, et al. 2019. Neural Approaches to Conversational AI. *Foundations and Trends in Information Retrieval* 13, 2-3 (2019), 127–298.
- [7] Kristian Gingstad, Øyvind Jekteberg, and Krisztian Balog. 2020. ArXivDigest: A Living Lab for Personalized Scientific Literature Recommendation. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (CIKM '20)*. 3393–3396.
- [8] Javeria Habib, Shuo Zhang, and Krisztian Balog. 2020. IAI MovieBot: A Conversational Movie Recommender System. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (CIKM '20)*. 3405–3408.
- [9] Claudia Hauff, Julia Kiseleva, Mark Sanderson, Hamed Zamani, and Yongfeng Zhang. 2021. Conversational Search and Recommendation: Introduction to the Special Issue. *ACM Transactions on Information Systems* 39, 4 (2021), 1–6.
- [10] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. 2021. A Survey on Conversational Recommender Systems. *Comput. Surveys* 54, 5 (2021), 36 pages.
- [11] Ivica Kostric, Krisztian Balog, and Filip Radlinski. 2021. Soliciting User Preferences in Conversational Recommender Systems via Usage-Related Questions. In *Fifteenth ACM Conference on Recommender Systems (RecSys '21)*. 724–729.
- [12] Wenqiang Lei, Xiangnan He, Yisong Miao, Qingyun Wu, Richang Hong, Min-Yen Kan, and Tat-Seng Chua. 2020. Estimation-Action-Reflection: Towards Deep Interaction Between Conversational and Recommender Systems. In *Proceedings of the 13th International Conference on Web Search and Data Mining (WSDM '20)*. 304–312.
- [13] Changheng Li, Yongjing Hao, Pengpeng Zhao, Fuzhen Zhuang, Yanchi Liu, and Victor S. Sheng. 2021. Tell Me Where to Go Next: Improving POI Recommendation via Conversation. In *Proceedings of the 26th International Conference on Database Systems for Advanced Applications (DASEAA '21)*. 211–227.
- [14] Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018. Towards Deep Conversational Recommendations. In *Advances in Neural Information Processing Systems 31 (NeurIPS '18)*.
- [15] Shuchen Li, Peter Brusilovsky, Sen Su, and Xiang Cheng. 2018. Conference Paper Recommendation for Academic Conferences. *IEEE Access* 6 (2018), 17153–17164.
- [16] Pablo Sánchez and Alejandro Bellogin. 2022. Point-of-Interest Recommender Systems Based on Location-Based Social Networks: A Survey from an Experimental Perspective. *Comput. Surveys* (2022).
- [17] Marilyn Walker, Steve Whittaker, Amanda Stent, Preetam Maloor, Johanna Moore, Michael Johnston, and Gunaranjan Vasireddy. 2004. Generation and Evaluation of User Tailored Responses in Multimodal Dialogue. *Cognitive Science* 28, 5 (2004), 811–840.
- [18] Jinfeng Yi, Qi Lei, Junchi Yan, and Wei Sun. 2018. Session Expert: a Lightweight Conference Session Recommender System. In *2018 IEEE International Conference on Big Data (Big Data) (BigData '18)*. 1677–1682.
- [19] Hamed Zamani and Nick Craswell. 2020. Macaw: An Extensible Conversational Information Seeking Platform. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*. 2193–2196.
- [20] Shuo Zhang and Krisztian Balog. 2020. Evaluating Conversational Recommender Systems via User Simulation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '20)*. 1512–1520.
- [21] Kun Zhou, Xiaolei Wang, Yuanhang Zhou, Chenzhan Shang, Yuan Cheng, Wayne Xin Zhao, Yaliang Li, and Ji-Rong Wen. 2021. CRSLab: An Open-Source Toolkit for Building Conversational Recommender System. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations (IJCNLP '21)*. 185–193.